Proceedings of the ASME 2017 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference
IDETC/CIE 2017
August 6-9, 2017, Cleveland, Ohio, USA

DETC2017-68371

OPTIMIZATION UNDER UNCERTAINTY VERSUS ALGEBRAIC HEURISTICS: A RESEARCH METHOD FOR COMPARING COMPUTATIONAL DESIGN METHODS

William R. Binder
Georgia Institute of Technology
Atlanta, GA, 30332

Christiaan J.J. Paredis
Georgia Institute of Technology
Atlanta, GA, 30332

ABSTRACT

In this paper, we introduce a research method for comparing computational design methods. This research method addresses the challenge of measuring the difference in performance of different design methods in a way that is fair and unbiased with respect to differences in modeling abstraction, accuracy and uncertainty representation. The method can be used to identify the conditions under which each design method is most beneficial. To illustrate the research method, we compare two design methods for the design of a pressure vessel: 1) an algebraic approach, based on the ASME pressure vessel code, which accounts for uncertainty implicitly through safety factors, and 2) an optimization-based, expected-utility maximization approach which accounts for uncertainty explicitly. The computational experiments initially show that under some conditions the algebraic heuristic surprisingly outperforms the optimization-based approach. Further analysis reveals that an optimization-based approach does perform best as long as the designer applies good judgment during uncertainty elicitation. An ignorant or overly confident designer is better off using safety factors.

1 INTRODUCTION

Choosing a design method is an important step in any design. This choice ultimately influences both the design process and the final artifact. Ideally, a designer chooses a method that leads to the most desired design outcome. This outcome not only depends on the value of the final artifact but also on the time and cost invested during the design process [1]. Because achieving higher artifact value typically requires more design effort, designers must choose a design method that best balances artifact value and design process resources for a particular design situation. The situation is important because design methods perform differently in different contextual situations. For example, an electrical engineering design method is not likely to be as valuable as a chemical engineering design method in a chemical engineering situation. Because the value of the design method may be substantial, it is important to choose the most appropriate design method.

To choose the best design method, we must first be able to compare design methods. Real-world examples are quite limited for a variety of reasons. First, solving the same design problem with multiple design methods can take months or years and is thus prohibitively expensive. Second, comparing the value of two different final artifacts is difficult, again because of costs, but also because the two cannot fairly be placed in the same market environment without affecting each other. As a consequence, very few real-world comparisons are published in the literature, and, because context matters, any existing comparisons offer little help in valuing the design methods.

Instead, researchers may use computational models and simulations to compare design methods. Simulation is relatively inexpensive and allows thus for a broad comparison across many design problems and situations. However, comparing design methods with simulations still presents many challenges. One challenge is to evaluate the results of those different designs methods, the design artifacts and design process costs, in a fair and unbiased way. How should one compare two design methods when one suggests a vessel fails, while the other suggests the vessel does not fail? Clearly, if one design method is used as the predictive analysis model then this will bias the comparison. Instead, the design methods could be compared using an unbiased third design method. This is the basis for the Design Decision Framing Model (DDFM), which is used in the research method to compare design methods. The introduction of the DDFM is the primary contribution of this paper. It is further described in the methodology section of this paper and can be used to best select a design method for a given context.

To illustrate the DDFM, we compare two different design methods for the design of a pressure vessel, 1) an algebraic approach, based on the ASME pressure vessel code, which accounts for uncertainty implicitly through safety factors, and 2) an optimization-based, expected-utility maximization approach which accounts for uncertainty explicitly. The design problem and corresponding results are provided in the motivating example section. Finally, we present the results and conclusions of the research method and the motivating example.

2 RELATED WORK

Previous work can be divided into two areas: uncertainty analysis in design and value of information theory in design.

In design, uncertainty may be accounted for in many different ways. Most commonly, uncertainty is addressed implicitly using safety factors. Such safety factors may result from regulation or standardization, increasing with the size of the uncertainty and severity of the consequences of failure [2]. The challenge with safety factors is that they often lead to over-design because the safety factors must be prescribed conservatively to cover a broad set of situations.

Uncertainty may also be captured explicitly using probability theory [3]. Doing so has the advantage that it allows the designer to frame the design problem explicitly as a decision problem under uncertainty. A rational decision maker, who aims to act in a fashion that is consistent with his beliefs and preferences, would then choose the design alternative that maximizes his expected utility [4]. For instance, one could use an optimization algorithm to maximize an estimate of the expected utility determined using Monte Carlo simulation [5]. Compared to a safety-factor approach, modeling and solving such an optimization problem requires a lot more time and resources.

Rather than framing the design problem as expected utility maximization, Reliability Based Design Optimization (RBDO) frames it as a maximization of performance while meeting a specified reliability. In general, RBDO suffers from similar computational complexity challenges as expected utility maximization. Researchers have therefore proposed simplifications, such as the first-order reliability method [6] and the second-order reliability method [7], which define performance functions and approximate those performance functions using first and second order Taylor expansions, respectively. In these approaches, accuracy of the analysis is sacrificed in favor of reduced computation time. Care must be used when the failure region is not well approximated by a first or second order equation.

Some researchers have argued that probability theory does not apply in design when insufficient data is available to characterize the uncertainties [8-11]. However, such a conclusion is usually based on an outdated frequentist interpretation of probability, and no advantage over Bayesian, subjectivist probability theory, has been demonstrated.

In this paper, to compare the quality of different design methods, we build on an extension of decision theory, namely, value of information theory [12, 13]. Its aim is to determine the economic value of information and to provide guidance regarding the price one should be willing to pay to consult a source of information. Information is valuable only if it may change the decision maker's choice to a more valuable alternative. In design, the performance of an artifact is often predicted using models. These models serve as sources of information that help inform the designer in his decisions. Treated this way, the concepts of value of information theory can also apply to engineering models [14]. Valuing information enables decision makers to rationally choose whether to gather additional information until more valuable alternatives present themselves. To decide which actions to take, in [15], a method for conceptual design is prescribed in which the expected value of a refinement of the design space is compared with a value of information approach to specifying evaluation functions. Value of information theory has also been used to compare artifact refinement and analysis in design methods [16], allowing the decision maker to choose the next synthesis or analysis action based on the currently available information.

