Joslyn, Cliff A; Buehring, W; Kaplan, PG; and Powell, D: (2004) "Critical Infrastructure Protection Decision Support System (CIP/DSS): Addressing Uncertainty and Risk", Los Alamos National Laboratory Technical Report LAUR 04-6720

Critical Infrastructure Protection Decision Support System (CIP/DSS): Addressing Uncertainty and Risk

Cliff A Joslyn[1], William A Buehring[2], Paul G Kaplan[3], and Dennis R Powell[4]

Several sources of uncertainty are present in the CIP/DSS models, for example uncertain input parameters that are drivers for estimating consequences in the individual infrastructure models. CIP/DSS should be capable of assisting DHS in analyzing risk-informed[5] strategies for implementing preventive, protective, mitigation, response, and recovery measures. Thus DHS expects the level of uncertainty to be explicitly addressed in our CIP/DSS results. Although the CIP/DSS project is not charged with estimating the probability of terrorist attack, the CIP/DSS analysis must be able to accommodate, without great difficulty, the relevant threat information available to DHS. Thus selecting an appropriate method for dealing with uncertainty should account not only for suitable theoretical underpinnings but also practical needs, such as providing users of CIP/DSS results with a reasonable understanding of the level of uncertainty in those results.

This paper overviews some technical issues about, and options for addressing, uncertainty and risk in CIP/DSS. Discussions on this topic were held at the June 30 – July 1, 2004, CIP/DSS meeting held at Argonne National Laboratory, where consensus was reached that sensitivity and uncertainty analysis should be part of the software package of CIP/DSS. The classes of methods discussed here were presented at that and other meetings in some depth. To date, the sampling approach to sensitivity analysis presented in Sec. 2.1 has been applied using native capabilities within the Vensim implementation. The other topics presented here are under active exploration, and after further experience with CIP/DSS modeling and use of CIP/DSS results, may be recommended for implementation in future versions.

This paper represents a work in progress, and we look forward to updating it as we gain experience and try different approaches. In particular, we will aim over the next six months to evaluate both the methods described in this paper and other available methods in the context of the biological scenarios defined by DHS.

1. HOW DOES CIP/DSS WORK?

A brief overview of how CIP/DSS works provides a basis for describing some options for treatment of uncertainty and risk. The performance of the 14 infrastructures and their primary

---

[1] Computer and Computational Sciences, Los Alamos National Laboratory
[2] Decision and Information Sciences, Argonne National Laboratory
[3] Sandia National Laboratory
[4] Decision Applications, Los Alamos National Laboratory
[5] Apostolakis, George E., "How Useful is Quantitative Risk Assessment?," *Risk Analysis*, 24:3, pp. 515-520 (June 2004).
[7] Also known as state variables, consequence metrics in this paper refer to infrastructure model outputs that are passed through a database to the decision analysis module of CIP/DSS.

interdependencies are examined through detailed national and metropolitan models. The models provide information with respect to physical and transactional infrastructure activity and several other consequence metrics,[7] which are provided to the decision analysis module, where consequence metrics are combined and converted into ***decision metrics***,[8] such as presented in Table 1. For a given set of conditions and assumptions (a ***case***), the goal is for the infrastructure models to calculate and output an appropriate set of consequence metrics to the decision analysis module; that is, the infrastructure models will provide outputs that are either decision metrics without conversion or are as close as possible to the desired decision metrics.
The primary building block of the CIP/DSS is called a case. A case consists of consists of two or more ***scenario pairs;*** each scenario pair is composed of a ***Readiness Scenario*** and an ***Incident Scenario***:

- Base Scenario Pair
  - Base Readiness Scenario
  - Base Incident Scenario
- One or more Alternate Scenario Pair(s)
  - Alternate Readiness Scenario
  - Alternate Incident Scenario

Comparison of alternate scenario pairs with base scenario pairs indicates the effects that various investments and strategies, labeled here as ***optional measures*** (which include hardware, processes and strategies related to prevention, protection, mitigation, response, and recovery), could have if they were implemented by DHS or other organizations. Each scenario requires a separate simulation over a period of time (defined by the case) with the detailed national and metropolitan models.

Table 1 Decision Metrics for CIP/DSS Proof-of-Concept

	Category	**Metric**	**Units**
1	Human health and safety	Total fatalities (public and occupational)	deaths
2	Economic damage	Initial direct costs (e.g., cleanup, repair)	$
3	National security	Delay in force deployment capability	days
4	Socio-political	Public confidence	index
5	Environmental impacts	Unusable city blocks	blocks
6	Environmental impacts	Unusable farmland	acres

Table 2 summarizes the scenarios that are associated with each case and that need to be simulated by the models.

Table 2 Scenarios to be Simulated with the Infrastructure Models for Each Case

Scenario Name	**Description**
Base Readiness	Business as usual conditions; consequences in the absence of terrorist events
Base Incident	Postulated event occurs with no additional optional measures

[8] The decision metrics listed in Table 1 represent a minimum set selected for the proof-of-concept. A more extensive set is under development for future versions.

	implemented, beyond what exists at the time
Alternate Readiness 1	A specific set of additional optional measures are in place; postulated event is not initiated
Alternate Incident 1	Alternate Readiness 1 optional measures are in place; postulated event occurs
. . .	. . .
Alternate Readiness N	A specific set, different than other Readiness Scenarios, of optional measures are in place; postulated event is not initiated
Alternate Incident N	Alternate Readiness N optional measures are in place; postulated event occurs

The models are intended to identify, for each scenario of a given case, a set of consequences given by a vector, C, consisting of values for decision metrics, such as listed in Table 1.[9] For the purposes of this paper, let us consider the vector C to have only the 6 components shown in Table 1 (assume the values are summed over time without discounting over time).

