Proceedings of IMECE2005:
2005 ASME International Mechanical Engineering Congress and Exposition
Orlando, Florida, USA, November 5-11, 2005

IMECE2005-81181

APPLYING INFORMATION ECONOMICS AND IMPRECISE PROBABILITIES TO DATA COLLECTION IN DESIGN

Jason Matthew Aughenbaugh
gtg224k@mail.gatech.edu

Jay Ling
gtg004v@mail.gatech.edu

Christiaan J. J. Paredis
chris.paredis@me.gatech.edu

Systems Realization Laboratory
G. W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0405
www.srl.gatech.edu

KEYWORDS: information economics, imprecise probabilities, value of information, decision-based design

ABSTRACT

One important aspect of the engineering design process is the sequence of design decisions, each consisting of a formulation phase and a solution phase. As part of the decision formulation, engineers must decide what information to use to support the decision. Since information comes at a cost, a cost-benefit trade-off must be made. Previous work has considered these trade-offs in cases in which all relevant probability distributions were precisely known. However, engineers frequently must estimate these distributions by gathering sample data during the information collection phase of the decision process. In this paper, we introduce principles of information economics to guide decisions on information collection. We present a method that enables designers to bound the value of information in the case of unknown distributions by using imprecise probabilities to characterize the current state of information. We illustrate this method with an example material strength characterization for a pressure vessel design problem, in which we explore the basic performance, subtleties, and limitations of the method.

INTRODUCTION

Engineering design is a sequential and iterative process, consisting of the phases of product planning, clarification of task, conceptual design, embodiment design, and detail design [1]. Decision-based design research recognizes a sequence of decisions in the design process and emphasizes the importance of these decisions to the success of the design [2-4]. Each decision has two phases—*problem formulation* and *problem solution*. Part of the decision formulation phase is a sub-decision problem regarding how much information to collect—the subject of this paper.

By expending resources to create, collect, and analyze information, engineers can support decision-making throughout the design process. This information may provide value to the designers by leading to a better final design, but this benefit is uncertain until resources are spent and the information is actually collected. In this paper, we have incorporated the management of this cost-benefit tradeoff into the design decision model. We define the principles of such information management as *information economics* [5].

In related research, Gupta et al. have demonstrated the importance of incorporating the cost (in terms of number of design alternatives considered) of decision-making into the overall design decision model [6], but they do not provide a method for estimating the value of information in actual design problems. Radhakrishnan and McAdams consider the cost-benefit trade-offs in selecting models of various levels of abstraction in engineering design [7]. They present a framework in which a designer can reason about model uncertainty, but they admit that the designer is left with little guidance in estimating the actual value of information from different models. Along similar lines, Bradley and Agogino develop a decision-analytic approach to assist designers in cost-benefit analysis of resource expenditures using precisely characterized probability distributions to guide and prioritize information collection [8], but they do not explain how to estimate these distributions.

In the simulation literature, statistical output analysis is commonly performed to assess whether a sufficient number of simulation replicates have been performed to obtain statistically significant conclusions [9]. However, since the analysis is usually performed based on accuracy requirements, one cannot easily formulate the trade-off with respect to the simulation cost. As in any kind of cost-benefit analysis [10], a common unit of measure is needed. This can be achieved here by using the *economic value of information* [11]. Although the economic value of information is clearly correlated with accuracy, they are not equivalent. For example, a very accurate and expensive model used to distinguish between two alternatives that differ significantly in performance has less value than a simpler, less expensive model that could have made the same distinction at a much lower cost.

In this paper, we present a method that uses imprecise probabilities, as described by Walley [12], to extend the applicability of information economics to cases in which probability distributions are not perfectly known to designers. By using imprecise probabilities, this approach gives the designer a method by which information economics can be applied to the management of statistical data collection in support of engineering design decisions. We present this approach in the context of a pressure vessel design and material strength characterization example, including a computational experiment to illustrate how this approach performs under various conditions. Our primary contribution is the development of a method for bounding the value of information using imprecise probabilities when true probability densities are unknown. This contribution potentially can have significant impact on engineering design by opening more problem classes to formal cost-benefit analysis during the problem formulation phase and information collection tasks of design decisions.

INFORMATION ECONOMICS

In this section, we introduce the basics of *information economics*. *Economics* is the *study of choice under conditions of scarcity* [13]. Extending this definition, *information economics* is the study of choice in information collection and management when resources to expend on information are scarce. Because designers face a scarcity of resources, such as time and money, in the design process, the principles of information economics should be applied to the information collection process in engineering design.

The area of information economics grew out of statistical decision theory in the 1950s when Marschak published a series of papers on the economics of information and organization [5]. Recently, with the explosion of new information technologies, information economics has regained attention within the broader context of information management. Current areas of research focus on corporate finance and industry policy, such as intellectual property rights, industry regulation, and fostering innovation [14, 15], or on the infusion of information technology into a corporation [16]. Within engineering, the focus of information management has been primarily on data exchange, interoperability, and visualization to support collaborative design (for an overview of these areas, refer to the following review articles [17-21]). In a more general sense, information economics presents principles by which the cost-benefit tradeoffs of information collection can be managed in engineering design. Many of these concepts have been developed and employed previously in standard micro-economics and the theory of the economic value of information, pioneered by Marschak [5] and summarized by Lawrence [11].

Our goal here is to relate these concepts directly to the management of information and uncertainty in engineering design. A substantial difference between our area of application and that of Marschak and Lawrence is the availability of perfectly known probability distributions—knowledge that Marschak and Lawrence assume but engineers often lack. Specifically, we explore the decision of whether to continue to collect random data samples in the context of estimating the material strength for a pressure vessel design, as described in the next section. Since the distribution of the material strength is not known *a priori*, a direct application of existing methods is impossible.

EXAMPLE PROBLEM

Throughout the remainder of this paper, we discuss the application of information economics in the context of a pressure vessel design example. This example has been used previously to demonstrate the value of explicitly representing the imprecision in the available characterization of uncertainties for the design decision by using imprecise probabilities [22]. We now extend this experiment to explore the decision of how much information to collect in order to support design decisions.

In the example problem, we design a pressure vessel to meet certain requirements while maximizing payoff. The complication is that the pressure vessel is to be built using a material whose yield strength is unknown. The designers believe that the yield strength is well modeled as a normally distributed random variable, but they do not know the parameters of the normal distribution. They can perform yield strength tests, thus sampling the distribution at a cost *c* per test.

In our experiment, each yield strength test represents one sample of the true material strength distribution. The samples are drawn from a random normal distribution whose parameters are unknown. Specifically, the material strength is a random variable X such that:

$$X \sim N(\mu, \sigma^2)\ .$$

The mean μ and variance σ^2 are unknown, and the goal of the information collection is to determine these parameters such that a good design decision can be made. The experiment consists of drawing a set of n samples $\{x_i\}_{i=1}^{n}$ from X. Each sample x_i that is drawn from the distribution is a piece of information. This piece of information can be used to help characterize the true nature of the uncertainty. The more samples drawn, the more information a designer has about the truth. However, every sample drawn costs the designer resources, for example time and money. Unless the designer has infinite resources, he or she cannot collect the infinite number of samples necessary for a perfect characterization of the distribution. Thus, the designer needs to determine when to stop collecting information—in this case, data samples. For this, we turn to a cost-benefit analysis.

COST-BENEFIT ANALYSIS

As a designer collects data samples x_i, the marginal benefit of an additional sample decreases. For example, if the designer has only 10 samples, an 11^{th} sample will usually be quite valuable; in contrast, if the designer has 1000 samples, the 1001^{th} sample will be considerably less valuable. In this sense, information displays diminishing returns. At some point, the cost of gathering additional information will outweigh the benefit. Thus, the value of a sample is not merely inherent in the sample; rather, the value is measured as viewed from the perspective of the designer. A fundamental principle of information economics is that a decision maker (DM) should continue to collect information only as long as there is an information source available whose net value is positive. Putting things into more standard micro-economic terms, a rational designer will stop taking data samples at the point where the marginal benefit of the next sample is less than or equal to the marginal cost of acquiring it.

A formalization of the basic cost-benefit analysis noted above has been summarized in the context of information by Lawrence [11]. In his work, the measure by which information can be managed is *value*. We primarily adopt Lawrence's definitions, but there are some fundamental differences between our approaches, stemming from differing assumptions of initial information. Lawrence assumes that the DM has perfect knowledge of the probability distributions that describe uncertain events. In our engineering design example, the estimation of these distributions is the indirect goal of information collection, with the direct goal being a good design decision. Since perfect knowledge about the distribution is not available, we develop a method for estimating the value of information using imprecise probabilities.

