Risk Measures in Engineering Design under Uncertainty

R. Tyrrell Rockafellar
Emeritus Professor, Dept. of Mathematics, Univ. of Washington, Seattle, USA

Johannes O. Royset
Associate Professor, Operations Research Dept., Naval Postgrad. School, Monterey, USA

ABSTRACT: Engineering decisions are made under substantial uncertainty about current and future system cost and response. A risk-neutral decision maker would rely on expected values when comparing designs, while a risk-averse decision maker might adopt nonlinear utility functions or failure probability criteria. The paper shows that these models for making decisions fall within a framework of risk measures that includes many other possibilities. The paper provides an overview of the framework, highlights benefits derived from certain risk measures, and gives a truss design example.

In design and optimization of structures, engineers are faced with the challenge of assessing the adequacy of a system with uncertain performance or selecting the best among several uncertain candidate systems. We consider the situations where the uncertain performance is captured by a random variable whose distribution is estimated by probabilistic models. A risk-neutral decision maker would make assessments and ranking on the basis of expected values of such random variables. Traditionally, a risk-averse decision maker would rely on expected utility theory, with a nonlinear utility function, or consider the probability of exceeding a threshold, i.e., a failure probability. In this paper, we outline a framework based on *risk measures* for risk-averse decision making that encapsulates these approaches, but also offers new possibilities. Various alternatives are illustrated in a design optimization problem for a truss structure.

Since preferences of decision makers are highly situational dependent, we avoid a discussion about whether risk neutrality or risk averseness is more appropriate; see Rockafellar and Royset (2015) and references therein. Here, we provide tools for handling risk averseness regardless of its source and motivation.

Risk measures as the basis for decision making are supported by a well-developed theory and extensive use in financial engineering and increasingly in other fields; see Dowd (2005); Commander et al. (2007); Rockafellar and Royset (2010); Minguez et al. (2013). Their connections with classical expected utility theory and risk-neutral decision making are revealing as discussed below. For other models of decision making, we refer to references in Rockafellar and Royset (2015).

We proceed in Section 2 with a description and examples of risk measures. Sections 3 and 4 make connections with expected utility theory and risk neutrality under distributional uncertainty, respectively. Section 5 discusses optimization of random variables that depend on design parameters. Sections 2-5 summarize the material by Rockafellar and Royset (2015). Section 6 illustrates the framework with a numerical example.

1. Risk Measures

A broad class of decision models that encapsulates essentially all reasonable approaches rely on measures of risk as defined next:

> A *measure of risk* is a functional $\mathscr{R}$ that assigns to a random variable Y a number $\mathscr{R}(Y)$, which could be infinity, as a quantification of the risk in Y.

The answer to the question of how "risky" is Y, is now simply defined to be $\mathscr{R}(Y)$. The comparison of two choices Y and Y' then reduces to comparing $\mathscr{R}(Y)$ and $\mathscr{R}(Y')$. A requirement that Y should be "adequately" $\leq b$ is interpreted as having $\mathscr{R}(Y) \leq b$. The minimizing of $\mathscr{R}(Y)$ over a set of candidate random variables Y then amounts to finding the lowest b such that there is a Y "adequately" $\leq b$. Here and throughout the paper we assume that high values of Y are undesirable. For example, Y might be the (life-cycle) cost or response amplitude of a system. For technical reasons and convenience, we limit the scope to random variables with finite second moments. We illustrate the breadth of possibilities with examples.

Expectation. The choice $\mathscr{R}(Y) = E[Y]$, the expected value, is simple, but not sensitive to the possibility of high values. Obviously, this choice incorporates no level of risk averseness.

Worst-case. The choice $\mathscr{R}(Y) = \sup Y$, the smallest value that Y exceeds only with probability zero, is conservative, usually overly so as it is infinite for distributions such as the normal. In fact, the corresponding decision model ignores all the information in the distribution of Y except its highest "possible" realization. Still, in some applications there may be thresholds that should not be exceeded.