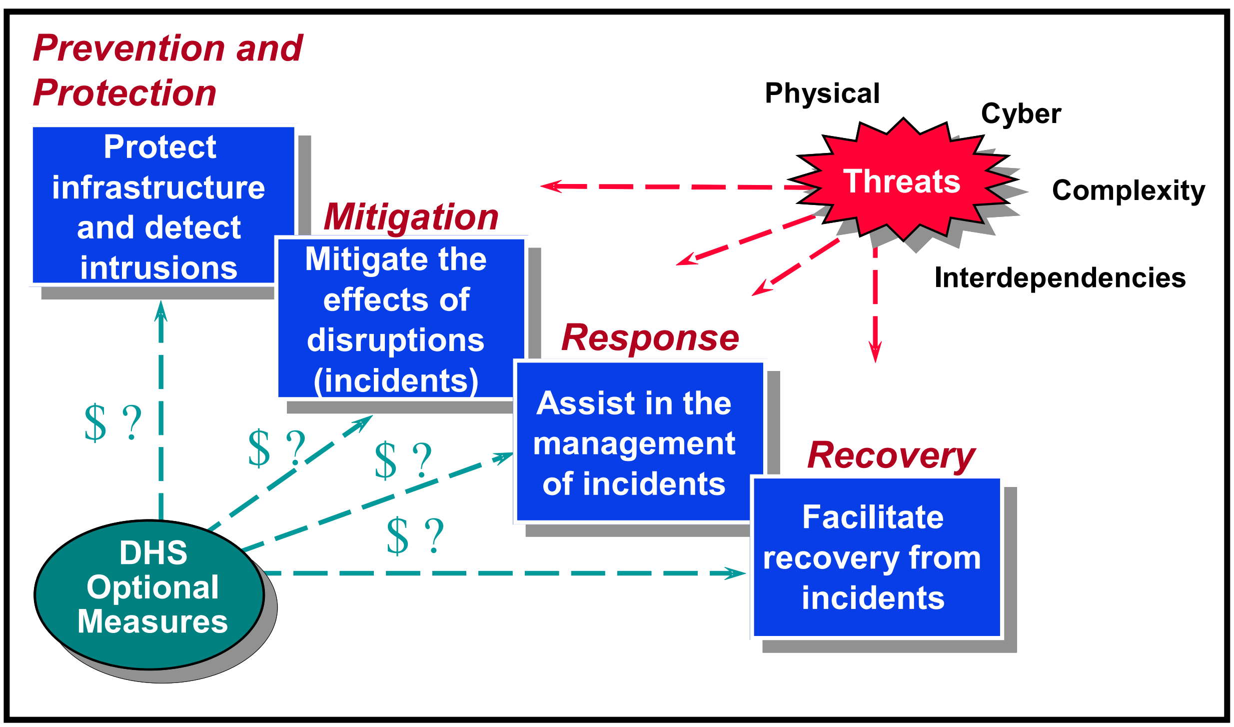

Figure 1 Analysis Framework for DHS Optional Measures with CIP/DSS

[9] The vector C is an array of values over the different consequence metrics and over time, where the time interval depends on the particular case.

Figure 1 demonstrates the overall analysis framework as CIP/DSS assists DHS in analyzing alternative optional measures that include prevention, protection, mitigation, response, and recovery.

2. UNCERTAINTIES

A brief examination of what is meant by *uncertainty* is in order before proceeding to recommended approaches. Kaplan[10] has presented a historical view of uncertainty analysis that relates modern views of uncertainty with ancient foundations and provides a sense of the difficulty in arriving at a precise description of uncertainty. For the purposes of the current effort, the definition of uncertainty from the draft ASME report[11] on risk analysis and management should suffice:

> *A measure of knowledge incompleteness due to inherent deficiencies in acquired knowledge. Also, a characterization of the degree to which the state of a system is unsettled or in doubt, such as the uncertainty of an outcome. In a quantified risk assessment, uncertainty is a representation of the confidence in the state of knowledge about the models and parameter values used.*

Although point values for consequence metrics are acceptable, decisions that are aided by CIP/DSS results would be more informed if the results reflected uncertainties, such as those introduced by CIP/DSS parameters and inputs (e.g., weather conditions at the time of an attack simulated by the Base Incident Scenario). Presenting the ranges for important CIP/DSS results may be appropriate to give users an appreciation of the extent of uncertainty. In addition, the decision analysis module allows CIP/DSS to account for risk attitudes other than risk neutral. DHS may wish to examine the implications for appropriate spending on protective measures if the government were to behave in a risk averse fashion rather than risk neutral for some or all decision metrics. Our society's relatively greater concern over a low probability and high consequence event versus a high probability and low consequence event with the same expected value for consequences is an indication that risk neutrality does not always match up well with public concerns.

Uncertain results from CIP/DSS could always be reduced to point values for simplicity or for users who do not care to address uncertainties. Thus, adding the capability to address uncertainty in CIP/DSS does not necessarily mean increasing complexity of presenting results.

There are a number of sources of uncertainties present in CIP/DSS analysis:

1. ***Structural***. Uncertainty is introduced into the consequence metrics because of assumed functional relationships within and between (i.e., interdependencies) infrastructure

---

[10] Kaplan, Paul G., *From "The Art of Conjecture" to "The Science of Inference:" A Historical Tour of Uncertainty Analysis*, Sandia National Laboratories, presented at PMC 2002.

[11] American Society of Mechanical Engineers (ASME), *Guidance on Risk Analysis and Management for Critical Asset Protection,* draft report submitted to the U.S. Department of Homeland Security, April 2, 2004. We note that a July 30, 2004 draft is now available.

models. An example is the quantity of electricity produced from an electrical generating unit, given the load, weather, availability of other generating units, and other important conditions. That is, the actual amount of electricity from that generating unit for those conditions is only approximated by mathematical relationships used in the CIP/DSS energy infrastructure model. A related example involving an interdependency is the quantity of natural gas used by that generating unit, if it were a natural-gas-fired generating unit. Such structural uncertainties in the models can have a major impact on consequences.

2. ***Inherent Randomness***. Uncertainty is introduced into consequence metrics because of inherent randomness in the physical world and CIP/DSS input parameters. An example is the forced outage of an electrical generating unit, which affects the quantity of generation from other generating units, which in turn affects some consequence metrics, such as electricity cost and environmental impacts associated with electricity generation.
3. ***Lack of Knowledge.*** Uncertainty is introduced into consequence metrics because of lack of knowledge about input parameters, such as in relationships used to estimate consequences. An example is the dose/response relationship that estimates the number of fatal cancer cases for a given integrated population exposure to radiation. Many dose/response relationships remain the subject of ongoing research and are imprecise. Educated estimates of remaining uncertainty are often available.
4. ***Estimating Future Conditions***. Uncertainty is introduced into consequence metrics because of the need to estimate model input parameters at a future point in time. Examples include the average premium for property/casualty insurance and the telephone call overload shape function (overload level vs. time). The historical information often gives an excellent starting point, but, presumably, future conditions and characteristics of the scenario could affect these parameters. The importance of this category of uncertainty increases as the time horizon of the CIP/DSS study increases.
5. ***Threat Characterization***. Uncertainty is also introduced as a result of the probability of attack and threat scenarios. This category of uncertainty is addressed in the discussion of risk in Sec. 3.