MATHEMATICAL PROBLEM FORMULATION

The value of information is measured by how the information changes the design decision. In this section, we illustrate a framework for determining this value. We start by examining how new information changes the DM's probability estimates over the state space. Next, we review the payoff of a decision (a metric for measuring the value of a decision), and show how payoff is used to arrive at an optimal decision under uncertainty. Then, we provide general definitions of information and information sources. Finally, we unite these concepts in the value of information framework and illustrate its application with an example problem.

Specifying probabilities over the state space

For every decision problem, a DM has a set of available *actions* $\boldsymbol{A}=\{a\}$ from which to chose one action. All possible states of the world form a *state space* $\boldsymbol{X}=\{x\}$. In our example problem, the action a consists of a set of design variables that specify the pressure vessel dimensions, and the state of the world is the actual material strength x of the material used in a particular pressure vessel. The material strength, or state, is assumed to be normally distributed with associated probability density function $p(x)$, which is unknown to the designer. The state of the world is outside the DM's control, so the DM can at best estimate the probabilities of the different individual states, thus forming the estimated distribution $\tilde{p}(x)$.

From among the many possible interpretations of probability [12, 23-26], we interpret the DM's estimate $\tilde{p}(x)$ according to a subjective interpretation of probability. Under a *subjective* interpretation, probabilities are an expression of belief based on an individual's willingness to bet [23, 25, 27]. While there may be random variables that assume outcomes according to true relative frequencies, we avoid a *frequentist* interpretation, under which a probability represents the ratio of times that one outcome occurs compared to the total number of outcomes in a series of identical, repeatable, and possibly random trials. We choose the subjective interpretation because the true relative frequencies cannot be determined with any finite number of data samples, and because a subjective interpretation is applicable to a broader class of problems, as it is not limited to repeatable events. Naturally, subjective probabilities should be consistent with available information, including knowledge about observed relative frequencies and the DM's actual beliefs; such probabilities can be considered *rationalist* subjective probabilities [12].

The payoff of a decision

The outcome of a decision problem can be represented by a *payoff function*, $\pi(x,a)$, that depends on both the chosen action a and the realized state of the world x. The payoff function used in our example problem is shown in Figure 1. Note that for a given yield strength and design, the failure cost will either be zero (no failure) or a constant (the cost of the damage, lost productivity, etc. when the pressure vessel fails).

Making an optimal decision

Because of uncertainty in the state of the world x, the DM cannot know the payoff of any action with certainty. We assume that the DM seeks to maximize the *expected payoff*, given by $E_x[\pi(a,x)]$.

The expectation is taken over all states x because that is what the DM is modeling as random. Note that the expectation is taken with respect to the true distribution $p(x)$. In our

$$\pi(a,x) = P_{selling} - C_{material}(a) - C_{failure}(a,x),$$

where:

$P_{selling}$ ≡ selling price = \$200

$C_{material}$ ≡ material cost per volume = \$8500/m³

x ≡ true yield strength of pressure vessel

a ≡ action or design variables (radius, thickness, length)

$$C_{failure}(a,x) \equiv \begin{cases} 0 \text{ if } x \geq \text{maximum stress in pressure vessel} \\ \$1{,}000{,}000 \text{ otherwise} \end{cases}$$

Figure 1: Payoff function

example (and in most real world design scenarios), the DM does not know the true distribution $p(x)$, and must use his or her subjective distribution $\tilde{p}(x)$. The DM thus chooses optimal a such that

$$a^* = \underset{a}{\operatorname{argmax}}\ (E_{\tilde{p}(x)}[\pi(a,x)]) .$$

We deviate slightly from standard notation and write $E_{\tilde{p}(x)}$ to emphasize that the DM maximizes the expectation, as calculated using his or her subjective probability density function $\tilde{p}(x)$.

It is important to distinguish two expected payoffs. The true expected payoff of a design is calculated using the true $p(x)$ that is unknown to the designer. We write the true expected payoff as:

$$\pi_{true} = E_x[\pi(a^*,x)] .$$

The estimated expected payoff according to the designer's subjective distribution is

$$\pi^*_{\tilde{p}(x)} = E_{\tilde{p}(x)}[\pi(a^*,x)] .$$

This payoff $\pi^*_{\tilde{p}(x)}$ will in general differ from the true payoff π_{true} because the designer's estimated distribution $\tilde{p}(x)$ usually will not match the true distribution $p(x)$. Lawrence does not make this distinction in his work, but the distinction is crucial in cases in which the designer has only imprecise information about the random distribution.

Information and information sources

The definition of information varies significantly by subject and application. In this paper, we modify Lawrence's definition [11] and define *information* as any stimulus that changes the recipient's subjective probability distribution $\tilde{p}(x)$ over a well-described set of states.

An *information source* is anything that provides information. This information arrives in the form of a *message y*. In our example, the information source is the yield strength testing process, and a message is the result of a single yield strength test—that is, one observation of material strength. Information economics studies whether it is valuable to pay an information source for a message. Before the message is received, a DM does not know what information that message contains, and therefore the DM does not know exactly how it will change the subjective probability distribution $\tilde{p}(x)$ over the state space. In turn, the DM does not know how it will affect the decision a^* and its payoff. Thus, a DM should apply the principles of information economics to arrive at a formal metric for determining if the benefit of a message outweighs the cost of acquiring it—the *value of information*.

The value of information

We begin by considering two possible decisions: the first decision is made using the current state of information, and the other is made after the receipt of message *y*. In the first case, assume the DM's subjective probability distribution of the states is represented as $\tilde{p}(x)$. These are the *prior* probabilities, and decision a_0^* is the optimal prior decision, given by

$$a_0^* = \underset{a}{\operatorname{argmax}}\ (E_{\tilde{p}(x)}[\pi(a,x)]) . \quad (1)$$

After the message y is received, the DM has an updated *posterior* probability distribution $\tilde{p}(x \mid y)$. The corresponding optimal decision a_y^* is given by

$$a_y^* = \underset{a}{\operatorname{argmax}}\ (E_{\tilde{p}(x|y)}[\pi(a,x)]) . \quad (2)$$

How can we compare these two decisions? If we wait until the true state of the world *x* is revealed, we can calculate the *ex-post value of the message y* for the particular realized state *x* as:

$$\upsilon(\mathrm{y|x}) = \pi(a_y^*,x) - \pi(a_0^*,x) .$$

This represents the amount that the receipt of message y (and the incorporation of its information into the decision) changed the decision maker's payoff, given the particular outcome x of the state.

The term *value* is used throughout this paper in a *marginal* sense, that is, in terms of differences. The ex-post value of a message y is the *marginal payoff* of acquiring that message—the difference between the payoff of the decision with and without the information from message y. This value can be positive, negative, or zero. It will be positive if the message leads the DM to chose action a_y^* that has a higher payoff under realized state *x* than action a_0^*. It will be negative if the message in someway misled the DM into choosing an action a_y^* that has a lower payoff than the prior decision a_0^*. If the message did not change the choice of action, such that a_y^* is the same as a_0^*, then the value is zero.

The previously defined ex-post value is not useful for determining the potential change in payoff of receiving a message because it measures the actual benefit, which can only be known *after* the decision is made and the truth realized. It is common knowledge that a good decision can lead to a bad outcome, especially if a very rare, adverse state of the world is realized—a situation referred to in the vernacular as *bad luck*. Conversely, a bad decision can lead to a good outcome—a case of *good luck*.

Rather than assessing the value of a message for a particular state x, a DM is really interested in the expected value over all the possible states of the world. The *gross value*—where *gross* implies *before* factoring in cost—of the message *y* is defined as the expected difference in the payoff with and without the message, such that:

$$\upsilon(y) = \text{gross value(y)} = E_x[\pi(a_y^*,x) - \pi(a_0^*,x)] . \quad (3)$$

Calculating the true gross value of a message requires the expectation over the true distribution $p(x)$, which the decision-maker does not have. Even this is not the whole story or the complete challenge. Equation (3) is valid for analysis of the value of a particular message y only after it is received. However, when the DM needs to decide whether or not to purchase a message, the content of the message—that is the particular message y from the set $\mathbf{Y}$ of all possible messages—is also unknown. When purchasing a message, it is as if the DM is purchasing a sealed envelope; he or she does not know what is inside until he or she buys and opens the envelope. The DM must therefore consider the value of a particular information source $\mathbf{I}$ that can provide messages from the set $\mathbf{Y}$, instead of a single message.

If the DM had access to the true probability of the messages over the set $\mathbf{Y}$, he or she could calculate the gross value of an information source $\mathbf{I}$:

(4) $$\text{gross value}(\mathbf{I}) = E_y E_x[\pi(a_y^*, x) - \pi(a_0^*, x)].$$

Because the DM does not have access to the true probability distribution of the messages $p(y)$ or of the states $p(x)$, Equation (3) cannot be used directly to estimate the value of an information source. In this paper, we develop a method for bounding the value of information that incorporates the imprecision of the DM's information state.

