A Probabilistic Treatment of Expert Knowledge and Epistemic Uncertainty in NESSUS®

Luc Huyse*, Ben H. Thacker† and David S. Riha*
Reliability & Materials Integrity, Southwest Research Institute, San Antonio, TX

Simeon H.K. Fitch‡
Mustard Seed Software, Charlottesville, VA

Jason Pepin§ and Edward A. Rodriguez**
Engineering Sciences and application – Los Alamos National Lab, Los Alamos, NM

The risk assessment in many engineering applications is hampered by a lack of hard data. Under these conditions the selection of probability density function (PDF) seems arbitrary. Quite often the data are not only sparse but also vague expert knowledge or conflicting. Several non-probabilistic methods have been proposed in the literature to perform a risk assessment under these conditions. We propose to use probabilistic techniques using uncertain PDFs. The uncertainty on the PDF is characterized by treating the parameters in the PDF as random variables. We expand the classical Bayesian updating scheme to make use of vague or imprecise interval data. Each expert is considered to be a sample from a parent distribution of experts. Consequently, a conflict between experts is accounted for through the likelihood function. The uncertain PDFs can be used in both simulation-based and MPP-based reliability methods. Because of the uncertainty on the PDF of the random variables, the risk or reliability index itself will be a random variable. Design decisions are made on the basis of the risk assessment and an incorrect risk assessment increases the total cost of the design. Since a cost can be associated with either an over or underestimation of the risk, an optimal reliability index can be determined, which minimizes this cost. The probabilistic framework we present in this paper establishes a direct link between the amount and quality of the available data and the optimal reliability estimate. This link allows the decision maker to weigh expected value of additional data collection efforts against the expected optimal reliability index improvement.

I. Background

THERE is growing interest in the technical community to develop uncertainty analysis methods that are more suitable for representing subjective (or epistemic) uncertainties. This paper describes a method that treats epistemic uncertainties within a probabilistic framework and will be a new feature in the NESSUS® probabilistic analysis software. NESSUS is a general-purpose advanced reliability tool that computes failure probabilities as well as probabilistic importance and sensitivity factors for both component and system problems. The upcoming version

* Principal Engineer, Reliability & Materials Integrity, Southwest Research Institute, San Antonio, TX, Member AIAA.
† Director R&D, Materials Engineering, Southwest Research Institute, San Antonio, TX, Associate Fellow AIAA.
‡ Owner, Mustard Seed Software, Charlottesville, VA
§ Technical Staff Member, Engineering Sciences and Applications – Weapon Reliability, Los Alamos National Lab, Los Alamos, NM, Member AIAA.
** DynEx Deputy Project Director, Engineering Sciences and Applications – Weapon Reliability, Los Alamos National Lab, Los Alamos, NM, Member AIAA.

8.3 of the NESSUS software will allow users to fit imprecise probability densities on the basis of interval data and will compute the effect of these epistemic uncertainties on the reliability index or failure probability.

Epistemic uncertainties arise when there is little or no data available to assign traditional probability distributions, or when information is non-specific, ambiguous, or conflicting, as is often the case with expert knowledge. Such data are often available only as an interval estimate, without detailed knowledge about the relative likelihood of specific values within the interval.

In this paper we discuss a probabilistic framework that incorporates expert knowledge. The desire to have a probabilistic approach is motivated by the availability of proven analysis techniques and software implementations. Over the years considerable experience and familiarity has been gained with these techniques, which improves the credibility of the risk assessment. We demonstrate how existing Bayesian estimation techniques can be expanded to work with expert knowledge, given in the form of an interval estimate instead of a single, precise data point. The technique interprets each expert as a sample from a pool of experts and uses a likelihood function to update the prior distribution of the PDF parameters. The Bayesian estimation method can be used for both crisp (or precise) and interval data.

Because the PDF parameters are estimated on the basis of limited data, considerable statistical uncertainty may be associated with the PDF parameters. The paper subsequently illustrates how the uncertain PDFs, obtained using the presented Bayesian updating technique, can be integrated in existing, efficient probabilistic computational tools. A practical example using the First-Order Reliability Method (FORM) is given.

