

Decision analysis using targets instead of utility functions

Robert Bordley¹, Marco LiCalzi²

¹ Knowledge Network, GM Research Development Center
e-mail: robert.bordley@gm.com

² Department of Applied Mathematics, University “Ca’ Foscari” of Venice
e-mail: licalzi@unive.it

Received: 12 December 1999 / Accepted: 4 March 2000

Abstract. A common precept of decision analysis under uncertainty is the choice of an action which maximizes the expected value of a utility function. Savage’s (1954) axioms for subjective expected utility provide a normative foundation for this principle of choice. This paper shows that the same set of axioms implies that one should select an action which maximizes the probability of meeting an uncertain target. This suggests a new perspective and an alternate target-based language for decision analysis. We explore the implications and the advantages of this target-based approach for both individual and group decision-making.

Mathematics Subject Classification (2000): 91B06, 91B10, 90B50

Journal of Economic Literature Classification: D81

1. Introduction

Two of the most important approaches to decision making are optimizing and satisficing. In a world with no uncertainty, where each feasible action d leads to a (known) consequence x_d , the optimizing approach postulates that the agent should choose an action d which maximizes his value function $v(d)$. The satisficing approach, instead, explicitly recognizes the cost or the practical impossibility of searching among all possible actions for the

We thank Erio Castagnoli, Stefano Della Vigna, Massimo Marinacci and the seminar audiences at the universities of Calabria, Michigan and Paris IX (Dauphine) for helpful comments. This paper was mostly written while the second author was enjoying the warm hospitality of the Department of Economics at Purdue University.

optimal one; see Simon (1955). Therefore, it suggests that the agent should establish some target t and then pick the *first* action d which meets the target.

Rational decision making is based on the optimizing principle. However, the satisficing approach has some appealing features because thinking about targets is very natural in many problems. Moreover, as it is often easier to formulate targets and monitor their achievement, they are prominent in the decision processes of bureaucracies and organizations.

The purpose of this paper is to bring a target-based language to decision making under uncertainty, while maintaining the optimizing principle as a guide to rational decisions. The satisficing approach is sufficient but not necessary to make target-based decisions.

When there is uncertainty, an action d may lead to different outcomes, usually summarized in a random consequence X_d . Given a target t , then, the agent can only assess the probability $P(X_d \geq t)$ that the action d leads to a consequence that meets the target. According to the optimizing principle, the agent should choose an action d which maximizes the probability $v(d) = P(X_d \geq t)$ of meeting the target t ; see Manski (1988).

This simple target-based model is not complete because there may be uncertainty about the target itself. For example, many businesses trying to meet customer requirements set these as their target. But, while government contractors frequently have exact specifications to satisfy, commercial firms must usually view their customers' requirements as uncertain. The target may not even be known in individual contexts. Many individuals have the goal of "being successful"; but only a very few know precisely which combinations of money, leisure time, culture, etc. must be attained to achieve this goal.

Hence, we relax the assumption of a known target and replace t with a random consequence T . The ensuing target-based decision model prescribes that the agent should choose an action d which maximizes the probability $v(d) = P(X_d \geq T)$ of meeting an uncertain target. Borch (1968) develops a closely related approach.

To assess the normative appeal of this model for decision analysis, we consider the subjective expected utility theory perfected in Savage (1954). It provides an axiomatic foundation which implies that the agent should choose an action d which maximizes his expected utility $v(d) = EU(X_d)$ with respect to a subjective probability distribution.

Quite unexpectedly, the target-based model satisfies the Savage axioms! As we prove in the appendix, there is no way to tell if an agent abides by Savage's axioms is maximizing his expected utility or is maximizing the probability of meeting his uncertain target. A similar result for the von Neumann and Morgenstern's axiomatization is given in Castagnoli and LiCalzi (1996).

This implies that Savage's axioms allow us to use two different languages in formulating our decision models – the language of utilities or the language

of targets. As it turns out, after a trivial normalization, translating between the two languages is simply a matter of setting $U(x) = P(x \geq T)$, that is, interpreting the “utility” of a consequence x as the subjective probability that x will meet the target. The choice between a utility-based language or a target-based language for decision theory ultimately comes down to a question of convenience.