Quantile. For $\alpha \in (0,1)$, the α-quantile of a random variable Y, $q_\alpha(Y)$, is simply $F_Y^{-1}(\alpha)$ when the cumulative distribution function F_Y of Y is strictly increasing, with a slightly more complex formula in the general case. The choice of risk measure $\mathscr{R}(Y) = q_\alpha(Y)$ is widely used in financial engineering under the name "value-at-risk" with typically an α of nearly one, and is equivalent to the failure probability. We recall that the probability of failure $p(Y) = \text{prob}(Y > 0)$, where we assume that positive realizations of Y are considered "failure." The failure probability is widely used in reliability analysis. It is clear that

$$p(Y) \leq 1-\alpha \text{ if and only if } q_\alpha(Y) \leq 0. \quad (1)$$

Consequently, the choice of a quantile as risk measure is equivalent to adopting a failure probability criterion. There are two immediate concerns with these approaches. First, there may be two design with the same failure probability, but their distributions could be different, especially in the critical upper tail. In fact, the failure probability is insensitive to the tail of the distribution and an exclusive focus on the corresponding decision models may hide significant risks. The second concern when using the failure probability is its lack of convexity and smoothness as a function of the design parameters. These deficiencies dramatically increase the difficulty of solving design optimization problems involving failure probability terms; see Rockafellar and Royset (2010) and Rockafellar and Royset (2015) for details.

Superquantile. The α-superquantile of Y at probability $\alpha \in (0,1)$ is given by

$$\bar{q}_\alpha(Y) = \frac{1}{1-\alpha}\int_\alpha^1 q_\beta(Y)d\beta, \quad (2)$$

i.e., an α-superquantile is an average of quantiles for probability levels $\alpha < \beta < 1$. When the cumulative distribution function of Y has no discontinuity at the realization $y = q_\alpha(Y)$, we have the equivalent formula $\bar{q}_\alpha(Y) = E[Y \mid Y \geq q_\alpha(Y)]$, i.e., the α-superquantile is simply the conditional expectation of Y above the α-quantile. Despite its somewhat complicated definition, convenient expressions facilitate the computation of superquantiles making them almost as accessible as an expectation. If Y is normally distributed with mean μ and standard deviation σ, then $\bar{q}_\alpha(Y) = \mu + \sigma\phi(\Phi^{-1}(\alpha))/(1-\alpha)$, where ϕ and Φ are the probability density (pdf) and cumulative distribution functions for a standard normal random variable. If Y follows a discrete distribution with realizations $y_1 < y_2 < \ldots < y_n$ and corresponding probabilities $p_1, p_2, \ldots, p_n$, then

$$\bar{q}_\alpha(Y) = \begin{cases} \sum_{j=1}^n p_j y_j & \text{for } \alpha = 0 \\ \frac{1}{1-\alpha}\left[\left(\sum_{j=1}^i p_j - \alpha\right) y_i + \sum_{j=i+1}^n p_j y_j\right] & \\ \quad \text{for } \sum_{j=1}^{i-1} p_j < \alpha \leq \sum_{j=1}^i p_j < 1 & \\ y_n & \text{for } \alpha > 1 - p_n. \end{cases}$$

We note that the realizations are sorted, without loss of generality, to simplify the formula. Generally,

$$\bar{q}_\alpha(Y) = \min_c c + \frac{1}{1-\alpha} E[\max\{0, Y-c\}]$$

i.e., a superquantile is the minimum value of a one-dimensional optimization problem with variable c. A risk measure that focuses primarily on the important upper tail of the distribution of Y is then the superquantile risk measure $\mathscr{R}(Y) = \bar{q}_\alpha(Y)$; also called conditional value-at-risk. A superquantile risk measure depends on the parameter α that represents the degree of risk averseness of the decision maker. For $\alpha = 0$, $\bar{q}_\alpha(Y) = E[Y]$ and therefore corresponds to the risk-neutral situation. An $\alpha = 1$ gives $\bar{q}_\alpha(Y) = \sup Y$ and therefore corresponds to the ultimate risk-averse decision maker. The superquantile risk measure leads to a number of benefits in subsequent analysis. For example, if Y depends on a set of design parameters, then the risk remains convex in the parameters as long as the parameterization is convex; see Rockafellar and Royset (2015) for a detailed discussion. The correspondence between a failure probability constraint $p(Y) \leq 1-\alpha$ and the quantile condition $q_\alpha(Y) \leq 0$ is given in (1). Analogously, a superquantile condition $\bar{q}_\alpha(Y) \leq 0$ corresponds to the condition $\bar{p}(Y) \leq 1-\alpha$, where $\bar{p}(Y)$ is the buffered failure probability of Y defined as the probability $1-\alpha$ that satisfies $\bar{q}_\alpha(Y) = 0$. We refer to Rockafellar and Royset (2010) for a discussion of the advantages that emerge from replacing a failure probability by a buffered failure probability.