Because of the epistemic or statistical uncertainty associated with the PDF, the risk measure (such as the reliability index) is an uncertain variable itself. The effect of the epistemic uncertainty (e.g. statistical uncertainty on the PDF shape or parameters) on the reliability can be quantified through confidence bounds (see Figure 1). These confidence bounds play an important role in the decision making process. Consider for instance the case of two competing designs with slightly different reliability curves. If the confidence bounds overlap considerably, it would be hard to argue that one design represents a significant improvement over the other, under the current state of knowledge. At least in theory, it is, however, possible to shrink the confidence bounds due to model uncertainty, either through additional data collection or physical modeling efforts. A comparison of the cost of these additional experiments or analyses against the potential gain that can be made by shrinking the confidence bounds leads to more rational decision-making under imperfect states of knowledge.

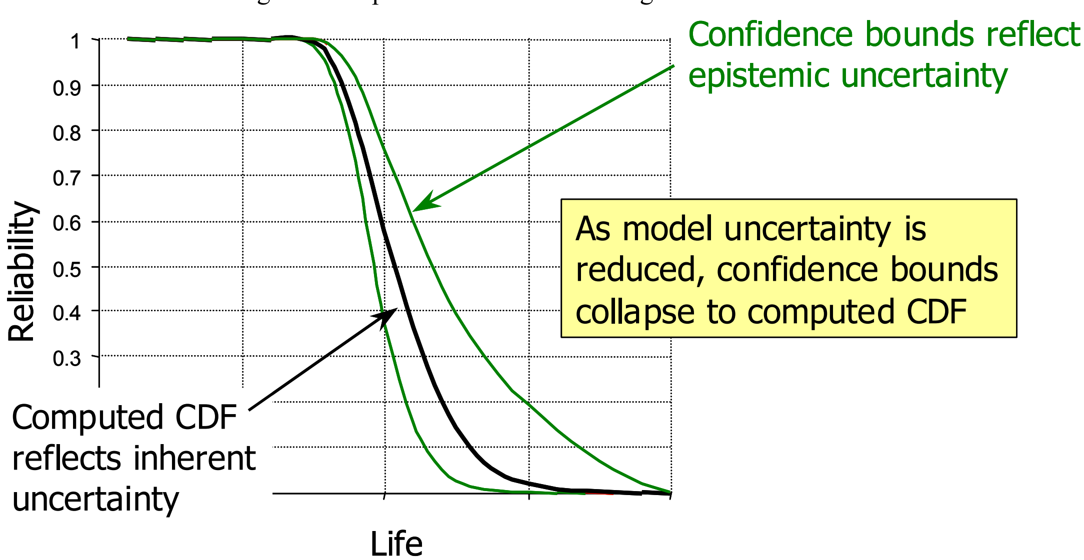

Figure 1: Effect of Inherent and Epistemic Uncertainty

The last part of the paper connects the reliability computations to the decision making process during design. Overestimation of the reliability may cause severe economic or societal loss whereas an underestimation of the reliability may unnecessarily increase the operational or manufacturing cost. A specific penalty can therefore be associated with a mischaracterization of the reliability index. In such a cost-based approach, the optimal reliability index is associated with the minimum total cost. An example demonstrates how the uncertainty associated with expert knowledge affects this minimum penalty reliability index.

II. Selection of the Probability Density Function

In the absence of data the objectivity of the probabilistic risk assessment is challenged in two ways. The first one is the selection of the PDF (PDF shape). When insufficient data exists to justify the choice of a particular PDF, we recommend the use of a parametric family of PDFs. A PDF distribution family or system contains shape parameters in addition to the location and scale parameters. In this approach, the shape of the PDF becomes a statistic that can be estimated directly from the data. Details about this approach can be found in a companion paper (Huyse and Thacker, 2004). The second challenge is the use of probabilistic methods to determine the PDF parameters when insufficient data are available. In this paper we assume that the PDF type is given and the focus is on the incorporation of vague, conflicting data into the probabilistic assessment.

III. Probabilistic Analysis with Interval Data

A. Methodology: Bayes' Theorem

We will demonstrate that Bayesian estimation methods can be used successfully to estimate the parameters in the PDF only interval data are given for the random variable Y. In Bayes' theorem prior knowledge about the parameter θ is combined with additional data $\mathbf{y}$ to compute the posterior density of θ (Box and Tiao, 1992).

Assume that the random variable Y is modeled using a PDF family with parameters θ: $f_Y(y \mid \theta)$. The Bayesian estimation approach treats the parameters θ within the selected PDF family as random variables and assigns a PDF $f_\Theta(\theta)$ to them. The prior beliefs about θ are updated with the data $\mathbf{y}$ and combined in a posterior density $f_\Theta(\theta \mid \mathbf{y})$. Several point estimates can be derived for θ from the posterior density $f_\Theta(\theta \mid \mathbf{y})$: posterior mean, median and mode are commonly used. It should be noted that the posterior mode will coincide with the MLE estimator if and only if the prior density $f_\Theta(\theta)$ is uniform.

