

DETC2005-85354

THE VALUE OF USING IMPRECISE PROBABILITIES IN ENGINEERING DESIGN

Jason Matthew Aughenbaugh
gtg224k@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

Imprecision, imprecise probabilities, epistemic uncertainty, aleatory uncertainty, engineering design, probability bounds analysis, PBA

ABSTRACT

Engineering design decisions inherently are made under uncertainty. In this paper, we consider imprecise probabilities (i.e. intervals of probabilities) to express explicitly the precision with which something is known. Imprecision can arise from fundamental indeterminacy in the available evidence or from incomplete characterizations of the available evidence and designer's beliefs. Our hypothesis is that, in engineering design decisions, it is valuable to explicitly represent this imprecision by using imprecise probabilities. We support this hypothesis with a computational experiment in which a pressure vessel is designed using two approaches, both variations of utility-based decision making. In the first approach, the designer uses a purely probabilistic, precise best-fit normal distribution to represent uncertainty. In the second approach, the designer explicitly expresses the imprecision in the available information using a probability box, or p-box. When the imprecision is large, this p-box approach on average results in designs with expected utilities that are greater than those for designs created with the purely probabilistic approach. In the context of decision theory, this suggests that there are design problems for which it is valuable to use imprecise probabilities.

INTRODUCTION

Of the many challenges in engineering design, one of the greatest is uncertainty. During the design process, engineers must make decisions without being certain of the outcomes of the decisions. Prior to making a decision, engineers can remove some of this uncertainty by expending resources to acquire more information, for example, by modeling the

product or by performing experiments. Engineers also reduce uncertainty by making decisions and moving to a more precise product specification. However, even once a final design is reached, uncertainty about its performance remains due to variations in the manufacturing process and the product's environment of use.

Engineers have developed or adopted various methods to support design decisions under uncertainty, such as safety factors [1], utility theory [2-5], probabilistic risk assessments [6], reliability based design optimization [7], and robust design [8, 9]. We believe that current methods are still limited in their ability to clearly and quantitatively reflect the uncertainty encountered in engineering design because they do not accommodate imprecise characterizations of the uncertainty.

Imprecision can result from fundamental indeterminacy in the available evidence or from incomplete characterizations of the available evidence and designer beliefs. Common practice is to ignore this imprecision and to represent uncertainty using precise probabilities. Our hypothesis is that, in engineering design decisions, it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities (i.e. intervals of probabilities).

In the context of decision theory and utility theory, one design method is preferred to another if on average it yields designs with higher expected utilities. In this paper, we present an example pressure vessel design problem and computational experiment that demonstrates, in comparison to an expected utility maximization design method that uses precise probabilities, the value of a design method that explicitly

represents the imprecision in the available characterization of uncertainties with imprecise probabilities.

UNCERTAINTY IN ENGINEERING DESIGN

Definitions of uncertainty

We view uncertainty in the context of decision theory [10, 11]. Following Nikolaidis [12], we define uncertainty indirectly from the definition of certainty. Nikolaidis defines *certainty* as the condition of knowing everything necessary to choose the course of action whose outcome is most preferred. We define a decision-maker's uncertainty as the gap, shown in Figure 1(a), between certainty and the information the decision-maker currently has available for decision making, which is the *present state* of information, a slight refinement of Nikolaidis's definition.

We define a state of *precise information* as the state of having acquired all information available at any price. This is the most information that can be collected, even given infinite time and resources. However, this is not necessarily the same as a state of *certainty*. For example, a particular outcome of a truly random event cannot be known until it happens, regardless of how many resources are spent studying the event. Even if the random distribution describing the process is known precisely, the next outcome remains uncertain.

This uncertainty, and hence the gap between a state of precise information and certainty in Figure 1(b), is defined as *irreducible uncertainty*. Other authors use terms such as *variability* or *aleatory uncertainty* [13-16] to describe similar aspects of uncertainty. The actual existence of such irreducible uncertainty is a philosophical issue about which many people disagree. Most authors [13-16], including skeptics of a fundamental distinction between types of uncertainty [17], are willing to admit that it is useful in practice to accept that some uncertainties, such as machining errors, are the result of truly random processes. Even if the process is not random at the level of fundamental physics, engineers may choose to assume it is for practical reasons, much as they make other assumptions and simplifications when modeling real systems.

Philosophical arguments about the existence of irreducible

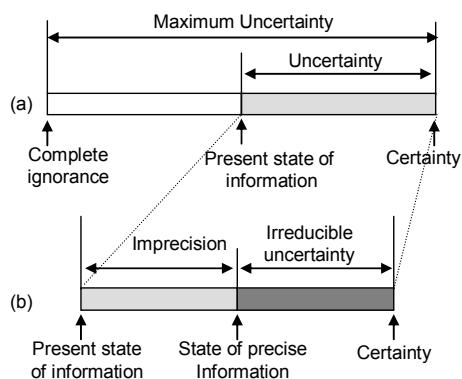


Figure 1: Characteristics of Uncertainty (adapted from [12])

uncertainty are tangential to our main point. We focus on the gap between the present state of information and a state of precise information, defined in Figure 1(b) as *imprecision*. Other authors use terms such as *epistemic uncertainty*, *ignorance*, or *reducible uncertainty* [13-16] to describe similar aspects of uncertainty. We specifically avoid the terms *aleatory* and *epistemic* because they have been used primarily in an attempt to differentiate the inherent nature of different uncertainties, rather than focusing on how human decision-makers should manage uncertainty in engineering design. In the following sections, we motivate the representation of uncertainty using imprecise probabilities, as suggested and formalized by Walley [18]. We begin with a discussion of interpretations of probability.

Interpretations of probability

Most engineers are familiar with the mathematics of probability, but many have not been formally exposed to the competing interpretations of probability. The philosophical arguments for or against different interpretations can be quite passionate. The interpretations most commonly adopted in engineering design are variations of the frequentist and subjective interpretations [19]. We briefly introduce both and indicate how imprecise probabilities are relevant under either interpretation. For more complete discussions of interpretations of probability, see for example [18-22].

The *frequentist* interpretation is based on the notion of relative frequencies of outcomes. Under a frequentist interpretation, 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. In engineering design, events are not always repeatable. Even assuming some events are essentially repeatable and data can be collected, there is no guarantee that a particular sample is representative of the true relative frequency. Although in theory the relative sample frequency approaches the true relative frequency as the sample size goes to infinity, an infinite sample size is impossible in practice. Consequently, engineers will always face *imprecision* in their characterizations of the frequentist probabilities.

Proponents of a *subjective* interpretation of probability assert that there is no such thing as a true or objective probability, but rather probabilities are an expression of belief based on an individual's willingness to bet [17, 19, 20]. One of the subjectivists' primary arguments against a frequentist perspective is the absence of truly repeatable events, especially in practical problems. For example, the probability that Team A beats Team B in a basketball game has no real meaning under a frequentist interpretation, because that event—that particular game—will occur exactly once. The notion of a long term frequency, and even random events, is meaningless [20]. Nevertheless, many people are willing to express their belief of who will win in terms of bets. When framed appropriately, such bets can be taken as subjective probabilities [20].