In summary, different design methods analyze and account for uncertainty differently. These methods lead to different design choices, but also require more or less time and effort to be applied, so that it is not clear when one method is superior to another. Using value of information theory, we propose in this paper a research method that allows us to compare different design methods fairly, so that researchers can focus on further improving the best methods for specific contexts.

3 METHODOLOGY

Design methods are rarely compared directly. This leads to the questions: If multiple design methods are available but suggest different actions, which one should one use? How should design methods be compared for a given context? These are the questions this paper attempts to answer. In this section, we describe a research method for comparing different design methods. The intended use of this research method is to characterize the performance of different design methods rigorously and to collect evidence in support of claims regarding the performance of design methods.

3.1 A Metric for Comparing Design Methods

The focus in this paper is on computational design methods that support design decision making. Each design method corresponds then to a particular way of framing a design decision. Tversky and Kahneman define a decision frame as "the decision-maker's conception of the acts, outcomes, and contingences associated with a particular choice" [17]. We define the design frame of a design method as consisting of a specific design space, set of modeling assumptions, and a search strategy used for decision making. The design space consists of the set of alternatives that will be evaluated using the set of modeling assumptions, as instructed by the search strategy. The outcomes of a design method are a chosen artifact and associated design process costs.

To compare design methods, we must also define the metric for comparison. This metric should be related to the goal of design, as described by Herbert Simon in his seminal work on the Sciences of the Artificial [18]: "Everyone designs who devises courses of action aimed at changing existing situations into preferred ones." This improvement in situations can be characterized mathematically by a value function. We adopt the perspective of a designer whose value function is profit, and more is better. The designer receives profit by designing an artifact and extracting value, perhaps by selling the artifact. However, the designer must first expend time and effort, design process costs, to design the artifact, whose value may not be realized for a considerable time. To account for the time difference between the value streams, a designer may use Net Present Value (NPV) [19, 20]. Further, because the value streams occur in the unknown future, the designer may be uncertain about the precise value of the NPV. This uncertainty presents itself as risk to the designer. To account for risk preferences, rational decision makers consider the expectation, E, of utility as their metric of value [4]. In this case, utility, U, is the value to the decision maker derives from the sale of an artifact. The objective of a design is then to maximize the expected utility of the NPV of the artifact and design process costs:

$$\mathcal{A}: \max_{a \in A} \mathrm{E}\left[U\left(NPV\left(a, C_{DP}(\mathcal{A})\right)\right)\right] \tag{1}$$

where a is an artifact from the set of all possible artifacts, A, and C_{DP} is the design process cost. However, we need to reframe Equation (1) to consider the effect of design methods. A design method is a sequence of design actions for selecting the artifact. Designers search for the best design method, which will have an appropriate tradeoff between the value of the artifact and the design process costs:

$$\mathcal{P}: \max_{p \in P} \mathrm{E}\left[U\left(NPV\left(a(p), C_{DP}(p)\right)\right)\right] \tag{2}$$

where P is the set of all possible design methods. We now have a metric, the expected utility of the NPV, for valuing design methods for a specific contextual situation. This metric, and the value of a design method, depends on how much the design method benefits the designer. That is, the value of a design method depends on how much better it performs as compared to the alternative design methods, the value is relative.

So far we have ignored the fact that the value of a design method depends on other factors that are included in the contextual situation. We need a metric for comparing across a range of different contextual situations, a context. A context may include properties of the artifact not specified by a, information about a competitor's product, beliefs concerning material properties, etc. Design methods may perform differently in different contextual situations if either design method's outcomes change as a result. For example, a design method that has many steps and requires a large quantity of time and effort may be inappropriate if there is limited time as the design method will be too costly or yield an artifact of low value. Our metric for comparing design methods must then also change as a range of contextual situations is considered. Assuming each contextual situation is equally likely, the metric for comparing design methods then becomes the average expected utility. If a design method has the greatest average expected utility for a given context, it is preferred for this context. We call such a context the applicability context, as this design method should be used if a designer's contextual situation is within this context.

3.2 A Method to Compute the Value of Design Methods

To overcome the challenges of evaluating design methods in the real world, we propose to use computational experiments to evaluate design methods. Instead, researchers may use computational experiments to evaluate design methods. Unlike real world experiments, computational experiments can repeat the exact same contextual situation, allowing for the expected utility to be calculated. Since the cost of computational experiments is low compared to real world experiments, the entire context can be explored to determine the average expected utility. Furthermore, computational experiments can evaluate design methods without the design methods mutually influencing each other.

However, computational experiments introduce their own challenges. While real world experiments are evaluated by actual outcomes, computational experiments must use models of reality to predict the performance of a design method. But this presents a problem if design methods utilize different assumptions or models. For example, consider two design methods that recommend two different minimum thicknesses of a pressure vessel to avoid failure under the same load. For a given contextual situation, there is only one minimum thickness to avoid failure. Therefore, one or both of the design methods must be incorrect in determining the minimum thickness to avoid failure.

In order to evaluate design methods fairly, we must be able to evaluate their outcomes using similar assumptions. However, if one of the design methods were used to evaluate the outcomes, there would be a clear bias towards this design method. This problem arises in [21], where the concept of imprecise probabilities is considered and compared to a probabilistic characterization of uncertainty. A motivating example of a pressure vessel is used to show how both uncertainty representations affect the value of a decision. To compare the two methods, we introduce an "omniscient supervisor" who knows the (artificially generated) truth, controls how much of this truth is revealed to the designer, and determines the value of artifacts and process costs resulting from each design method, so that the methods can be objectively compared. The omniscient supervisor removes the bias from choosing a particular design method's computation of the artifact value.

The above ideas form the basis of the Design Decision Framing Method (DDFM). Figure 1 shows the structure of the DDFM. To evaluate design methods, the DDFM considers the context, allows decision makers to use the design methods in developing the final artifacts, uses the omniscient supervisor to evaluate the "true" value of each artifact, and then combines this value with the associated costs of the design process to determine the value of the design process for each design method. Because no one can truly know the "truth," researchers can then further reduce bias by investigating different values of the truth. In

Figure 1, this is equivalent to considering different contextual situations. The value of the design method is then averaged over these different contextual situations, including different truths. The DDFM thus provides a fair and unbiased computational framework for evaluating different design methods.

