Operations Research

Decision Making with Multiattribute Performance Targets: The Impact of Changes in Performance and Target Distributions

Ilia Tsetlin, Robert L. Winkler,

Please scroll down for article—it is on subsequent pages

Decision Making with Multiattribute Performance Targets: The Impact of Changes in Performance and Target Distributions

Ilia Tsetlin
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, ilia.tsetlin@insead.edu

Robert L. Winkler
Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120, rwinkler@duke.edu

In many situations, performance on several attributes is important. Moreover, a decision maker's utility may depend not on the absolute level of performance on each attribute, but rather on whether that level of performance meets a target, in which case the decision maker is said to be target oriented. For example, typical attributes in new product development include cost, quality, and features, and the corresponding targets might be the best performance on these attributes by competing products. Targets and performance levels typically are uncertain and often are dependent. We develop a model that allows for uncertain dependent targets and uncertain dependent performance levels, and we study implications for decision making in this general multiattribute target-oriented setting. We consider the impact on expected utility of modifying key characteristics of performance (or target) distributions: location, spread, and degree of dependence. In particular, we show that explicit consideration of dependence is important, and we establish when increasing or decreasing dependence is beneficial. We illustrate the results numerically with a normal model and discuss some extensions and implications.

1. Introduction

In many decision-making situations, objectives can be expressed in terms of meeting targets. These targets might be externally imposed (e.g., sales quotas) or driven by a decision maker's internal preferences (e.g., aspirations such as wanting to be in a list of top-performing mutual fund managers). They can be fixed, as in the case of a fixed quota, or uncertain, as in the case of a contest where the target will depend on the performance of competitors. They can involve a single attribute (e.g., a limit on emissions set by regulation) or multiple attributes (e.g., developing a product that performs better than competing products on attributes such as cost, reliability, ease of use, and number of features). As indicated in Hymowitz (2005a, b), the use of targets, even short-term targets involving daily or weekly sales, revenues, and costs, is widespread in business. Targets set by upper management, as well as higher-level targets related to stockholders' expectations about various measures such as a firm's quarterly earnings per share, provide incentive for good performance, but can also create stress and perhaps lead to the types of ethical problems that have come to light in recent years.

Borch (1968) may have been the first to present the concept of target-based utility, using different terminology related to the probability of ruin. Castagnoli and LiCalzi (1996) show that expected utility can be expressed in terms of "expected probability," with the utility function for performance interpreted as a cumulative distribution function (c.d.f.) in the case of a single attribute (see also Bordley and LiCalzi 2000). In maximizing expected utility, a decision maker behaves as if maximizing the probability that performance is greater than or equal to a target, whether the target is real or just a convenient interpretation. Note that if a target is fixed, its c.d.f. simplifies to a step function with a single step at the target. Abbas and Matheson (2005) model target setting in organizations. They define "aspiration equivalents" for the alternatives under consideration based on the organization's utility function, drawing an analogy with Simon's (1955) notion of satisficing by seeking an alternative that meets or exceeds an aspiration level, and show that these aspiration equivalents can be used as targets.

Bordley and Kirkwood (2004) (hereafter BK) consider situations in which a target-oriented approach is natural and define a target-oriented decision maker for a single attribute as one with a utility that depends only on whether a target for that attribute is achieved. They extend this definition to targets for multiple attributes, requiring that the decision maker's utility for a multidimensional outcome depend only on the subset of attributes for which targets are met, and

they develop a target-oriented approach to assess a multiattribute preference function. Abbas and Howard (2005) introduce a class of multiattribute utility functions called attribute dominance utility functions that can be manipulated like joint probability distributions and allow the use of probability assessment methods in utility elicitation.

In this paper, we develop a model for investigating decision-making strategies in the presence of multiattribute performance targets. In general, both performance levels and targets might be uncertain. Thus, our model allows for uncertain performance levels, uncertain targets, and dependence both among the performance levels and among the targets. We consider the impact on expected utility of changes in the location and spread of distributions of performance levels or targets, as well as changes in dependence among performance levels or among targets for different attributes. In comparison to the previous work noted above, our primary focus is on how decision making is influenced by model parameters in a target-oriented setting rather than on connections to utility theory and assessment.

This paper is organized as follows. In §2, we discuss multiattribute target-oriented utility functions, clarify notions of independence, and argue for explicit consideration of dependence. In §3, we present our model, develop the main results, and discuss implications for decision making. In particular, we establish when increasing or decreasing dependence increases expected utility and also show the impact of changes in location and spread. Section 4 contains a brief summary and discussion.

2. Multiattribute Target-Oriented Utility Functions

The consequences in decision making often involve multiple attributes (Keeney and Raiffa 1976), and BK give a nice overview of multiattribute situations where a target-oriented approach is natural. Suppose that n attributes of performance are of interest, that a decision maker's levels of performance on the attributes are represented by $\mathbf{x} = (x_1, x_2, \ldots, x_n)$, and that targets for the attributes are represented by $\mathbf{t} = (t_1, t_2, \ldots, t_n)$. If the decision maker cares only about meeting targets, her utility function should reflect that. Following BK, we say that a decision maker is target oriented if her utility depends only on the subset of attributes for which the targets are met (i.e., for which $x_i \geqslant t_i$). Let $\mathbf{I} = (I_1, \ldots, I_n)$, where $I_i = 1$ if $x_i \geqslant t_i$ and zero otherwise. A target-oriented decision maker has a utility function $U_I(\mathbf{I})$ assigning utilities to the 2^n possible values of $\mathbf{I}$. Let $U_I(\mathbf{I}) = u_A$, where A is the set of indices $\{i \mid I_i = 1\}$ corresponding to the attributes in $\mathbf{I}$ for which the targets are met. For example, $U_I(1, 0, \ldots, 0) = u_1$, $U_I(0, 1, 1, 0, \ldots, 0) = u_{2,3}$, and so on. If $A_1 \subseteq A_2$, then $u_{A_1} \leqslant u_{A_2}$; utility can never be reduced by meeting additional targets. Also, $0 \leqslant u_A \leqslant 1$ for all A, with $u_\varnothing = U_I(0, \ldots, 0) = 0$ and $u_{1,2,\ldots,n} = 1$, leaving $2^n - 2$ utilities u_A to be assessed.

If $n = 2$,

$$U_I(\mathbf{I}) = u_1 I_1 + u_2 I_2 + (1 - u_1 - u_2) I_1 I_2. \tag{1}$$

The sign of $1 - u_1 - u_2$ has important implications here. For example, if $u_1 = u_2 = 0.9$, the decision maker wants to achieve at least one target, whereas if $u_1 = u_2 = 0.1$, the decision maker wants to achieve both targets.