A measure of risk $\mathscr{R}$ is *regular* if it satisfies

$\mathscr{R}(Y) = c$ when $Y \equiv c$ (constant equivalence);

$\mathscr{R}((1-\tau)Y + \tau Y') \leq (1-\tau)\mathscr{R}(Y) + \tau\mathscr{R}(Y')$ for all Y, Y' and $\tau \in (0,1)$ (convexity);

$\{Y \mid \mathscr{R}(Y) \leq c\}$ is a closed set for every constant c (closedness);

$\mathscr{R}(Y) > E[Y]$ for nonconstant Y (averseness).

The first condition is natural as it simply asserts that a random variable that always takes on the same value, has risk equal to that value. The second condition insists that a linear combination of two random variables has a risk that is no larger than the linear combination of the individual risks. This condition is also natural as it promotes diversification. The third condition is mostly technical as it simply asserts that a risk measure should have a certain continuity property. The last condition asserts that the risk should be greater than the expectation of a random variable as long as the random variable is not a deterministic constant. The choice $\mathscr{R}(Y) = E[Y]$ is therefore not regular, which is reasonable as it does not capture any degree of risk averseness. Of the examples above, the worst-case risk, and the superquantile risk measures satisfy the conditions.

2. RISK MEASURES AND UTILITY FUNCTIONS

Although a utility function u from classical expected utility theory by von Neumann and Morgenstern leads to a "quantification" $E[u(Y)]$ of a random variable Y, it is not natural to call this quantity a measure of risk. First, the orientation is flipped, with high values preferred to low ones. Second, the utility function distorts even a deterministic constant and therefore regularity cannot be achieved except in trivial cases. Still, important connections exist as we see next.

To avoid the awkward inconsistency between our orientation concerned with high values of Y and that of utility theory, concerned with low values, we define an analogous concept to a utility function.

A *measure of regret* is a functional $\mathscr{V}$ that assigns to a random variable Y a number $\mathscr{V}(Y)$, which may be infinity, as a quantification of the displeasure with the mix of possible realizations of Y. It could correspond to a utility function u through

$$\mathscr{V}(Y) = -E[u(-Y)], \tag{3}$$

but we ensure that it is anchored at zero. Hence, we insist that $\mathscr{V}(0) = 0$ and $\mathscr{V}(Y) > E[Y]$ when Y is not the constant zero. The correspondence is therefore with *relative* utility. Analogously to the regularity of risk measures, we say that a measure of regret is regular if it satisfies the closedness, convexity, and the two above conditions. If the random variable is not discrete, an additional technical condition might also be required; see Rockafellar and Uryasev (2013) for details. An example of a measure of regret is $\mathscr{V}(Y) = \frac{1}{1-\alpha}E[\max\{0,Y\}]$, with $\alpha \in (0,1)$, where negative realizations of Y are assigned zero regret, but positive realizations are viewed increasingly "regretable," with the increase being linear. This expression corresponds to a piecewise linear utility function with a kink at zero.

Major advantages derive from the following fact (see Rockafellar and Uryasev (2013)): A regular measure of risk $\mathscr{R}$ can be constructed from a regular measure of regret $\mathscr{V}$ through the one-dimensional optimization problem

$$\mathscr{R}(Y) = \min_c c + \mathscr{V}(Y - c). \quad (4)$$

For example, $\mathscr{R}(Y) = \bar{q}_\alpha(Y)$ derives from the measure of regret $\mathscr{V}(Y) = \frac{1}{1-\alpha} E[\max\{0, Y\}]$, which leads to the already claimed expression (1). A large number of other measures of risk can be constructed in a similar manner; see Rockafellar and Uryasev (2013). With the connections between regret and relative utility, this implies that every utility function u, with $u(0) = 0$ and $u(y) > y$ for $y \neq 0$, is in correspondence with a regular measure of risk through (3) and (4).