The DDFM is limited by the assumptions that it relies upon. For one, future decisions need to be either ignored or straightforward such that they can be evaluated by the omniscient supervisor. If the decisions are complex, the research method may require substantial computational resources. The DDFM also assumes that the values and model of the truth reasonably approximate reality. These assumptions are a reasonable approximation to applying the design methods in practice while being executed at a comparably smaller cost that is suitable for academic research.

3.2.1 Computational Considerations

Computational tools are considered to make the application of the DDFM computationally tractable. These tools are used in a computational experiment to compare different design methods, and further described in the motivating example section. The experiment focuses on comparing the two design methods over a range of contextual situations to determine which is preferred.

A design method is preferred for a context if its average expected utility is greater than the alternative. For a range of contextual situations, the omniscient supervisor can determine average expected utility, $\overline{U^\dagger}$, of the NPV based on the outcomes of the decision maker's search for an artifact, p, and the particular contextual situation, x, from the set of possible contextual situations, X, and with a probability density function, $f(x)$:

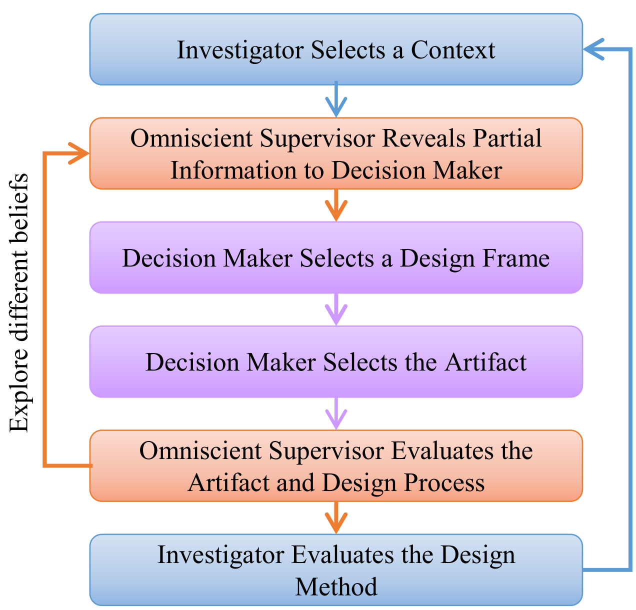

Figure 1. The Design Decision Framing Method. The method provides an unbiased estimate of the value of design methods.

$$\overline{U^\dagger} = \int_{x \in X} f(x) U^\dagger(p, x)\, dx \tag{3}$$

where the † denotes that the function is evaluated by the omniscient observer, as compared to a decision maker. Equation (3) requires that we know truth, and therefore the future. However, the future is unknowable, and we therefore do not know the truth. In Equation (2), this is accounted for in the expected utility since the truth is considered a part of the contextual situation, thus considering a range of truths as well.

The average utility could then be approximated by computing the expectation with the Monte Carlo Method [5, 22]. This approximation discretizes Equation (3) and the accuracy depends on the number of samples, m:

$$\overline{U^\dagger} \cong \frac{1}{m} \sum_{i=1}^{m} U^\dagger(p, x_i) \tag{4}$$

where we replace the probability density function with $1/m$ since we are using a simple average. But we are not only interested in computing one average utility, as we are comparing potentially many different design methods.

To improve the accuracy of the expectation estimate, we use common random numbers [23, 24]. By using the same samples to determine both estimates, the variance of the difference in estimates is reduced when the estimates are correlated [22, 24]. The difference in expected utility, $\Delta\overline{U}$, is thus estimated as:

$$\Delta\overline{U} \cong \frac{1}{m} \sum_{i=1}^{m} U^\dagger(p_1, x_i) - U^\dagger(p_2, x_i) \tag{5}$$

Where p_1 and p_2 refer to the first and second design methods, respectively. For this paper, p_1 refers to the optimization design method, and p_2 refers to the algebraic design method. When $\Delta\overline{U}$ is positive, the optimization heuristic is preferred, when negative, the algebraic heuristic. These computational tools allow for a computationally efficient determination of which design method is preferred for a given context.

To characterize the design methods, we investigate sets of contextual situations. By comparing the design methods over a range, we can identify how much the variables influence the preference of each design method. The range of values of the variables will be the context for the motivating example, in addition to variables that characterize the unknown truth. In this way we can analyze, for example, how a designer's beliefs of the truth can affect the most preferred design method.

4 MOTIVATING EXAMPLE

To explore the DDFM, an example study of a pressure vessel is investigated, based on the problem introduced in [16]. The example is centered on a designer who must select a value for the wall thickness of a pressure vessel with otherwise predetermined geometry and dimensions. In this case, the decision maker is a seller of pressure vessels, and receives revenue for each pressure vessel, but incurs a cost for pressure vessels that fail prematurely as well as the manufacturing costs of the pressure vessels. The business averages approximately \$3.3 million in revenues per year. The nominal pressure, p, for this vessel operates at 1.4 MPa, with a radius, r, of 0.2286 m, and a length, L, of 1.2 m. A selling price, P_s, of \$415 per pressure vessel was assumed for this nominal case, which correlates to the number of vessels sold, $n = 8{,}000$ vessels. The decision maker must deal with uncertainty in the material's ultimate strength, choosing the thickness that maximizes his profit. It is assumed that each pressure vessel's material is randomly drawn from a normal distribution, modeling different qualities of material one would expect from a vendor's batch. To select the thickness, two design methods are considered. To simplify the comparison between the two design methods, we restrict our focus to one year of revenues. The first design method is an algebraic method, based on the ASME pressure vessel code for thin walled pressure vessels. The second design method is an optimization method, based on value-driven design, utility theory, and thin walled pressure vessel assumptions.