From U_I, the induced $U(\mathbf{x})$ can be found for any distribution of $(\tilde{\mathbf{t}}, \tilde{\mathbf{x}})$ with $\tilde{\mathbf{t}}$ independent of $\tilde{\mathbf{x}}$. Recalling that I_i depends on whether $x_i \geqslant t_i$, and integrating out the uncertainty about $\tilde{\mathbf{t}}$,

$$U(\mathbf{x}) = u_1 F_{t_1}(x_1) + u_2 F_{t_2}(x_2) + (1 - u_1 - u_2) F_{t_1, t_2}(x_1, x_2), \tag{2}$$

where F_{t_i} is the c.d.f. of $\tilde{t}_i$ and F_{t_1, t_2} is the joint c.d.f. of $\tilde{\mathbf{t}}$. This looks almost multilinear (i.e., linear in F_{t_1} and F_{t_2} separately), and it is if $F_{t_1, t_2}(x_1, x_2) = F_{t_1}(x_1) F_{t_2}(x_2)$, i.e., if $\tilde{t}_1$ and $\tilde{t}_2$ are independent.

Proposition 1 extends these results to the case of n targets. As in (1), U_I can be expressed as a weighted average of the products of I_i terms for the $2^n - 1$ combinations of such terms, not including the empty set. Then, as in (2), $U(\mathbf{x})$ is a weighted average of corresponding c.d.f. terms,

$$U(\mathbf{x}) = \sum_A w_A F_{\{t_i \mid i \in A\}}(\{x_i \mid i \in A\}), \tag{3}$$

with $\sum_A w_A = 1$. Note that w_A is a linear combination of u_B terms, with $w_A = u_A$ as a special case. In (2), for example, $w_i = u_i$ for $i = 1, 2$ but $w_{1,2} = u_{1,2} - u_1 - u_2$.

PROPOSITION 1. *A target-oriented utility function U_I is consistent with a multilinear $U(\mathbf{x})$ if and only if $\{\tilde{t}_i \mid i \in A\}$ are independent for all A such that $w_A \neq 0$. In particular, if $w_{1,2,\ldots,n} \neq 0$, then a target-oriented U_I is consistent with a multilinear $U(\mathbf{x})$ if and only if all n targets are independent.*

PROOF. From (3), $U(\mathbf{x})$ is a weighted average of corresponding $F_{\{t_i \mid i \in A\}}(\{x_i \mid i \in A\})$ terms. Thus, $U(\mathbf{x})$ is multilinear if and only if for each A such that $w_A \neq 0$, the term $F_{\{t_i \mid i \in A\}}(\{x_i \mid i \in A\})$ can be factored into a product $\prod_{i \in A} F_{t_i}(x_i)$, which can be done if and only if $\{\tilde{t}_i \mid i \in A\}$ are independent. $\square$

Similarly, we can find an induced $U(\mathbf{t})$ for the targets. Starting with (1) and integrating out the uncertainty about $\tilde{\mathbf{x}}$ yields $U(\mathbf{t}) = u_1 G_{x_1}(t_1) + u_2 G_{x_2}(t_2) + (1 - u_1 - u_2) G_{x_1, x_2}(t_1, t_2)$, where $G_{x_i}(t_i)$ and $G_{x_1, x_2}(t_1, t_2)$ are the right-hand c.d.f.s $P(\tilde{x}_i \geqslant t_i)$ and $P(\tilde{x}_1 \geqslant t_1, \tilde{x}_2 \geqslant t_2)$. Extending this to n targets as in (3), we get $U(\mathbf{t}) = \sum_A w_A G_{\{x_i \mid i \in A\}}(\{t_i \mid i \in A\})$, with $\sum_A w_A = 1$.

Proposition 1 requires that the targets $\tilde{t}_1, \ldots, \tilde{t}_n$ be independent. Many results in BK are based on such independence, which implies independence of $\tilde{I}_1, \ldots, \tilde{I}_n$ given $\mathbf{x}$, as represented in their Equation (5) (2004, p. 826). However,

this probabilistic independence does not follow from their Definition 3, which says that targets are independent if the decision maker's probability of achieving the target on any attribute i depends only on the value of x_i. What needs to be added at the end of the definition is the qualifier "and not on whether targets for other attributes are achieved." Without some such qualifier, the terminology "independent targets" is not appropriate.

Details of definitions aside, results that rely on the probabilistic independence of $\tilde{t}_1, \ldots, \tilde{t}_n$ are restrictive because such independence, although it can lead to tractable results, is often not realistic. In fact, dependence among the targets is very natural in some of the excellent examples given in BK. In their product development example, for instance, negative dependence between targets for different features can occur due to trade-offs among the attributes. Suppose that a firm has one major competitor, so that the firm's targets are the performance levels achieved by that competitor. An improvement in the competitor's distribution of performance on one attribute leaves fewer resources available for other attributes. Also, improvements in the quality of the competitor's new product are likely to increase the competitor's price.

On the other hand, positive dependence among targets can occur if some common factor influences all of the targets in a similar way. For example, targets for a salesperson in a sales contest for multiple products could be influenced in a similar way by economic conditions. Targets related to a firm's stockholders' expectations about measures such as revenues, earnings per share, and stock price are also likely to be positively correlated.

It might be possible to redefine attributes to achieve independence among the targets. In the sales example, the attributes could be redefined relative to a base set by the state of economic conditions. In the product development example, performance on each attribute could be defined as performance relative to the competitor's performance on that attribute, implying that the targets would all be zero. However, this most likely would complicate the modeling of dependence among performance attributes.

Moreover, because the target on the ith attribute is achieved if $\tilde{t}_i - \tilde{x}_i \leqslant 0$, we are interested in whether $\tilde{t}_1 - \tilde{x}_1, \ldots, \tilde{t}_n - \tilde{x}_n$ are independent, not just in whether the targets are independent. Even if $\tilde{t}_1, \ldots, \tilde{t}_n$ are independent, we can have dependence among $\tilde{t}_1 - \tilde{x}_1, \ldots, \tilde{t}_n - \tilde{x}_n$ through dependence among $\tilde{x}_1, \ldots, \tilde{x}_n$. As discussed above, that dependence is likely to be negative if there are trade-offs among attributes and positive if there are some common factors influencing different attributes in the same direction. For instance, in BK's example of monthly sales quotas, a salesperson facing sales quotas for multiple products has to make choices about allocation of effort. Negative dependence is likely because greater effort on certain products tends to increase sales of those products and take effort away from other products, thereby tending to reduce sales on the other products. A firm's performance levels on

attributes such as revenues, earnings per share, and stock price are likely to be influenced in the same direction by economic conditions and therefore positively correlated.