One final definition that ties our notion of value back to the fundamental concept of cost-benefit analysis in information economics is *net value*. In our example, a message *y* must be purchased at some cost; resources need to be expended in order to acquire more information. Denoting this as $\text{cost}(y)$, the *net value* of a message is defined as

$$\text{net value(y)} = E_x[\pi(a_y^*, x) - \pi(a_0^*, x)] - \text{cost}(y).$$

Similarly the net value of an information source is:

(5) $$\text{net value}(\mathbf{I}) = E_y E_x[\pi(a_y^*, x) - \pi(a_0^*, x)] - \text{cost}(\mathbf{I}),$$

where $\text{cost}(\mathbf{I})$ is the cost of receiving one message y from information source $\mathbf{I}$.

If we revisit a designer's goal of making a reasonable cost-benefit tradeoff during information collection, we can now state the information economic principle that a designer should purchase a message from an information source if the net value of that information is positive. According to Equation (5), this requires the calculation of expectations across the distributions $p(x)$ and $p(y)$, which in general are not known to a designer. We will return to this problem after illustrating the simpler case of known probabilities.

Example with known probabilities

In this section, we present an example to illustrate the calculation of value of information in the case of known probabilities. We later extend this example to the case of unknown probabilities. While the information used in this example is not available to a DM, it is useful for illustration of the basic method, shown in Figure 2.

We assume that there is an omniscient supervisor overseeing the experiment. This supervisor knows the true distributions and can perform the actions shown in the gray boxes. These actions are normally not available to the DM. In this method, the DM begins with the observed set of samples $\Sigma = \{x_i\}_{i=1}^n$. The DM first uses this set of samples to construct a best-fit distribution $\tilde{p}(x)$, and then to choose an optimal design a_0^*, as shown on the left side of the figure. The DM then receives a hypothetical additional sample y_j from the supervisor. The DM constructs a new best fit distribution $\tilde{p}(x \mid y_j)$ and makes a new decision $a_{y_j}^*$. The difference in expected payoffs of the two decisions is then calculated by the supervisor to determine the true expected gross value $\upsilon(y_j)$ of the particular message y_j. This process of calculating the value of an additional sample is repeated over many y_j to calculate the average value of the next sample for a particular starting set of samples, which we will denote as $V(n+1 \mid \Sigma)$.

Recall that the net value of the next piece of information depends on the prior decision a_0^*, which in turn is dependent on the existing data samples. For example, the net value of the purchasing an 11th sample from the information source will depend on the first 10 samples. If the initial 10 samples just happen—by chance—to yield very good estimates of the distribution parameters, then the net value of the 11th sample will be small, but if they yield bad estimates of the distribution

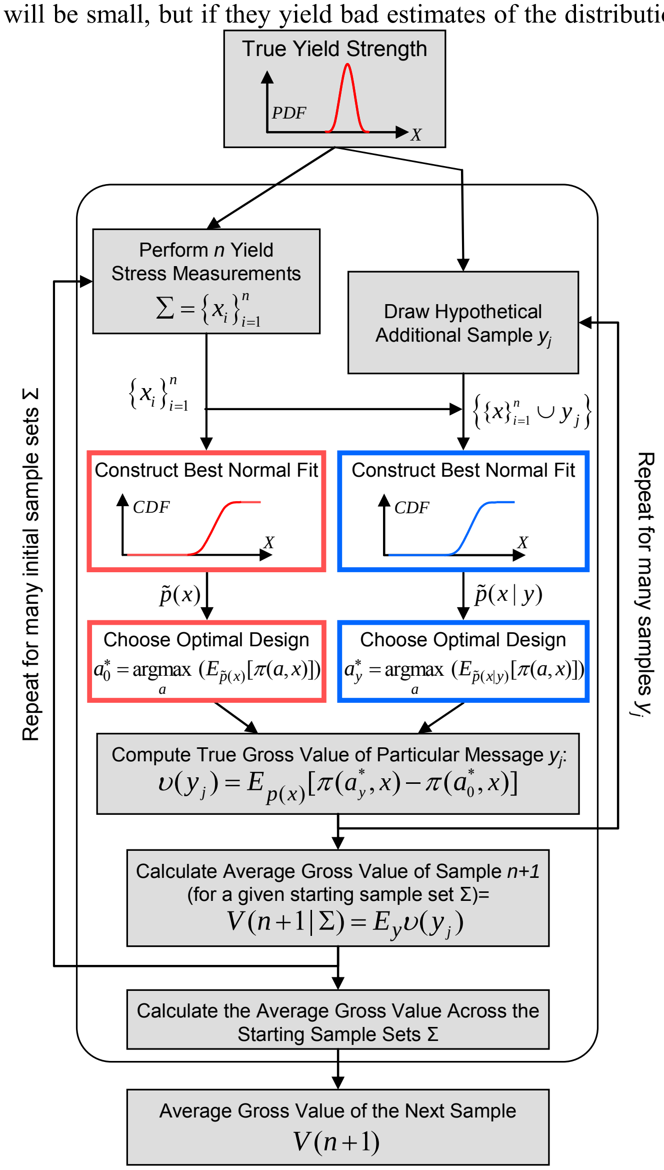

Figure 2: Calculating the value of information with known probabilities

parameters, then the net value of the 11th sample could be large. Consequently, the next step is to repeat the experiment over many initial sample sets Σ, which gives the average gross value of the next sample, denoted $V(n+1)$.

The final step is to repeat the process for different initial samples set sizes. By repeating the calculation over many initial sample sizes, we can construct a curve of the average net value of an additional sample at different sample sizes, as shown in Figure 3. This figure can be interpreted with the following example. At a prior sample size n=32, the average net value of an additional sample (the 33rd sample in this case), is about $2. The net value of the 52nd sample, starting from 51 samples, is negative, but the net value of the 51st sample is positive. This means that the 52nd sample is the first sample whose average net value is negative. This graph suggests that stopping at 51 samples is on average the optimal. Note that this conclusion is drawn using the true $p(y)$ and $p(x)$, which are not available to the DM.

The results can also be interpreted by considering the net expected payoff, the payoff of design that would have been realized if no additional information were collected, less the cost of the already collected n samples:

(6) $$\text{net expected payoff} = E_{p(x)}[\pi(a_0^*, x)] - n * \text{cost}(y)$$

The results are shown for different sample sizes in Figure 4. Again, because the actual observed samples affect the payoff, the payoff of the design is averaged over many initial sample sets. The relationship between this result and the net value of additional samples should be clear. The maximum net expected payoff occurs at the same sample size that the net value of an additional sample first becomes negative. Recalling that the net value is defined in a marginal sense, moving from 51 samples to 52 samples means a decrease in total payoff of the decision, as revealed in both plots.

In the preceding analysis, it appears simple to determine the optimal number of samples to collect. However, this simplicity is due to the assumption that the DM has precise knowledge of the true distributions $p(y)$ and $p(x)$. In our example problem, these two distributions are unknown and identical—they both describe the unknown true material strength. The characterization of $p(x)$ is the DM's indirect goal for data collection—the DM wants to characterize $p(x)$ well enough that the design based on the estimated $\tilde{p}(x)$ is acceptable.

In order to evaluate the value of information during the actual design process, the DM needs a method by which he or she can estimate the net value of additional data samples when $p(y)$ and $p(x)$ are unknown. We propose a method that uses imprecise probabilities to calculate an interval of net value for an additional sample.

What performance characteristics should we expect or demand of this method? Insight can be gained by examining the distribution of the net payoffs about the expected value curve of Figure 4. Box plots for sample sizes of 50, 100, and 150 are shown in Figure 5. The plots are constructed with the whiskers at the 5% and 95% quantiles, and the boxes from 25% to 75%. The box plots clearly reveal that both the variance and the chance of a very bad result decrease as the sample size increases, though at a penalty in expected net value as shown in Figure 4. The behavior shown in Figure 5 suggests that a reasonable estimation of the optimal number of samples will often be well beyond 51 samples, because by stopping at 51, a DM faces a very large downside risk, due to the long tail of the distribution. Consideration of this distribution and tradeoff is important when constructing a method for determining the value of additional samples because, in practice, an engineer will not have the information in Figure 4 or Figure 5 available for decision-making. We return to this problem after introducing the concept of imprecise probabilities.

IMPRECISE PROBABILITIES

Imprecise probabilities have been formalized by Walley [12], and the value of using imprecise probabilities in certain engineering design decisions has been demonstrated previously [22]. We extend this work to estimate the value of information through the application of information economics and imprecise probabilities.