The trade-off formula (4) provides important interpretations of a regular measure of risk as the result of a two-stage decision process involving a regular measure of regret (and therefore also a relative utility function). As an example, suppose that Y gives the damage cost of a system and the measure of regret $\mathscr{V}(Y)$ quantifies our displeasure with the possible damage costs. In (4), view c as the money put aside today to cover future damage costs and $Y - c$ as the net damage cost in the future. Then, $c + \mathscr{V}(Y - c)$ becomes the total cost consisting of the sum of the money put aside today plus the current displeasure with future damage costs. The risk $\mathscr{R}(Y)$ is then the smallest possible total cost one can obtain by selecting the amount to put aside today in the best possible manner. Consequently, a risk measure probes deeper than a measure of regret as it considers how one can mitigate displeasure.

With the close connection between regret and risk, one may be led to believe that a decision model based on regret would be equivalent to one based on the corresponding risk measure. Section 6 shows that this conclusion is incorrect.

3. RISK NEUTRALITY AND UNCERTAINTY

Regular measures of risk have alternative "dual" expressions; see Rockafellar and Uryasev (2013). Specifically, every regular measure of risk that is positively homogeneous (i.e., $\mathscr{R}(\lambda Y) = \lambda \mathscr{R}(Y)$ for all $\lambda \geq 0$, which implies scale invariance) can be expressed in the form

$$\mathscr{R}(Y) = \text{the max of } E[YQ] \text{ across all } Q \in \mathscr{Q}, \quad (5)$$

where Q is a random variable that is taken from a set $\mathscr{Q}$ of random variables called a risk envelope associated with the risk measure. For example, if $\mathscr{R}(Y) = \bar{q}_\alpha(Y)$, then $\mathscr{Q}$ consists of those random variables with realizations between zero and $1/(1-\alpha)$ and that has expectation one. An example illustrates the formula.

We consider the simple situation where the random variable Y of a system takes the value 1 with probability 0.1 and the value 0 with probability 0.9, with expected value 0.1. A risk-neutral decision maker centered on the expectation would use 0.1 in numerical comparisons with other systems and requirements. Next, we consider a risk-averse decision maker that has adopted the superquantile risk measure with $\alpha = 0.8$. Since $q_\beta(Y) = 0$ for $\beta \leq 0.9$ and $q_\beta(Y) = 1$ for $\beta > 0.9$, the formula (2) gives that $\mathscr{R}(Y) = 0.5$. A risk-averse decision maker with this decision model would use 0.5 in comparison with other designs. We now consider the dual expression. In this case, with the scaling $1/(1-\alpha) = 5$, (5) simplifies to

$$\begin{aligned}\mathscr{R}(Y) = \; & \text{maximum value of } 0.9 \cdot 0 \cdot q_1 + 0.1 \cdot 1 \cdot q_2 \\ & \text{such that } 0 \leq q_1, q_2 \leq 5, 0.9q_1 + 0.1q_2 = 1,\end{aligned}$$

which has the optimal solution $q_1 = 5/9$ and $q_2 = 5$. The maximum value then becomes $0.9 \cdot 0 \cdot 5/9 + 0.1 \cdot 1 \cdot 5 = 0.5$ that confirms the previous calculation of $\mathscr{R}(Y)$. More interestingly however, the expression can be interpreted as the assessment made by a risk-neutral decision maker that has a nominal distribution with probabilities 0.9 and 0.1 for the realizations 0 and 1, respectively, but that is uncertain about the validity of this distribution. To compensate, she allows the probabilities to be scaled up with a factor of at most 5, while still making sure that they sum to one, in a manner that is the *least* favorable. This risk-neutral decision maker then makes the exact same assessment of the situation as the risk-averse decision maker.

The alternative formula (5) helps explain a source of risk averseness: lack of trust in probabilistic models. In fact, this insight can help quantify the exact benefit of better probabilistic models.