The process of eliciting and assessing an individual's beliefs, or willingness to bet, is resource intensive. Even assuming that

precise beliefs—and hence precise probabilities—exist, it will often be impractical to fully characterize them due to constraints such as bounded rationality, time, and computational ability [18]. Consequently, only a partial—and therefore imprecise—characterization of the subjective probabilities is normally available. We address additional motivations for the use of imprecise probabilities in the next section, including discussing whether precise beliefs always exist.

The Motivation for Imprecise Probabilities

The general motivation for imprecise probabilities is that the more evidence on which a probability estimate is based, the more confidence a decision maker can have in it. Thus, the imprecision in the probabilities should be expressed explicitly in order to signal the appropriate level of confidence to ascribe to them.

For example, adapting an example from Walley [18], consider the toss of a thumbtack. The goal of the exercise is to determine the probability that the tack lands pin-up. Three experimenters perform this exercise, as follows:

1. Experimenter A is in a hurry and does not even look at the thumbtack. Experimenter A employs a non-informative prior distribution, in this case the principle of indifference or insufficient reason, and assumes that the probability of the tack landing pin-up is equal to the probability it doesn't land pin-up, thus ascribing a probability of 0.5 to both possibilities.
2. Experimenter B tosses the thumbtack 10 times and gets 6 pin-ups. Experimenter B's estimated precise probability of the tack landing pin-up is thus 0.6.
3. Experimenter C tosses the thumbtack 1000 times and gets 400 pin-ups. Experimenter C's estimated precise probability of the tack landing pin-up is thus 0.4.

If we think of the three experimenters as three analysts that could provide us with information, which analyst would we prefer to hire? Because Experimenter C's estimate was based on more data, it is more precise than Experimenter B's estimate. Experimenter A's estimate was based on no data, so it does not seem reasonable to place much confidence in it. Nevertheless, the precise probability estimates of 0.5, 0.6, and 0.4 appear equally credible. By not expressing the imprecision in these estimates, one is arbitrarily eliminating it, thus assuming precision that has no justification in the available evidence. This problem can be overcome by allowing analysts to state imprecise probabilities.

A full discussion of Walley's formalization of imprecise probabilities [18] is well outside the scope of this paper, but we summarize a few essentials. Walley defines upper and lower probabilities, which are special cases of upper and lower previsions. The lower prevision represents a price at which you are sure you would buy a bet, and the upper prevision represents a rate at which you are sure you would buy the opposite of the bet. If the two are equal, then they jointly represent your *fair price* for the bet, the price at which you are

willing to take either side of the bet. At a price between your upper and lower previsions, you are not willing to enter the bet on either side, at least not without collecting more information and updating your previsions. This connects directly with a common argument against imprecise probabilities.

One objection to the use of imprecise probabilities is that they can lead to indeterminacy of action during decision making. That is, given imprecise probabilities, there is no single clear "best" solution according to standard decision theories. We counter this by arguing that if the available evidence does not clearly suggest a particular course of action, then the representation of this evidence should not arbitrarily pretend that it does. An approach that demands precise probabilities necessitates an arbitrary resolution of the indeterminacy, a resolution that may have no basis in the available evidence. Such an approach does not differentiate well-grounded probabilities (such as Experimenter C's) from arbitrary ones (such as Experimenter A's). By admitting imprecise probabilities, one can abstain from these arbitrary judgments during analysis and support better decision making, as follows.

It is true that in order for a single solution to be chosen, the indeterminacy will need to be resolved. However, the explicit representation of imprecision in the characterization of uncertainty allows for the direct management of the imprecision in the context of a decision. For example, if a decision is not sensitive to the current level of imprecision, a robust decision can be made using rules such as satisficing, minimax, or arbitrary choice [18]. On the other hand, if the decision is sensitive to the existing imprecision, a decision can be made to collect more information (such as more tack tosses in our earlier example).

Another frequently levied objection to imprecise probabilities is that they are irrational. It is generally accepted that in order to be rational, one's probabilities must avoid a sure loss, a situation often referred to as a Dutch Book [18, 20, 21]. The general idea is that if your probabilities violate certain rules, a group of bets can be constructed, all of which you are willing to accept, but the combination of which results in a sure loss; you will lose money under any outcome. This argument is often presented in favor of precise probabilities and the axioms of Kolmogorov [23] or de Finetti [20]. However, Walley [18] presents axioms of coherence for imprecise probabilities that also avoid sure loss.

Once again, the details of Walley's formalization of imprecise probabilities are beyond the scope of this paper, but he begins with the fundamental notion that rationality requires you to avoid a sure loss. His axioms of coherence assure that if your imprecise probabilities satisfy them, then you are not subject to a sure loss. His primary deviation from axioms of precise probabilities such as de Finetti's is the allowance for a range of indeterminacy—prices at which you will not enter a bet as either a buyer or a seller. Thus, if your imprecise probabilities are coherent, you do not enter into a situation in which you suffer a sure loss. If you do spend the time and effort to collect

infinite evidence and fully elicit your beliefs, then your imprecise probabilities will collapse into precise probabilities.

The only concrete way to determine the value of using imprecise probabilities in engineering design is in the context of decision theory. Informally, one method of representing uncertainty is better than another if it allows designers to make better decisions. With this in mind, we discuss decision theory in the following section.

DECISION-MAKING UNDER UNCERTAINTY

Engineering design decisions are hampered by the uncertainty inherent in predicting the outcomes and payoffs of different actions. In this section, we discuss a safety-factor approach and traditional statistical decision theory.

Design with safety factors

One way engineers deal with uncertainty is by using safety factors. For example, if engineers are building a pressure vessel with material yield strength σ_y , the requirement to avoid failure is that σ_y exceeds the maximum stress in the pressure vessel, σ_{\max} . That is: $\sigma_y > \sigma_{\max}$. Engineers rarely, if ever, know these parameters with certainty. In a safety factor approach, engineers employ a safety factor $SF > 1$ with their best point estimates $\tilde{\sigma}_y$ and $\tilde{\sigma}_{\max}$ of the true σ_y and σ_{\max} , respectively, and design the pressure vessel such that

$$\tilde{\sigma}_y > SF \times \tilde{\sigma}_{\max}.$$

The safety factor is an attempt to wrap all uncertainty into one number. The engineers hope that by designing the pressure vessel with a safety margin around their estimates, the true yield strength σ_y (which may be much less than $\tilde{\sigma}_y$) will exceed the true σ_{\max} (which may be much larger than $\tilde{\sigma}_{\max}$).

The simplicity of employing safety factors is its biggest advantage, but the unanswered question in this approach is how large of a safety factor is necessary to meet reliability and performance requirements given the existing uncertainty. In the presence of large data histories, safety factors actually can be linked to probabilistic characterizations of uncertainty and reliability [1]. However, this link assumes precise probabilities are available.