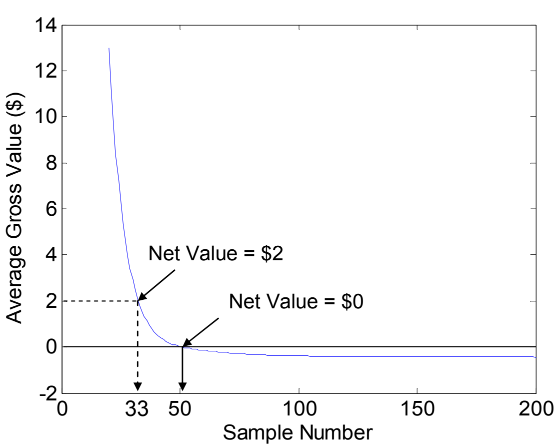

Figure 3: Net gain in payoff per sample

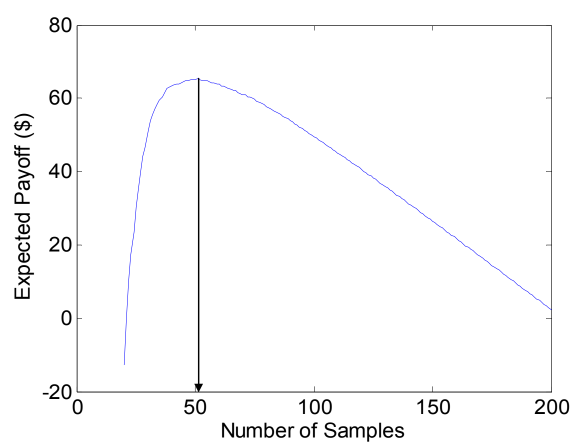

Figure 4: Net expected payoff of the design

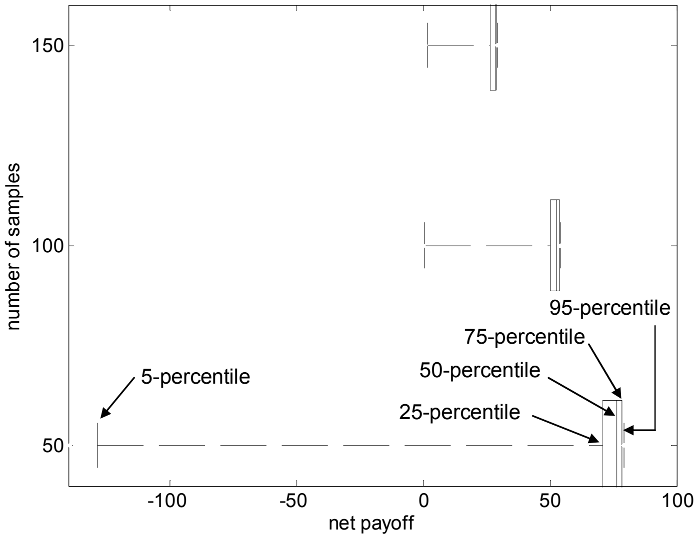

Figure 5: Box plots for various sample sizes

In the context of this paper, one can think of imprecise probabilities as intervals of probabilities. As noted previously, under a subjective interpretation, a probability represents an individual's willingness to enter a bet. Every bet has a price associated with it, and one can either buy or sell a bet at that price. The use of precise probabilities presumes that a DM can determine exactly the price at which he or she is indifferent between buying and selling the bet, the DM's so-called *fair price* [28]. The use of imprecise probabilities allows for a range of prices at which a DM would neither buy nor sell the bet, because he or she is not sure how betting at these prices will affect his or her expected payoff.

Imprecise probabilities can in theory be reduced to precise probabilities by collecting infinite evidence and expending infinite effort to elicit the DM's beliefs. In this example problem, we are explicitly assuming a finite amount of evidence (in the form of data samples), so precise probabilities are unattainable. We therefore use imprecise probabilities to capture the DM's current state of information, which in our example corresponds to his or her probability assessments.

We use a probability-box [29], or p-box, to represent imprecise cumulative probability distributions. A p-box incorporates both imprecision and probabilistic characterizations by expressing *interval* bounds on the cumulative *probability* distribution function (CDF) for a random variable. More formally, the bounds on a p-box, such as shown in Figure 6(a), are given by two CDFs (F_1 and F_2) that enclose a set of CDFs that are consistent with the current state of information and the DM's beliefs. The p-box shown in Figure 6(a) is for a random variable Z with known variance $\sigma^2 = 1$ and mean bounded by the interval μ=[0,1]. Thus extending the notation of probability, we can write $Z \sim N([0,1],1)$.

The true CDF is unknown, and any of the infinite number of normal CDFs with $\sigma^2 = 1$ inside the p-box could be the true one, such as those shown in Figure 6(b).

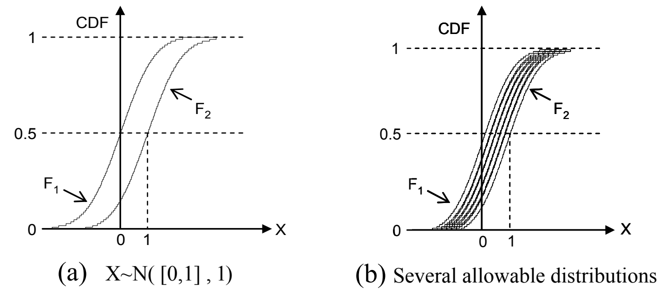

Figure 6: Example P-box and distributions

While there are several ways to construct p-boxes [30, 31], we choose a practical method based on traditional confidence intervals on the mean and variance, as described by Equation (A2) and Equation (A4) respectively in the Appendix. By choosing a particular confidence level for the mean and variance intervals, a DM is essentially stating that he or she is comfortable assuming that the true distribution lies entirely in the resultant p-box. This assumption is similar to accepting the p-box as a model of the truth. This distinction becomes important in our method for estimating the value of information, as explained in the following section.

ESTIMATING THE VALUE OF INFORMATION

In this section, we explain the method that we have created to bound the gross value of the next information message. We start by describing how design decisions are made. We then motivate the use of imprecise probabilities, describe our methods of estimating the value of information, and present a computational experiment that illustrates the results of our method.

Design decision policy

The DM optimizes the expected payoff of a design using his or her estimate $\tilde{p}(x)$ of the true distribution $p(x)$. This best guess is derived by assuming that the material strength is normally distributed and using the sample mean and variance of the observed samples as estimates of the true mean and variance. In previous work, we presented a decision policy that uses imprecise probabilities [22]. In this paper, we chose a decision policy based on precise probabilities in order to isolate the effect of using imprecise probabilities for estimating the value of information.

Motivation for imprecise probabilities

One motivation for using imprecise probabilities is that the use of precise probabilities does not enable useful estimates of value. The necessity of an alternative to precise probabilities is illustrated in the following example. Assume that the DM represents his or her state of information $\tilde{p}(x)$ precisely. Using this information, the DM chooses an optimal design a_0^* according to Equation (1), using $\tilde{p}(x)$ when evaluating the expectation.

Now assume that the DM acquired an additional data sample y. With this information, the DM can create a new subjective distribution $\tilde{p}(x \mid y)$, where in general $\tilde{p}(x \mid y) \neq \tilde{p}(x)$. The DM would then choose an optimal design a_y^* according to Equation (2), using $\tilde{p}(x \mid y)$ when evaluating the expectation.

If the DM wanted to calculate the gross value of this message y, he or she would use Equation (3), repeated here for convenience:

$$(7) \qquad \upsilon(y) = \text{gross value}(y) = E_x[\pi(a_y^*, x) - \pi(a_0^*, x)] .$$

Ideally the expectation E_x would be taken over the true $p(x)$, but this distribution is unknown. The DM's two best options for approximating $p(x)$ are $\tilde{p}(x)$ and $\tilde{p}(x \mid y)$.

If the DM uses $\tilde{p}(x)$ as the best estimate of $p(x)$, we can adopt our notation from earlier and rewrite Equation (7) as:

$$(8) \qquad \upsilon(y) = E_{\tilde{p}(x)}[\pi(a_y^*, x) - \pi(a_0^*, x)] .$$

This can be rewritten by distributing the expectation as:

$$\upsilon(y) = E_{\tilde{p}(x)}[\pi(a_y^*, x)] - E_{\tilde{p}(x)}[\pi(a_0^*, x)] .$$

According to Equation (1), the design decision a_0^* maximizes $E_{\tilde{p}(x)}[\pi(a, x)]$, thus

$$(9) \qquad E_{\tilde{p}(x)}[\pi(a_y^*, x)] \leq E_{\tilde{p}(x)}[\pi(a_0^*, x)] .$$

This means that the gross value of message y is always estimated to be zero or negative, no matter how much information is available. Yet intuitively, the gross value of additional information should often be positive—acquiring information should improve the DM's ability to make a good decision on average.