The structural uncertainties in the models need investigation before significant effort is expended addressing the effect of parameter uncertainties. In particular, the potential effect of structural uncertainties on the key consequence metrics needs to be estimated. If the level of uncertainty associated with a particular assumed structural relationship is high, consideration should be given to defining a new input parameter that specifically addresses that uncertainty. If the uncertainty in consequence metrics introduced by structural uncertainties can be shown to be very small compared to the uncertainty in consequence metrics introduced by uncertain input parameters, it may be possible to avoid treating structural uncertainties explicitly.

Given that a mathematical model is a formal statement of assumptions about a relationship among known inputs $\boldsymbol{x}$ and predicted outputs $\boldsymbol{y}$, the structure of a model defines how the characteristics $\boldsymbol{y}$ are determined from $\boldsymbol{x}$. It is usually a mathematical algorithm or set of computational rules that map $\boldsymbol{x}$ into $\boldsymbol{y}$. This set of algorithms or rules can combine elements from random number streams in stochastic models (or not and result in a deterministic model). Structural uncertainty arises from variation in $\boldsymbol{y}$ observed from plausible alternative model structures. Let $\mathcal{M} = \{ \mathrm{m}(x; z; \theta) \mid \theta \in \Theta \}$ define the set of alternative plausible models. In general θ is an enumeration of models or is a means to index a family of models. In practice, alternative models are typically explored, albeit mainly to improve the accuracy, runtime, or clarity of the model. It is relatively rare to explore the uncertainty induced on the output by variations in the structure. The recommended approach for this project is to acknowledge the existence of uncertainty arising from the structure of the model, but to admit the impracticality of exploring that uncertainty due to resource constraints.

2.1 Sampling Approaches to Addressing Uncertainty in Input Parameters

Estimating the uncertainty in consequence metrics is a desirable feature for CIP/DSS. The uncertainty in input parameters may stem from several sources, including inherent randomness, model structure, lack of knowledge, or the need to estimate future conditions. Sampling-based approaches to uncertainty and sensitivity analysis are both effective and widely used[12]. Analyses of this type involve the construction and exploration of a mapping from uncertain analysis inputs to uncertain analysis results. The basic idea is that analysis results $\mathbf{y} = \mathbf{m}(\mathbf{x}) = [y_1(\mathbf{x}), y_2(\mathbf{x}), \ldots, y_{nk}(\mathbf{x})]$ are functions of uncertain analysis inputs $\mathbf{x} = [x_1, x_2, \ldots, x_{nk}]$. In turn, uncertainty in input $\mathbf{x}$ results in a corresponding uncertainty in output $\mathbf{y}$. This leads to two questions:

a) What is the uncertainty in $\mathbf{y}(\mathbf{x})$ given the uncertainty in $\mathbf{x}$?, and
b) How important are the individual elements of $\mathbf{x}$ with respect to the uncertainty in $\mathbf{y}$?

Uncertainty analysis addresses the first question, and sensitivity analysis addresses the second. In practice, the process of performing an uncertainty analysis is very similar conceptually and computationally to doing sensitivity analysis.

Implementation of a sampling-based uncertainty and sensitivity analysis involves five components:

1. Definition of distributions $D_1, D_2, \ldots, D_{nk}$ that characterize the uncertainty in the components $x_1, x_2, \ldots, x_{nk}$ of $\mathbf{x}$,
2. Generation of a sample $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{nS}$ from the $\mathbf{x}$'s in consistency with the distributions $D_1, D_2, \ldots, D_{nk}$,
3. Propagation of the sample through the analysis to produce a mapping $[\mathbf{x}_j, \mathbf{y}_j]$, $j = 1, 2, \ldots, nS$, from analysis inputs to analysis results,
4. Presentation of uncertainty analysis results (i.e., approximations to the distributions of the elements of $\mathbf{y}$ constructed from the corresponding elements of $\mathbf{y}_j$, $j = 1, 2, \ldots, nS$), and
5. Determination of sensitivity analysis results (i.e., exploration of the mapping $[\mathbf{x}_k, \mathbf{y}_j]$, $j = 1, 2, \ldots, nS$).

Definition of the distributions $D_1, D_2, \ldots, D_{nk}$ that characterize the uncertainty in the components $x_1, x_2, \ldots, x_{nk}$ of $\mathbf{x}$ is the most important part of a sampling-based uncertainty and sensitivity analysis as these distributions determine both the uncertainty in $\mathbf{y}$ and the sensitivity of $\mathbf{y}$ to the elements of $\mathbf{x}$. The distributions $D_1, D_2, \ldots, D_{nk}$ are typically defined through an expert review process (or examination of appropriate data), and their development can constitute a major analysis cost. A possible analysis strategy is to perform an initial exploratory analysis with rather

[12] Helton, J.C. and F.J. Davis. 2000. "Sampling-Based Methods for Uncertainty and Sensitivity Analysis," *Sensitivity Analysis*. A. Saltelli, K. Chan, and E.M. Scott (eds). New York, NY: Wiley. pp. 101-153.

crude definitions for D_1, D_2, …, D_{nk} and use sensitivity analysis to identify the most important analysis inputs; then, resources can be concentrated on characterizing the uncertainty in these inputs and further analysis can be carried out with these improved uncertainty characterizations.

Several sampling strategies are available, including random sampling, importance sampling, and Latin hypercube sampling. The latter is very popular for use with computationally demanding models because its efficient stratification properties allow for the extraction of a large amount of uncertainty and sensitivity information with a relatively small sample size. In addition, effective correlation control procedures are available for use with Latin hypercube sampling. The drawback of Latin hypercube sampling is that it can lead to biased estimates.