4. DESIGN OPTIMIZATION

In design, the random variable of interest is parameterized by a vector $\mathbf{x} = (x_1, ..., x_n)$ of design variables. The goal might then be to select $\mathbf{x}$ such that the risk of the random variable is minimized, usually subject to constraints on $\mathbf{x}$. This leads to the design optimization problem

$$\text{minimize } \mathscr{R}(Y(\mathbf{x})) \text{ subject to } \mathbf{x} \in \mathscr{X},$$

where $\mathscr{R}$ is a regular risk measure applied to a response or cost random variable $Y(\mathbf{x})$ depending on the design vector $\mathbf{x}$. For example, $Y(\mathbf{x}) = g(\mathbf{x}, \mathbf{V})$, with g a (limit-state) function parameterized by the design vector $\mathbf{x}$ and a random vector $\mathbf{V}$. The set $\mathscr{X}$ specifies constraints on $\mathbf{x}$, which we for simplicity drop below. This formulation can be expanded to include multiple random variables and multiple measures of risk with few complications.

A key property of regular measures of risk is that the canonical formulation is a convex optimization problem whenever $Y(\mathbf{x})$ is an affine function of $\mathbf{x}$ for every realization, possibly except for an event with probability zero, and $\mathscr{X}$ is a convex set. If $\mathscr{R}$ is monotone, i.e., $\mathscr{R}(Y) \leq \mathscr{R}(Y')$ whenever $Y \leq Y'$ with probability one, then linearity can be relaxed to convexity. The value of convexity of an optimization problem cannot be overestimated as it dramatically improves the ability of algorithms to obtain globally optimal solutions efficiently. In the absence of convexity, a globally optimal solution is usually inaccessible unless $\mathbf{x}$ only involves a small number of variables and a huge computational effort is employed.

The trade-off formula (4) allows a simplification of the canonical formulation into the following equivalent form: minimize $c_0 + \mathscr{V}(Y(\mathbf{x}) - c_0)$, where $\mathscr{V}$ is a regular measures of risk corresponding to the regular risk measure $\mathscr{R}$ through (4) and c_0 is an auxiliary design variable to be optimized unconstrained. This equivalent form is computationally beneficial as expressions for regret are usually simpler than those for risk. For example, if $\mathscr{R}(Y(\mathbf{x})) = \bar{q}_\alpha(Y(\mathbf{x}))$, i.e., using a superquantile risk measure, then $\mathscr{V}(Y(\mathbf{x})) = \frac{1}{1-\alpha} E[\max\{0, Y(\mathbf{x})\}]$ and the design optimization problem takes the following equivalent form

$$\text{minimize } c_0 + \frac{1}{1-\alpha} E[\max\{0, Y(\mathbf{x}) - c_0\}]$$

which simply involves an expectation. If $Y(\mathbf{x}) = g(\mathbf{x}, \mathbf{V})$ for some function g and the distribution of $\mathbf{V}$ is discrete with realizations $\mathbf{v}^1, ..., \mathbf{v}^J$ and probabilities $\gamma^1, ..., \gamma^J$, then the formulation simplifies further to

$$\begin{aligned} &\text{minimize } c_0 + \frac{1}{1-\alpha} \sum_{j=1}^{J} \gamma^j c_j \\ &\text{subject to } g(\mathbf{x}, \mathbf{v}^j) - c_0 \leq c_j, \text{ for all } j = 1, ..., J \\ &\qquad 0 \leq c_j, \text{ for all } j = 1, ..., J, \end{aligned}$$

with c_j, $j = 1, ..., J$, being auxiliary design variables. The reformulation involves additional constraints and variables, but this is outweighed by the removal of all complicating expressions with the exception of the unavoidable function g. In fact, the formulation resembles the corresponding one in the absence of uncertainty. Consequently, design optimization under uncertainty using a superquantile risk measure is in some sense only marginally harder than the corresponding design optimization problem *without* uncertainty.