If the DM instead used the posterior distribution $\tilde{p}(x \mid y)$, we can rewrite Equation (7) as:

$$\upsilon(y) = E_{\tilde{p}(x|y)}[\pi(a_y^*, x) - \pi(a_0^*, x)] .$$

Expanding the expectation we find

$$\upsilon(y) = E_{\tilde{p}(x|y)}[\pi(a_y^*, x)] - E_{\tilde{p}(x|y)}[\pi(a_0^*, x)] .$$

According to design Equation (2), design decision a_y^* maximizes $E_{\tilde{p}(x|y)}[\pi(a, x)]$, thus we find:

$$E_{\tilde{p}(x|y)}[\pi(a_y^*, x)] \geq E_{\tilde{p}(x|y)}[\pi(a_0^*, x)] .$$

In this case, the gross value is always calculated to be positive, which is also unreasonable; there will always be "unlucky" samples, or messages, that lead to a worse design. Another objection to using the precise $\tilde{p}(x \mid y)$ is that it has no use in decision making, because $\tilde{p}(x \mid y)$ is only available after the information message y is collected.

This exercise illustrates that an information collection policy based upon an assumption of precisely characterized distributions is not useful. The principles of information economics cannot be applied meaningfully while using precise probabilities, but they can be implemented using a method based on imprecise probabilities that provides useful bounds on the value of information, as described in the next section.

Methodology for bounding the value of information

An overview of our method is shown in Figure 7. The DM begins with the actually observed set of data samples $\Sigma = \{x_i\}_{i=1}^n$. The DM first uses this sample to construct a best-fit normal distribution and choose an optimal design a_0^*, the

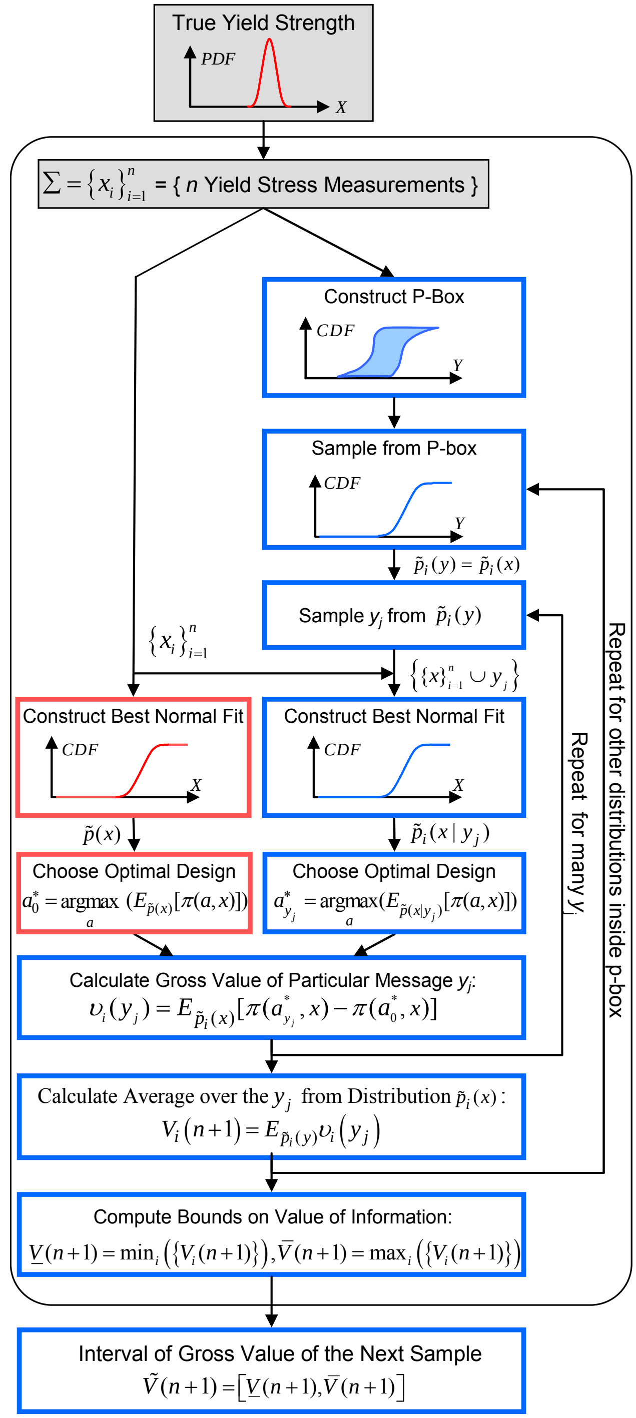

Figure 7: Overview of method using imprecise probabilities to bound the value of information

left side of the figure. The DM then uses the observed samples to construct a p-box by the principles mentioned previously (confidence intervals on the mean and variance of the normal distribution).

Recall that by assumption, this model of information—the p-box—contains the CDF of the true distribution $p(x)$. The DM discretizes the p-box, as described below, and selects one normal distribution from the p-box to represent both $\tilde{p}_i(x)$ and $\tilde{p}_i(y)$. In our example problem, these two distributions are identical since they both describe the unknown true material strength. This selected distribution is used to estimate the gross value of collecting an additional piece of information. If the DM repeated this for *every* normal distribution inside the p-box, one of the calculated values will be the true gross value of the next piece of information. Clearly the DM cannot try every distribution, so we propose the following procedure.

The DM can partition the p-box into a finite, representative set of distributions. This is done by discretizing the confidence intervals on the mean and variance The DM selects a finite set of distributions from the infinitely many possible distributions by pairing all the means and variance combinations, resulting in a set of distributions such as shown in Figure 8. Future work will explore more efficient methods for this partitioning, such as concepts from design of experiments, direct manipulation of the p-boxes, or random sampling across the confidence intervals, but this method suffices for illustration of the concept.

The DM proceeds by selecting one distribution, say $\tilde{p}_i(x)$, from this finite set and assuming that this distribution is the true distribution ($\tilde{p}_i(x) = p(x)$). The DM then calculate the gross value of taking the $(n+1)^{th}$ sample, denoted $V_i(n+1)$, via a Monte Carlo simulation, as follows.

Given the assumed message distribution $\tilde{p}_i(y)$, the DM can draw hypothetical next samples y_j from this distribution. Each of these samples is used, along with the actually observed samples $\{x_i\}_{i=1}^{n}$, to estimate a new posterior distribution $\tilde{p}_i(x \mid y_j)$. The DM uses this distribution to choose optimal design $a^*_{y_j}$, for the given distribution and hypothetical sample. The DM then evaluates the expected payoff of this design using the assumed $\tilde{p}_i(x)$, and calculates the gross value $\upsilon_i(y_j)$ of that particular y_j. The DM repeats this for many different y_j's drawn from $\tilde{p}_i(y) = \tilde{p}_i(x)$ and calculates the average, or expected, gross value of the next message $V_i(n+1)$ with distribution $\tilde{p}_i(x)$ assumed to be the true distribution. Finally, the DM repeats this process for all $\tilde{p}_i(x)$ in the chosen set (from the p-box). This results in a set of gross values $\{V_i(n+1)\}$.

Recall that if the p-box had been infinitely sampled, then one of the values $V_i(n+1)$ in this set would be the true gross value of the $(n+1)^{th}$ sample, given the previously observed n samples. The set $\{V_i(n+1)\}$ would then form an interval $V(n+1) = [\underline{V}(n+1), \overline{V}(n+1)]$. In our method, the p-box is only finitely sampled, so the set of values $\{V_i(n+1)\}$ only gives an approximate interval, $\tilde{V}(n+1) = [\underline{V}(n+1), \overline{V}(n+1)]$, with the

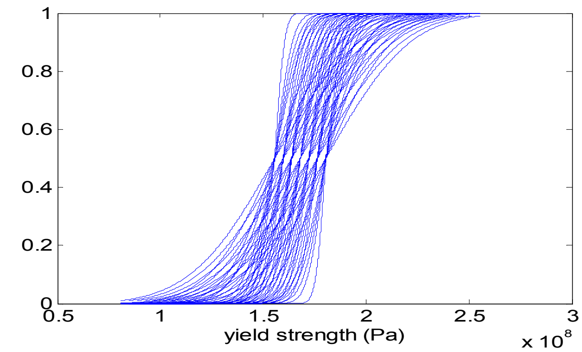

Figure 8: Various distributions in the P-box

lower and upper-bounds defined as $\underline{V}(n+1) = \min_i(\{V_i(n+1)\})$ and $\overline{V}(n+1) = \max_i(\{V_i(n+1)\})$. Clearly the accuracy of this estimated interval improves as the density of sampling from the p-box increases.