Propagation of the sample through the analysis to produce the mapping $[\mathbf{x}_k, \mathbf{y}_k]$, k = 1, 2, …, nS, from analysis inputs to analysis results is often the most computationally demanding part of a sampling-based uncertainty and sensitivity analysis. The details of this propagation are analysis specific and can range from very simple for analyses that involve a single model to very complicated for large analyses that involve complex systems of linked models.

Presentation of uncertainty analysis results is generally straightforward and involves little more than displaying the results associated with the already calculated mapping $[\mathbf{x}_k, \mathbf{y}_k]$, k = 1, 2, …, nS. Presentation possibilities include means and standard deviations, density functions, cumulative distribution functions (CDFs), complementary cumulative distribution functions (CCDFs), and box plots. Presentation formats such as CDFs and box plots are usually preferable to means and standard deviations because of the large amount of information that is lost in the calculation of means and standard deviations.

Determination of sensitivity analysis results is usually more demanding than the presentation of uncertainty analysis results due to the need to actually explore the mapping $[\mathbf{x}_k, \mathbf{y}_k]$, k = 1, 2, …, nS, to assess the effects of individual components of $\mathbf{x}$ on the components of $\mathbf{y}$. Available sensitivity analysis procedures include examination of scatterplots, regression analysis, correlation and partial correlation analysis, and other standard statistical techniques.

Sampling-based uncertainty and sensitivity analysis is widely used, and as a result, is a fairly mature area of study. However, a significant challenge is the education of potential users of uncertainty and sensitivity analysis about such issues as:

- The importance of such analyses and their role in both large and small analyses;
- The need for an appropriate separation in the conceptual and computational implementation of analyses of complex systems of *aleatory* (irreducible uncertainty resulting from inherent variability of a quantity) from *epistemic* (theoretically reducible uncertainty resulting our lack of knowledge about a quantity) uncertainty
- The need for a clear conceptual view of what an analysis is intended to represent and a computational design that is consistent with that view, and
- The importance of avoiding deliberately conservative assumptions if meaningful uncertainty and sensitivity analysis results are to be obtained.

The Vensim modeling approach offers a powerful, convenient, but somewhat limited way to examine model and input uncertainties. Monte Carlo simulation can be automated and conducted conveniently in each infrastructure model, such that the effect of variation in important input parameters on key consequence metrics can be examined. Model inputs can be varied according to a specific list of values, range of values, or according to specified probability distributions. Model outputs are specified and the effect of the variation of the inputs can be seen on these variables. In the figure below, hospital bed utilization, a metric of the public health model, is examined as the major inputs to the public health model are assigned a uniform random distribution over their individual minimum and maximum values. A significant assumption for these runs is that all infrastructure systems are fully functional.

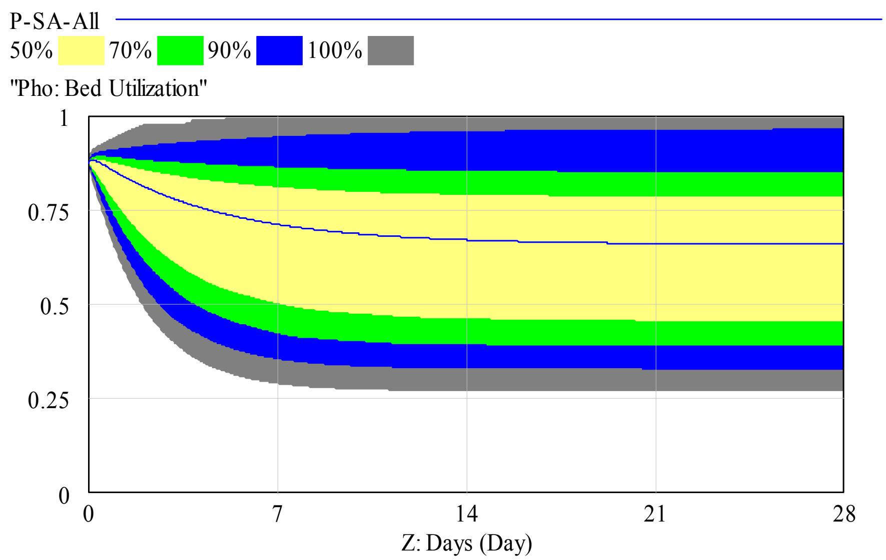

The chart shows the observed variation in bed utilization over a 28 day period by examining the output of 500 runs of the model. In each run, a new set of inputs was sampled from the probability distributions of the inputs; the model was executed, and the output values stored. Latin Hypercube sampling (LHS) is a stratified sampling technique where the random variable distributions are divided into equal probability intervals. A probability is randomly selected from within each interval for each sample. This distributes input samples over equi-probable bins, making it more likely that a representative range of the input space is sampled. In the figure above, the green and yellow central regions represent those simulation runs that accounted for 70% of the observed values. The blue line splitting the yellow region is the value of a single run based on the default input values to the model. The utilization for this (blue) run tends to stabilize near 0.7. The width of the 70% variation region on day 28 is about 0.50 for a ratio of approximately 0.70 (0.50/0.70).

This approach to inspecting the sensitivity of the sector models and subsector elements was applied to the Metropolitan model and is discussed in the documentation for those sectors[13]. In these analyses, a significant assumption was that the inputs were uncorrelated, as each sector model was exercised for sensitivity independent of the inputs used to initialize the remaining

[13] Bush, Brian, et al., "Documentation for the Metropolitan Model of the Critical Infrastructure Protection Decision Support System," Los Alamos National Laboratory report LA-UR-xxxx-04, in development.

sector models. Due to the lack of adequate data, it was infeasible to derive an estimate of the correlation of the metropolitan inputs. The sensitivity analysis was nonetheless useful. It identified several shortcomings in the models; enabled an initial assessment of potential model uncertainty; and helped build confidence that the models were exhibiting expected behaviors.

While this sensitivity approach is useful for visual inspection of model output variation due to uncertainties in inputs, it lacks the ability to identify the relative importance of input parameters. Adoption of techniques for more precise statistical estimation of uncertainty, sensitivity, and importance, as discussed in Saltelli[14] is desired.