Assume that $\mathbf{y}$ is a vector of n observations of the random variable Y whose PDF $f_Y(y)$ depends on the value of the parameters θ within the selected PDF family. Assume also that the parameters θ themselves have a probability distribution $f_\Theta(\theta)$. Bayesian estimation computes not just a point estimate of θ, but rather the entire PDF.

From the theorem of total probability, it follows that:

$$f_Y(\mathbf{y} \mid \theta).f_\Theta(\theta) = f_{Y,\Theta}(\mathbf{y},\theta) = f_\Theta(\theta \mid \mathbf{y}).f_Y(\mathbf{y}) \tag{1}$$

Given the data $\mathbf{y}$, $f_Y(\mathbf{y} \mid \theta)$ may be regarded as a function of θ instead of $\mathbf{y}$. In this case it is called the likelihood of θ for given data $\mathbf{y}$ and is usually written as $l(\theta \mid \mathbf{y})$. Since $f_Y(\mathbf{y})$ does not depend on θ, it can be omitted from Eq. 1 and the non-normalized posterior density of θ, given the data $\mathbf{y}$, becomes:

$$f_\Theta(\theta \mid \mathbf{y}) \propto \ell(\theta \mid \mathbf{y}).f_\Theta(\theta) \tag{2}$$

Bayes' theorem combines prior knowledge about the parameter θ with additional data $\mathbf{y}$ to compute the posterior density of θ:

$$f_\Theta(\theta \mid \mathbf{y}) \propto \frac{\ell(\theta \mid \mathbf{y}).f_\Theta(\theta)}{\int_\Theta \ell(\theta \mid \mathbf{y}).f_\Theta(\theta)d\theta} \tag{3}$$

The posterior distribution $f_\Theta(\theta \mid x)$ is proportional to the product of the likelihood $l(\theta \mid \mathbf{y})$ and the prior distribution $f_\Theta(\theta)$. The integral in the denominator of the expression normalizes the density so the total probability for θ remains equal to 1. Since Bayes' theorem is sequential in nature, the parameter θ can be continually updated using Eq.(3) as more observations $\mathbf{y}$ are taken.

The likelihood function follows directly from the selected PDF family whereas the prior distribution may include subjective probabilities. The selection of prior distributions is discussed at length in a subsequent section.

B. Use of Interval Data

Consider now the case where not a precise point value y is observed but rather an interval $[y_1, y_2]$. Only a single scalar interval observation $[y_1, y_2]$ is considered in order not to overload the notations. The formulation is easily extended to multiple dimensions. This interval can be thought of as representing an expert's knowledge about the value of the parameter Y in the model. Alternatively, the interval $[y_1, y_2]$ may represent the measurement uncertainty of an observation. The probability of observing this interval realization for a given value of θ is:

$$f([y_1, y_2]|\theta) = \int_{y_1}^{y_2} f(y|\theta)\,dy \tag{4}$$

where $f(y \mid \theta)$ represents the probability density. When interpreted as a function of θ for given values of y_i, the value $f([y_1, y_2] \mid \theta)$ can be interpreted as the likelihood of θ for the given interval observation $[y_1, y_2]$. Application of the Bayes' theorem to this result gives:

$$f(\theta|[y_1, y_2]) = \frac{f(\theta)\int_{y_1}^{y_2} f(y|\theta)\,dy}{\int_{\Theta} f(\theta)\int_{y_1}^{y_2} f(y|\theta)\,dy\,d\theta} \tag{5}$$

This equation indicates that Bayes' theorem can be applied equally well to imprecise interval observations as it is to single-point observations. Since Bayes' theorem is sequential, interval data and precise single-point observations can readily be combined in the estimation of the parameter θ.

IV. Conflicting data

At times, experts may disagree with each other and give (partially) conflicting information. In the proposed Bayesian approach no direct assessment is made regarding conflicts in the interval data. It is assumed that discrepancies in the expert opinion are not due to misinterpretation of the questions. The various interval data, which represent multiple expert opinions, are treated like sample observations of the random variable *Y*. Therefore conflicting interval estimates do not pose a numerical problem.