COMPUTATIONAL EXPERIMENT

We will now apply the principles of information economics to the design of the pressure vessel as explained in the example problem section. The experiment proceeds according to the method shown in Figure 7. The method is repeated for many sample sizes n. This generates intervals on the gross value for one particular sample trace, that is, one particular sequence of random samples $\{x_i\}$. This experiment is then repeated many times to generate multiple sample traces.

RESULTS

Using the method and experiment described above, we can bound the gross value of the next piece of information. A graph of these bounds for a particular sequences of samples—a particular sample *trace*—is shown in Figure 9. A trace represents the bounds on gross value of the n^{th} sample, given a particular set of $n-1$ previously observed samples $\{x_i\}_{i=1}^{n-1}$. Figure 10 shows the upper-bound, lower-bound, and midpoint for two additional traces in the vicinity of their crossing of the cost line—the zero net value point. The curves in the two figures reveal several interesting characteristics, as discussed in the following sections.

Small sample sizes yield large value intervals

Figure 9 shows that the potential value for very small sample sizes covers a very large range, from largely negative to largely positive and skewed towards the positive side. The bounds on the gross value of the n^{th} sample are the lower-bound and upper-bound at that point. For example, the gross value of the 16th sample is in the interval [-10, 940] for the trace shown in Figure 9. We will discuss decision policies for resolving this interval in the discussion section, but for now assume one extreme—the DM stops collecting data when the upper-bound on the gross value is less than the cost—that is, when the upper-bound on the net value is negative. This is a so-called Γ-*maximax* policy [30]. At the accepted confidence level, the true value is assumed to lie in the interval, so this represents the

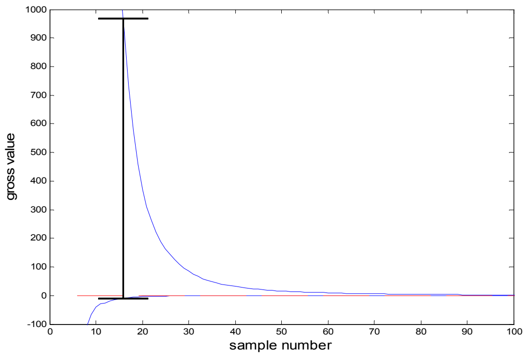

Figure 9: Example high-level behavior of gross value

point at which the true net value cannot exceed zero (since the upper-bound is negative), and no rational DM would take an additional sample.

The bounds on value are not monotonic

In a general sense, it is reasonable to expect that the value of additional samples would decrease as n increases. However, each trace represents one sequence of actually observed samples. Thus, the gross value of the i^{th} sample is estimated based on the first $i-1$ samples. Once the i^{th} sample is collected, the value of the $(i+1)^{th}$ sample is calculated using all i acquired samples. If the actually acquired i^{th} sample is really "lucky" or "unlucky", the gross value of the next sample can change significantly, potentially yielding non-monotonic bounds. An example of such non-monotonicity is labeled in Figure 10. Non-monotonicity can result in multiple cost line crossings, but these crossings were never observed to be more than a few sample sizes apart. Because the bounds are already estimates, a deviation of a few samples is not significant.

The lower-bound is always non-positive

It is worth noting that the lower-bound on the interval will always be non-positive (i.e. given the available information, it is always possible that the gross value of the next piece of information will be less than or equal to zero). This happens because the best fit distribution $\tilde{p}(x)$ on which we base our design decision is always contained in our set of distribution samples from the p-box—it is a candidate for the truth in our method. This means that during the calculation of the interval on gross value, this distribution will be considered as the truth at some point, yielding the situation described in Equation (8) —if the DM's estimate already is the true distribution, which is possible though rare, then no information can make the estimate any better. It will in fact make the DM worse off.

Examining the net value

The next point to note is the relationship between the gross value and cost. In practice, there is a relationship between the number of pressure vessels being designed and the cost, because the cost of information collection is amortized over all the pressure vessels. In this example, we are assuming that each yield strength test on a material sample costs $0.50 per pressure vessel, and we proceed to discuss the design of one pressure vessel. Other cost functions could be used without adding significant complexity. With the cost fixed at $0.50, an experiment following Trace A and using the upper-bound decision rule will stop with the 147th sample, because the upper-bound on the gross value of the 148th sample is less than the cost. By Equation (5), this implies a negative net value.

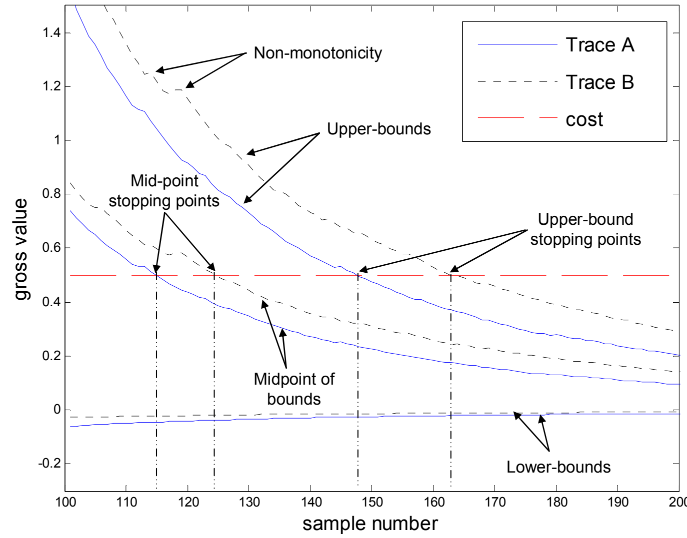

Figure 10: Two example traces of gross value

The same logic can be applied to Trace B. In this case, a DM would collect 162 samples. At this point, the upper-bound on the gross value of the 163rd sample is less than the cost, so the net value is negative.

In this section we have presented two representative results. Our overall results consist of many sample traces that can be analyzed in the same way as the examples above. In the next section, we discuss more general aspects of our method and results.

DISCUSSION

With basic results presented in the previous section, we now move to more general discussion. We examine the realized payoffs based on the described method and discuss alternative decision policies for resolving ambiguity.

Comparison of realized payoffs

Using the true material strength distribution $p(x)$ —which is *not* known by the designers—an omniscient supervisor can evaluate the actual expected payoff of the designs that would be created using the samples in the traces. The results for Trace A from Figure 10 are shown in Figure 11. Each point represents the true expected net payoff (y-axis) of a design chosen based on the current number of samples (x-axis) and found using Equation (6). Because a DM would never create all of these designs and does not have access to $p(x)$, this is a hypothetical exercise that only the omniscient supervisor can perform using the true distribution $p(x)$.

The results in Figure 11 indicate that, given the actual observed sequence of samples, the DM would have been best-off stopping earlier (at 5 samples) than our method shows (at 147 samples). In this example, the DM loses about 60% of the payoff by collecting the additional 142 samples. Is this a result of the stopping policy? One alternative policy would be to use the midpoint of the bounds. Using this stopping rule, we find from Figure 10 that we would collect 114 samples. This still results in a loss of 50% payoff from the optimal. Is such a loss in payoff justified?

In the discussion surrounding the distribution of payoffs and Figure 5, we conclude the DM may wish to go beyond the average "optimal" stopping point due to the large downside risk of stopping too soon. The actual expected net payoffs for trace B from Figure 10 are shown in Figure 12. In this example trace, it turns out that given the actually observed samples, it would have been much worse to stop after 80 samples as compared to 100. According to Figure 10, the mid-point stopping rule would be at 124 samples for this trace. While this is still about 50% below the optimal, it yields a significantly better result than a policy that would have stopped at 80 samples. Before discussing stopping policies in more detail in the next subsection, we present one last trace in Figure 13**.**

For the trace shown in Figure 13, the optimal stopping point would have been at 110 samples. In this case, the mid-point stopping rule solution of 130 samples is relatively close, though still resulting in some payoff loss . What causes the optimal stopping point to be so high in this case? One contributing factor is that the first five actual samples were 192 MPa, 200 MPa, 194 MPa, 197 MPa, and 181 MPa. These are all above the true mean of 180 MPa. This initial "unlucky" bias leads to a severe over-estimate of the material strength, which in turn leads to a severe under-design of the pressure vessel. Consequently, the pressure vessel fails much more often than expected, leading to a significantly higher average failure cost. This example indicates how sensitive the design can be to sample data, and why a large number of samples may be needed to reach a stable result.

Before reaching a conclusion on the effectiveness of this method of bounding the value of information, we must emphasize that an actual designer would not have the actual expected payoff curves available. Therefore, a designer does not know if he or she is in an example similar to that of Figure 11, Figure 12, Figure 13, or something else altogether. A conservative policy therefore leads the designer to keep taking samples until it is reasonably assured that there is no chance of a large negative payoff; that is, samples are taken until the downside risk is acceptable.