2.2 Generalized Approaches to Addressing Uncertainty in Input Parameters

The sampling approach outlined above works well when uncertainty about input parameters can be precisely stated, for example as coming from a particular probability distribution D_k for input variable x_k. But probabilistic characterizations can give the appearance of more knowledge than is really present, and alternative representations for uncertainty merit consideration for their potential to represent uncertainty in situations where little information is available.

Generally, it might be the case that only limited or qualitative information about the uncertainty of a parameter x_k is known, for example its shape, or its minimum or maximum values, or its moments (mean and variance), or perhaps just its parametric family (for example being normal or beta). Methods based on "random intervals" and "probability boxes" (p-boxes)[15] are available for handling these situations. These methods are closely related to alternative uncertainty quantification techniques including fuzzy systems, possibility theory, Dempster-Shafer evidence theory, random sets, and imprecise probabilities[16]. Some of these will also be considered under Risk in Section 3 below, and we can collectively identify these as **Generalized Information Theory (GIT)**[17] approaches.

GIT techniques do work not with *single* probability distributions, but rather with *collections* of them. As an example to be used in this subsection, we will consider the Metropolitan Food sector, for which an initial analytical treatment is currently underway[18]. Given the following distributional assumptions about certain internal and external variables:

Water demand rate currently supplied, commercial: $W = Beta(4,1)$
Electrical fraction available, industrial: $E = Beta(25,1.3)$

---

[14] Saltelli, A., Tarantola, S., Campolongo, F., and Ratto, M, *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*, Wiley & Sons, New York, NY, 2004.
[15] Ferson, Scott; Kreinovich, V; Ginzburg, LM; D.S. Myers and K. Sentz: (2002) "Constructing Probability Boxes and Dempster-Shafer Structures", SAND Report 2002-4015, Sandia National Lab, Albuquerque NM, http://www.sandia.gov/epistemic/Reports/SAND2002-4015.pdf
[16] Joslyn, Cliff and Booker, Jane: ``Generalized Information Theory for Engineering Modeling and Simulation", in: *Engineering Design Reliability Handbook*, ed. E Nikolaidis and D Ghiocel, CRC Press, in press
[17] Klir, George and Wierman, Mark J: (1999) *Uncertainty-Based Information Elements of Generalized Information Theory*, Springer-Verlag,Berlin
[18] Joslyn, C and Bettencourt, L (2004): "Analytical Representation and Uncertainty Quantification of the Food Sector", in preparation

Transportation fractional free flow: *T = Beta(1.5,.6)*
Time from Raw to Processed: *TRP = Normal(3,.25)*
Time from Processed to Retail: *TPR = Normal(7,1)*

we can derive the following sensitivity analysis of the amount of uncontaminated beef processed from sampling:

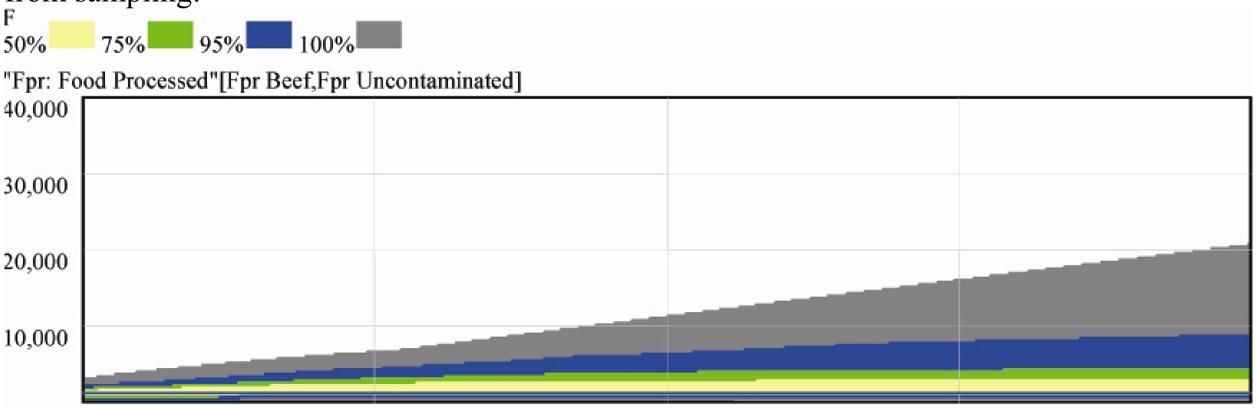

Again, here we presumed that *TPR* is normally distributed with mean 7 and standard deviation 1, which is shown below as `norm` in black as the particular cumulative distribution Normal(7, 1). But instead, we may know that:

- *TPR* is normally distributed, but all we know about the mean is that it is in the interval [6,8], and all we know about the standard deviation is that it's in the interval [.5,1.5]. This can be characterized as a GIT structure known as a p-box, which is a statistical representation of a random interval, also known as a Dempster-Shafer uncertainty structure on the line[19]. This p-box is shown in the figure below as `normint` in red.
- Alternatively, we may still know that *TPR* has a precise mean 7 and standard deviation 1, but that is *all* we know; in particular, *TPR* could be drawn from *any* parametric family of distributions, not just normals. This can, in fact, be represented, and indeed is shown as the p-box `meanst` in blue.

[19] Joslyn, Cliff and Ferson, Scott: (2004) ``Approximate Representations of Random Intervals for Hybrid Uncertainty Quantification", in: *4th Int. Conf. on Sensitivity Analysis of Model Output (SAMO 04),* ftp://ftp.c3.lanl.gov/pub/users/joslyn/samo04.pdf