As an example, consider a normal probability model with uniform prior for both μ and σ. The posterior densities are shown for the cases with overlapping (Figure 1a) and non-overlapping data (Figure 1b). With overlapping expert opinion, the overlapping point becomes the posterior mode for the mean (and σ=0). For a uniform prior, the posterior distribution essentially reflects the likelihood of realizing joint observations that fall within each of the expert opinion intervals. Consequently, if the intervals overlap by more than a single point, the posterior mode extends across the entire range of overlap.

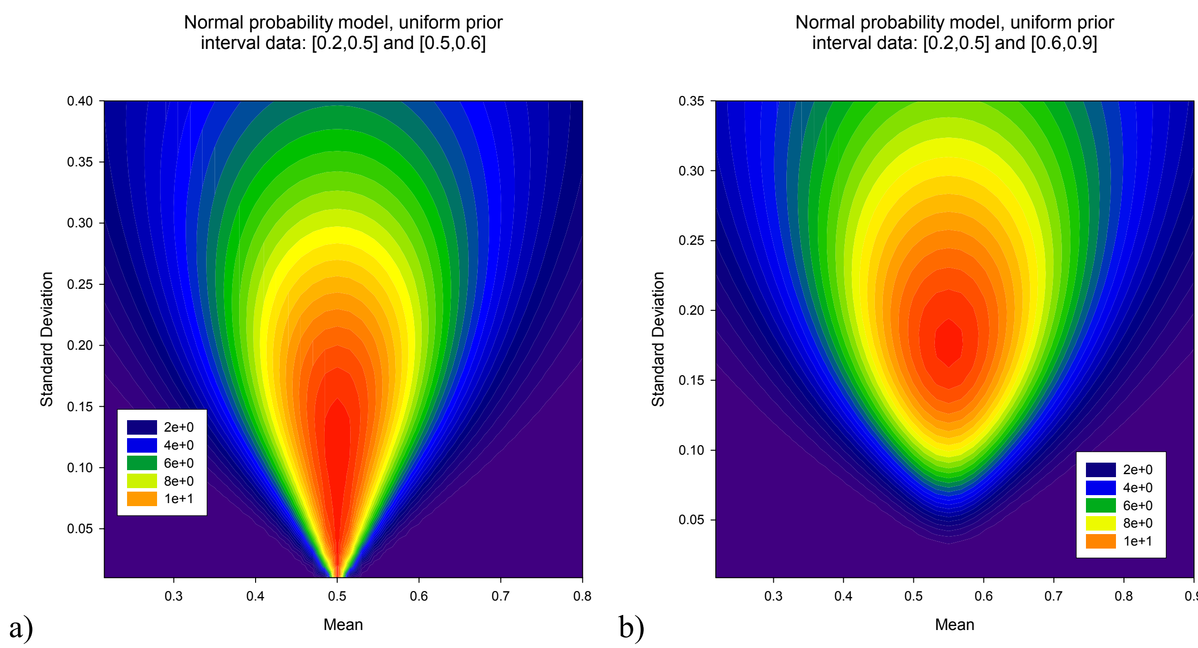

Figure 2: Posterior PDF for (μ, σ) when (a) the expert knowledge from two sources overlaps and (b) when the expert knowledge is in conflict

When the experts are in conflict the posterior mode of the mean is given by a weighted mean of the two intervals. The posterior mode of the standard deviation depends on the degree of conflict and width of individual intervals. The Bayesian estimation tries to maximize the odds of observing data in either the interval [0.2,0.5] or

[0.6,0.9]. Since both intervals in Figure 3 are of equal length, the posterior distribution is symmetric with respect to the mean and the model of the posterior mean is given by the mid-point between the intervals ($\mu = 5.5$).
Because of the conflict between the two experts (non-intersecting intervals), the total uncertainty on the parameters is larger.

Additionally, conventional statistical techniques can be applied to detect outliers. The relative weight of each expert opinion is not postulated by the analyst but follows directly from the data. The Bayes' factors (Box and Tiao, 1993) measure to what extent the new expert opinion is in conflict or in agreement with the available evidence to date.

V. Risk in Decision-Making Context

For each application a target reliability index β_{target} can be determined on the basis of societal criteria (such as risk averseness, severity of consequence…). Design decisions are made on the basis of the risk assessment. If the risk assessment indicates that the target risk is not achieved the design needs to be improved, which increases costs. If the risk assessment erroneously suggested that the target risk is not achieved, then these costs would amount to a waste of resources. A similar argument can be made when the risk is underestimated. Consequently, a penalty function $p(\beta)$ – i.e. loss of utility or cost – can be associated with an over or underestimation of the reliability index. The most appropriate choice for the reliability index is the one that minimizes the total expected cost. The lower bound of the confidence interval (safe or conservative value) may not be the most appropriate choice nor practically feasible or economically viable.