5. DESIGN EXAMPLE

We consider the simply supported truss in Figure 1. Let V_k be the yield stress of member k, $k = 1, 2, ..., 7$. Members 1 and 2 have lognormally distributed yield stresses with mean 100 N/mm^2 and standard deviation 20 N/mm^2. The other members have lognormally distributed yield stresses with mean 200 N/mm^2 and standard deviation 40 N/mm^2. The yield stresses of members 1 and 2 are correlated with correlation coefficients 0.8. However, their correlation coefficients with the other

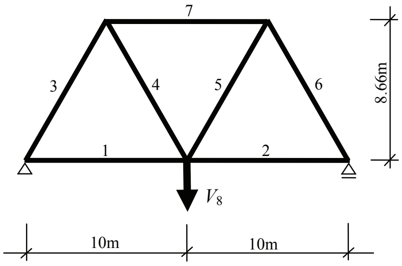

Figure 1: Design of Truss

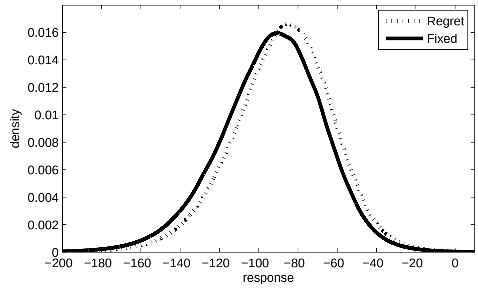

Figure 2: Pdf for fixed and regret-optimized designs.

yield stresses are 0.5. Similarly, the yield stresses of members 3-7 are correlated with correlation coefficients 0.8. The truss is subject to a random load V_8 in its mid-span. V_8 is lognormally distributed with mean 1000 kN and standard deviation 400 kN. The load V_8 is independent of the yield stresses. We use a joint lognormal distribution and the above correlation coefficients to approximate the joint distribution of $\mathbf{V} = (V_1, V_2, ..., V_8)$.

The design vector $\mathbf{x} = (x_1, x_2, ..., x_7)$, where x_k is the cross-section area (in 1000 mm^2) of member k. The truss is constrained by the set $\mathscr{X} = \{\mathbf{x} \mid 0.5 \le x_k \le 2,\ k = 1,2,...,7,\ x_1 + x_2 + ... + x_7 \le 9\}$, where the first restriction limits each member to be between 500 mm^2 and 2000 mm^2 and the last restriction limits the total cross-section area.

For each member, we compare load effect with capacity through $g_k(\mathbf{x}, \mathbf{v}) = v_8/\zeta_k - v_k x_k,\ k = 1,2,...,7$, where ζ_k is a factor given by the geometry and loading of the truss. From Figure 1, we determine that $\zeta_k = 1/(2\sqrt{3})$ for $k = 1,2$, and $\zeta_k = 1/\sqrt{3}$ for $k = 3,4,...,7$. If $g_k(\mathbf{x}, \mathbf{v})$ is positive the load effect is larger than the capacity of the member. A

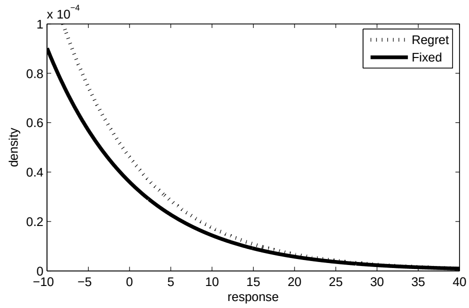

Figure 3: Tails of pdf for fixed and regret-optimized designs.

random variable of concern might then be the *response* $Y(\mathbf{x}) = \max_{k=1,...,7} g_k(\mathbf{x}, \mathbf{V})$, which gives the highest difference between load effect and capacity across all the members. In the following, we approximate the distribution of $\mathbf{V}$ by the empirical distribution generated by an independent sample of size 100,000. This approximation facilitates computations.

We initially consider the "fixed" design $\mathbf{x} = (9,9,...,9)/7$ that exactly satisfies the cross-section area budget. The third row labelled "Fixed" in Table 1 gives the mean $E[Y(\mathbf{x})]$ and probability of failure $p(Y(\mathbf{x})) = \text{prob}(Y(\mathbf{x}) > 0)$. The full pdf of $Y(\mathbf{x})$ is given by the solid line in Figure 2. Here, and below, we smooth the discrete data using a logconcave exponential epi-spline (see Royset et al. (2013)) to construct estimates of the pdf of $Y(\mathbf{x})$. Figure 3 highlights the corresponding upper tail. Although, the response is typically negative, high values might occur.