4.1 Algebraic Design Method

The algebraic design method is a simplified form of [25], which calculates two thicknesses and selects the more conservative one, i.e., the thicker one:

$$S = \sigma_{ts}/N \tag{6}$$

$$t_a = p * r/(S * E - 0.6 * p) \tag{7}$$

$$t_b = p * r/(2 * S * E + 0.4 * p) \tag{8}$$

$$t_{req} = max(t_a, t_b) \tag{9}$$

where σ_{ts} is the ultimate tensile strength of the material, N is the factor of safety, t_a is the minimum required thickness at longitudinal seam welds, t_b is the minimum required thickness at circular seam welds, p is the internal pressure, r is the radius of the spherical ends of the pressure vessel, E is the weld efficiency of the seams, and t_{req} is the minimum required thickness for the pressure vessel. A weld efficiency of 1 was assumed for all calculations. A factor of safety of 3.5 was used as per ASME standards when using the ultimate strength to determine the minimum required thickness [25].

4.2 Optimization Design Method

The optimization design method chooses the thickness that maximizes the expected utility of the decision maker. The decision maker's utility depends on the profitability of the business, $Profit(n_f, t)$, which depends on the number of vessels sold, the material cost, and the failure cost:

$$V(t) = \frac{4}{3}\pi * (r^3 - (r - t)^3) + \pi * L * (r^2 - (r - t)^2) \tag{10}$$

$$C_m(t) = P_m * V(t) \tag{11}$$

$$C_f(n_f) = P_f * n_f \tag{12}$$

$$Profit(n_f, t) = n * (P_s - C_m(t)) - C_f(n_f). \tag{13}$$

where V is the volume of material per vessel, C_m is the cost of materials per vessel, P_m is the material cost, t is the thickness of the material, C_f is the cost incurred from failed vessels, P_f is the per unit failure cost, n_f is the number of vessels that fail, n is the number of pressure vessels sold, and P_s is the price of each pressure vessel sold. To determine the probability of a particular pressure vessel failing, Pr_f, the decision maker expresses his beliefs about the ultimate strength.

The decision maker forms his beliefs using strength tests of the material. The decision maker evaluates the mean, $\bar{x}_{\sigma_{ts}}$, sample standard deviation, $S_{\sigma_{ts}}$, and degrees of freedom, ν, of the observed ultimate strength values and characterizes his beliefs using the Student's t-distribution. Then, the probability of a particular pressure vessel failing is the probability that the ultimate strength is less than the peak stress: the t-distribution's cumulative distribution function (CDF), F_ν, evaluated at the peak stress. The decision maker assumes a thin walled pressure vessel, and so evaluates the CDF at the hoop stress, σ_h [26]:

$$\sigma_h(t) = \frac{p * r}{t} \tag{14}$$

$$Pr_f(t) = F_\nu(\frac{\sigma_h(t) - \bar{x}_{\sigma_{ts}}}{S_{\sigma_{ts}}}) \tag{15}$$

The probability of a given number of failures is described by the probability mass function (pmf) of the binomial distribution:

$$B(n_f; n, Pr_f(t)) = \binom{n}{n_f} Pr_f^{\,n_f}(1 - Pr_f)^{n-n_f} \tag{16}$$

Given the probability of each possible failure, the expected profit, $\mathrm{E}[P(t)]$, is equal to:

$$\mathrm{E}[P(t)] = \sum_{i=0}^{n} B(i; n, Pr_f(t)) * Profit(i, t) \tag{17}$$

Since the decision maker is risk averse, he does not want to make a decision based on the expected profit, but rather choose the thickness which maximizes his expected utility. Although any monotonically increasing function can be used, we use the following equation for utility, which assumes a constant risk tolerance, R:

$$Utility(n_f, t) = R * \left(1 - e^{\frac{-Profit(n_f,t)}{R}}\right). \tag{18}$$

For businesses, a ratio of risk tolerance to sales of 0.064 has been commonly measured [27]. Given that sales are approximately \$3.3 million, a risk tolerance of \$212,000 is assumed.

Risk preferences should be considered when a design contains uncertainty. In this case, the decision maker is uncertain about the material's ultimate strength. Thus, the design method prescribes that the expected utility be maximized:

$$\mathrm{E}[U(t)] = \sum_{i=0}^{n} B\big(i; n, Pr_f(t)\big) * Utility(i,t) \tag{19}$$

with optimal thickness, t^*:

$$t^* = \arg\max_{t \in T} \mathrm{E}[U(t)] \tag{20}$$

where T contains all positive real numbers less than the radius.

4.2.1 Expert Decision Maker

Experts are able to update their beliefs based on past experience. Equation (15) gives the probability of failure assuming that the stress is adequately described by the thin-walled pressure vessel assumptions. However, given prior experience, an expert decision maker may recognize that Equation (15) does not properly reflect the probability of failure. In order to more accurately represent the probability of failure, the expert decision maker introduces the random variable, δ, which is used as a multiplicative term on the hoop stress. Then, the decision maker's probability of a particular pressure vessel failing is the expected value of the probability of failure, including the probability density function (PDF), f, of δ:

$$Pr_f(t) = \int_{-\infty}^{\infty} F_v\Big(\frac{\sigma_h(t) * \delta - \bar{x}_{\sigma_{ts}}}{S_{\sigma_{ts}}}\Big) * f(\delta) d\delta. \tag{21}$$

To compare the effect of experience, we will test the performance of the optimization design method using both Equations (15) and (21). The non-expert optimization design method uses Equations (10)-(20) while the expert optimization design method uses Equations (10)-(14) and (16)-(21) to select the thickness of the pressure vessels.