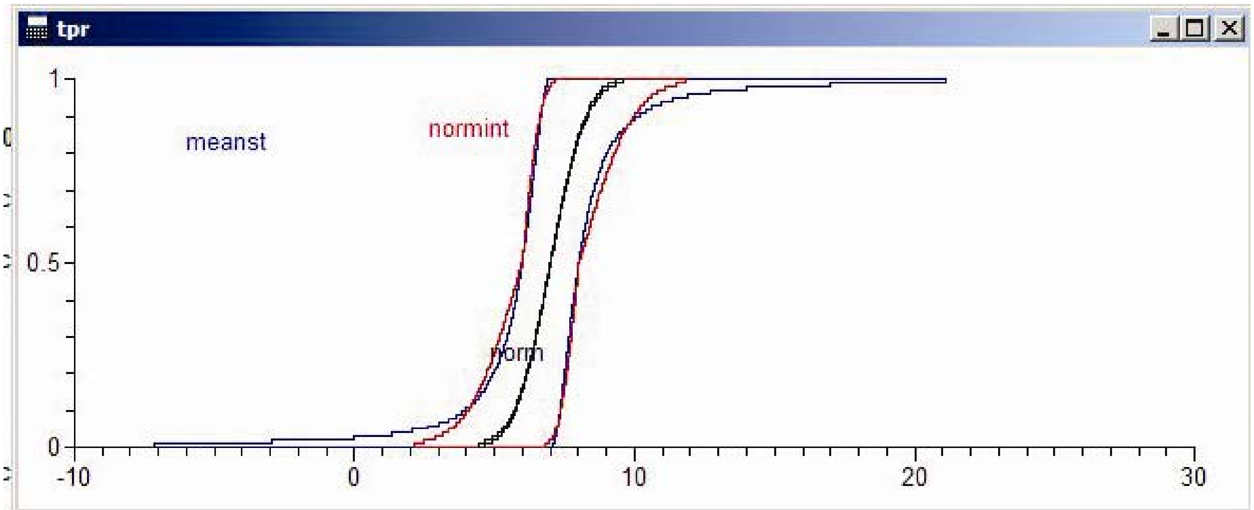

This discussion is merely exemplary, and in any particular situation, determination of the proper *way* to quantify information about a particular variable, in addition to the actual *numbers* to use, depends on the interaction between the model builder and the relevant data resources, including content experts. In some cases, the model builders may only have access to the most qualitative information, such as a variable being between two landmark values, or even linguistic information such as "very large". While methods to accommodate all of these cases is part of a wider GIT program involving fuzzy systems and possibility theory[16], here we will consider those cases which allow characterization of the uncertainty of variables in terms of p-boxes.

In any event, returning to our example, each of the p-boxes `meanst` and `normint` are the regions containing all the cumulative distributions satisfying their respective criteria. Note that while they are similar, they are also subtly different, although in each case, the single distribution Normal(7, 1) is included as one possible single cumulative distribution within the p-box in question. This reflects the great constraint present in specifying a particular distribution, as distinguished from the relative lack of constraint present in the actual information available to characterize the variable. Thus given either of these alternate scenarios, the p-box representations are more accurate, reflecting the greater range of uncertainty present, and most importantly not *forcing* the investigator to select *exactly one particular* distribution.

Risk and reliability approaches using GIT treatments of uncertainty are in active development[20]. Given a GIT representation of uncertainty, a complete analysis can then proceed in multiple ways.

First, an analytical approach is possible. In the Metropolitan Food sector, Joslyn and Bettencourt[18] have made an initial determination that the long-term qualitative behavior of the amount of uncontaminated beef processed is sensitive to the sign of the quantity $\omega = 4\alpha\beta - \beta^2$, where $\alpha = WE/TRP$ and $\beta = TE/TPR$. The first figure below shows the four single cumulative distributions for *W, E, TRP*, and *TPR,* while the second figure shows the p-boxes for α, β, and ω derived by convolution of the algebraic equations under the assumption of independence. Note

---

[20] Helton, JC and Oberkampf, WL eds.: (2004) "Special Issue: Alternative Representations of Epistemic Uncertainty", *Reliability Engineering and Systems Safety*, v 95:1-3

that these intermediate quantities in the lower figure have substantial "width" to them, and thus greater uncertainty than can be captured by the requirement of determining individual distributions. In particular, note that the p-box for ω spans zero, indicating that there is enough uncertainty present in the input parameters to make the long-term behavior difficult to predict even at the *qualitative* level.