The above examples illustrate why it is important to model dependence among $\tilde{x}_1, \ldots, \tilde{x}_n$ and among $\tilde{t}_1, \ldots, \tilde{t}_n$ explicitly instead of relying on independence assumptions. Next, in §3, we develop a model of multiattribute targets allowing dependence among the performance levels and among the targets. We study decision-making strategies in this target-oriented framework, with particular attention given to how these strategies change as a result of dependence (and other distributional parameters) and how they vary with different utility assumptions.

3. Interdependent Multiattribute Targets

The $n = 2$ case is easier to work with than the more general n-target case, and it provides useful insights into the impact of characteristics such as dependence on the decision maker's expected utility. We provide some general results for two targets in §3.1, give some numerical results with a normal model in §3.2, and discuss extensions to $n > 2$ in §3.3.

3.1. Two Targets

We assume that the decision maker is target oriented, so that her utility is given by (1). Let F_{d_1, d_2} represent the joint c.d.f. of $\tilde{\mathbf{d}} = \tilde{\mathbf{t}} - \tilde{\mathbf{x}}$, which can be found from the distributions for $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$; let F_{d_1} and F_{d_2} represent the corresponding marginal c.d.f.s; and let $F_{d_1 \mid d_2}$ represent the conditional distribution of $\tilde{d}_1$ given $\tilde{d}_2$. From (1), the decision maker's expected utility is

$$\mathrm{EU}(F_{d_1, d_2}) = u_1 F_{d_1}(0) + u_2 F_{d_2}(0)$$
$$+ (1 - u_1 - u_2) F_{d_1, d_2}(0, 0). \qquad (4)$$

The decision maker may be able to manipulate aspects of the distribution of performance, so the impact on EU of changes in these characteristics is of interest. Proposition 2 and its Corollary 1, Proposition 3, and Proposition 4 consider the impact of changes in location, dependence, and spread, respectively. As in §2, we assume that $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$ are independent, and we rule out the cases with $u_i = 1$ and $u_{3-i} = 0$, $i = 1, 2$, which reduce to a single target on attribute i and make the other attribute irrelevant.

PROPOSITION 2. *If $F_{d_2}^*(d_2) = F_{d_2}(d_2)$ and $F_{d_1 \mid d_2}^*(0 \mid d_2) \geqslant F_{d_1 \mid d_2}(0 \mid d_2)$ for all d_2, then $\mathrm{EU}(F_{d_1, d_2}^*) \geqslant \mathrm{EU}(F_{d_1, d_2})$.*

PROOF. From (4), $\mathrm{EU}(F_{d_1, d_2}) = u_1[F_{d_1}(0) - F_{d_1, d_2}(0, 0)] + u_2 F_{d_2}(0) + (1 - u_2)F_{d_1, d_2}(0, 0)$. We can write

$$F_{d_1}(0) - F_{d_1, d_2}(0, 0)$$
$$= P(\tilde{d}_1 \leqslant 0, \tilde{d}_2 > 0) = \int_0^\infty F_{d_1 \mid d_2}(0 \mid d_2) \, dF_{d_2}(d_2).$$

Similarly, $F_{d_1, d_2}(0, 0) = \int_{-\infty}^0 F_{d_1 \mid d_2}(0 \mid d_2) \, dF_{d_2}(d_2)$. Thus, $F_{d_1}^*(0) - F_{d_1, d_2}^*(0, 0) \geqslant F_{d_1}(0) - F_{d_1, d_2}(0, 0)$ and $F_{d_1, d_2}^*(0, 0) \geqslant F_{d_1, d_2}(0, 0)$ imply $\mathrm{EU}(F_{d_1, d_2}^*) \geqslant \mathrm{EU}(F_{d_1, d_2})$. □

COROLLARY 1. *If* $F^*_{d_1, d_2}(d_1, d_2) = F_{d_1, d_2}(d_1 + m, d_2)$ *(or* $F^*_{d_1, d_2}(d_1, d_2) = F_{d_1, d_2}(d_1, d_2 + m)$*) for all* (d_1, d_2)*, where* $m > 0$*, then* $\mathrm{EU}(F^*_{d_1, d_2}) \geqslant \mathrm{EU}(F_{d_1, d_2})$*, with a strict inequality unless* $F_{d_1}(0) = F_{d_1}(m)$ *and* $F_{d_1, d_2}(m, 0) = F_{d_1, d_2}(0, 0)$ *(*$F_{d_2}(0) = F_{d_2}(m)$ *and* $F_{d_1, d_2}(0, 0) = F_{d_1, d_2}(0, m)$*).*

The results of Proposition 2 are intuitively reasonable. Moving from F to F^* increases the probability of making the target on Attribute 1, and thus increases EU via a first-order stochastic-dominance type of result. By interchanging the subscripts in Proposition 2, the same result holds with an increase in the probability of making the target on Attribute 2.

A special case of the type of stochastic dominance in Proposition 2 is a location shift, as given in Corollary 1. Moving from F_{d_1, d_2} to $F^*_{d_1, d_2}(d_1, d_2) = F_{d_1, d_2}(d_1 + m, d_2)$ amounts to a negative location shift of size m in $\tilde{t}_1 - \tilde{x}_1$. Such a shift can be caused by a positive location shift in $\tilde{x}_1$ (making higher performance on Attribute 1 more likely), a negative location shift in $\tilde{t}_1$ (making lower targets for Attribute 1 more likely), or some combination of the two. In terms of strategies for a decision maker who has no control over the distribution of targets, a positive location shift in $\tilde{x}_1$ increases EU. Of course, a *positive* location shift in the distribution of $\tilde{t}_i - \tilde{x}_i$ (the conditions in Corollary 1 with $m < 0$) decreases the probability of meeting the target on that attribute and thus decreases EU. Note that Corollary 1 does not require any assumptions on F_{d_1, d_2}. In contrast, the impact of a change in spread depends on other conditions, as will be seen in Proposition 4 and §3.2.

Next, we consider the impact of dependence, using the following definition based on Yanagimoto and Okamoto (1969), which generalizes Lehmann's (1966) notion of positive quadrant dependence.