The quality and risk of the fixed design is assessed using various quantifiers. First, we consider the regret $\mathscr{V}(Y(\mathbf{x})) = (1/(1-\alpha))E[\max\{Y(\mathbf{x}), 0\}]$, with $\alpha = 0$. This expression gives the average capacity exceedance. It corresponds to a piecewise linear utility function. The expression can also be interpreted as the expected cost under the assumption that no load exceedance has a zero cost and a load exceedance has a cost proportional to the degree of exceedance. The last column, second row of Table 2 gives the regret of $Y(\mathbf{x})$ as $\mathscr{V}(Y(\mathbf{x})) = 0.01156$.

Second, we consider the measures of risk $\mathscr{R}(Y(\mathbf{x})) = \bar{q}_\alpha(Y(\mathbf{x}))$ for $\alpha = 0, 0.5, 0.9, 0.99, 0.999$; see the last column of

Design	Size of member (in mm^2)							Mean	$p(Y(\mathbf{x}))$
	1	2	3	4	5	6	7		
Fixed	1286	1286	1286	1286	1286	1286	1286	-93.5	0.00044
Regret	1220	1231	1316	1297	1315	1298	1323	-87.8	0.00053
$\alpha = 0$	1804	1805	1079	1077	1078	1080	1078	-125.9	0.00162
$\alpha = 0.5$	1755	1756	1099	1097	1096	1100	1097	-125.5	0.00145
$\alpha = 0.9$	1650	1649	1142	1138	1139	1141	1141	-122.1	0.00114
$\alpha = 0.99$	1465	1467	1212	1222	1210	1219	1204	-109.7	0.00060
$\alpha = 0.999$	1292	1288	1290	1263	1281	1296	1289	-93.9	0.00043

Table 1: Designs of truss

Risk measure	Optimized	Fixed
Regret	0.01118	0.01156
$\alpha = 0$	-125.9	-93.49
$\alpha = 0.5$	-98.21	-72.82
$\alpha = 0.9$	-65.41	-49.82
$\alpha = 0.99$	-28.89	-24.17
$\alpha = 0.999$	8.232	8.308

Table 2: Regret and risk in optimized and fixed designs

Table 2. We recall that $\bar{q}_\alpha(Y(\mathbf{x}))$ is essentially the conditional expectation of $Y(\mathbf{x})$ given $Y(\mathbf{x})$ is no smaller than its α-quantile. Consequently, $\bar{q}_\alpha(Y(\mathbf{x}))$ gives the average of the $(1-\alpha)100\%$ worst responses. All these averages are well below zero except for $\alpha = 0.999$; the average of the 0.1% worst responses exceeds zero. In comparison, the average load exceedance is only slightly above zero at 0.01156. The choice of α depends on the degree of risk averseness of the decision maker.

We next turn to optimization of the design. First, we minimize the regret $\mathscr{V}(Y(\mathbf{x}))$ subject to the constraint $\mathbf{x} \in \mathscr{X}$. The resulting optimization problem is a linear program solved in the General Algebraic Modeling System (GAMS) Distribution 24.1.3, with the CPLEX 12.5.1 solver, on a laptop computer with 4 GB of RAM and 2.6 GHz processor running Windows 7. The solver time is 0.78 seconds. The optimal design, only marginally different than the previous design, is given in row four of Table 1. We note that both the mean and probability of capacity exceedance are worst for the optimized design relative to the fixed design. However, the regret is 0.01118 and slightly better; see the second row of Table 2. Although similar, the pdf of the regret-optimized response is different than that for the fixed design as seen by comparing the dotted and solid lines in Figures 2 and 3. It is interesting to note that the optimized design gives up average response and worsens the probability of capacity exceedance to ensure slightly lower likelihood for high realizations and therefore a slightly improved regret. The reduction in likelihood is too small to be visible in Figure 3.

Using the same computational platform, we second minimize the superquantile risk $\bar{q}_\alpha(Y(\mathbf{x}))$ under the same constraints and obtain the designs of rows 5-9 in Table 1 using $\alpha = 0, 0.5, 0.9, 0.99, 0.999$. The solver times vary between 4 and 218 seconds. As α increases, the mean response worsens, but the probability of capacity exceedance decreases. The resulting pdf are given in Figure 4, with upper tails given in Figure 5, where we leave out the case $\alpha = 0.5$, which is similar to that with $\alpha = 0$. We see from Table 2 that the optimized designs have, usually, substantially lower risk than those of the fixed design. The effect of optimization diminishes as α increases simply to due to the fact that the fixed design happens to be a better design in those cases.