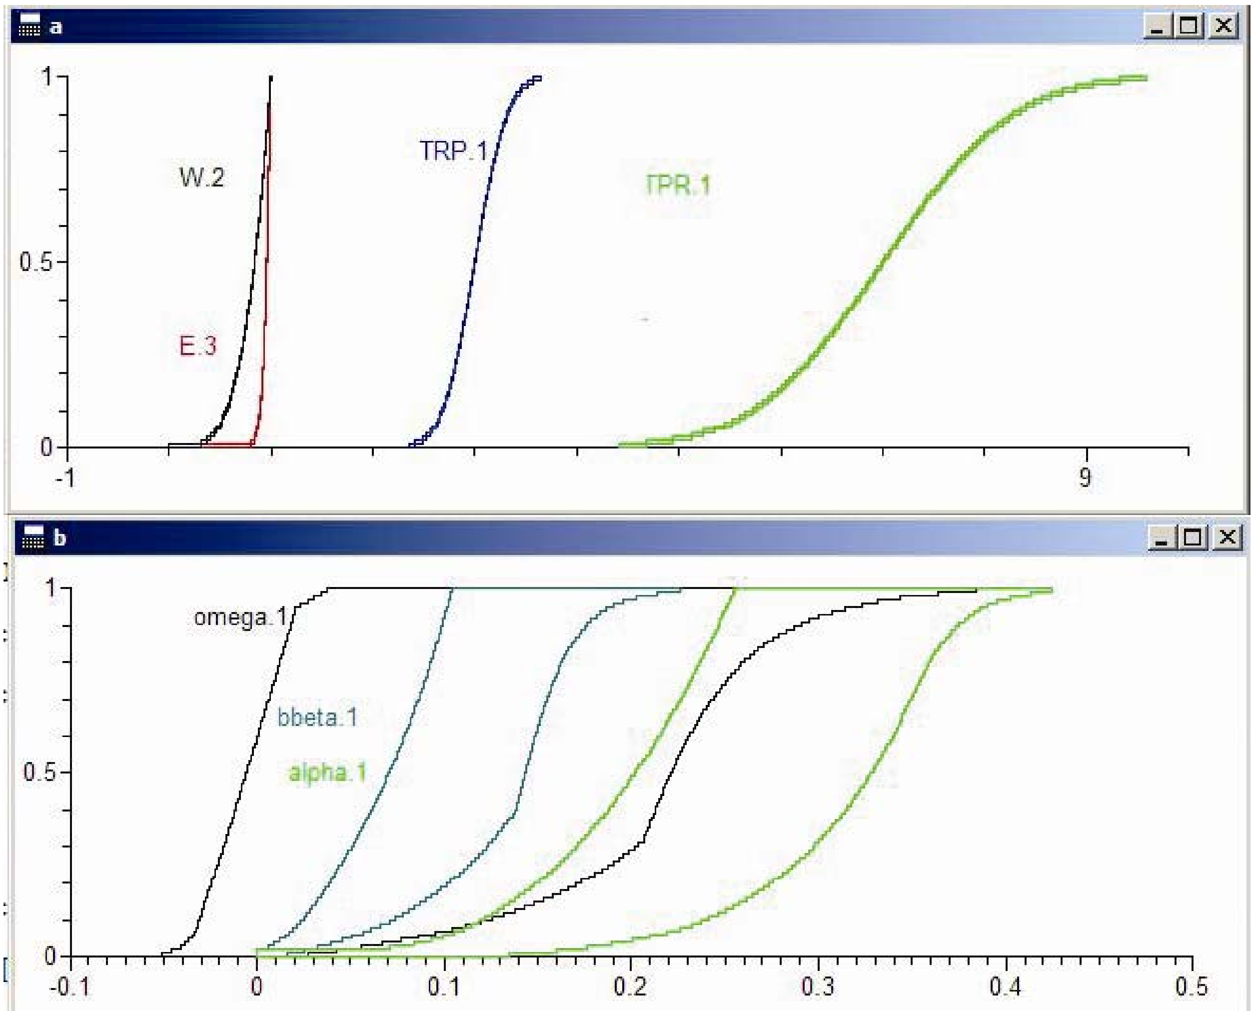

However, there are many reasons why purely analytical approaches to CIP/DSS models are not generally feasible, and thus the greater accuracy provided by GIT representations of uncertainty must be balanced against the greater difficulty of working with them (for example, Vensim does not support GIT methods). Instead, sampling approaches to GIT treatments of risk problems are in active development[21,22], including early work on their use specifically for sensitivity analyses[23]

[21] JC Helton, JD Johnson, and WL Oberkampf, "An Exploration of Alternative Approaches to the Representation of Uncertainty in Model Predictions", *Reliability Engineering and Systems Safety*, v. 95:1-3, pp. 39-72

[22] Joslyn, Cliff and Kreinovich, Vladik: (2004) "Convergence Properties of an Interval Probabilistic Approach to System Reliability Estimation", *Int. J. General Systems,* in press

[23] J. C. Helton, J.D. Johnson, W.L. Oberkampf, and C.J. Sallaberry, "Sensitivity Analysis in Conjunction with Evidence Theory Representations of Epistemic Uncertainty", *Reliability Engineering and Systems Safety*, in press

This subsection is intended only to indicate future direction of these classes of approaches. All of the subjects mentioned here are areas of current research.

2.3 Using Monte Carlo Output in the Decision Analysis Module

An important characteristic of the decision analysis methodology being applied in CIP/DSS is that it is highly suited for Monte Carlo simulation. For each individual run of the infrastructure models, a set of decision metrics is produced. The utility function operates on each individual run to produce a utility value. The variation in utility values across the numerous runs of the infrastructure is an important, observable characteristic. The distribution of utility values for the runs allows us to compute a mean and standard deviation, as well as ultimate ranges. The principles of utility theory allow us to use ***expected utility*** to determine preferred courses of action.

When examining two different (but structurally identical) scenarios (e.g., base incident and alternate incident 1) within the same case, starting with the same random number seed for all Monte Carlo runs offers the potential to compute a distribution of *utility differences* between the scenarios that is, in a sense, not dependent on the sequence of random numbers generated in the simulation (this will be especially helpful in comparing results based on 2 or more different decision-maker profiles). The mean value for the difference in utility can be computed along with statistical characteristics. In addition, when the utility values for the two scenarios are close, we can also compute the probability that the utility of one scenario is greater than the other. This approach also may serve to reduce some of the inherent uncertainties (e.g., a consequence metric that tends to be greater when the weather is hot will be consistent between the two scenarios because the Monte Carlo drawing for the weather is the same in each scenario).

It should be noted that the expected utility of uncertain decision metrics (e.g., 1,000 sets of the decision metrics for each of 1,000 model runs using Monte Carlo simulation) is generally not equal to the utility of the expected values of the decision metrics. That is, if we eliminate the uncertainty by using expected values for each decision metric and then operate on those single values with the utility function, we will generally obtain a different utility than when we compute the individual utility of the decision metrics produced from each Monte Carlo run and then determine the expected value for that array of utility values. The preference assessment to obtain the utility function is conducted under conditions of uncertainty, and therefore accounts for risk attitudes (risk averse, risk neutral, or risk prone) of the decision maker. If the decision-maker profile being used is not completely risk neutral, the difference in the calculation described above is critical.

The loss-of-load probability (LOLP) for an electrical generating system with a large quantity of hydroelectric generation can be used to illustrate this point. An LOLP equal to or less than 1 day in ten years (2.7×10^{-4}) is considered to be very good for an electrical generating system. An LOLP greater than 1 day per year (2.7×10^{-3}) is considered to be not so good. Suppose the hydroelectric generation varies considerably from year-to-year, and the probabilities of those types of years, based on extensive historical data, can be characterized as shown in Table 3.

Table 3 Hydroelectric Generation Characteristics and Calculation of LOLP

Type of Hydro Year	Probability of Occurrence	Hydro Generation (Megawatt-hours [MWh])	Loss-of-Load Probability (LOLP)
Very wet	0.10	10,000	0.0
Wet	0.25	8,000	0.0
Normal	0.30	6,000	0.0
Dry	0.25	4,000	5.4×10^{-3}
Very dry	0.10	2,000	13.7×10^{-3}

Further suppose the LOLP is 0.0 (no failures) whenever hydroelectric generation is 6,000 MWh or more, and the LOLP is two days per year (5.4×10^{-3}) at 4,000 MWh and is 5 days per year (13.7×10^{-3}) at 2,000 MWh. Expected annual hydro generation is 6,000 MWh: therefore, the LOLP of the expected value for hydro generation is 0.0 (no outages; perfect performance). However, the expected LOLP (computed by weighting the LOLP for each type of hydro year by its probability) is 2.7×10^{-3} (not so good). Two completely different answers are obtained based on when the uncertainty was addressed. The calculation of expected utility rather than utility of expected consequences exhibits similar characteristics.