Consider a linear penalty function $p(\beta)$:

$$p(\beta) = \begin{cases} a(\beta - \beta_{\text{target}}) & \text{if } \beta_{\text{target}} \le \beta \\ k\,a(\beta - \beta_{\text{target}}) & \text{if } \beta < \beta_{\text{target}} \end{cases} \tag{6}$$

where a is a scale factor, and k is a measure for the asymmetry of the cost function (see Figure 6a). The cost function $p(\beta)$ is seldom symmetric in an engineering context; an overestimation of the reliability may cause a catastrophic loss of life whereas an underestimation of the reliability leads to overly costly structures. The parameter k is usually greater than one in engineering applications.

The minimum-penalty reliability index β_{mp}, which minimizes the cost function $p(\beta)$, is an appropriate point estimator, which reflects the impact of the epistemic uncertainty. For linear penalty functions, β_{mp} is given by (see Der Kiureghian, 1989):

$$\beta_{mp} = F_B^{-1}\left(\frac{1}{k+1}\right) \tag{7}$$

where F_B is the CDF of the reliability index. When β is assumed normally distributed, the minimum penalty index in Eq.(7) can be approximated by:

$$\beta_{mp,N} = \mu_B - \Phi^1\left(\frac{1}{k+1}\right).\sigma_B \text{ where } \begin{cases} \mu_\beta = \beta\big|_{E(\boldsymbol{\Theta})} \\ \sigma_\beta = \sqrt{\nabla_{\boldsymbol{\Theta}}^T \beta.\Sigma_{\boldsymbol{\Theta\Theta}}.\nabla_{\boldsymbol{\Theta}}\beta} \end{cases} \tag{8}$$

VI. Application

Consider a deflection limit state for the simply supported beam in Figure 2 with uncertain load P and Young's modulus E. All other variables are considered deterministic. It is assumed that only expert opinion is available for the Normal variables P and E. Although one would generally expect better prior knowledge, a non-informative prior will be assumed. Two different scenarios are considered: in the first scenario only 7 expert opinion intervals are available (some experts are in conflict), in the second scenario it is assumed that 140 experts are available (this data set is generated by repeating the original data set 20 times).

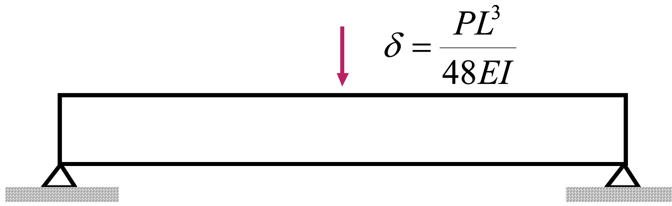

Figure 3: Simply supported beam

It can be concluded from the summary statistics in Table 1 that increasing the number of interval data generally does reduce the uncertainty on the parameter estimates. However, because of the conflict between several experts, the large standard error on the mean value estimates persists even as the data set is increased.

Table 1: Summary statistics for load *P* and modulus *E*.

Variable	Statistic	7 interval data		140 intvl data	
		Mean	StDev	Mean	StDev
Load *P*	Mode	3.11	0.89	3.12	1.00
	Exp. Value	3.18	1.17	3.15	1.00
	Std Error	1.53	0.39	1.11	0.07
	Correlation	0.22		0.29	
Modulus *E*	Mode	2.60	0.35	2.60	0.45
	Exp. Value	2.62	0.61	2.62	0.47
	Std Error	1.32	0.29	1.06	0.05
	Correlation	-0.03		-0.16	

For both cases the reliability index was computed using the expected values of the parameters (mean and standard deviation for *P* and *E*) and Eq. (8). Because of the lower estimate value for the standard deviation of both *P* and *E* (the mean value estimates remain unaffected), the reliability estimate is considerably higher for the cases with 140 data.
Subsequently, the sensitivity derivatives of β with respect to the epistemic parameters were computed. The results are given in Table 2. As anticipated, only $\partial\beta/d\mu_E$ is positive. The largest sensitivity derivatives are with respect to the standard deviation of the modulus.
Finally, an first-order second-moment estimate for the reliability index was obtained using Eq.(8). It can be seen that the standard error of the reliability index is not much smaller when many more data are present. The large value for σ_β is mostly due to the large standard error on the mean value of the mean value of the Young's modulus.