As seen in Table 1, the minimization of regret yields a rather different design than the minimization of risk. (We note that the choice of α has no bearing on the minimum-regret design as the factor $1/(1-\alpha)$ simply scales the regret.) Hence, although the measure of regret $\mathscr{V}$ is the foundation of the measures of risk $\bar{q}_\alpha$, the latter measures examine "deeper" the random variable in question by

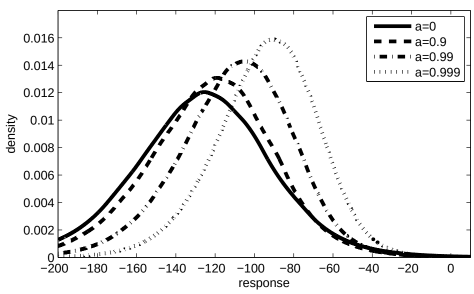

Figure 4: Pdf for risk optimized designs

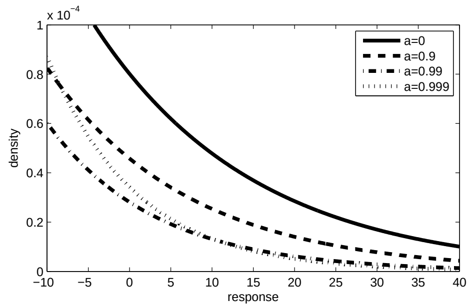

Figure 5: Tails of pdf for risk optimized designs

also considering how to best mitigate the displeasure of high responses. For example, the minimum-regret design in Table 1 is a substantially inferior design compared to the minimum-risk design, as measure by the corresponding $\bar{q}_\alpha$. The minimum-risk design is better in this sense as it more easily allows mitigation of risk through an intelligent choice of c in the trade-off formula (1).

We note that a risk neutral decision maker would select the design in row 5 of Table 1. In view of the discussion in Section 4, a risk-neutral decision maker that is uncertain about the underlying probability distribution might select one of the designs in the lower rows of that table. As stated above, we approximate $\mathbf{V}$ by a discrete random variable with 100,000 possible realizations, each with probability 10^{-5}. In this case, the discussion of Section 4 takes the following form. If the decision maker believes that she could have estimated the probabilities incorrectly with a factor of $1/(1-\alpha) = 1/(1-0.99) = 100$, i.e., the probability of each realization can be any number between 0 and $10^{-5} \cdot 100 = 10^{-3}$, then she would have selected the design of row 8 of Table 1. This results in a design that is significantly worse on average (-109.7 vs -125.9). Hence, the average worsening of the response due to incomplete information about the distribution is $-109.7 - (-125.9) = 16.2$. Analysis of this kind might help justifying efforts to improve probabilistic models.

ACKNOWLEDGEMENT

The authors acknowledges financial supported from AFOSR and the second author also from DARPA.

6. REFERENCES

Commander, C., Pardalos, P., Ryabchenko, V., Uryasev, S., and Zrazhevsky, G. (2007). "The wireless network jamming problem." *Journal of Combinatorial Optimization*, 14, 481–498.

Dowd, K. (2005). *Measuring market risk*. Wiley, New York, NY, 2. edition.

Minguez, R., Castillo, E., and Lara, J. (2013). "Iterative scenario reduction technique to solve reliability based optimization problems using the buffered failure probability." *Proceedings of ICOSSAR*, G. Deodatis, ed.

Rockafellar, R. and Royset, J. (2010). "On buffered failure probability in design and optimization of structures." *Reliability Engineering & System Safety*, 95, 499–510.

Rockafellar, R. and Royset, J. (2015). "Engineering decisions under risk-averseness." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*.

Rockafellar, R. and Uryasev, S. (2013). "The fundamental risk quadrangle in risk management, optimization and statistical estimation." *Surveys in Operations Research and Management Science*, 18, 33–53.

Royset, J., Sukumar, N., and Wets, R. J.-B. (2013). "Uncertainty quantification using exponential epi-splines." *Proceedings of ICOSSAR*, G. Deodatis, ed.