We explore the implications of this point of view in the rest of this paper, which is organized as follows. Section 2 compares the target-based model with the utility-based model for decision analysis under uncertainty. The following three sections suggest a few applications where the target-based model may have a comparative advantage over the utility-based formulation. More precisely, Section 3 considers individual decision making while Sect. 4 looks at group decision making; finally, Section 5 examines multiattribute decision making, both at the individual and at the group level. Section 6 summarizes and offers our conclusions.

2. The target-based model

Suppose that the agent has to rank several possible decisions. Assume for simplicity that the set C of consequences is finite and completely ordered by a preference relation \succeq . Denote by P_d his probability distribution for the random consequence X_d associated¹ with a decision d . The expected utility model suggests that the ranking be obtained by using the value function

$$v(d) = EU(X_d) = \sum_x U(x)P_d(x), \quad (1)$$

where $U(x)$ is a von Neumann and Morgenstern (NM-)utility function over consequences.

The target-based model, instead, suggests using the value function

$$v(d) = P(X_d \geq T) = \sum_x P(x \geq T)P_d(x), \quad (2)$$

where $P(x \geq T)$ is the cumulative distribution function of the uncertain target and T is stochastically independent of X_d .

Both models are linear functionals over the probability distributions associated with a decision. This formal similarity may be carried one step further. Note that U is bounded and increasing. After having normalized its range to $[0, 1]$, U has all the properties of a cumulative distribution function (c.d.f.) over consequences. By a standard probability-theoretic argument,

¹ Let P be the agent’s subjective probability distribution on the state space S . The probability distribution P_d is induced by the act $d : S \rightarrow C$ through the equality $P_d(X_d = x) = P(d(s) = x)$.

we can associate to the c.d.f. U a random consequence T stochastically independent of X_d ; see the appendix. We can thus view $U(x)$ as the probability that the target T is below x ; that is, $U(x) = P(x \geq T)$. This definition turns the NM-utility function U into a probability distribution and makes (1) and (2) formally identical.

The idea that the NM-utility function U should be interpreted as a probability distribution may appear unusual but, in fact, NM-utilities *are* probabilities. Consider how the agent assesses the utility $U(x)$ of a consequence x . For simplicity, assume that there is a best consequence b and a worst consequence w , so that $b \succeq x \succeq w$. Then $U(x)$ is the probability p that makes the agent indifferent between getting x for sure or playing a lottery where he has probability p of getting b and probability $(1 - p)$ of getting w ; that is, $U(x) = p$.

It is the utility-based language that baptizes the probability p as the NM-utility of x , as some decision theorists have noted long since. For instance, Pratt, Raiffa and Schlaifer (1995) equate utility with the “indifference probability”. Similarly, both Behn and Vaupel (1982) and Howard (1992) speak of utility as a “preference probability” and remark that the manipulation of preference probabilities is in many respects similar to the treatment of real probabilities. Nonetheless, they are very careful in keeping preference probabilities distinct from real probabilities because a preference probability is not associated with a potentially observable event.

The target-based model accepts $U(x)$ as a real probability only after having introduced the notion of an uncertain target T : the potentially observable event associated with the probability $U(x)$ is that x does meet the target.

As a result, the target-based model can avoid the notion of a cardinal utility function $U(x)$ over consequences. This has a major pedagogical advantage. While both target-based and utility-based decision making demand an understanding of probabilistic choice, the utility-based model additionally requires a comprehension of cardinal utility assessments. For target-based reasoning, the agent must only be able to handle probability judgements.

This has a potential for simplifying the assessment procedure for $U(x)$. For instance, assume that consequences are monetary outcomes described as changes to the current endowment. A simple way to elicit the distribution of his target is to let the agent draw a probability density $u(x)$ over the real line and then estimate his c.d.f. as

$$U(x) = P(x \geq T) = \int_{-\infty}^x u(t) dt. \quad (3)$$

In the utility-based language, the probability density $u(x)$ corresponds to the marginal utility associated with the NM-utility function U . It may be interesting to compare the difficulty of eliciting a probability density $u(x)$

versus the task of assessing the marginal utility of U . Berhold (1973) exploits a similar insight to estimate certainty equivalents.