3. RISK

The draft ASME report defines risk as:

> *The combination of the probability and consequences of an adverse event (i.e., threat). When the probability and consequences are expressed numerically, risk is the product of those values, which are combined considering uncertainties. Since probability and consequences are usually expressed as ranges rather than point values (to account for uncertainties associated with their computation), risk can generally be expressed in terms of a mean value and upper and lower bounds (or range). When expressed qualitatively, risk may be represented by the probability and consequences as indicated by their relative positions on a risk matrix. However, it is usually desirable to express qualitatively derived probabilities and consequences numerically so risk values can be compared to the results of other analyses.*

Analysis conducted with CIP/DSS is intended to assist in addressing risk and investment strategies, such as:

- Incorporating consequence, vulnerability, and threat information into an overall risk assessment, what are the highest risk areas?

- What investment strategies can be made that will have the most impact in reducing overall risk?

DHS needs to demonstrate a rational approach to funding preventive and protective measures and consequence mitigation, response, and recovery strategies. Given the relatively large resources dedicated to DHS, scrutiny is to be expected, such as the *Time* article entitled "How We Got Homeland Security Wrong."[24] The article stated that the vast majority of 2003 federal agency grants totaling $13.1 billion to states and localities for homeland-security purposes "was distributed with no regard for the threats, vulnerabilities and potential consequences faced by each region." Although federal agencies must operate in a political environment and obtain Congressional approval for spending patterns, analysis of funding options with CIP/DSS should assist DHS in making risk reduction one of the important factors in such funding decisions and in presenting risk-informed funding recommendations to Congress.

Risk is measured as a combination of probability of attack and consequences given an attack, each of which may be uncertain. As described in Sec. 2, various options are available for estimating the uncertainty in decision metrics. Uncertainties with respect to type, nature, and frequency of threat are much greater than most other problems that have been historically studied. For this reason, various advanced methods, including possibility theory, fuzzy sets, and approximate reasoning are being investigated to help address threat.[25] Probabilistic techniques are considered less appropriate because terrorist actions are intentional, not random, and detailed statistical analysis of past data is insufficient given the dynamic nature of the problem.[26]

Although it is not within the CIP/DSS work scope to estimate probability of attack, CIP/DSS must be able to accommodate the form of threat information available to DHS. At this time, it appears that reduction of threat information to a probability of attack, with a range of uncertainty, is feasible. However, further investigation and work with DHS is needed to verify this approach.

If probability of attack data is available, with an uncertainty range, the resultant risk can be portrayed using the combined uncertainties from the CIP/DSS decision metrics and from the threat. A typical application will involve comparing a base scenario (no additional investments beyond existing measures) with an alternate scenario that includes investments for various optional measures. Since the investment in optional measures is made before the attack, the comparison ultimately comes down to risk reduction (measured by comparing outputs from the CIP/DSS decision analysis module for the base scenario and for the alternate scenario combined with the probability of attack) versus expenditures for optional measures. By examining different investment strategies with CIP/DSS (e.g., different optional measures or different levels of implementation or different locations for implementation), DHS can gather information on what appropriate investment strategies might be, given the threat information.

[24] Ripley, Amanda, "How We Got Homeland Security Wrong: The Fortification of Wyoming, and Other Strange Tales from the New Front Line," *Time*, March 29, 2004 (pp. 32 -37).

[25] Darby, John, *Evaluating Terrorist Risk Using Possibility Theory, Fuzzy Sets, and Approximate Reasoning*, Los Alamos National Laboratory.

[26] Darby, J., B. Bush, S. Eisenhawer, and T. Bott, *Methodology for Optimizing Allocation of Resources to Protect Infrastructure against Acts of Terrorism,* Los Alamos National Laboratory, LA-UR 04-0590 (Dec. 8, 2003).

It should be noted that DHS investment strategies in an alternate scenario may change the probability of attack and/or its uncertainty when compared to a base scenario. For example, the threat of airplane hijacking has presumably been reduced by introduction of new security procedures, barriers preventing cockpit entry, and wider use of air marshals. Thus, we are currently living the alternate scenario in which those investments have been made and the probability of hijacking is less than it would have been in the base scenario (without those measures). CIP/DSS needs a framework to allow adjustment of probability of attack based on what optional measures are implemented in the alternate scenario.

Another way in which CIP/DSS information can be useful to DHS is to estimate the minimum probability of attack that justifies implementing a particular set of optional measures. That is, a base scenario and an alternate scenario are analyzed as usual with CIP/DSS. The reduction in consequences, as measured by CIP/DSS decision metrics, between the alternate scenario and the base scenario is then compared with the investment costs to determine the minimum probability of attack at which these particular investments are justified. Any probability of attack equal to or greater than this minimum would indicate that the alternate scenario investments are appropriate.

Considering the uncertainty in decision metrics, an upper and lower bound for the minimum probability of attack that justifies implementing a particular set of optional measures can be estimated using CIP/DSS. This range for minimum probability will help CIP/DSS users realize that, although extensive calculations lie behind the CIP/DSS outputs, considerable uncertainty in results will often be the case.

Of course, there are numerous threats (chemical, biological, explosives, natural disasters, etc.) and DHS investment strategies to investigate. To the extent that a representative basis set of attacks can be defined by DHS, CIP/DSS could be used to examine the entire set and indicate what types of optional measures look best overall, given the current thinking on probabilities of attack for the various types of threats. Some optional measures are likely to be good investments even though they are not justified for any single type of threat because they serve to reduce risk for more than one type of threat. Once CIP/DSS has been set up to examine a certain type of threat, the capability to change the probability of attack for any threat (e.g., to reflect new information) and examine the appropriate scenarios quickly will be a genuine strength of the analytic procedure.

4. ACKNOWLEDGEMENTS

GIT calculations were performed in, and screen shots taken from, RAMAS RiskCalc 4.0[27,28] We would also like to acknowledge the assistance of many other project members from the three laboratories, and especially Luis Bettencourt, Brian Bush, and Michael Ebinger.

[27] http://www.ramas.com/riskcalc.htm
[28] Ferson, Scott: (2002) *RAMAS Rick Calc 4.0 Software: Risk Assessment with Uncertain Numbers*, CRC Press