In practice, particularly for novel design tasks, engineers do not have precise probabilities and often must resort to an *ad hoc* choice of a safety factor. Because engineers do not want structures to fail, engineers usually choose safety factors that are much larger than necessary. This frequently results in costly over-design of the product. At the same time, the *ad hoc* method fails to provide any guarantee of reliability. Consequently, engineers have pursued more formal methods for dealing with uncertainty, such as traditional statistical decision theory.

Traditional Statistical Decision Theory

In traditional statistical decision theory [24], utility analysis, as originally proposed by von Neumann and Morgenstern [25], is

used for making decisions under uncertainty. von Neumann and Morgenstern assume a frequentist interpretation of probability in their theory, but Savage [22] presents an alternative axiomization of utility based on subjective probabilities. In general, *utility* expresses preference—more preferred decision outcomes are assigned higher utility values. If chosen correctly, utilities reflect the decision-maker's preferences, even under uncertainty. By applying the expected value operator, the decision-maker weights all possible outcomes according to their likelihood of occurring, and then chooses the action that maximizes the expected utility. Utility theory has been studied extensively by economists and decision theorists, and there continues to be an increasing interest in applying utility methods to engineering problems, as in [2-5].

In traditional statistical decision theory, one characterizes the uncertainty with a precise probability density function. Consequently, a designer is forced to either eliminate or ignore imprecision. Elimination requires the designer to expend resources to acquire more information, thus increasing the costs of the design process. Ignoring imprecision involves overstating the true current state of information by making assumptions. Since both of these approaches have inherent problems, it appears that an extension to existing statistical decision theory is needed.

We investigate the hypothesis that by generalizing the statistical formalism to allow imprecise probabilities, better decisions can be obtained. For this purpose, we introduce probability bounds analysis in the following section.

PROBABILITY BOUNDS ANALYSIS

In order to consider imprecision explicitly and distinctly from irreducible uncertainty, we use a recently developed formalism that extends traditional probability theory—called probability bounds analysis (PBA) [26, 27]. Other methods are available, such as Dempster-Shafer Evidence Theory [28] and double-loop or joint Monte Carlo sampling [29], but PBA has additional appealing properties, such as its ability to propagate uncertainty in a computationally efficient fashion, as described in the *Discussion and Future Work* section. PBA is also better starting point for illustration because it is more closely aligned with familiar concepts from probability, statistics, and interval analysis.

PBA basics

PBA expresses uncertainty in a structure called a *probability-box*, or *p-box*. The p-box incorporates both imprecision and probabilistic characterizations by expressing interval bounds on the cumulative probability distribution function (CDF) for a random variable. This marriage of probability theory and interval analysis is captured in the analogy that a p-box is a stretched-out distribution function in the same way that an interval is a stretched out scalar. More formally, the bounds on a p-box, such as shown in Figure 2(a), are given by two CDFs (F_1 and F_2) that enclose a set of CDFs that are, under some interpretation, consistent with the current state of information.

Interpreting a p-box

While there are several ways to construct p-boxes [31, 32], we choose a practical method based on traditional confidence intervals, as explained in Appendix A. The method by which a p-box is constructed influences its exact interpretation. For illustration of meaning, it is easiest to consider a p-box as constructed at the 100% confidence level. In this case, the p-box expresses the range of all CDFs that are still deemed possible based on existing information. For example, assume that for all practical purposes, X is a random variable, and engineers have strong theoretical information that X is normally distributed with known variance $\sigma^2=1$. However, the engineers can only characterize the mean imprecisely, bounding it in the interval $\mu=[0,1]$. Extending the notation of probability, we can write

$$X \sim N(\mu, \sigma^2) = N([0,1], 1),$$

such that X is distributed normally with mean μ and variance σ^2 . The corresponding p-box is shown in Figure 2(a). In this case, the bounds on the p-box are defined by the two distributions, $F_1 \sim N(0,1)$ and $F_2 \sim N(1,1)$. 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 2(b). However, any distribution that falls partially or entirely outside of the p-box is definitely inconsistent with the present state of information.

Vertical slices of the p-box yield intervals on the CDF for a particular realization. For example, a vertical slice at zero yields the interval for the CDF of [0.1587, 0.5]. This means that the probability that X is less than zero is between 0.1587 and 0.5, but one does not have enough information to specify a precise probability within that interval. Horizontal slices of a p-box result in intervals on the quantiles of the CDF. For example, a slice at the median (CDF=0.5) gives the interval [0,1] for the median.

In summary, a p-box is a more expressive generalization of both traditional probability distributions and interval representations, as is illustrated in Figure 3. A general p-box explicitly expresses both probability (represented by the shapes of the boundary CDFs) and imprecision (represented by the separation between the upper and lower bounds).

MEASURING THE VALUE OF USING IMPRECISE PROBABILITIES

In this section, a pressure vessel is designed under incomplete information using two design approaches. Descriptions of the design scenario and the computational experiment follow.

Design scenario

One needs to design a pressure vessel that is to contain 0.15 m^3 of gas under 7 MPa of pressure. Due to space limitations, certain maximum dimensions are imposed. The goal is to determine the dimensions (radius R , wall thickness t , and length L) of the vessel, shown in Figure 4, for which the overall utility,

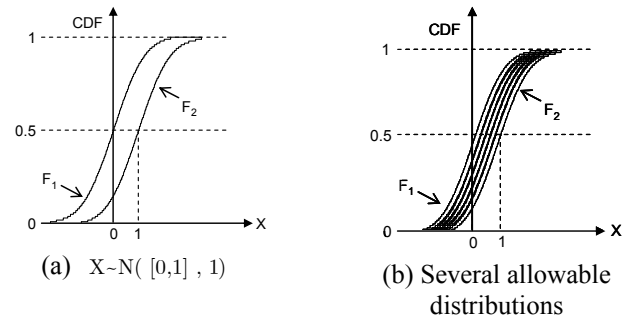


Figure 2: Example P-box and distributions

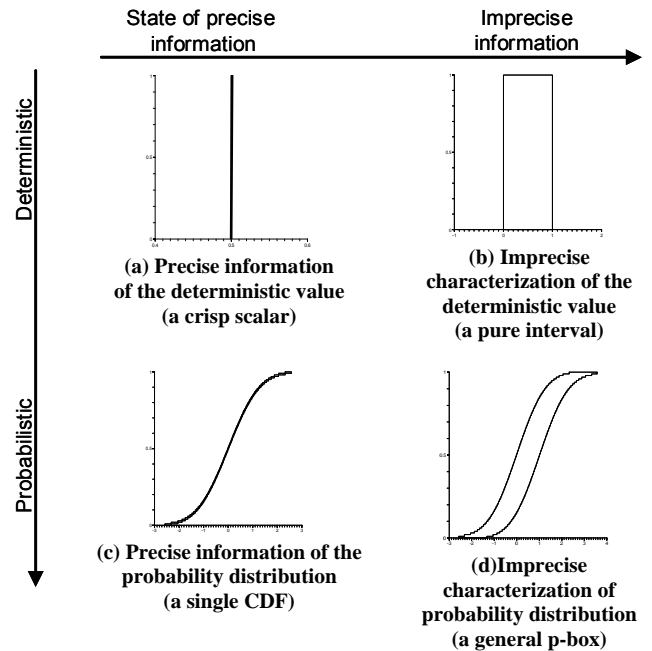


Figure 3: Dimensions of uncertainty

defined in the next section, is maximized. Since the vessel will be used in a human-occupied location, the cost of failure is weighted heavily. The vessel will be made of a new type of steel for which the yield strength is not well characterized. The material production process produces variations in the material properties such that the material yield strength is well modeled by a random variable. Because the material is new and testing is relatively expensive, variations in yield strength have only been measured in a set Σ of n independent tension tests, where n is a relatively small number due cost considerations.