We have drawn two probability densities for T in Fig. 1. On the left, there is a decreasing probability density $u_1(x)$ that suggests a conservative assessment of the target: the most likely targets to meet are the worst outcomes. On the right, there is a symmetric unimodal probability density $u_2(x)$ that suggests an uncertain target very close to the *status quo*, represented by the mode of the distribution and corresponding to a zero change in the endowment.




Fig. 1. Two probability densities

Figure 2 shows the cumulative distribution functions of the targets associated with these probability densities. By (3), the conservative assessment on the left of Fig. 1 leads to the concave c.d.f. $U_1(x)$ on the left of Fig. 2, which is equivalent to a concave NM-utility function. Therefore, in using a target-based model, global risk aversion stems from a conservative assessment of the target.




Fig. 2. Two c.d.f.'s for the target

Looking at the right-hand side of Fig. 1, the unimodal probability density for the target corresponds to the *S*-shaped c.d.f. $U_2(x)$. This is equivalent to the *S*-shaped utility function of prospect theory, according to which people tend to be risk averse over gains and risk seeking over losses; see Kahneman

and Tversky (1979). In the target-based language, this amounts to saying that perceiving the uncertain target as (more or less) symmetrically distributed around the *status quo* leads to risk averting choices when gains are involved and to risk seeking decisions when losses are at stake.

Many variations on this theme are possible, but we report only one more. Suppose that the probability density for the target is unimodal and positively skewed (e.g., a lognormal distribution). Again, the mode represents the *status quo*. Consistent with prospect theory, the resulting distribution function $U(x)$ will not only be concave over gains and convex over losses, but steeper for losses than for gains. In the target-based language, a unimodal probability density for the target with fatter right tails implies a behavior that is less risk averse over gains than it is risk seeking over losses.

3. Individual decision making

The previous section and the appendix show that there is a formal equivalence between the target-based model and the utility-based model. Therefore, the decision analyst may wish to transfer insights from one model to the other. This section considers a few examples concerning the relationships between expected utility models and target-based reasoning in individual decision making.

3.1. State-dependence

Suppose that an agent is interested in maximizing his expected utility. By our result of formal equivalence, his decision process can be modelled as if he is maximizing the probability of meeting an uncertain target T with c.d.f. $U(x)$, provided that the target is stochastically independent of the random consequences to be evaluated. Since it is always possible to make sure that the target satisfies this assumption, there are no restrictions in going from the utility-based model to the target-based formulation.

More care, instead, is needed for going in the opposite direction because the target-based model allows for higher generality. Suppose that the agent is interested in maximizing the probability of meeting an uncertain target T . If the target T is stochastically independent of X_d , we know that

$$P(X_d = x, x \geq T) = P(x \geq T)P_d(x). \quad (4)$$

Therefore, the value function is

$$P(X_d \geq T) = \sum_x P(X_d = x, x \geq T) = \sum_x P(x \geq T)P_d(x),$$

which is equivalent to an expected utility model if we define $U(x) = P(x \geq T)$. But when the assumption of stochastic independence does not hold, (4) may fail and we have to resort to more general formulations.

There are situations in which the assumption of stochastic independence is too strong. Suppose that there are two different decisions d_1 and d_2 both resulting in x when the states are respectively s_1 and $s_2 \neq s_1$. For example, assume that X_1 and X_2 are the (nominal) rates of return on two stock portfolios and that there are two states $s_1 = H$ and $s_2 = L$ representing respectively high and low inflation. Imagine that $X_1 = x$ in state H , while $X_2 = x$ in state L . The otherwise identical nominal rate of return x should probably be weighed differently if it obtains in a state of high or low inflation. The assumption of stochastic independence, instead, entails

$$\begin{aligned} P(x \geq T | s = H) &= P(x \geq T | X_1 = x) \\ &= P(x \geq T | X_2 = x) = P(x \geq T | s = L), \end{aligned}$$

so that the probability that the nominal rate x meets the target cannot depend on the level of inflation.