4.3 Omniscient Supervisor

The omniscient observer must evaluate the value of the two design methods by evaluating the pressure vessels. Therefore, the omniscient observer requires its own set of assumptions to determine failure: its own design method. To minimize bias, the omniscient supervisor should not use either of the design methods being investigated. The omniscient observer's design method is similar to the optimization design method, but is made more conservative by calculating the Von-Mises stress from the tangential, radial, and longitudinal stress to reduce potential bias. The equations for the tangential, radial, and longitudinal stress are calculated in the more general thick walled pressure vessel case [26]:

$$\sigma_t = \frac{p * (r^2 + (r-t)^2)}{r^2 - (r-t)^2} \tag{22}$$

$$\sigma_r = -p \tag{23}$$

$$\sigma_z = \frac{p * (r-t)^2}{r^2 - (r-t)^2} \tag{24}$$

$$\sigma_v = \sqrt{{\sigma_t}^2 + {\sigma_r}^2 + {\sigma_z}^2}. \tag{25}$$

where σ_t, σ_r, σ_z, and σ_v are the tangential stress, radial stress, longitudinal stress, and Von-Mises stress, respectively. In this case, the omniscient supervisor compares the Von-Mises stress to the ultimate strength to determine the true probability of failure, ${Pr_f}^\dagger$, based on the ultimate strength's true mean, $\mu_{\sigma_{ts}}$, and variance, $var_{\sigma_{ts}}$:

$${Pr_f}^\dagger(t) = F_v\Big(\frac{\sigma_v(t) - \mu_{\sigma_{ts}}}{\sqrt{var_{\sigma_{ts}}}}\Big). \tag{26}$$

Here the ultimate strength's true mean and standard deviation are used, as the omniscient supervisor has no uncertainty concerning the distribution of ultimate strengths. The omniscient supervisor evaluates the probability of failure based on the normal distribution's CDF, which reflects the true distribution of ultimate strengths.

In addition to determining the probability of failure, the omniscient supervisor also accounts for the design process costs. The omniscient supervisor explicitly considers the additional computational cost, C_O, based on the amount of time, $time$, necessary to determine t^* for a given contextual situation, x_i, to determine the omniscient supervisor's evaluation of profit:

$$C_O = P_O * time(x_i) \tag{27}$$

$$Profit^\dagger(n_f, t) = n\big(P_s - C_m(t)\big) - C_f(n_f) - C_O. \tag{28}$$

Thus, for a particular contextual situation and a given thickness, the omniscient supervisor can determine the expected utility, $E[U^\dagger(t)]$:

$$Utility^\dagger(n_f, t) = R * \left(1 - e^{\frac{-Profit^\dagger(n_f,t)}{R}}\right). \tag{29}$$

$$\mathrm{E}[U^\dagger(t)] = \sum_{i=0}^{n} B\big(i; n, {Pr_f}^\dagger(t)\big) * Utility^\dagger(i,t). \tag{30}$$

The omniscient supervisor uses Equations (22)-(30) to evaluate each of the design methods: the algebraic, non-expert optimization, and expert optimization. By comparing the expected utility for each design method, the most preferred design method can be determined for a given contextual situation. However, we are concerned about more than one particular contextual situation.

In this computational experiment, we aim to determine which method is best across a given context—the range of contextual situations in which the design methods are more preferred than the other. Table 1 shows the variables and their bounds delineating the context. The variable, r, is the radius of the spherical ends of the cylindrical pressure vessel. The rated pressure, p, is the specified internal pressure of the pressure vessel relative to the external pressure. The length, L, is the length of the cylindrical portion of the pressure vessel, excluding the spherical caps. To allow for a fair comparison between the methods, the market price is assumed independent of the method, but varying across the contextual situations. It is determined by applying a profit margin, PM, applied to the costs associated with the optimal thickness, $t^\dagger$, as determined by the omniscient supervisor under perfect information:

$$P_s = \frac{{Pr_f}^\dagger(t^\dagger)P_f + C_m(t^\dagger)}{1 - PM} \tag{31}$$

Because only material and failure costs are considered, the lower and upper bounds of the profit margin are chosen to be quite high so both design methods result in reasonable profits. The per unit material cost, P_m, is the cost of a cubic meter of steel. The per unit failure cost, P_f, includes payment for expected damage to property and nearby individuals. The per unit computing cost, P_O, is included to account for the additional cost of the optimization design method over the algebraic design method. The costs for model and code development are ignored for both design methods, and the time to run the algebraic design method is considered negligible. The optimization cost is based on the average cost of using Amazon's EC2 m3.large and m3.2xlarge on demand computing services using Windows [28]. The true mean of the ultimate strength, $\mu_{\sigma_{ts}}$, and the true standard deviation of the ultimate strength, $\sigma_{\sigma_{ts}}$, are characteristics of the distribution of ultimate strengths in delivered steel. The values, and choice of normal distribution, are those of delivered high grade streel [29]. The decision maker is given information about this distribution in the form of strength tests that reveal the ultimate strength of a plate of steel, randomly selected from the true distribution. The number of strength tests, n_{st}, reflects the amount of information available to the decision maker. The strength tests are samples of the ultimate strength from the true distribution. The decision maker then uses those samples and determines the sample mean, $\bar{x}_{\sigma_{ts}}$, sample standard deviation, $S_{\sigma_{ts}}$, and the number of degrees of freedom, ν, to use as his beliefs in Equation (15) for the non-expert optimization design method, and Equation (16) for the expert optimization design method. The algebraic design method uses only the sample mean as the ultimate strength in Equation (6).

Table 1. Variables in the context

	Symbol	Lower Bound	Upper Bound
Pressure vessel radius [m]	r	0.183	0.274
Pressure vessel length [m]	L	0.96	1.44
Rated pressure [MPa]	p	11.2	16.8
Profit margin [%]	PM	64.8	73.4
Material cost [\$/m^3]	P_m	4,040	6,060
Failure cost [\$/vessel]	P_f	80,000	120,000
Computing cost [\$/hour]	P_O	0.518	0.777
True mean of the ultimate strength [MPa]	$\mu_{\sigma_{ts}}$	340	510
True variance of the ultimate strength [MPa]	$var_{\sigma_{ts}}$	16.2	24.4
Number of strength tests [samples]	n_{st}	24	36

To gain a deeper understanding of the difference in performance between the two design methods, we performed a 10-level full factorial across the number of strength tests and the failure cost. We focus on these two variables because they emphasize the difference between the two design methods: implicitly versus explicitly accounting for uncertainty and risk. The number of strength tests determines the amount of information available to the decision maker and hence his or her uncertainty about the material properties while the failure cost strongly influences the risk faced by the designer.

Finally, the experiments are performed for both a risk neutral and a risk averse decision maker to investigate the effect of risk preferences on the choice of design methods.