DEFINITION 1. Assume that c.d.f.s $F(d_1, d_2)$ and $F^*(d_1, d_2)$ have common marginal distributions. Then, F^* has strictly larger quadrant dependence than F if and only if $F^*(d_1, d_2) \geqslant F(d_1, d_2)$ for all d_1 and d_2, with the inequality strict for some (d_1, d_2).

PROPOSITION 3. *If* $F^*_{d_1, d_2}$ *has strictly larger quadrant dependence than* F_{d_1, d_2} *and* $F^*_{d_1, d_2}(0, 0) > F_{d_1, d_2}(0, 0)$*, then* $\mathrm{EU}(F^*_{d_1, d_2}) > (=, <)\mathrm{EU}(F_{d_1, d_2})$ *if and only if* $u_1 + u_2 < (=, >)1$.

PROOF. From (4) and the common marginal distributions,

$$\mathrm{EU}(F^*_{d_1, d_2}) - \mathrm{EU}(F_{d_1, d_2})$$
$$= (1 - u_1 - u_2)\big[F^*_{d_1, d_2}(0, 0) - F_{d_1, d_2}(0, 0)\big] > (=, <)0$$

if and only if $u_1 + u_2 < (=, >)1$. $\square$

The assumption of common marginal distributions in the definition of strictly larger quadrant dependence allows us to isolate the impact of dependence on EU. Proposition 3 shows that an increase in the dependence between $\tilde{t}_1 - \tilde{x}_1$ and $\tilde{t}_2 - \tilde{x}_2$ increases EU if $u_1 + u_2 < 1$ and decreases EU if $u_1 + u_2 > 1$. Increased dependence makes it more likely that both targets or neither target will be met, and less likely that one target will be met. Thus, the decision maker is more likely to get a utility of one or zero and less likely to get u_1 or u_2. This trade-off is favorable if $u_1 + u_2 < 1$ and unfavorable if $u_1 + u_2 > 1$. For example, if $u_1 = u_2 = 0.1$ (0.9), increasing (decreasing) dependence is preferable because the decision maker wants to achieve both targets (at least one target). If $u_1 = 0.9$ and $u_2 = 0.2$, decreasing dependence is preferable, but in this case Target 1 is more important than Target 2 and thus deserves more effort.

In the special case with $u_1 = u_2 = u$, the decision maker views the targets symmetrically in the sense that the only thing that matters is the number of targets that are met, not which specific targets they are. From Proposition 3, the impact of dependence on EU depends on u. High dependence is desirable for the decision maker when $u < 0.5$, and low dependence is preferred when $u > 0.5$. If $u = 0$, for example, then the only goal is to make both targets; making only one target has no utility. Thus, high correlations are clearly preferred. At the other extreme, when $u = 1$, all that matters is making at least one target, and low correlations are better. If $u = 0.5$, then the final term in (4) is zero and EU does not depend on the correlations.

We might think of u in the symmetric two-attribute case as a measure of risk aversion with respect to the number of targets that are met. If $u < (=, >)0.5$, utility as a function of the number of targets that are met is convex (linear, concave), implying risk-taking (risk-neutral, risk-avoiding) behavior. For a target-oriented decision maker, this type of risk aversion dictates attitudes toward dependence, not toward spread. In the asymmetric case with $u_1 \neq u_2$, the same notion applies with $(u_1 + u_2)/2$ playing the role of u.

What about a decision maker considering changes in the dependence between the two performance levels, $\tilde{x}_1$ and $\tilde{x}_2$? An increase (decrease) in quadrant dependence for the distribution of $\tilde{\mathbf{x}}$ implies an increase (decrease) in quadrant dependence for the distribution of $\tilde{\mathbf{d}} = \tilde{\mathbf{t}} - \tilde{\mathbf{x}}$. Therefore, increasing (decreasing) dependence between $\tilde{x}_1$ and $\tilde{x}_2$ is preferable if $u_1 + u_2 < (>)1$. The same is true for dependence between $\tilde{t}_1$ and $\tilde{t}_2$.

Finally, we consider shifts in the spread of a distribution, using the following definition.

DEFINITION 2. The c.d.f. F^* has greater spread than the c.d.f. F (F has less spread than F^*) in the sense of median-preserving spread if $F^*(d) > (=, <)F(d)$ for all d such that $F(d) < (=, >)0.50$.

REMARK. Median-preserving spread is similar in spirit to the mean-preserving spread of Rothschild and Stiglitz (1970). For both, greater spread involves some shift of

probability from the center of the distribution to the tails, while holding a measure of location constant. In many decision-making applications, the mean of the distribution is the important measure of location. However, in a target-oriented situation, the key issue is whether the probability of meeting a target is less than or greater than one-half, so one tail is of particular interest and the median is the appropriate measure of location. Hereafter, any mention of changes in spread refers to median-preserving spread.

PROPOSITION 4. *If $\tilde{d}_1$ and $\tilde{d}_2$ are independent, $F_{d_i}^*$ has greater spread than F_{d_i}, $F_{d_{3-i}}^* = F_{d_{3-i}}$, and $0 < F_{d_{3-i}}(0) < 1$, $i = 1, 2$, then $\mathrm{EU}(F_{d_1, d_2}^*) > (=, <)\mathrm{EU}(F_{d_1, d_2})$ if $F_{d_i}(0) < (=, >)0.5$.*

PROOF. Suppose that $F_{d_1}^*$ has greater spread than F_{d_1} and $F_{d_2}^* = F_{d_2}$. From (4) and the independence of $\tilde{d}_1$ and $\tilde{d}_2$,

$$\mathrm{EU}(F_{d_1, d_2}^*) - \mathrm{EU}(F_{d_1, d_2})$$
$$= u_1[F_{d_1}^*(0) - F_{d_1}(0)]$$
$$\quad + (1 - u_1 - u_2)F_{d_2}(0)\big[F_{d_1}^*(0) - F_{d_1}(0)\big]$$
$$= [F_{d_1}^*(0) - F_{d_1}(0)]\big\{u_1[1 - F_{d_2}(0)] + (1 - u_2)F_{d_2}(0)\big\}$$
$$> (=, <)0$$

if $F_{d_1}^*(0) > (=, <)F_{d_1}(0)$, which holds if $F_{d_1}(0) < (=, >)$ 0.5 because $F_{d_1}^*$ has greater spread than F_{d_1}. The proof is identical if the roles of $F_{d_1}^*$ and $F_{d_2}^*$ are reversed. $\square$