In this example, the assumption of stochastic independence fails because the evaluation of x should be state-dependent. Adopting a utility-based language requires that we analyze the situation by introducing a state-dependent utility function $U(x, s)$ for the consequence x in state s ; see Karni (1985) or Schervisch, Seidenfeld and Kadane (1990). However, this is not an easy task because the eliciting of $U(x, s)$ requires the agent to compare the cardinal utility of x across different states. Moreover, since there may be different ways in which the stochastic independence of the target may fail, there is no easy general recipe to convert a target-based model into a utility-based model. It is possible that target-based reasoning may offer a more convenient approach.

Continuing in the example, the agent might, for instance, assess a target T_H for high inflation and a (possibly different) target T_L for low inflation. Different targets for different states of the world can be easily understood and formulated. Conditional on the states of the world, the agent may view the nominal rate of return X_d on the stock portfolio d as a random variable. Thus, he would have to use the value function

$$\begin{aligned} v(d) &= P(X_d \geq T) \\ &= P(X_d \geq T_H | H)P(H) + P(X_d \geq T_L | L)P(L), \end{aligned} \tag{5}$$

where T represents the state-dependent target and $P(H)$ and $P(L)$ are the (subjective) probabilities for the two states.

Note that (5) has a hierarchical structure which makes it very easy to tackle the decision problem in steps. First, the agent assesses which are the relevant states and assesses their probability. Second, for each relevant state

the agent estimates his uncertain target and the (conditional) distribution of X_d . Third, he computes the probability that X_d would meet his state-dependent target T^s in state s . Fourth, he aggregates these assessments using the laws of probability theory. The target-based model makes the decision problem easy to *divide and conquer* even if stochastic independence does not hold.

3.2. Untying correlations

There are some important cases where the lack of stochastic independence in a target-based model may be overcome by recoding the problem. In general, the uncertain target T and an arbitrary random consequence X to be evaluated may be stochastically dependent. However, assume that there exist an increasing transformation $g : C \rightarrow \mathbb{R}$ and a random number Y such that $g(T) + Y$ and $g(X) + Y$ are stochastically independent random numbers. Note that $Y = Y(X)$ may depend on the random number X that we are considering.

Define a new target $T^* = g(T) + Y$ and a new random number $X^* = g(X) + Y$. By construction, T^* and X^* are stochastically independent. Moreover, since $X \geq T$ if and only if $g(X) \geq g(T)$ and $g(X) + Y \geq g(T) + Y$ if and only if $X^* \geq T^*$, we have

$$\begin{aligned} P(X \geq T) &= P(g(X) \geq g(T)) = P(g(X) + Y \geq g(T) + Y) \\ &= P(X^* \geq T^*), \end{aligned}$$

where X^* and T^* are stochastically independent. In this situation, we can transform T and X and obtain new random numbers X^* and T^* which satisfy stochastic independence.

One special case of this model corresponds to a situation where X and T are nominal rates of return and the only source of correlation is the underlying (random) rate of inflation J . If we let $Y = -\log(1 + J)$ and $g(x) = \log(1 + x)$, then $X^* = \log[(1 + X)/(1 + J)]$ and $T^* = \log[(1 + T)/(1 + J)]$ are defined in real terms and may be taken to be stochastically independent. An appropriate recoding of targets and payoffs may eliminate some sources of correlation and recover stochastic independence. It is obvious that stochastic independence still holds if, for reasons of clarity, one introduces the transformation $h(x) = e^x - 1$ and makes reference to the (random) real rates of return $X^{**} = h(X^*)$ and $T^{**} = h(T^*)$.

Another situation where the source of correlation may be eliminated is the following. Given a random consequence X , assume that the agent evaluates X by using an *a posteriori* target T which is a function of his *a priori* target T^* and X itself. This may correspond to a situation where the agent “adapts” his target to the random consequence he is evaluating. This kind of effect

is observed quite commonly in real life. The agent who has access to some very promising random consequence X raises his expectations. The agent who is offered a poor prospect lowers his target.