5 RESULTS

We first focus on the comparison for the non-expert optimization design method and the algebraic design method. Figure 2 shows the difference in expected profit between the non-expert optimization and algebraic design methods for the risk neutral and risk averse cases. It is important to note that positive values refer to the non-expert optimization design method being preferred. The algebraic design method is only preferred in two regions for the risk neutral case. The first region is the area at the top of the figure, where the non-expert is well informed about the material strength. Counterintuitively, this suggests that a non-expert decision maker with available information may actually be worse off by using the optimization design method. As available information increases, uncertainty about the material strength decreases. Thus, to maintain a similar probability of failure, the thickness of the pressure vessel can be decreased. However, the non-expert optimization design method incorrectly characterizes the probability of failure as compared to the omniscient supervisor. By assuming that Equation (14) is a fair approximation of the max stress in the pressure vessel, the non-expert is too aggressive in sizing the pressure vessel, recommending overly thin pressure vessels. Aggressive sizing of the pressure vessels also occurs for the risk averse case, but the risk aversion causes the design method to be more conservative even in the high available information region. For this region, the uncertainty in the material strength is dwarfed by the error in judgment introduced by Equation (14). Thus, even with perfect information, the algebraic design method may be preferred for certain per unit failure costs. If the optimization design method more accurately modeled the probability of failure, additional information would benefit the optimization design method. This is confirmed by Figure 3, which shows the comparisons of the expert optimization design method and the algebraic design method for the risk neutral and risk averse cases. Now, the expert optimization design method is always more preferred as additional information becomes available.

The results of Figure 3 also show that the expert optimization design method is preferred more as the per unit failure cost decreases. This is true even when considering risk aversion. The slope of the difference in expected profit as a function of available information decreases as compared to the risk neutral case. Conversely, the slope increases as a function of the per unit failure cost., is a maximum at the lowest per unit failure cost, where available information has little influence on the difference in expected profit. In the extreme, if the per unit failure cost was zero, then the problem becomes a minimization of material cost: a minimization of thickness, where knowledge of the material strength is irrelevant.

The algebraic design method recommends a thickness independent of information or failure cost. With sufficient Monte Carlo samples, the average thickness recommended by the algebraic design method is related to the true mean strength of the material, and thus approximately constant at 28 mm. The algebraic design method implicitly accounts for the probability of failure and the per unit failure cost in Equations (6) through (9), by using safety factors. If these implicit assumptions do not match the contextual situation well, the design method may not perform well, as is the general case. Typically, the algebraic design method is too conservative, recommending a pressure vessel that is too thick. Figure 4 shows the average thickness recommended by the expert optimization design method for the risk neutral and risk averse cases. Except in the region to the lower right, the thickness is less than the algebraic design method's thickness of 28 mm. While being excessively conservative decreases the average performance of the design method by increasing the material cost, it prevents extremely poor performance. This allows the algebraic design method to outperform the non-expert and expert optimization design methods when available information is limited and the per unit failure cost is high.

In general, the optimization design method performs better than the algebraic design method, but can suffer from occasional poor performance. This is readily seen from Figure 3 in the lower right region. The cause for the poor performance is twofold: the optimization design method is overly conservative, and a series of failed pressure vessels increases the average failure cost. The optimization design method recommends on average an overly thick pressure vessel, but will occasionally recommend an overly thin pressure vessel. First, it can be easily seen from Figure 4 that the average thickness in the lower right is greater than that of the algebraic design method for both the risk neutral and risk averse cases. This occurs because the amount of information is so limited that the decision maker in the optimization design method must consider extremely low material strengths as reasonably likely. In turn, the optimization design method recommends very thick pressure vessels to reduce the probability of failure. The expected thickness tends to decrease with increasing strength test samples as the decision maker becomes more confident in the shape of the underlying distribution. Said another way, because the number of degrees of freedom increases, the sample standard deviation tends to decrease, and the decision maker no longer needs to consider very low ultimate strengths as reasonably likely. Second, Figure 5 shows the average failure cost for the expert optimization design method, with a quickly rising peak occurring in the lower right region. This occurs because of the relatively limited number of strength test samples. While on average the expected thickness is greatest in this region, there are cases where a cluster of high strength test samples misleads the decision maker into believing the ultimate strength has a comparatively higher mean and lower standard deviation than the true distribution would suggest. This belief will mislead a rational decision maker into choosing a very thin pressure vessel, and result in a high number of failures. The fewer strength test samples are taken, the more likely this type of event can occur. Thus, despite the thickness being greater on average, the occasional thin pressure vessel greatly increases the expected costs associated with failure. This phenomenon also

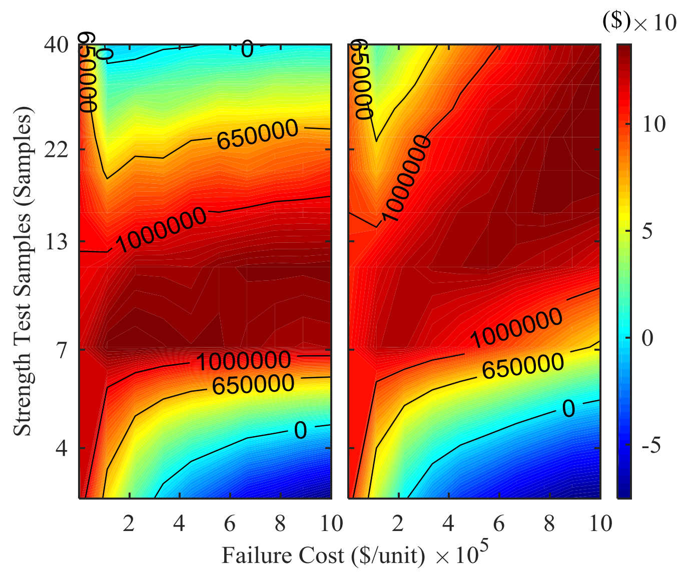

Figure 2. Contour plots of the difference in expected profits between the non-expert optimization and algebraic design methods for the risk neutral (left) and risk averse (right) cases. Positive values indicate preference for the non-expert optimization design method in the risk neutral case.