These tests can at best give an estimate of the true distribution, so in addition to inherent randomness (irreducible uncertainty), engineers also face imprecision—they cannot characterize the parameters of the random variable precisely. The number of samples n can varied to explore different levels of imprecision, just as the number of tosses was varied in the thumbtack example. A detailed description of the design assumptions, requirements, and parameters is given in Appendix B.




Figure 4: Pressure vessel schematic and design variables

The computational experiment

The goal of the experiment is to compare the utility of the design solutions that result from different approaches for representing uncertainty applied to the same design problem. The comparison is made possible in this experiment because we assume that overseeing the experiment is a supervisor who is in a state of precise information about the steel's material properties. From the supervisor's perspective, only irreducible uncertainty about the yield strength of the material exists—uncertainty that is precisely characterized by a normal distribution with a mean of 180 MPa and a standard deviation of 15 MPa. The supervisor can therefore determine precisely the dimensions of the pressure vessel that result in the maximum expected utility. This optimal design under the precise information is the benchmark for comparison of the other design approaches.

The general layout of the experiment is shown in Figure 5. The experiment consists of two designers: one using a single best-fit normal distribution (approach A), and the other (approach B) using a p-box to represent the uncertainty about the yield strength. The details of these approaches are explained in the sections following this general overview of the experiment.

Both approaches start with the same information about the uncertain yield strength. This information is a set Σ of n random samples (torsion test results) from the true normal distribution:

$$\Sigma = \{\sigma_{y_i}\}_{i=1}^n.$$

Both designers assume that the true distribution is normal, but they use their own approach to represent their uncertainty about the distribution's parameters. According to decision policies explained in the following sections, each designer then selects an optimal design, denoted as

$$a_{\Sigma}^* = \{R_A^*, t_A^*, L_A^*\} \text{ and } b_{\Sigma}^* = \{R_B^*, t_B^*, L_B^*\}$$

respectively for approach A and approach B.

The supervisor compares each of the design solutions to determine the *expected utility evaluated under precise information* for each solution. For approach A this is written as

$$E_{\sigma_y|\Sigma}[U(\sigma_y, a_{\Sigma}^*)],$$

and for approach B as

$$E_{\sigma_y|\Sigma}[U(\sigma_y, b_{\Sigma}^*)].$$

In order to compare the value of the two approaches, the supervisor, who has access to precise information, computes the difference in expected utility. Using the fact that the two

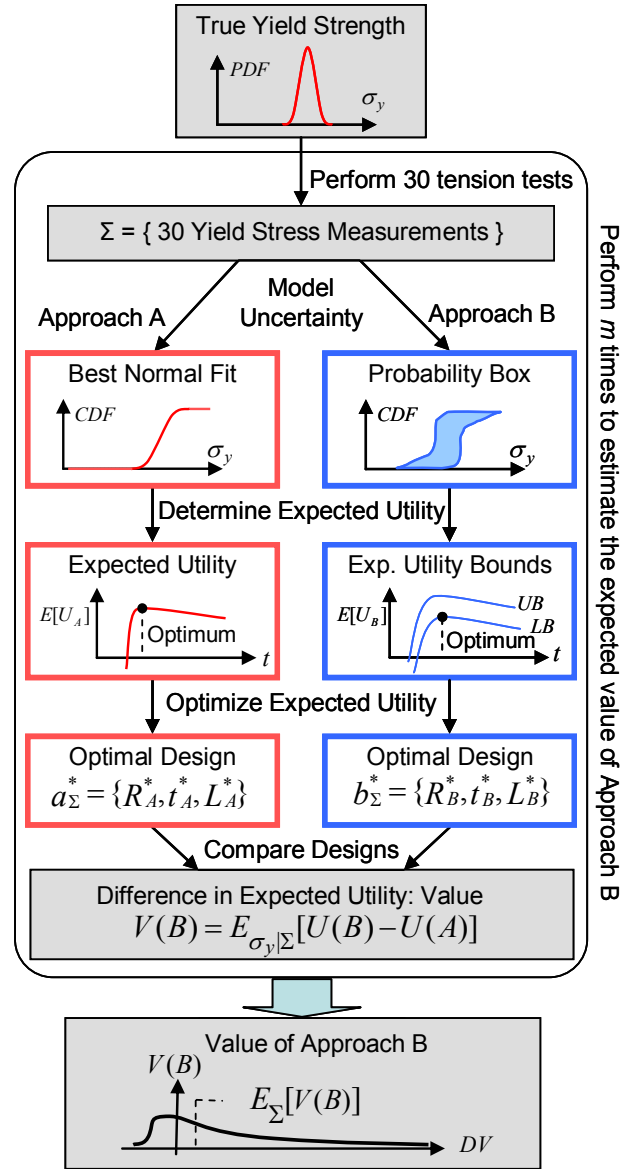


Figure 5: A computational experiment for determining the value of using imprecise probabilities (shaded actions are performed by a supervisor)

approaches start with the same information (sample Σ), the value of approach B over approach A can then be expressed as

$$V(B) = E_{\sigma_y|\Sigma}[U(\sigma_y, b_{\Sigma}^*) - U(\sigma_y, a_{\Sigma}^*)].$$

It is necessary to note that this value was for only one sample Σ —the yield strength measurements with which each designer starts. Due to the randomness in Σ , one trial is not sufficient to judge the relative value of each approach; the supervisor needs to repeat the above experiment many times in order to determine which design approach performs best *on average*, over m different sample sets Σ . Mathematically, the expectation must be taken with respect to Σ in order to calculate the *average expected value of approach B over A*, written

$$E_{\Sigma}[V(B)] = E_{\Sigma}[E_{\sigma_y|\Sigma}[U(\sigma_y, b_{\Sigma}^*) - U(\sigma_y, a_{\Sigma}^*)]].$$

The addition of the word *average* emphasizes that this quantity is the expectation over the samples of the expected utility of particular design solutions.

Both design approaches are based on utility theory and use the same preference structure and utility function given by:

$$U(\sigma_y, DV) = P_{\text{selling}} - C_{\text{material}}(DV) - C_{\text{failure}}(\sigma_y, DV),$$

where:

$$P_{\text{selling}} \equiv \text{selling price} = \$200$$

$$C_{\text{material}} \equiv \text{material cost per volume} = \$8500/\text{m}^3$$

$$\sigma_y \equiv \text{true yield strength of pressure vessel}$$

$$DV \equiv \text{design variables (radius, thickness, length)}$$

$$C_{\text{failure}}(\sigma_y, DV) \equiv \begin{cases} 0 & \text{if } \sigma_y \geq \sigma_{\text{max}}(DV) \\ \$1,000,000 & \text{otherwise} \end{cases}$$

As defined in the previous line, for 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).