In particular, suppose that the *a priori* target T^* is stochastically independent of X and that there exists some increasing function $g : C \rightarrow \mathbb{R}$ such that $g(T) = ag(X) + (1 - a)g(T^*)$ with $a = a(X) < 1$. Since $X \succeq T$ if and only if $g(X) \geq g(T)$, we have

$$P(X \succeq T) = P(g(X) \geq g(T)) = P(g(X) \geq g(T^*)),$$

where $g(X)$ and $g(T^*)$ are stochastically independent. Recoding the problem in terms of the *a priori* target eliminates the source of correlation.

For these cases where stochastic independence may be recovered, we can convert the target-based formulation into a standard expected utility model. We think that this conversion may help the agent to make better decisions in some cases. For example, suppose that X_d represents the (nominal) wage that the agent expects to earn in the next year if he accepts job d while T represents the (real) income that he feels to be necessary to maintain a satisfactory standard of living. The agent, moreover, ignores what the rate of inflation J is going to be over the next year; therefore, T is a random variable.

As shown in Shafir, Diamond and Tversky (1997), the agent is often confused about nominal and real values. Framing this problem in a utility-based language may lead him to a poor decision because the agent ends up assessing the utility of X_d without discounting away the inflationary effect and does not know how to take inflation into account. Using target-based reasoning, the agent may be asked to describe his target T in real values and his wage expectations X_d in nominal values, if this makes him feel more confident about their assessment. When the probability distributions for X_d and T have been elicited, it should not be difficult to show him how to correct his estimates and account for inflation. This procedure may be useful in many situations where the agent faces similar cognitive limitations: a firm which sets performance targets related to the growth of the market, a government that must impose different social security contributions based on changes in the demographic trends or environmental regulations based on trends in pollution, etc.

3.3. Projecting future preferences

Applied decision analysis assesses the current preferences of an agent and uses them as predictors of what his preferences will be in the future. In particular, with some possible discounting to account for time preferences, the current NM-utility function is often used as the best estimate of the

NM-utility function in the future. Target-based decision theory suggests an alternative way of modelling the relationship between today's preferences and tomorrow's preferences.

The target-based model interprets the current utility function as a description of the agent's uncertainty about his current target. Similarly, it must interpret his future utility function as describing the uncertainty about his target in the future from today's viewpoint. In most cases, it is reasonable to expect greater uncertainty about the future target than about the current target. Hence, the target-based approach indicates that today's estimate of the utility function in the future might differ systematically from today's utility function.

There are many ways to model this systematic difference. For concreteness, we consider only one example which should however suffice to convey the main point. Suppose that the consequences in C are real numbers and that the current utility function $U(x)$ corresponds to a normal c.d.f. for the uncertain current target. Once $U(x)$ is assessed, then, we may compute its mean m_0 and its variance v_0 . For convenience, we can think of m_0 as the current *status quo*. As discussed in Section 2, a normal distribution for the target generates risk-averse preferences for $x \geq m_0$ and risk-seeking preferences for $x \leq m_0$.

Now, suppose that we expect the *status quo* to evolve over time according to a random walk, that is, that there exists a sequence $\{Y_k\}$ of independent and identically distributed zero-mean random variables such that $m_t = m_0 + \sum_{k=1}^t Y_k$. If we denote the variance of any Y_k by σ^2 , this implies that today's estimate of the variance of m_t is $t\sigma^2$. Hence, from today's viewpoint, the uncertainty about the future target T_t consists of two parts: the uncertainty about how the target varies around its mean, which is measured by v_0 ; and the uncertainty about how the mean of the target varies, which is measured by $t\sigma^2$.

The uncertainty about the future target, then, is described by a normal c.d.f. $U_t(x)$ with a mean of m_0 and a variance of $v_0 + t\sigma^2$. Reinterpreting $U_t(x)$ as today's utility function for the agent's future preferences at time t , we see that this results in less risk-aversion for $x \geq m_0$ and less risk-seeking for $x \leq m_0$ than would be implied by today's utility function for the agent's current preferences. That is, in this example the increased uncertainty about the future favors a more aggressive behavior over gains and a more conservative attitude over losses. However, depending on the details of the stochastic process underlying the evolution of m_t , we may obtain a variety of conclusions.