Proposition 4 shows that greater spread in F_{d_i} is preferred if the probability of meeting the ith target is less than 0.5, and less spread is preferred if that probability is greater than 0.5. In the former case, the important effect of increasing the spread is to move probability into the lower tail, specifically below $\tilde{d}_i = 0$, where the target is met. As the spread gets larger, the probability of meeting the target increases, approaching 0.5. In the latter case, decreasing the spread moves probability out of the upper tail, specifically above $\tilde{d}_i = 0$, where the target is not met. Although Proposition 4 involves a shift in spread for only one attribute, a shift in the same direction for both attributes will yield the same result. Note that the preference for an increase or decrease in spread is not related to risk aversion; it is simply driven by a desire to increase the probability of meeting the target, a motivation similar to increasing the probability of winning a contest (Gaba et al. 2004). Finally, the spread considered in Proposition 4 is that of $\tilde{d}_i = \tilde{t}_i - \tilde{x}_i$. Because $\tilde{t}$ and $\tilde{x}$ are independent, this spread can be increased (decreased) through an increase (decrease) in the spread of either $\tilde{t}_i$ or $\tilde{x}_i$. If the distribution of targets is indeed out of the decision maker's control, any adjustment would have to involve the spread of $\tilde{x}_i$.

Of the results in §3.1, only Proposition 4 imposes the condition of independence of $\tilde{d}_1$ and $\tilde{d}_2$. Without this independence, the impact of a change in spread affects the joint cumulative probability $F_{d_1, d_2}(0, 0)$ in a more complicated fashion than just through the product of the marginal probabilities. We return to this point in the context of a normal model in §3.2.

3.2. Two Targets, Normal Distributions

For additional insight into the impact of shifts in location, spread, and dependence, we consider a normal model and generate some numerical results. Assume that the decision maker's performance $\tilde{\mathbf{x}}$ is normally distributed with means μ_1 and μ_2, standard deviations σ_1 and σ_2, and correlation ρ_x. The vector of targets $\tilde{\mathbf{t}}$ is independent of $\tilde{\mathbf{x}}$ and is also normally distributed, with means λ_1 and λ_2, standard deviations τ_1 and τ_2, and correlation ρ_t. Then the distribution of $\tilde{\mathbf{t}} - \tilde{\mathbf{x}}$ is normal with means $\lambda_i - \mu_i$ and standard deviations $\gamma_i = \sqrt{\sigma_i^2 + \tau_i^2}$ for $i = 1, 2$ and correlation $r = (\sigma_1\sigma_2\rho_x + \tau_1\tau_2\rho_t)/\gamma_1\gamma_2$. From (4), the decision maker's expected utility is

$$\mathrm{EU}(F_{d_1, d_2}) = u_1\Phi_1(z_1) + u_2\Phi_1(z_2)$$
$$\quad + (1 - u_1 - u_2)\Phi_2(z_1, z_2, r), \qquad (5)$$

where $z_i = (\mu_i - \lambda_i)/\gamma_i$ and Φ_n is the c.d.f. of the standard n-variate normal distribution. Note that Φ_1 is increasing and Φ_2 is increasing in all three arguments.

In terms of decision making, we assume that the parameters λ_1, λ_2, τ_1, τ_2, and ρ_t of the distribution of targets are out of the decision maker's control. However, the decision maker may be able to manipulate μ_1, μ_2, σ_1, σ_2, and ρ_x, the parameters of the performance distribution, so the impact on $\mathrm{EU}(F_{d_1, d_2})$ of changes in these parameters is of interest.

PROPOSITION 5. *Consider the two-attribute target-oriented situation with bivariate normal distributions for $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$.*

(a) $\mathrm{EU}(F_{d_1, d_2})$ *is increasing in μ_1 and μ_2.*

(b) $\mathrm{EU}(F_{d_1, d_2})$ *is increasing (constant, decreasing) in ρ_x and ρ_t if $u_1 + u_2 < (=, >)1$.*

(c) *If $\rho_x = \rho_t = 0$, $\mathrm{EU}(F_{d_1, d_2})$ is increasing (constant, decreasing) in σ_i if $\mu_i - \lambda_i < (=, >)0$.*

PROOF. Part (a) follows from Proposition 2 because an increase in μ_i ($i = 1, 2$) is a positive location shift in $\tilde{x}_i$, and a larger such shift results in a larger increase in $\mathrm{EU}(F_{d_1, d_2})$ for the strictly increasing normal c.d.f. Part (b) follows from Proposition 3, the fact that a larger correlation r implies strictly larger quadrant dependence, and the fact that r is increasing in ρ_x and ρ_t. Part (c) follows from Proposition 4 because an increased (decreased) σ_i means greater (less) spread. These results can also be verified by differentiating $\mathrm{EU}(F_{d_1, d_2})$ with respect to its parameters. $\square$

The results of Proposition 5 are consistent with the discussion of Propositions 2–4 in §3.1. Numerical examples can provide further insight into the nature and extent of changes in EU as parameters are varied. For instance, Figure 1 shows EU as a function of r for $u_1 + u_2 = 0(0.5)2$ when $\mu_i - \lambda_i = 0$, $i = 1, 2$. The bottom curve corresponds to the case in which the decision maker cares only about making both targets; it gives the probability of making both targets, which increases from 0 to 0.5 as r goes from -1 to 1. The top curve, which is a mirror image of the bottom curve, applies when all that matters is making at least one target.

Figure 1. Expected utility as a function of correlation r for $u_1 + u_2 = 0(0.5)2$ when $\mu_i - \lambda_i = 0$, $i = 1, 2$.
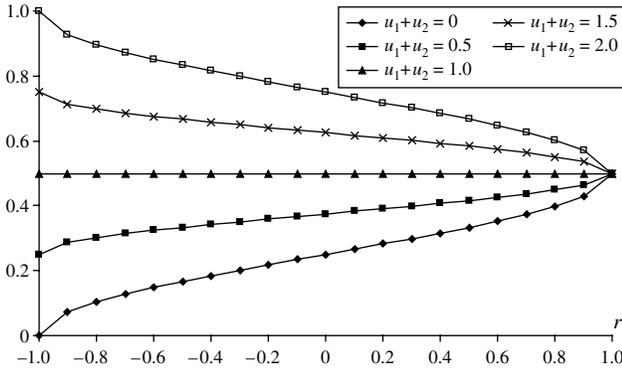
As $u_1 + u_2$ moves from these extremes of zero and two toward one, EU becomes less sensitive to r. The correlation ρ_x affects EU through r, with $\partial r/\partial \rho_x = \sigma_1 \sigma_2/\gamma_1 \gamma_2$. Thus, $0 < \partial r/\partial \rho_x < 1$, and the impact of a change in ρ_x is greater (less) when σ_1 and σ_2 are larger (smaller) relative to τ_1 and τ_2.