FUTURE WORK

This paper lays a strong foundation for applying information economics to engineering design decisions involving unknown probability distributions, but there is still significant room for improvement and additional exploration.

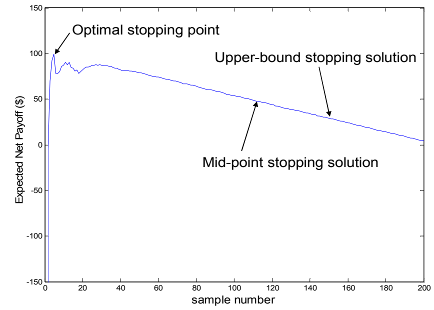

Figure 11: Actual expected net payoffs for Trace A

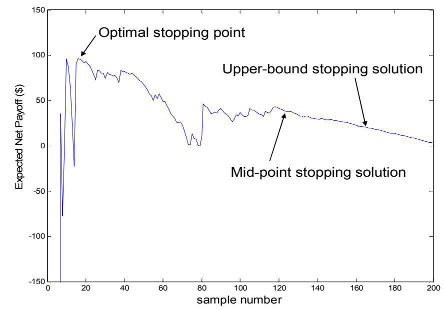

Figure 12: Actual expected net payoffs for Trace B

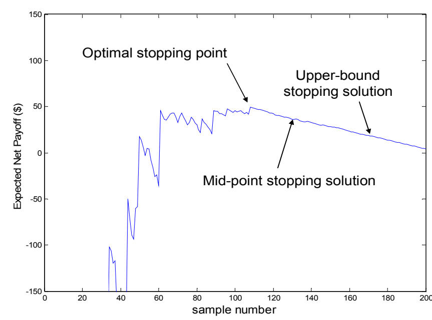

Figure 13: Actual expected net payoffs, additional trace

Decision policies for gathering information

In this paper, we have presented a method for bounding the value of additional data samples. These bounds need to be resolved according to some policy in order to make a decision. Given just the bounds, *any* policy that selects a point between the bounds is rational, because the true value is only known to be somewhere between them. However, given meta-data about the design problem and information at hand, one policy may actually perform better on average than another.

We have so far discussed two policies—an upper-bound Γ-maximax policy and a midpoint policy. Using the Γ-maximax policy, a DM compares the upper-bound on gross value to cost, thus deciding to stop collecting data when the upper-bound on the net value becomes negative. Because the DM stops collecting information only when he or she is absolutely sure that the value of the next piece of information is less than its cost, this extreme decision policy will often be overly conservative. Since it is rare that the true gross value will be the upper-bound, the DM almost always takes too many samples,.

An opposite extreme is to use the lower-bound of the interval of gross value, a so-called Γ-*maximin* policy [30, 31]. As noted earlier, the lower-bound is always below the cost line. Consequently, the net value of all additional samples is negative, and the DM never collects additional information. This extreme policy is defendable according to the basics of the method, but it seems to be hard to defend in engineering design applications. There may be some applications (such as deterministic data) in which stopping after some initial number of samples is optimal, but in many cases, including our design example, a decision made without collecting additional information will be a bad one. Therefore, in engineering design, this decision policy is clearly unacceptable.

Of the infinite possible policies between these two extreme policies, we have already illustrated one based on the midpoint of the bounds. A midpoint policy is a special case of using the Hurwicz criterion [32], a policy that picks a point somewhere between the upper and lower-bounds according to an index of pessimism and optimism. In our example problem, we can use the supervisor's knowledge to determine that a midpoint policy is clearly better than either a Γ-maximax (upper-bound) or Γ-maximin (lower-bound) decision policy. However, this will not necessarily be true in general. Our example problem involved a very skewed payoff function—the payoff when the vessel fails is very negative (minus $1 million), yet the payoff when it succeeds is only slightly positive (the selling price of $200).

In this section we have mentioned three possible decision policies for using the value of information when collecting sequential data. Other policies for reducing the set of possible decisions include enforcing conditions of interval dominance, E-admissibility, or maximality [30]. However, these policies will not always result in a unique decision; they are merely means for managing the candidate solution space.

Clearly the selection of a stopping rule is not a trivial matter. What we seek is a decision policy that *usually* results in good designs. In order to determine the most valuable policy, an experiment similar to [22] could be conducted to determine which decision policies yield better decisions for this example problem. In such an experiment, an omniscient supervisor with perfect knowledge judges the net payoffs of the decisions made using the different decision policies. Based on these payoffs, we can judge which decision policies do better and for which conditions they do better. Such information would help a designer select a policy suitable for a given decision problem.

In this paper, we loosely compare the Γ-maximax and midpoint policies, and find that the midpoint policy almost always performs better. If this type of analysis can be extended to other types of design problems, the method for bounding the value of information presented in this paper would have an even larger impact of engineering design. Whether such meta-data can be captured is an open research question.

Design decision policies

As mentioned in the Estimating the Value of Information section, we choose a design decision policy based on precise probabilities so that we can focus on the effectiveness of using imprecise probabilities to estimate the value of information. Previous work has shown the value of incorporating the DM's imprecise probabilities directly into the design decision [22]. Using imprecise probabilities for both the design decision and the prediction of the value of information would complicate the application of our method for determining the value of future information and confound the effects of the two policies. However, it appears to be a more realistic representation of a typical design problem and could possibly lead to additional insightful results.

P-box construction

For this experiment, we construct the p-box using 95% confidence intervals on our mean and variance estimates as explained in the Appendix. We then proceed to assume that this p-box contains the truth. Further experiments should be conducted to evaluate how the design method performs: how often is this valid, and what confidence levels are appropriate? It is impractical to use 100% confidence intervals because these will be infinite. It may be that alternative methods for constructing p-boxes are needed [33]. Similar to the decision policy work noted above, it may be that different p-box construction policies work better for different design problems. Any such relationships need to be explored before this method can be put to practical use.

Computational cost

The method presented in this paper requires a double-loop Monte Carlo simulation for every sample size. For this experiment, the calculation of bounds on the value of the next sample takes about 5 minutes with a high number of p-box and message samples, though results for runs as short as 30 seconds appear nearly as good . These times are on a single 2.6 GHz Pentium 4 processor system with 512 MB of RAM. This computation time is perfectly reasonable, but the computational complexity can be expected to increase substantially for more complicated design problems. Future work investigating how to compute and simulate directly with p-boxes in order to avoid multiple-loop Monte Carlo simulations would be useful.

Design problems

The pressure vessel example used in this paper is deliberately simple, in order to illustrate the concept. To assess the general applicability of our method, the design example needs to be varied. Specifically, several variations of the design problem could be investigated. First, the true distribution could be varied. Second, a situation in which there are multiple uncertain parameters could be explored. Third, the payoff function could be varied to consider different levels of risk-preference. All of these steps would lead to more general conclusions about the applicability of our method.

SUMMARY

In this paper we have introduced the principles of information economics and related them to engineering design problems in which statistical distributions are not fully known. The main contribution of this work is the formulation of a method by which the bounds on the value of information can be calculated by a designer during the information collection process through the use of imprecise probabilities. An open question is how to make a decision given these bounds on value. We have explored several example situations and decision policies, described the limitations, and identified areas for future work.

ACKNOWLEDGEMENTS

Jason Aughenbaugh and Jay Ling are each supported by National Science Foundation Graduate Research Fellowships. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. This research is further sponsored in part by NASA Ames Research Center under cooperative agreement NNA04CK40A. Additional support is provided by the G.W. Woodruff School of Mechanical Engineering at Georgia Tech.

APPENDIX: CONSTRUCTING A P-BOX FROM LIMITED DATA

PBA is a recently developed method, and most publications have focused on introducing the concept of PBA, explaining its theoretical value, and introducing methods for computing with p-boxes [34-36]. Only more recently have researchers addressed the construction of p-boxes from sample data [33].

While several approaches exist [33], we choose a pragmatic approach based on standard statistical confidence intervals. In this example, the designers assume the true yield strength (X) is normally distributed, but with unknown mean and variance:

$$X \sim Normal(\mu, \sigma^2) .$$

One basis of reference for the true but unknown μ and σ^2 are the minimum variance unbiased point estimates:

$$\hat{\mu} = \bar{x} = \tfrac{1}{n}\sum_{i=1}^{n} x_i \text{ and } \hat{\sigma}^2 = s^2 = \tfrac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2 ,$$

where the x_i's are the sample observations, and n is the sample size. These quantities are called respectively the sample mean and sample variance and are commonly used in pure probabilistic approaches. In order to construct a p-box, we broaden these point estimates to confidence intervals. In this experiment, a 95% confidence level is used.