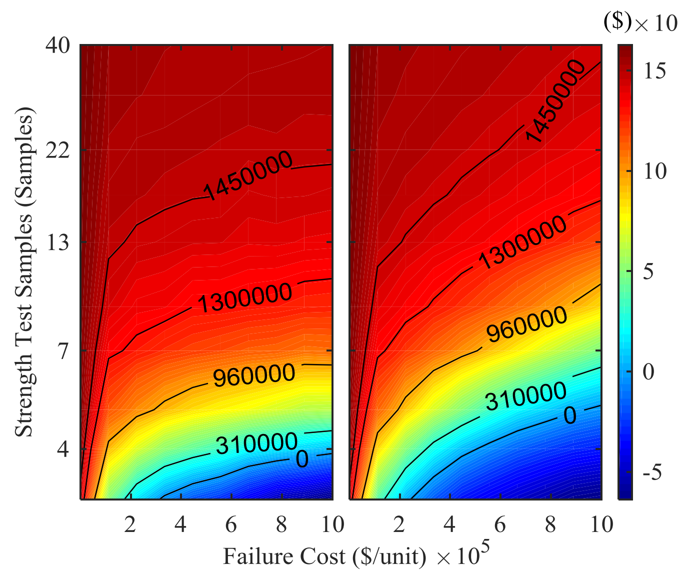

Figure 3. Contour plots of the difference in expected profits between the expert optimization and algebraic design methods for the risk neutral (left) and risk averse (right) cases. Positive values indicate preference for the non-expert optimization design method in the risk neutral case.

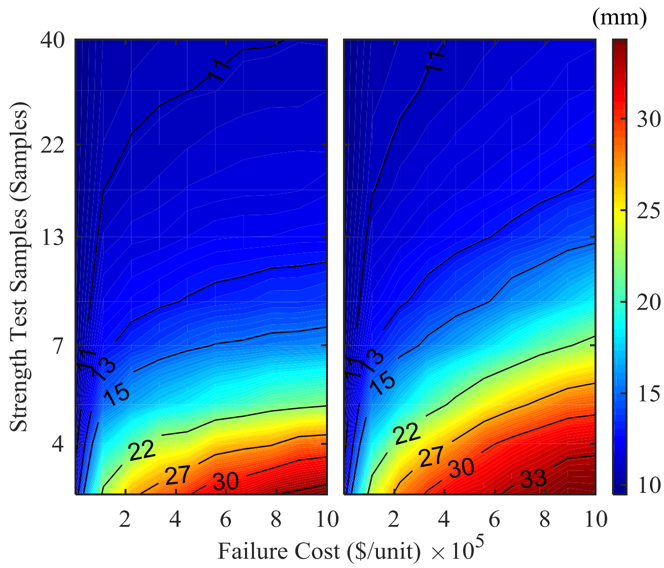

Figure 4. Contour plots of the expected thickness for the expert optimization design method for the risk neutral (left) and risk averse (right) cases.

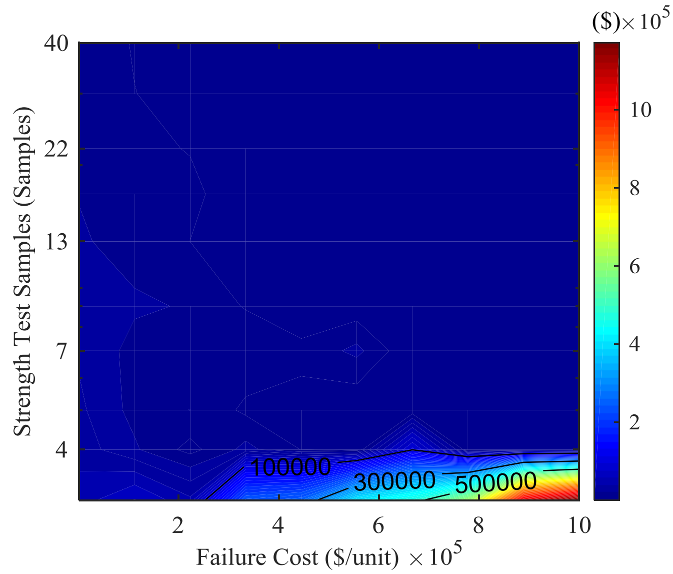

Figure 5. A contour plot of the expected total costs associated with failure for the expert optimization design method in the risk averse case.

appears for the algebraic design method, although not to the same degree because of the safety factor. As a result, expected costs associated with failure for the algebraic design method are extremely close to zero for all contextual situations investigated. Clustering never benefits the decision maker, as even when strength test samples group on the lower end of the strength tests, the decision maker will be more conservative and choose a greater thickness than is truly necessary. Any deviation between the decision maker's beliefs and the true distribution of the ultimate strength has a negative effect on profitability for both design methods, but especially so for the optimization design method, and even more so in cases where the ultimate strength is overestimated.

6 CONCLUSIONS

Two design methods for the design of pressure vessels have been investigated. The design methods were compared using a design decision framing method to determine each design method's relative value for different contextual situations. The design decision framing method uses an omniscient supervisor who provides a fair and unbiased comparison of the design methods. Comparing the design methods, the expert optimization design method was preferred over a larger applicability context. However, the algebraic design method was preferred in some regions, in particular when compared to the non-expert design method.

The expert optimization design method outperformed the non-expert optimization design method by using good judgment. In this case, the expert used his experience to update his beliefs about the probability of failure to approximate reality more accurately. Designers that ensure their models and methods are accurate representations of reality benefit over those who do not. However, even expert designers who exercise good judgment may be misled in special cases where available information is limited, as was the case for limited available information and a high failure cost. This emphasizes the need for good designers to gain sufficient information to be confident in their knowledge and to use applicable methods.

A design method should only be used in its applicability context, and care must be exercised when the applicability context is not explicit. One example is the algebraic design method, which makes assumptions for the failure cost and the material cost. In regions of low failure cost, the algebraic design method performed very poorly as compared to both the non-expert and expert optimization design methods.