Table 2: Summary statistics for load *P* and modulus *E*.

Parameter	7 interval data	140 interval data
β	2.944	3.763
$\partial\beta/d\mu_P$	-0.33	-0.42
$\partial\beta/d\sigma_P$	-0.37	-0.66
$\partial\beta/d\mu_E$	1.53	1.94
$\partial\beta/d\sigma_E$	-4.11	-6.52
σ_β	2.444	2.195

It can be concluded from this example that the conflicting expert opinion regarding the Young's modulus causes a large uncertainty on the reliability index. Because of the sensitivity of the failure probability to the Young's modulus, the minimum penalty reliability index will not increase as more data are added. The conflict must be resolved instead.
Analysis of the relative contributions of σ_β reveal that the conflicting opinion regarding the Young's modulus weighs more heavily on the overall reliability index than the conflict about the load level does. The large uncertainty on the reliability index cannot be reduced unless the conflict regarding the Young's modulus is resolved.

VII. Conclusion

In this paper we have demonstrated that probabilistic methods can successfully be used in applications where only vague interval measurements are available. The approach is characterized by:

- The epistemic uncertainty is modeled by treating parameters as random variables.
- PDF parameters are estimated using Bayesian techniques, which are expanded to accommodate imprecise interval data. Non-informative priors can be used to characterize the prior lack of knowledge.
- Each of the interval data is interpreted as a random interval sample of a parent PDF. Consequently, a partial conflict between experts is automatically accounted for through the likelihood function.
- First-order approximations of the epistemic uncertainty on the reliability index are readily computed using results of only a single reliability analysis.
- The sensitivity of the reliability index to the model uncertainty parameters or to individual data points can readily be assessed. This is very powerful for experiment planning.

Note that, in the absence of specific knowledge about the variable, the selection of a specific PDF for a variable can be avoided altogether. A hierarchical model of a continuous family of PDF's can be used instead as described in a companion paper (Huyse and Thacker, 2004). This approach establishes a comprehensive framework for probabilistic analysis in absence of hard data.

References

1. Box, G. and G. Tiao (1992). *Bayesian Inference in Statistical Analysis.* John Wiley & Sons.
2. Der Kiureghian, A. and P.-L. Liu (1986). Structural Reliability under Incomplete Probability Information. *ASCE Journal of Engineering Mechanics 112*(1), 85-104.
3. Der Kiureghian, A. (1989). Measures of Structural Safety under Imperfect States of Knowledge. *ASCE Journal of Structural Engineering 115*(5), 1119-1140.
4. Ditlevsen, O. (1979). Generalized Second Moment Reliability Index. *ASCE Journal of Structural Mechanics 7*(4), 435-451.
5. Gelman, A., J. Carlin, H. Stern, and D. Rubin (1995). *Bayesian Data Analysis.* London, UK: Chapman & Hall.
6. Huyse L. and B.H. Thacker (2004), A Framework to Estimate Uncertain Random Variables, *45th AIAA/ASME/ASCE/AHS/ ASC Structures, Structural Dynamics and Materials Conference*, Palm Springs, CA, CD-ROM.
7. Huyse L. and B.H. Thacker (2005), Decision Making under Uncertainty on the Basis of Expert Knowledge, 9th *ICOSSAR Conference*, Rome, Italy, CD-ROM.
8. Oberkampf, W.L., T.G. Trucano and C. Hirsch, *Verification, Validation and Predictive Capability in Computational Engineering and Physics*, SAND2003-3769, February 2003.
9. Oberkampf, W. and J. Helton (2002). *Investigation of Evidence Theory in Engineering Applications*, AIAA-Paper 2002-1569. In 43rd AIAA/ASME/ASCE/AHS/ASC
10. Pratt, J., H. Raiffa, and R. Schlaifer (1995). *Introduction to Statistical Decision Theory.* Cambridge, MA: MIT Press.
11. Thacker, B.H. and E.A. Rodriguez, 2002, *Verification and Validation Plan for Dynamic Experimentation Containment Vessels*, Los Alamos National Laboratory Draft Report, December.
12. Thacker, B.H., 2003, *The Role of Nondeterminism in Verification and Validation of Computational Solid Mechanics Models*, Reliability & Robust Design in Automotive Engineering, SAE Int., SP-1736, No. 2003-01-1353, SAE 2003 World Congress, Detroit, MI, 3-6 March.
13. Southwest Research Institute, 2006, NESSUS User's and Theory manuals, Version 8.3, also http://www.nessus.swri.org