Relaxing the independence assumptions in Proposition 5(c) is of particular interest so that we can better understand the implications of modifying σ_i. If either ρ_x or ρ_t is nonzero, then $r \neq 0$, which influences the impact of σ_i on EU through the last term of (5): $\partial \Phi_2/\partial r > 0$ (as illustrated in the bottom curve in Figure 1), and $\partial r/\partial \sigma_i > (=, <)0$ if $\sigma_i \tau_{3-i} \rho_t < (=, >) \sigma_{3-i} \tau_i \rho_x$. In Figure 2, we show the probability of making both targets (EU with $u_1 = u_2 = 0$) as a function of σ_1 for five choices of (ρ_x, ρ_t). The results are analogous for σ_2. As in Figure 1, $\mu_i - \lambda_i = 0$, $i = 1, 2$, which enables us to isolate the impact of changes in σ_1 related to ρ_x and ρ_t because varying these parameters will not change $z_1 = z_2 = 0$. We set $\sigma_2 = \tau_1 = \tau_2 = 1$ without much loss of generality because other values will just rescale the curves.

Figure 2 illustrates the five types of curves that can occur in this setting:

(1) If $\rho_x = \rho_t = 0$, EU $= 0.25$ for all σ_1, a horizontal line.

(2) If $\rho_x \geqslant 0$ and $\rho_t \leqslant 0$ with one or both inequalities strict, EU is strictly increasing in σ_1.

(3) If $\rho_x \leqslant 0$ and $\rho_t \geqslant 0$ with one or both inequalities strict, EU is strictly decreasing in σ_1.

(4) If $\rho_x > 0$ and $\rho_t > 0$, EU is increasing in σ_1 at first and then decreasing, with a maximum at $\sigma_1 = (\sigma_2 \tau_1/\tau_2)(\rho_x/\rho_t)$, which equals $\rho_x/\rho_t = 0.7/0.7 = 1$ in Figure 2 because $\sigma_2 \tau_1/\tau_2 = 1$.

(5) If $\rho_x < 0$ and $\rho_t < 0$, EU is decreasing in σ_1 at first and then increasing, with a minimum at $\sigma_1 = (\sigma_2 \tau_1/\tau_2)(\rho_x/\rho_t)$, which equals $\rho_x/\rho_t = -0.7/(-0.7) = 1$ in Figure 2 because $\sigma_2 \tau_1/\tau_2 = 1$.

Thus, depending on ρ_x and ρ_t, σ_1 might not matter or the decision maker might want σ_1 to be as large as possible,

Figure 2. Probability of making both targets as a function of σ_1 for selected values of ρ_x and ρ_t when $\mu_i - \lambda_i = 0$, $i = 1, 2$, and $\sigma_2 = \tau_1 = \tau_2 = 1$.
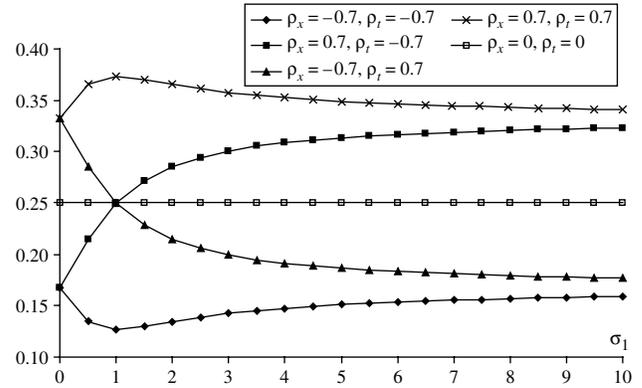

as small as possible, or at some intermediate value, with nontrivial variations in EU as σ_1 varies.

The parameters in Figure 3 are identical to those in Figure 2 except that $\mu_1 - \lambda_1 = 0.5$. If $\rho_x = \rho_t = 0$, Proposition 5(c) indicates that EU is decreasing in σ_1 when $\mu_1 - \lambda_1 > 0$. The curve for $\rho_x = \rho_t = 0$ in Figure 3 is indeed decreasing, as are two other curves, implying that lower values of σ_1 are preferred. However, the curves for $\rho_x = \rho_t = 0.7$ and $(\rho_x, \rho_t) = (0.7, -0.7)$ increase first before decreasing, peaking at $\sigma_1 = 0.5$ and 8.5, respectively, given the increments of 0.5 used for σ_1 in the figure (the latter maximum is difficult to see because the curve is quite flat). The increasing portions of these curves represent segments where increases associated with the correlations, as in Figure 2, outweigh the decreases due to the positive $\mu_1 - \lambda_1$. In this sense, the balancing of the two effects can temper the primary effect via $\mu_1 - \lambda_1$ of wanting to make σ_1 as small as possible, sometimes even leading to a relatively large preferred value of σ_1. Results (not presented) for negative $\mu_1 - \lambda_1$ are similar, with EU increasing in σ_1 if $\rho_x = \rho_t = 0$.

The curves in Figures 2 and 3 represent EU if $u_1 = u_2 = 0$, as noted above. As u_1 and u_2 increase with $u_1 + u_2 < 1$,

Figure 3. Probability of making both targets as a function of σ_1 for selected values of ρ_x and ρ_t when $\mu_1 - \lambda_1 = 0.5$, $\mu_2 - \lambda_2 = 0$, and $\sigma_2 = \tau_1 = \tau_2 = 1$.
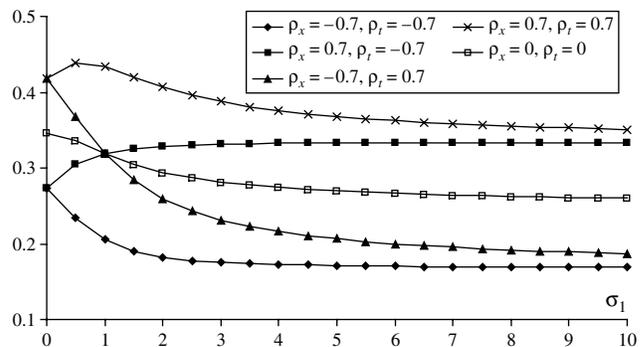
the curves retain the same general shapes and are closer together. When $u_1 + u_2 > 1$, the curves "flip;" for example, the curve with positive correlations moves from bottom to top and the curve with negative correlations moves from top to bottom. The curves are closer together as $u_1 + u_2 \to 1$ and further apart as $u_1 + u_2 \to 2$.