When tailored to a specific situation, this approach may help the agent appreciate the implications of a target-based approach; see Cyert and deGroot (1975) for a similar approach using a utility-based language. Note

that, as time unfolds and new information obtains, the estimate of m_t will change. Therefore, the estimate of U_t today will differ from the estimate of U_t tomorrow. This has two implications: on the one hand, the change in U_t is likely to result in temporal inconsistency; on the other hand, the agent is now given a way to take into account how the future arrival of information will impact his willingness to implement today's choices and to predict how his preferences may change. This should be particularly important in a situation where the agent must decide today which options he wants to leave open in the future.

The option-based approach to decision making usually presumes that the agent's uncertainty about his prospects increases as his time horizon becomes longer. Options are valuable because they increase the flexibility of the agent's future decision set and thus work to reduce the downside of his uncertainty. Hence, the value of an option also tends to increase with the length of the time horizon. On the other hand, the example suggests that in some cases the agent's uncertainty about his requirements (as summarized by T) will also increase with time, making him more willing to take risks over gains. This reduces the attractiveness of an option taking values in the domain of gains. A longer time horizon may increase both the uncertainty about the prospects (enhancing the value of the option) and the uncertainty about the target (curtailing the value of the option). In general, the net impact of longer time horizons on option values may not be obvious.

4. Group decision making

In many applications of decision analysis, a group of agents is trying to reach an agreement on some decision problem under uncertainty. These situations may involve club memberships, corporate bureaucracies, academic faculties, parliamentary assemblies or judiciary committees.

From the viewpoint of utility-based models, there are two major approaches to the analysis of this kind of problem. The first is theoretically oriented: it assumes that there are n agents and that each agent i has a known individual NM-utility function $U_i(x)$ and considers the problem of aggregating the n individual utility functions into a group utility function $U_g(x)$; see Harsanyi (1955). The second approach is process oriented: it has the members of the group share their utility functions and focus on identifying alternatives that appear beneficial from the perspective of each agent; see Keeney (1992).

In this section, we present some of the possible contributions of the target-based model to either approach. As is customary in the utility-based model, we assume that agents do not misrepresent their preferences or their probability assessments. Consistent with the previous discussion, we will

interpret $U_i(x)$ as the c.d.f. of an uncertain target. The target-based language helps to raise a preliminary issue that is often neglected.

4.1. Common target

In some situations, there is a group target T_g that the agents would like to pursue collectively. For example, the members of an academic faculty may share the common target of achieving excellence for their department. To make the target of excellence operational, we can define it as being listed among the top five schools in a given national ranking. Even when all members of the department share this objective, they may still disagree about how likely a specific outcome x is to meet the common target. For instance, they may have different opinions about how many publications and positive teaching evaluations are necessary to achieve their shared target of excellence.

In other situations, each agent i has his own individual target T_i and the group tries to choose a course of action which takes into account these (possibly conflicting) targets. This may be the case if each member of the faculty has the target of promoting his own reputation. Advancing the excellence of the department, then, would not necessarily be a commonly shared target. These situations are known as problems of social choice.

While both the situations with one common target and those with many individual targets fall into the realm of group decision making, they should be kept distinct. Formally, let $U_i(x)$ represents agent i 's uncertainty about the target. When the agents share a common target T_g and agent i offers his own honest assessment of the group target, then

$$U_i(x) = P_i(x \geq T_g),$$

where P_i denotes i 's subjective probability for T_g . On the other hand, if the agent is honestly reporting about his own target T_i , then

$$U_i(x) = P_i(x \geq T_i).$$

We must attach a different interpretation to U_i depending on whether the agents share a common target or have individual targets. Switching to the utility-based language, we see that this difference is easily lost unless we do not qualify $U_i(x)$ separately: in the case of a common target, $U_i(x)$ summarizes agent i 's own view of what should be the group utility function; in the case of individual targets, $U_i(x)$ describes his individual preferences. It seems, then, that the utility-based notation is not sufficiently rich to capture this important difference.