Confidence interval for the mean:

Since x_1, x_2, …, x_n is a random sample from a normal distribution with unknown mean μ and unknown variance σ^2, the sampling distribution of the statistic

$$\text{(A1)} \qquad t = \frac{\hat{\mu} - \mu}{s/\sqrt{n}}$$

is the t distribution with $n-1$ degrees of freedom. Letting $t_{\alpha/2,n-1}$ be the upper $\alpha/2$ percentage point of the t distribution with $n-1$ degrees of freedom, it can be shown that

$$P\left\{-t_{\alpha/2,n-1} \le t \le t_{\alpha/2,n-1}\right\} = 1-\alpha .$$

Substituting for t in Equation (A1) and solving for the mean μ, we arrive at a $(1-\alpha)100\%$ confidence interval for the mean

$$\text{(A2)} \qquad [\underline{\mu}, \overline{\mu}] = [\hat{\mu} - t_{\alpha/2,n-1}\, s/\sqrt{n} \, , \hat{\mu} + t_{\alpha/2,n-1}\, s/\sqrt{n}] .$$

Confidence interval for the variance:

Since x_1, x_2, …, x_n is a random sample from a normal distribution with unknown mean μ and unknown variance σ^2, it can be shown that the sampling distribution of

$$\text{(A3)} \qquad \chi^2 = \frac{(n-1)s^2}{\sigma^2}$$

is chi-square with $n-1$ degrees of freedom, where n is the sample size and s^2 is the sample variance [37]. To develop the confidence interval, we note that

$$P\left\{\chi^2_{1-\alpha/2,n-1} \le \chi^2 \le \chi^2_{\alpha/2,n-1}\right\} = 1-\alpha .$$

Substituting for χ^2 in Equation (A3) and solving for the variance σ^2, we arrive at a $(1-\alpha)100\%$ confidence interval for the variance

$$\text{(A4)} \qquad \left[\underline{\sigma}^2, \overline{\sigma}^2\right] = \left[\frac{(n-1)s^2}{\chi^2_{\alpha/2,n-1}}, \frac{(n-1)s^2}{\chi^2_{1-\alpha/2,n-1}}\right].$$

A table of t and χ^2 values is found in most probability and statistic books, such as [37].

REFERENCES

[1] Pahl, G., and Beitz, W., 1996, *Engineering Design: A Systematic Approach*, Springer Publishing, London.

[2] Hazelrigg, G. A., 1996, *Systems Engineering: An Approach to Information-Based Design*, Prentice Hall, Englewood Cliffs, NJ.

[3] Hazelrigg, G. A., 1998, "A Framework for Decision-Based Design," *Journal of Mechanical Design*, **120**(4), pp. 653-658.

[4] Marston, M., Allen, J. K., and Mistree, F., 2000, "The Decision Support Problem Technique: Integrating Descriptive and Normative Approaches in Decision Based Design," *Engineering Valuation and Cost Analysis*, **3**, pp. 107-129.

[5] Marschak, J., 1974, *Economic Information, Decision, and Prediction: Selected Essays*, D. Reidel Publishing Company, Boston.

[6] Gupta, S. K., and Xu, C., 2002, "Estimating the Optimal Number of Alternatives to Be Explored in Large Design Spaces: A Step Towards Incorporation

Decision Making Cost in Design Decision Models," *ASME DETC Computers and Information in Engineering Conference*, Montreal, Canada, pp. DETC2002/CIE-34491.
[7] Radhakrishnan, R., and McAdams, D. A., 2005, "A Methodology for Model Selection in Engineering Design," *Journal of Mechanical Design*, **127**(May), pp. 378-387.
[8] Bradley, S. R., and Agogino, A. M., 1994, "An Intelligent Real Time Design Methodology for Component Selection: An Approach to Managing Uncertainty," *Journal of Mechanical Design*, **116**(December), pp. 980-988.
[9] Law, A. M., and Kelton, D. W., 2000, *Simulation Modeling and Analysis*, McGraw-Hill, St. Louis.
[10] Layard, R., and Glaister, S., eds, 1994, *Cost-Benefit Analysis (2nd Edition)*, Cambridge University Press, Cambridge.
[11] Lawrence, D. B., 1999, *The Economic Value of Information*, Springer-Verlag, New York.
[12] Walley, P., 1991, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, New York.
[13] Lieberman, M., and Hall, R., 2000, *Introduction to Economics*, South-Western College Publishing.
[14] Rubin, M. R., 1983, *Information Economics and Policy in the U.S.*, Libraries Unltd., Littleton, CO, USA.
[15] Strassmann, P. A., 1999, *Information Productivity*, Information Economics Press.
[16] Strassmann, P. A., 2004, "The Foundations of Information Economics - Part I," *Information Economics Journal*, **March 2004**.
[17] Ciocoiu, M., Gruninger, M., and Nau, D., 2001, "Ontologies for Integrating Engineering Applications," *ASME Journal of Computing and Information Science in Engineering*, **1**(1).
[18] Urban, S. D., Dietrich, S. W., Saxena, A., and Sundermier, A., 2001, "Interconnection of Distributed Components: Current Middleware Solutions," *ASME Journal of Computing and Information Science in Engineering*, **1**(1).
[19] Jayaram, S., Vance, J., Gadh, R., Jayaraman, J., and Srinivasan, H., 2001, "Assessment of VR Technology and Its Applications to Engineering Problems," *ASME Journal of Computing and Information Science in Engineering*, **1**(1).
[20] Rangan, R. M., and Chadha, B., 2001, "Engineering Information Management to Support Enterprise Business Processes," *ASME Journal of Computing and Information Science in Engineering*, **1**(1), pp. 32-40.
[21] Szykman, S., Sriram, R., and Regli, W. C., 2001, "The Role of Knowledge in Next-Generation Product Development Systems," *ASME Journal of Computing and Information Science in Engineering*, **1**(1), pp. 3-11.
[22] Aughenbaugh, J. M., and Paredis, C. J. J., 2005, "The Value of Using Imprecise Probabilities in Engineering Design," *ASME DETC Design Theory and Methodology*, Long Beach, CA, USA, pp. DETC2005-85354.
[23] de Finetti, B., 1980, "Foresight. Its Logical Laws, Its Subjective Sources (Translated)," *Studies in Subjective Probability*, Kyburg, H. E. and Smokler, H. E. eds., E. Krieger Publishing Company.
[24] Hajek, A., 2003, Interpretations of Probability, http://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=probability-interpret.
[25] Joslyn, C. A., and Booker, J. M., 2005, "Generalized Information Theory for Engineering Modeling and Simulation," *Engineering Design Reliability Handbook*, Nikolaidis, E., Ghiocel, D. M., and Singhal, S. eds., CRC Press, New York, pp. 9.1-9.40.
[26] Savage, L. J., 1972, *The Foundation of Statistics*, Dover Publications, New York.
[27] Winkler, R. L., 1996, "Uncertainty in Probabilistic Risk Assessment," *Reliability Engineering & System Safety*, **54**(2-3), pp. 127-132.
[28] de Finetti, B., 1974, *Theory of Probability: A Critical Introductory Treatment*, Wiley, New York.
[29] Ferson, S., and Donald, S., 1998, "Probability Bounds Analysis," *International Conference on Probabilistic Safety Assessment and Management (PSAM4)*, Springer-Verlag, New York, NY, pp. 1203-1208.
[30] Troffaes, M. C. M., 2004, "Decision Making with Imprecise Probabilities: A Short Review," *The SIPTA Newsletter*, **2**(1), pp. 4-7.
[31] Berger, J. O., 1985, *Statistical Decision Theory and Bayesian Analysis*, Springer.
[32] Arrow, K., and Hurwicz, L., 1972, "An Optimality Criterion for Decision-Making under Uncertainty," *Uncertainty and Expectation in Economics: Essays in Honour of G. L. S. Shackle*, Carter, C. F. and Ford, J. L. eds., Blackwell.
[33] Ferson, S., Hajagos, J., Myers, D. S., and Tucker, W. T., 2005, "Constructor: Synthesizing Information About Uncertain Variables," Applied Biomathematics for Sandia National Laboratories, http://www.ramas.com/construct.pdf.
[34] Ferson, S., 2000, "Probability Bounds Analysis Solves the Problem of Incomplete Specification in Probabilistic Risk and Safety Assessments," *Ninth Conference on Risk-Based Decisionmaking in Water Resources*, Santa Barbara CA, USA.
[35] Ferson, S., 2002, "Ramas Risk Calc," 4.0 ed., Lewis Publishers, New York.
[36] Tucker, W. T., and Ferson, S., 2003, "Probability Bounds Analysis in Environmental Risk Assessment," Applied Biomathematics, Setauket, New York, http://www.ramas.com/pbawhite.pdf.
[37] Hines, W. W., Montgomery, D. C., Goldsman, D. M., and Borror, C. M., 2003, *Probability and Statistics in Engineering*, John Wiley & Sons, Inc., Hoboken, NJ.