In addition to approach A and approach B, the experiment's supervisor is able to create a design using precise information. The three approaches are described in the following sections.

Supervisor's design under precise information

The supervisor can create a design using the true distribution, since he or she is in a state of precise information. The supervisor therefore knows precisely that $\sigma_y \sim \text{Normal}(180\text{MPa}, (15\text{MPa})^2)$. If we define this as approach K (not shown in Figure 5), the supervisor then chooses the design variables

$$DV = k = \{R_K, t_K, L_K\}$$

such that the expected utility $E_{\sigma_y}[U(\sigma_y, k)]$ is maximized. This leads to the *optimal design under precise information*:

$$k^* = \underset{k}{\text{argmax}}(E_{\sigma_y}[U(\sigma_y, k)]).$$

This optimal design, with expected utility denoted $E[U(k^*)]$ for brevity, serves as the baseline for comparison because no other approach can yield an average higher expected utility across many repetitions m .

Design using approach A: Precise normal fit

Designer A does not have access to precise information, but instead only has access to the set Σ of n data samples. Because the designer does not know the true distribution of σ_y , he or she must make an approximation, denoted $\tilde{\sigma}_y(A, \Sigma)$. In this approximation, the representation of $\tilde{\sigma}_y$ depends on both the approach, in this case A, and the observed random sample Σ . Designer A represents the uncertainty as a normal distribution, using the sample mean and sample variance as unbiased estimates of the true mean and true variance, respectively. This yields the probabilistic model

$$\tilde{\sigma}_y(A, \Sigma) \sim \text{Normal}\left(\frac{1}{n} \sum_{i=1}^n \sigma_{y_i}, \frac{1}{n-1} \sum_{i=1}^n (\sigma_{y_i} - \frac{1}{n} \sum_{i=1}^n \sigma_{y_i})^2\right).$$

Designer A therefore chooses design variables

$$a_\Sigma = \{R_A, t_A, L_A\}$$

that maximize the *estimated* expected utility

$$E_{\tilde{\sigma}_y|\Sigma}[U(\tilde{\sigma}_y(A, \Sigma), a_\Sigma)],$$

given his or her information about the randomness. The expected utility is only *estimated* because designer A does not have access to a precise characterization of the random variable σ_y . The expected utility maximization results in the optimal design using approach A given samples Σ , denoted:

$$a_\Sigma^* = \underset{a_\Sigma}{\text{argmax}}(E_{\tilde{\sigma}_y|\Sigma}[U(\tilde{\sigma}_y(A, \Sigma), a_\Sigma)]).$$

Design using approach B: Imprecise probabilities

Designer B takes a different approach for capturing the uncertainty in the yield strength. Specifically, designer B represents the uncertainty in σ_y by $\tilde{\sigma}_y(B, \Sigma)$, where the nature of $\tilde{\sigma}_y(B, \Sigma)$ is expressed as a p-box rather than a pure normal distribution. The construction of this p-box is addressed in Appendix A. The estimated expected utility under design approach B is defined as

$$E_{\tilde{\sigma}_y|\Sigma}[U(\tilde{\sigma}_y(B, \Sigma), b_\Sigma)]$$

where:

$$b_\Sigma = \{R_B, t_B, L_B\}$$

is the designer's chosen design action given sample set Σ . Because the p-box expresses a range of possible distributions, the expected utility is no longer a crisp number but rather an interval defined by lower-bound \underline{E} and upper-bound \bar{E} , such that

$$E_{\tilde{\sigma}_y|\Sigma}[U(\tilde{\sigma}_y(B, \Sigma), b_\Sigma)] = [\underline{E}, \bar{E}].$$

Because the expected utility is now an interval, the designer cannot choose design variables b_Σ that maximize the expected utility in the traditional sense. Instead, a new decision rule is required. While other approaches exist, in this experiment, a conservative best-worst case, or maxi-min, rule is used. Designer B therefore chooses the design action b_Σ that has the highest lower bound \underline{E} on the expected utility. This results in an optimal design decision using approach B given the observed samples Σ , denoted:

$$b_\Sigma^* = \underset{b_\Sigma}{\text{argmax}}(\underline{E}) = \underset{b_\Sigma}{\text{argmax}}(\underline{E}_{\tilde{\sigma}_y|\Sigma}[U(\tilde{\sigma}_y(B, \Sigma), b_\Sigma)]).$$

EXPERIMENTAL RESULTS

The computation experiment was repeated for many different initial sample sets Σ of size n . For each level of n , the design process was repeated $m=100,000$ times in order to determine the average performance of the two methods. We first present the results for the particular sample size $n=30$, and then discuss the results over varying values of n , which represent different levels of imprecision.

Value of using imprecise probabilities for $n=30$ samples of the true yield strength

We first conduct the experiment with a sample set Σ of size $n=30$, meaning the designers are given the results of 30 independent yield stress tests. As measured by the supervisor using precise information and averaged over $m=100,000$ initial sets, approach B on average yields designs with greater expected utility than approach A. Specifically, using standard statistical analysis [33], we find the 95% confidence interval (CI) on the value of approach B over A to be:

$$95\% \text{ CI on } V(B) \text{ is } [\$16, \$22].$$

To put this result in perspective, the expected utility of the supervisor's design, which is the best possible because it is designed under a state of precise information, is $E[U(k^*)] = \$104$. Thus the CI on the expected value of approach B over A can also be expressed as [16%, 21%] of the optimal utility $E[U(k^*)]$. This is a substantial deviation that suggests that there is value in using the p-box approach for this design problem. We also note the average expected utilities realized under approach A and B :

$$\text{Approach A: } E_{\Sigma}[E_{\sigma_y \Sigma}[U(\sigma_y, a_{\Sigma}^*)]] = \$70,$$

$$\text{Approach B: } E_{\Sigma}[E_{\sigma_y \Sigma}[U(\sigma_y, b_{\Sigma}^*)]] = \$89.$$

The total deviations from optimal (\$34 for approach A and \$15 for approach B), coupled with the relative value of approach B over A, indicate that B is a better approach at this level of imprecision. In the next section, we explore how these results differ at varying levels of imprecision.

Variation of value with level of imprecision

The previous discussion dealt with a fixed sample set size of $n=30$ material strength tests. While those results demonstrated that it was valuable to use the p-box approach in that case, a more general result is desirable. By varying the number of material strength tests n , we can vary the imprecision of the characterization. The supervisor's design yields the best possible expected utility $E[U(k^*)]$, and hence the designs of designers A and B can at best equal it. In Figure 6, we plot (in log-log scale) the percent deviation from this best possible expected utility for approach A and approach B for different values of n .