When the agents share a common target T_g , $U_i(x)$ represents agent i 's subjective probability distribution for T_g . The problem of specifying a group

utility function U_g as a function of the individual utility functions can be reformulated as the problem of combining the probability distributions U_1, U_2, \dots, U_n into a single probability distribution U_g . This problem has a long history in probability theory and many solutions have been advanced; see Genest and Zidek (1986). For example, given a (sure) consequence x , let $U_i(x)$ be the probability that x will meet the target T_g according to agent i . If we follow the well-known linear opinion pooling rule studied and axiomatized in, among others, Bacharach (1975), deGroot and Mortera (1991), and Keeney (1976), then $U_g(x)$, the group probability for x meeting the target, would be an average of the individual probabilities.

Given a common target, the process-oriented approach to group decision making would have the members of the group exchange their probabilistic assessments of T_g and discuss the evidence or the information supporting these assessments, in order to facilitate the identification of critical uncertainties and the pooling of opinions. Let us consider an example. Suppose that the members of the executive board convene to discuss the firm's target. One executive might say that the minimum acceptable target for the pre-tax profit is between 8% and 12%. Another executive might place it between 9% and 11%. A third executive might add that it should be between 10% and 14%. At that point, there would be an extensive debate over what is acceptable or not. Why does the third executive consider it essential to have a profit of at least 10%? Why does the first executive think it might be acceptable to have a profit as low as 8%? While reaching a consensus cannot be guaranteed, this process should foster convergence.

4.2. Individual targets

Things change substantially when the agents do not share a common target, so that $U_i(x)$ represents agent i 's subjective probability distribution for his own target T_i . In the utility-based approach, this situation is described as a social choice problem with cardinal utilities; see d'Aspremont and Gevers (1977) or Roberts (1980). In spite of a few positive results along the lines of Harsanyi (1955), there is an ample literature about the impossibility of reasonable general-purpose aggregating procedures; see Hylland and Zeckhauser (1979) or Seidenfeld, Kadane and Schervish (1989). Thus, we must resort to developing robust solutions for specific classes of problems; see, for instance, Eliashberg and Winkler (1981) or Keeney and Kirkwood (1975).

Target-based reasoning may offer a systematic approach for some specific problems. Suppose that the n individual targets T_1, T_2, \dots, T_n are stochastically independent. This mildly restrictive assumption holds whenever the target of each agent is defined on a different probability space. In other cases, some appropriate transformations similar to those in Section 3.2

may untie the correlations and lead to a formulation in which the individual targets are stochastically independent.

The aggregation of the individual targets into a group objective function U_g should be based on a principle of fairness. One possible criterion is that a consequence be acceptable to the group only when it meets the individual target of each member of the group. The group should therefore aim to maximize the probability of meeting everybody's target. If we let $T_{(n)} = \max\{T_1, \dots, T_n\}$ denote the n -th order statistic, this would result in

$$U_g(x) = P(x \geq T_{(n)}) = \prod_{i=1}^n P_i(x \geq T_i) = \prod_{i=1}^n U_i(x) \quad (6)$$

being the probability that x meets everybody's target. The random variable T_g associated with U_g would therefore be used as the "group benchmark" to rank any feasible decision d by the probability $P(X_d \geq T_g)$ that it meets this benchmark. Note that, if we assume independence of T_g from the random consequences X_d , this is in turn equivalent to the maximization of the expected value of the group utility function U_g .

Other criteria, of course, are possible. For example, a different principle could be that a consequence is acceptable to the group whenever it is acceptable to at least one member of the group. If we let $T_{(1)} = \min\{T_1, \dots, T_n\}$ denote the first order statistic, the "group benchmark" T_g would be distributed according to the c.d.f.

$$\begin{aligned} U_g(x) &= 1 - P(x < T_{(1)}) = 1 - \prod_{i=1}^n [P_i(x < T_i)] \\ &= 1 - \prod_{i=1}^n [1 - U_i(x)] \end{aligned} \quad (7)$$

and the value function would be again $P(X_d \geq T_g)$. Similarly, we might follow the principle of deeming acceptable to the group any consequence that meets the target of at least k of the n members of the group. This would lead to a more complicated form for U_g .