There are limitations to the design decision framing method. While the DDFM is based on normative decision theory and information theory, additional tests and investigations are necessary to prove it is generally applicable to all design methods that recommend actions. Also, in our motivating example, we did not consider any future actions that would depend on the chosen alternative. Including such actions would greatly increase the computational complexity. It is also necessary for the compared design methods to be similar, that is, to recommend similar types of actions. In our motivating example, we only investigated design methods that recommended a pressure vessel's thickness. If the two design methods recommend dissimilar actions, they cannot be compared in a straightforward manner. One way to compare dissimilar design methods is to consider sets of design methods that perform similar functions when used together. In addition to the above limitations, the design decision framing method also requires a design method to act as the omniscient supervisor. If this design method is not a good model of reality, then the results will similarly not be appropriate.

In future work, we will compare design methods to analyze and design flexible systems, systems that may modify their configuration because of the changing environment. Currently these systems are analyzed using a real options analysis and many new design methods are being proposed. The design decision framing method may be able to characterize such design methods and identify the conditions under which a given design method should be used.

ACKNOWLEDGMENTS

This research was in part supported by Idaho National Laboratory under contract #00131206, and by the G.W. Woodruff School of Mechanical Engineering at Georgia Tech. The research presented in this paper was also supported in part by the National Science Foundation under Award CMMI-1645316. The United States Government retains, and by accepting the article for publication, the publisher acknowledges that the United States Government retains, a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for United States Government purposes.

REFERENCES

[1] Lee, B. D., and Paredis, C. J. J., 2014, "A Conceptual Framework for Value-Driven Design and Systems Engineering," *24th CIRP Design Conference*, Elsevier, Milan.

[2] Bristow, J. W., and Irving, P., 2007, "Safety Factors in Civil Aircraft Design Requirements," *Engineering Failure Analysis*, **14**(3), pp. 459-470.

[3] Durrett, R., 2010, *Probability: Theory and Examples*, Cambridge university press.

[4] von Neumann, J., and Morgenstern, O., 1980 (1944), *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, NJ.

[5] Rubinstein, R. Y., and Kroese, D. P., 2011, *Simulation and the Monte Carlo Method*, John Wiley & Sons.

[6] Der Kiureghian, A., 2005, "First-and Second-Order Reliability Methods," *Engineering design reliability handbook*, pp. 14-11.

[7] Der Kiureghian, A., Lin, H.-Z., and Hwang, S.-J., 1987, "Second-Order Reliability Approximations," *Journal of Engineering mechanics*, **113**(8), pp. 1208-1225.

[8] Ben-Haim, Y., 2001, "Common Knowledge, Coherent Uncertainties and Consensus," Technical Report ETR-2001-03, Israel Institute of Technology

[9] 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.

[10] Liu, B., 2007, "Uncertainty Theory," *Uncertainty Theory*, Springer, pp. 205-234.

[11] Bellman, R. E., and Zadeh, L. A., 1970, "Decision-Making in a Fuzzy Environment," *Management Science*, **17**(4), pp. B-141-B-164.

[12] Lawrence, D. B., 1999, *The Economic Value of Information* Springer.

[13] Howard, R. A., 1966, "Information Value Theory," *Systems Science and Cybernetics, IEEE Transactions on*, **2**(1), pp. 22-26.

[14] Nickerson, R. C., and Boyd, D. W., 1980, "The Use and Value of Models in Decision Analysis," *Operations Research*, **28**(1), pp. 139-155.

[15] Wood, W. H., and Agogino, A. M., 2005, "Decision-Based Conceptual Design: Modeling and Navigation of Heterogeneous Design Spaces," *Journal of Mechanical Design*, **127**(1), pp. 2-11.

[16] Thompson, S. C., and Paredis, C. J., 2010, "An Investigation into the Decision Analysis of Design Process Decisions," *Journal of Mechanical Design*, **132**(12), pp. 121009.

[17] Tversky, A., and Kahneman, D., 1981, "The Framing of Decisions and the Psychology of Choice," *Science*, **211**(4481), pp. 453-458.

[18] Simon, H. A., 1996, *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

[19] Coller, M., and Williams, M. B., 1999, "Eliciting Individual Discount Rates," *Experimental Economics*, **2**(2), pp. 107-127.

[20] Shrieves, R. E., and Wachowicz Jr, J. M., 2001, "Free Cash Flow (FCF), Economic Value Added (EVA™), and Net Present Value (NPV):. A Reconciliation of Variations of Discounted-Cash-Flow (DCF) Valuation," *The Engineering Economist*, **46**(1), pp. 33-52.

[21] Aughenbaugh, J. M., and Paredis, C. J. J., 2006, "The Value of Using Imprecise Probabilities in Engineering Design," *Journal of Mechanical Design*, **128**(4), pp. 969-979.

[22] Fishman, G., 2013, *Monte Carlo: Concepts, Algorithms, and Applications*, Springer Science & Business Media.

[23] Gal, S., Rubinstein, R. Y., and Ziv, A., 1984, "On the Optimality and Efficiency of Common Random Numbers," *Mathematics and Computers in Simulation*, **26**(6), pp. 502-512.

[24] Kleinman, N. L., Spall, J. C., and Naiman, D. Q., 1999, "Simulation-Based Optimization with Stochastic Approximation Using Common Random Numbers," *Management Science*, **45**(11), pp. 1570-1578.

[25] ASME, 2007, " Boiler and Pressure Vessel Code Section VIII, Division 1,"Rules for Construction of Pressure Vessels,"" *American Society of Mechanical Engineers, New York, New York*.

[26] Shigley, J. E., and Mischke, C. R., 2001, *Mechanical Engineering Design*, McGraw-Hill, New York, NY.

[27] Howard, R. A., 1988, "Decision Analysis: Practice and Promise," *Management Science*, **34**(6), pp. 679-695.

[28] Amazon, Amazon Ec2 Pricing, https://aws.amazon.com/ec2/pricing/.

[29] Hess, P. E., Bruchman, D., Assakkaf, I. A., and Ayyub, B. M., 2002, "Uncertainties in Material and Geometric Strength and Load Variables," *Naval engineers journal*, **114**(2), pp. 139-166.