When the imprecision is large, approach B performs significantly better than approach A. For example, at a sample size of 10, a 95% CI on the value of approach B over A is [540%, 580%] of $E[U(k^*)]$. There is no doubt that this value is significant. Around sample size 50, the two approaches yield similar results. Past this point, approach A performs better, but not by much. For example, for a sample size of 100, the 95% CI on the value of approach B is [-0.7%, -0.2%] of $E[U(k^*)]$. This difference appears insignificant, but we note there may be cases in which such small percentages matter.




Figure 6: Variation of value with imprecision




Figure 7: Histogram of value of p-box approach

In summary, as the imprecision increases, the value of approach B over approach A increases significantly. As the imprecision approaches zero, approach A becomes only slightly better than B. For different design problems, the designer will not necessarily know where the two curves cross. Thus, unless the designer is sure *a priori* that the consequences of the imprecision are insignificant, the results of this computational experiment suggest that it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities.

Explanation of results

In this section, we provide some insight into the results of the previous two sections. The results for 100,000 different sample sets Σ of sizes $n=10, 30$, and 100 are shown in histograms in Figure 7. The x-axis of the histogram in Figure 7 is the expected value of approach B over A, denoted $E_{\sigma_y \Sigma}[U(\sigma_y, b_{\Sigma}^*) - U(\sigma_y, a_{\Sigma}^*)]$.

In many cases, approach B yields a design with a lower utility than approach A. However, in some cases, approach B yields a design with a much higher expected utility. The tails of the distribution for $n=10$ and $n=30$ extend much farther than shown in the figure. For example, the maximum expected value of

approach B seen in any trial of $n=10$ was \$95,000. Results such as these skew the overall distribution such that on average, approach B yields a design with a higher expected utility than approach A yields. As the imprecision decreases, both the skewness and value of approach B over A also decrease.

The results can also be understood in terms of the expected utility curves used in the experiment. In Figure 8 and Figure 9, we illustrate the expected utility as a function of wall thickness t for two different samples, Σ_1 and Σ_2 respectively, both of size $n=30$. The four curves shown in each figure are:

1. Estimated expected utility under approach A: $E_{\sigma_y|\Sigma}[U(\tilde{\sigma}_y(A, \Sigma), t)]$
2. Estimated lower bound on expected utility under approach B: $E_{\sigma_y|\Sigma}[U(\tilde{\sigma}_y(B, \Sigma), t)]$
3. Estimated upper bound on expected utility under approach B: $E_{\sigma_y|\Sigma}[U(\tilde{\sigma}_y(B, \Sigma), t)]$
4. The true expected utility in a state of precise information: $E_{\sigma_y}[U(\sigma_y, t)]$

Designers A and B choose a thickness (t_A^* and t_B^* respectively) that is optimal according to their *estimated* expected utilities. Their estimates are in general *not* equivalent to the true expected utility function. Therefore, the expected utility actually realized by a particular design is not reflected in their estimates, but rather in the *true* curve $E_{\sigma_y}[U(\sigma_y, t)]$, which is known only to the supervisor.

In Figure 8 (based on sample set Σ_1), the true curve is between the curve from approach A and the upper bound from approach B. By noting t_A^* and t_B^* , we can read off the true expected utility evaluated under precise information for each approach from the truth, $E_{\sigma_y}[U(\sigma_y, t)]$. For this particular sample set Σ_1 , the expected utility realized from approach A is about \$20 higher than the expected utility realized from approach B. This means for sample Σ_1 the relative value of approach B is negative: $V(B) = -\$20$, indicating approach A performs better.

In Figure 9 (based on sample set Σ_2), the true curve is near the lower bound of approach B. For this particular sample set Σ_2 , the expected utility from approach B is about \$60 more than that from approach A, so the relative value of approach B for sample Σ_2 is $V(B) = \$60$.

The preceding results are for two representative cases. The overall results of the experiment, discussed previously, indicate that, on average, the latter case dominates. Approach A is more likely to overestimate the true material strength, and when it does, the consequences are disastrous — a high probability of failure. Approach B is more conservative, resulting in higher material costs, but, on average, these material costs are offset by the reduced failure costs. As the sample size increases, the best-fit normal distribution of approach A becomes, on average, closer to the true distribution, such as the example sample Σ_3 of size $n=200$ shown in Figure 10. For this sample, the optimal designs of the three approaches converge, and therefore yield




Figure 8: Example Expected Utility Functions, $n=30$

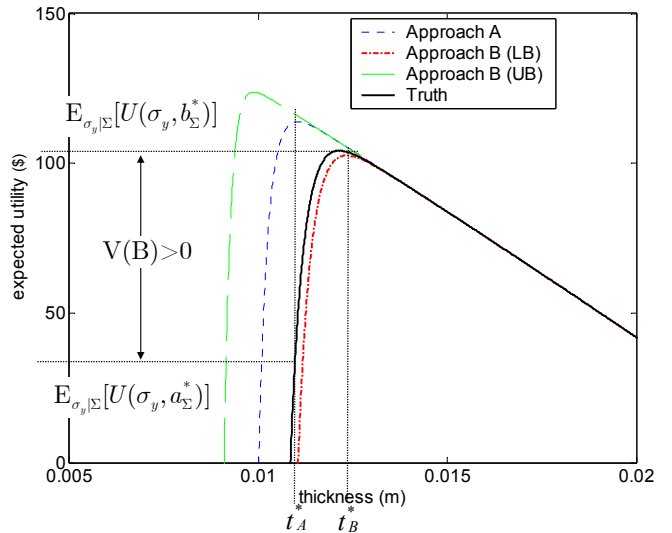


Figure 9: Example Expected Utility Functions, $n=30$




Figure 10: Example Expected Utility Functions, $n=200$

similar utilities. Thus we can see how when the imprecision is small, the value of approach B over A is near zero.

Summary of results

For this design problem, the experimental results indicate that when the imprecision is large, approach B (using imprecise probabilities) performs significantly better than approach A (using precise probabilities). When the imprecision is small, the difference between the two approaches is insignificant. This computational experiment has therefore demonstrated that there are scenarios in which it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities.

DISCUSSION AND FUTURE WORK

In this paper, we present a specific and simplified design problem and computational experiment that demonstrates the value of using imprecise probabilities in engineering design. The demonstration of the value in one design problem leads to the natural question: *what are the characteristics of a design problem that make the use imprecise probabilities valuable?* The specific use of PBA in this example leads to the question: *are there more effective approaches for making decisions when the characterization of uncertainty is imprecise?* These questions form the basis for future work.

Net value of using imprecise probabilities

In this paper, we have demonstrated the potential value of a design method that uses imprecise probabilities. What we have heretofore referred to as *value* is really the *gross value*. What designers should really care about is the *gain*, or *net value*, which is defined as the difference between gross value and cost [34].

One cost factor is computational cost. The probabilistic approach often requires Monte Carlo analysis or related methods in order to calculate expected values [35]. These methods have an inherently high computational cost, even when used in combination with surrogate modeling [36, 37] to reduce the computational burden. In this example, the design problem was simplified to the point that Monte Carlo analysis could be avoided. First, all of the distributions considered were normal, and second, the utility function was simple enough that the expected utility could be calculated analytically using the cumulative normal distribution, a commonly available statistic.