The implications for the decision maker who can modify her distribution of performance levels are straightforward. It is always desirable to increase the means, μ_1 and μ_2, and to increase (decrease) the correlation, ρ_x, if $u_1 + u_2 < (>)1$. The primary effect of modifying the standard deviations of $\tilde{x}_1$ and $\tilde{x}_2$ when correlations equal zero suggests increasing (decreasing) a standard deviation σ_i when the corresponding μ_i is less than (greater than) the expected target λ_i, but this effect can be tempered and even reversed if the correlations are nonzero.

3.3. Extension to n Targets

The two-target case can be extended formally to the case of targets on n attributes. From (3), for example,

$$\mathrm{EU}(F_{d_1,\ldots,d_n}) = \sum_A w_A F_{\{d_i \mid i \in A\}}(\mathbf{0}), \qquad (6)$$

where $\sum_A w_A = 1$ and $\mathbf{0}$ is a conformable vector of zeros. Such an extension is at the cost of greater complexity as the number of terms in $\mathrm{EU}(F_{d_1,\ldots,d_n})$ increases exponentially and the weights w_A become more complicated linear combinations of the $2^n - 1$ possible utilities u_B (excluding $u_\varnothing = 0$), making it difficult to get insight from (6) except in special cases.

Some results do generalize to the case of n targets, however. For instance, Corollary 1 and Proposition 5(a) extend directly to n attributes. EU is nondecreasing with a negative location shift in any $\tilde{d}_i$, which can be caused by a positive location shift in $\tilde{x}_i$ or a negative location shift in $\tilde{t}_i$. For the normal model, it is always desirable to increase any of the performance means $\mu_1,\ldots,\mu_n$ or decrease any of the target means $\lambda_1,\ldots,\lambda_n$.

Propositions 4 and 5(c) can be extended by induction to n attributes. If $\tilde{d}_1,\ldots,\tilde{d}_n$ are independent, an increase in the spread of $\tilde{d}_i$ for one or more $i \in \{1,\ldots,n\}$ while holding spreads constant for the other attributes will increase (decrease) EU if $F_{d_i}(0) < (>)0.5$ for all $\tilde{d}_i$ with increased spreads. If $\tilde{d}_1,\ldots,\tilde{d}_n$ are not independent, however, the impact of changes in spread can vary greatly depending on the many different combinations of dependence that can occur. With the normal model, for example, any of the $n(n-1)/2$ pairwise correlations can vary in sign and magnitude, subject only to the constraint that the correlation matrix be positive definite. Figure 3 and the discussion in §3.2 about relaxing the independence assumptions in Proposition 5(c) provide some indication of the effects of dependence on the impact of changes in spread and show that such effects may be difficult to untangle. Things get even more complex and more difficult to untangle as

n increases, although it is still feasible to generate numerical results.

Propositions 3 and 5(b) cannot be extended fully to n attributes because generalizations of quadrant dependence do not easily cover all possible cases. Extensions for some special cases are possible, however. For example, F^* has larger lower (upper) orthant dependence than F if $F^*(\mathbf{d}) \geqslant F(\mathbf{d})$ $[H^*(\mathbf{d}) \geqslant H(\mathbf{d})]$ for all $\mathbf{d}$, where $H(\mathbf{d}) = P_F(\tilde{\mathbf{d}} > \mathbf{d})$ (Joe 1997). If the only thing that matters is meeting all targets (i.e., if $u_A = 0$ for all A except $A = \{1,\ldots,n\}$), then $\mathrm{EU}(F) = F(\mathbf{0})$. Thus, $\mathrm{EU}(F^*) \geqslant \mathrm{EU}(F)$ if F^* has larger lower orthant dependence than F. At the other extreme, if the only thing that matters is making at least one target (i.e., $u_A = 1$ for all A except $A = \varnothing$), then $\mathrm{EU}(F) = 1 - H(\mathbf{0})$. In this case, $\mathrm{EU}(F^*) \geqslant \mathrm{EU}(F)$ if F^* has smaller upper orthant dependence than F.

4. Summary and Discussion

Target-oriented situations, in which objectives can be expressed in terms of meeting targets on one or more attributes, are common, as illustrated by the examples given in BK and in §§1–2 of this paper. In many cases, the targets are uncertain; for example, they might not be set in advance or they might be driven by competition. Furthermore, multiattribute targets are often dependent, as discussed in §2; for example, negative dependence can occur due to trade-offs among the attributes, and positive dependence can be caused by different attributes being influenced in a similar way by some common factor such as economic conditions.

In §3, we develop a model of multiattribute targets, allowing dependence among the performance levels $\tilde{x}_i$ and among the targets $\tilde{t}_i$. We focus primarily on the case of two attributes, which provides useful insights into the impact of dependence and other distributional characteristics on expected utility. We show that a positive location shift in any $\tilde{x}_i$ or a negative location shift in any $\tilde{t}_i$ increases EU; that increasing the degree of quadrant dependence of $\tilde{x}_1$ and $\tilde{x}_2$ (or of $\tilde{t}_1$ and $\tilde{t}_2$) increases (decreases) EU if the sum of the utilities for making just one target is less than (greater than) the utility for making both targets; and that increasing (decreasing) the spread of $\tilde{d}_i = \tilde{t}_i - \tilde{x}_i$ when $\tilde{d}_1,\ldots,\tilde{d}_n$ are independent increases EU if the probability of meeting the target is less than (greater than) 0.5.

For a better sense of the impact of shifts in location, spread, and dependence, we generate some numerical results for a model in which $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$ have bivariate normal distributions. Of particular interest here is the impact of changes in spread when the independence assumption in Propositions 4 and 5(c) is relaxed. When at least one of the correlations ρ_x and ρ_t is nonzero, EU can be strictly increasing, strictly decreasing, or nonmonotonic as the spread of $\tilde{x}_i$ increases, with the exact form of the curve depending on the difference between the performance mean and the target mean for attribute i as well as on the correlations.

The model can be extended to the n-target case at a cost of increased complexity because of the exponential growth of the possible number of combinations of targets

that can be met. However, some of the results obtained for two targets do generalize to n targets. A positive location shift in any performance level or a negative location shift in any target is always desirable, and increasing (decreasing) the spread of any $\tilde{d}_i$ when $\tilde{d}_1, \ldots, \tilde{d}_n$ are independent is desirable if the probability of meeting the target on the ith attribute is less than (greater than) 0.5.