This computational simplicity extended to the PBA approach in this example. The boundaries of the p-box were piece-wise normal CDFs. Thus, the expected utility interval could be calculated from the p-box using the cumulative normal distribution, too. Consequently, the computational cost of the two design methods was comparable. In this example, since the PBA approach yielded a greater expected value at a similar cost when the imprecision was large, the net value (gross value less cost) of moving to an approach based on PBA is positive, and is therefore worth further research. Particularly, we must consider more general example problems.

Propagation of uncertainty using p-boxes

In this experiment, there is no need to combine uncertainties from different sources because only one parameter is uncertain. However, most engineering design problems contain multiple sources of uncertainty and thus do require a means to propagate uncertainty. PBA includes algorithms with foundations in interval analysis for propagating p-boxes through calculations [30, 38]. Ferson and Ginzburg [13] claim that these methods are on average less computationally expensive for computing with imprecise probability distributions than double-loop Monte Carlo methods. However, it does not appear that these methods have been explored in detail for engineering design applications.

Another consideration is the flexibility of the algorithms. In Monte Carlo analysis, engineers often assume specific correlations between variables as a matter of convenience that enables the use of standard statistics. This understates the true imprecision in the available characterizations of uncertainty and can have serious consequences. PBA provides methods [30] that can propagate uncertainty under various conditions of dependence, including the extremes of fully known and completely unknown dependencies. Despite this promise, it is not yet clear how well PBA can propagate uncertainty from many sources, or how well PBA can be integrated with more complex design tools and models, such as discrete-event simulations and optimization methods. These issues need to be resolved before PBA can be put to use in general engineering design problems.

Variation in the distributions of random variables

In this experiment, we assumed a design attribute was well modeled as a normally distributed random variable. However, the nature of the randomness, such as its distribution type and parameters, may affect the value of various design approaches. For example, if the true distribution is not normal, then both of these approaches may perform much worse compared to the reference case of precise information. It is therefore important to explore different distributions.

An additional consideration is purely subjective probabilities, that is, probabilities for events that are not well modeled as inherently random events. The construction of p-boxes based on more subjective measures is an open question. Similarly, the combination of evidence and beliefs is a difficult problem that has been investigated, but not resolved, in many fields [39, 40]. Thus the combination of subjective probabilities is an area for future research.

Decision policies and preferences

In this experiment, the relatively high cost of vessel failure, in comparison to material costs, means it is very important to avoid failure. Because the PBA approach—including the maximin decision policy—used in this experiment is conservative compared to a normal fit approach, it may be that PBA provides more value only when there is a strong incentive to avoid failure. Therefore the performance of PBA should be explored for different preference structures. Similarly, given an

interval on expected utility, there are alternative (and less conservative) methods in interval theory [41] for making decisions that should be explored, including midpoint, satisficing, and maximum entropy policies [18, 41].

SUMMARY

Engineers inherently lack information during the engineering design process and therefore face uncertainty. Existing design methods do not explicitly represent the imprecision in the available characterizations of uncertainty. The pressure vessel design example in this paper demonstrates that when the designers only have access to a small set of sample data, a PBA approach that uses imprecise probabilities to model uncertainty in design decisions can lead on average to better designs than a purely probabilistic approach that requires precise probabilities. We can then conclude that at least in some design problems, it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities.

ACKNOWLEDGMENTS

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. 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. The authors would like to thank Rich Malak, Steve Rekuc, Jay Ling, Morgan Bruns, and Scott Duncan for their valuable insights and feedback.

APPENDIX A: 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 [27, 30, 42]. Only more recently have researchers addressed the construction of p-boxes from observed sample data [31].

While several approaches exist [31], we choose a pragmatic approach based on standard confidence intervals. In this example, the designers know the true yield strength is normally distributed, but with unknown mean and variance. For clarity here, we denote the random variable for yield strength as X to distinguish it from standard statistical notations that use σ for the standard deviation. Thus we have:

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

Because the true μ and σ^2 are unknown, the designer uses the unbiased point estimates

$$\hat{\mu} = \bar{X} = \frac{1}{n} \sum_{i=1}^n X_i \quad \text{and} \quad \hat{\sigma}^2 = s^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2$$

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. So far this is exactly what we have

defined as the pure probabilistic approach. In order to construct a p-box, we broaden these point estimates to confidence intervals. For the mean we find a $100(1 - \alpha)\%$ confidence interval as

$$[\underline{\mu}, \bar{\mu}] = [\hat{\mu} - z_{\alpha/2} \hat{\sigma}_{\bar{X}}, \hat{\mu} + z_{\alpha/2} \hat{\sigma}_{\bar{X}}] \quad \text{where} \quad \hat{\sigma}_{\bar{X}} = s/\sqrt{n}$$

where $z_{\alpha/2}$ is found from the inverse standard normal

$$z_{\alpha/2} = \Phi^{-1}(\alpha/2).$$

A table of these values is found in most probability and statistic books, such as Devore [33]. In this experiment, a 99% confidence interval is used, and thus $z_{\alpha/2} = z_{0.05/2} = 2.58$.

The $100(1 - \alpha)\%$ confidence interval on the variance estimate is a bit more complicated, but it can be found [43] to be

$$[\underline{\sigma}^2, \bar{\sigma}^2] = \left[s^2 - z_{\alpha/2} \sqrt{\text{var}(s^2)}, s^2 + z_{\alpha/2} \sqrt{\text{var}(s^2)} \right]$$

$$\text{where: } \text{var}(s^2) = \frac{(n-1)^3}{n^4} (s^2)^2.$$

APPENDIX B: DETAILS OF EXPERIMENT

As discussed in the body of the paper, the experiment consists of the design of a pressure vessel by different design approaches. In this appendix, we present the details of this experiment, including assumptions and numerical values.

Goal:

The goal of this problem is to design a pressure vessel with volume of 0.15 m^3 that can sustain an internal gas pressure of 7 MPa such that the expected utility of the design is maximized subject to the following constraints. As noted in the body of the paper, the utility function is given by

$$Utility = P_{\text{selling}} - C_{\text{material}} \cdot \text{volume} - C_{\text{failure}}.$$

In this utility function, material cost is directly proportional to the volume of material used. This assumption is a simple aggregation of total cost of manufacture. Because the pressure vessel will be used in a location occupied by humans, the cost of failure is relatively high when the vessel fails.