The impact of increasing dependence is trickier in the n-target case, although we can say that increased lower orthant dependence is desirable if the only thing that matters is meeting all targets and decreased upper orthant dependence is desirable if the only thing that matters is meeting at least one target. These situations belong to the class of symmetric n-target situations with $u_A = u_{|A|}$, where things are simplified a bit because the number of possibly nonzero utilities in the model is reduced from $2^n - 1$ to n. As in the two-target case, we can think of the decision maker as being risk taking (risk neutral, risk averse) in terms of the number of targets that are met if $u_{|A|}$ is convex (linear, concave) as a function of $|A|$.

Target-oriented utilities can be thought of in a primitive sense as a direct reflection of preferences for meeting targets. An alternative interpretation is to view such utilities as being driven by payoffs associated with meeting particular sets of targets. Suppose that there is a monetary payoff $\pi_A > 0$ for each specific combination A of targets that are met, and that the decision maker's interest in meeting targets is solely based on these monetary payoffs. Then, $u_A = U_\pi(\pi_A)$, where U_π is the decision maker's utility function for money. From Propositions 3 and 5(b), a decision maker in a two-target situation prefers higher (lower) correlations ρ_x and ρ_t if $u_1 + u_2 < (>)u_{1,2}$, which implies that $\pi_1 + \pi_2 < (>)\pi_{1,2}$ if the decision maker is risk neutral with respect to money. For any set of payoffs with $\pi_{1,2} > \max\{\pi_1, \pi_2\}$, a sufficiently concave (convex) U_π will yield $u_1 + u_2 > (<)u_{1,2}$. In that sense, a risk avoider is more likely to prefer negative correlations and a risk taker is more likely to prefer positive correlations.

We have assumed that $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$ are independent of each other. In some cases, $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$ may be dependent. For example, if attribute i represents a firm's revenues, $\tilde{x}_i$ and $\tilde{t}_i$ are likely to be positively correlated through a common factor, economic conditions. In a golf tournament, the performance of a given golfer and the target (the total score needed to win) will be positively correlated because all of the golfers are affected similarly by the grooming of the course and by weather conditions. All of the results of §3.1 hold if $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{t}}$ are dependent because the crucial distribution is that of $\tilde{\mathbf{d}} = \tilde{\mathbf{t}} - \tilde{\mathbf{x}}$. The results of §3.2 can be generalized by considering a normal distribution for $(\tilde{\mathbf{x}}, \tilde{\mathbf{t}})$ with the correlation of $\tilde{x}_i$ and $\tilde{t}_j$ denoted by $\rho_{x_i t_j}$. For $n = 2$, EU is as in (5) with $\gamma_i = \sqrt{\sigma_i^2 + \tau_i^2 - 2\sigma_i \tau_i \rho_{x_i t_i}}$ and $r = (\sigma_1 \sigma_2 \rho_x + \tau_1 \tau_2 \rho_t - \sigma_1 \tau_2 \rho_{x_1 t_2} - \sigma_2 \tau_1 \rho_{x_2 t_1})/\gamma_1 \gamma_2$. If all correlations are nonzero only within attributes and not between attributes (i.e., $\rho_x = \rho_t = \rho_{x_i t_j} = 0$ for $i \neq j$), which might often be a reasonable assumption, then $r = 0$, and the only impact of correlations is through the standard deviations γ_i, via $\rho_{x_i t_i}$.

A strict target-oriented approach assumes that whenever an additional target is met, there is a "step" in the utility function, and the utility function is flat between steps. In a less strict interpretation, which may often be appropriate, the step and the flatness may be a bit extreme, with utilities possibly showing some increase between steps and/or a steep but continuous increase around the target. This is particularly likely when the target is not externally generated with a specific prize (e.g., winning first prize in a tournament), but is internally generated and can be thought of as an aspiration level or a reference point. As noted in BK (2004, p. 823), Heath et al. (1999) suggest that the inflection point in Kahneman and Tversky's (1979) S-shaped value function can be interpreted as a target. We anticipate that as long as meeting the targets is viewed as quite important by the decision maker, the strategies implied by the strict interpretation should be quite reasonable, robust, and applicable.

Acknowledgments

References

Abbas, A. E., R. A. Howard. 2005. Attribute dominance utility. *Decision Anal.* **2**(4) 185–206.

Abbas, A. E., J. E. Matheson. 2005. Normative target-based decision making. *Managerial Decision Econom.* **26**(6) 373–385.

Borch, K. 1968. Decision rules depending on the probability of ruin. *Oxford Econom. Papers* **20**(1) 1–10.

Bordley, R. F., C. W. Kirkwood. 2004. Multiattribute preference analysis with performance targets. *Oper. Res.* **52**(6) 823–835.

Bordley, R. F., M. LiCalzi. 2000. Decision analysis using targets instead of utility functions. *Decisions Econom. Finance* **23**(1) 53–74.

Castagnoli, E., M. LiCalzi. 1996. Expected utility without utility. *Theory Decision* **41**(3) 281–301.

Gaba, A., I. Tsetlin, R. L. Winkler. 2004. Modifying variability and correlations in winner-take-all contests. *Oper. Res.* **52**(3) 384–395.

Heath, C., R. Larrick, G. Wu. 1999. Goals as reference points. *Cognitive Psych.* **38**(1) 79–109.

Hymowitz, C. 2005a. When meeting targets becomes the strategy, CEO is on wrong path. *Wall Street Journal* (Eastern ed., March 8) B.1.

Hymowitz, C. 2005b. Readers share tales of jobs where strategy became meeting target. *Wall Street Journal* (Eastern ed., March 22) B.1.

Joe, H. 1997. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London, UK.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2) 263–291.

Keeney, R. L., H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York.

Lehmann, E. L. 1966. Some concepts of dependence. *Ann. Math. Statist.* **37**(5) 1137–1153.

Rothschild, M., J. E. Stiglitz. 1970. Increasing risk: I. A definition. *J. Econom. Theory* **2**(3) 225–243.

Simon, H. A. 1955. A behavioral model of rational choice. *Quart. J. Econom.* **69**(1) 99–118.

Yanagimoto, T., M. Okamoto. 1969. Partial orderings of permutations and monotonicity of a rank correlation statistic. *Ann. Inst. Statist. Math.* **21** 489–506.