Constraints and assumptions

- Total length of pressure vessel less than 1.5 meters
- Total height of the pressure vessel less than 0.6 meters
- Vessel internal pressure $P=7\text{MPa}$.
- Manufacture of the pressure vessel results in a uniform thickness t , and a constant material strength throughout
- It is assumed the pressure vessel will fail due to the stress in the walls. Therefore, the critical point for failure is when the maximum stress in the material of the walls exceeds the yield strength, or $\sigma_{\text{max}} > \sigma_y$.
- The maximum stress $\sigma_{\text{max}} = \max(\sigma_s, \sigma_c)$, where σ_c is the stress in the cylindrical section of the pressure vessel and σ_s is the stress in the spherical section. These quantities are given by the following approximations [44], which are meaningful for $R \geq 5t$:

$$\sigma_s = P \cdot [(R + 0.5t)/2t]$$

$$\sigma_c = P \cdot [(R + t)^2 + R^2]/(2Rt + t^2)$$

REFERENCES

- [1] Elishakoff, I., 2004, *Safety Factors and Reliability: Friends or Foes?* Kluwer Academic Publishers, Dordrecht, Netherlands.
- [2] Hazelrigg, G. A., 1998, "A Framework for Decision-Based Design," *Journal of Mechanical Design*, **120**(4), pp. 653-658.
- [3] Fernández, M. C., Seepersad, C. C., Rosen, D. W., Allen, J. K., and Mistree, F., 2001, "Utility-Based Decision Support for Selection in Engineering Design," *ASME DETC Design Automation Conference*, Pittsburgh, PA, pp. DETC2001/DAC-21106.
- [4] Scott, M. J., 2004, "Utility Methods in Engineering Design," *Engineering Design Reliability Handbook*, CRC Press, New York.
- [5] Thurston, D. L., 1990, "Multiattribute Utility Analysis in Design Management," *IEEE Transactions on Engineering Management*, **37**(4), pp. 296-301.
- [6] Bedford, T., and Cooke, R. M., 2001, *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge University Press, New York.
- [7] Mourelatos, Z. P., and Liang, J., 2004, "An Efficient Unified Approach for Reliability and Robustness in Engineering Design," *NSF Workshop on Reliable Engineering Computing*, Savannah, Georgia.
- [8] Byrne, D. M., and Taguchi, S., 1987, "The Taguchi Approach to Parameter Design," *Quality Congress*.
- [9] Taguchi, G., translated by Tung, L. W., 1987, *System of Experimental Design: Engineering Methods to Optimize Quality and Minimize Cost*, UNIPUB/Kraus International Publications, Dearborn, Michigan.
- [10] Bunn, D., 1984, *Applied Decision Analysis*, McGraw-Hill, New York.
- [11] Resnik, M., 1987, *Choices: An Introduction to Decision Theory*, University of Minnesota Press, Minneapolis.
- [12] Nikolaidis, E., 2005, "Types of Uncertainty in Design Decision Making," *Engineering Design Reliability Handbook*, Nikolaidis, E., Ghiocel, D. M., and Singhal, S., CRC Press, New York.
- [13] Ferson, S., and Ginzburg, L., 1996, "Different Methods Are Needed to Propagate Ignorance and Variability," *Reliability Engineering and System Safety*, **54**(2-3).
- [14] Hofer, E., 1996, "When to Separate Uncertainties and When Not to Separate," *Reliability Engineering & System Safety*, **54**(2-3), pp. 113-118.
- [15] Oberkampf, W. L., Helton, J. C., Joslyn, C. A., Wojtkiewicz, S. F., and Ferson, S., 2001, "Challenge Problems: Uncertainty in System Response Given Uncertain Parameters," <http://www.sandia.gov/epistemic/>.
- [16] Parry, G. W., 1996, "The Characterization of Uncertainty in Probabilistic Risk Assessment of Complex Systems," *Reliability Engineering and System Safety*, **54**(2-3), pp. 119-126.
- [17] Winkler, R. L., 1996, "Uncertainty in Probabilistic Risk Assessment," *Reliability Engineering & System Safety*, **54**(2-3), pp. 127-132.
- [18] Walley, J., 1991, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, New York.
- [19] 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., CRC Press, New York, pp. 9.1-9.40.
- [20] de Finetti, B., 1980, "Foresight. Its Logical Laws, Its Subjective Sources (Translated)," *Studies in Subjective Probability*, Kyburg, H. E. and Smokler, H. E., E. Krieger Publishing Company.
- [21] Hajek, A., 2003, Interpretations of Probability, <http://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=probability-interpret>.
- [22] Savage, L. J., 1972, *The Foundation of Statistics*, Dover Publications, New York.
- [23] Kolmogorov, A. N., translated by Morrison, N., 1956 (1933), *Foundations of Probability*, Chelsea Publishing Company, New York.
- [24] Pratt, J. W., Raiffa, H., and Schlaifer, R., 1995, *Introduction to Statistical Decision Theory*, The MIT Press, Cambridge, MA.
- [25] von Neumann, J., and Morgenstern, O., 1980, *Theory of Games and Economic Behavior, Third Edition*, Princeton University Press, Princeton, NJ.
- [26] 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.
- [27] 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>.
- [28] Shafer, G., 1976, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ.
- [29] Hofer, E., Kloos, M., Krzykacz-Hausmann, B., Peschke, J., and Woltereck, M., 2002, "An Approximate Epistemic Uncertainty Analysis Approach in the Presence of Epistemic and Aleatory Uncertainties," *Reliability Engineering & System Safety*, **77**(3), pp. 229-238.
- [30] Ferson, S., 2002, "Ramas Risk Calc," 4.0 ed., Lewis Publishers, New York.
- [31] 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>.
- [32] Ferson, S., Kreinovich, V., Ginzburg, L., Myers, D. S., and Sentz, K., 2002, "Constructing Probability Boxes and Dempster-Shafer Structures," SAND2002-4015, Sandia National Laboratories, Albuquerque, New Mexico, <http://www.sandia.gov/epistemic/Reports/SAND2002-4015.pdf>.

- [33] Devore, J. L., 1995, *Probability and Statistics for Engineering and the Sciences*, Duxbury Press, New York.
- [34] Lawrence, D. B., 1999, *The Economic Value of Information*, Springer-Verlag, New York.
- [35] Robert, C. P., and Casella, G., 1999, *Monte Carlo Statistical Methods*, Springer-Verlag, New York, NY.
- [36] Myers, R. H., and Montgomery, D. C., 1995, *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, John Wiley and Sons, New York.
- [37] Simpson, T. W., Peplinski, J. D., Koch, P. N., and Allen, J. K., 2001, "Metamodels for Computer-Based Engineering Design," *Engineering with Computers*, **17**(2), pp. 129-150.
- [38] Williamson, R. C., and Downs, T., 1990, "Probabilistic Arithmetic I: Numerical Methods for Calculating Convolutions and Dependency Bounds," *International Journal of Approximate Reasoning*, **4**, pp. 89-158.
- [39] Oberkampf, W. L., and Helton, J. C., 2002, "Investigation of Evidence Theory for Engineering Applications," *Non-Deterministic Approaches Forum, 43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Denver, CO.
- [40] Yager, R. R., Kacprzyk, J., and Fedrizzi, M., 1994, *Advances in the Dempster-Shafer Theory of Evidence*, John Wiley and Sons, New York.
- [41] Moore, R., 1979, *Methods and Applications of Interval Analysis*, Society for Industrial and Applied Mathematics, Philadelphia.
- [42] 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.
- [43] Weisstein, E. W., 2005, Sample Variance Distribution. From Mathworld--a Wolfram Web Resource, <http://mathworld.wolfram.com/SampleVarianceDistribution.html>.
- [44] Roark, R., and Young, W., 1975, *Formulas for Stress and Strain*, McGraw-Hill Book Co.