This target-based viewpoint suggests many ways of aggregating individual expected utility functions U_i 's into an aggregate expected utility function U_g . The implications of these and other possible fairness principles, however, must still be explored in depth. As a first step in their study, we state one simple result. Note that the two formulations in (6) and (7) are dual. If we assume that U_i is absolutely continuous for all i , it is a simple exercise in differentiation to establish the following.

Suppose that U_i is convex for all i ; then the U_g implied by (6) will also be convex. Similarly, if U_i is concave for all i , then the U_g implied by (7) will

also be concave. That is, if we start with all the agents being risk-seeking, trying to meet everybody's target makes the group also risk-seeking. Analogously, if we start with risk-averse agents, trying to meet at least someone's target makes the group also risk-averse. One simple example will suffice to illustrate this point: suppose that each agent i has an exponential utility function $U_i(x) = 1 - e^{-\lambda_i x}$. According to (7), the group utility function would be

$$U_g(x) = 1 - e^{-(\sum_{i=1}^n \lambda_i)x},$$

so that the group would also be risk averse. Furthermore, since $(\sum_i \lambda_i)$ is its coefficient of risk aversion and $\sum_i \lambda_i > \max_i \lambda_i$, the group would exhibit more risk aversion than any of its members. This conforms with the group polarization effect empirically observed by Doise (1969) and Moscovici and Zavalloni (1969).

5. Multiattribute decision making

The focus of this section is multiattribute utility theory under uncertainty. Assume that objectives have been specified and that n appropriate attributes $\alpha_1, \alpha_2, \dots, \alpha_n$ have been identified. If x_i denotes a specific level for attribute α_i , a (sure) consequence is described by the vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$ reporting the level of each attribute.

An important problem of multiattribute utility theory is how to construct the overall utility function $U(\mathbf{x})$ as a function of the n single attributes x_1, x_2, \dots, x_n . Many assumptions are possible, but the most important is utility independence. This states that the conditional preferences over random consequences involving only α_i do not depend on the levels of the other attributes. This implies the existence of a single-attribute utility function $U_i(x_i)$ for each $i = 1, 2, \dots, n$.

According to Keeney and Raiffa (1976, p. 224), the role of utility independence in multiattribute theory is similar to stochastic independence in multivariate probability theory. In particular, it implies that the multiattribute utility function is a multilinear function of the utilities associated with the various attributes. Furthermore, the range of both the overall and the single-attribute utility functions can be normalized to be $[0, 1]$.

5.1. Multiattribute individual decisions

We can interpret this result by considering an individual multiattribute utility problem and switching to the target-based language. Suppose that reaching the overall target is a function of the levels x_1, x_2, \dots, x_n of the single attributes; that is, there exists a set A such that the overall target is attained

if and only if $\mathbf{x} \in A$. In particular, assume that A is decomposable as the cartesian product of n subsets A_i such that $\mathbf{x} \in A$ if and only if $x_i \in A_i$ for all $i = 1, 2, \dots, n$.

If we assume stochastic independence of the attributes, the multiattribute utility function can be written as

$$U(\mathbf{x}) = P(\mathbf{x} \in A) = \prod_i P(x_i \in A_i) = \prod_i U_i(x_i), \quad (8)$$

where $U_i(x_i) = P(x_i \in A_i)$ is interpreted as the probability that the target for the single attribute i is met. As in the standard multiattribute utility formulation, the overall probability of success $U(\mathbf{x})$ is a multilinear function of the probabilities $U_i(x_i)$ of meeting the single-attribute targets, for $i = 1, 2, \dots, n$.

A related problem is found in the statistical theory of reliability, which expresses the reliability of a system as a function of the survival probabilities of its various components. In particular, for coherent systems with independent components, it is a standard result that the survival probability of the system is a multilinear function of the survival probabilities of each component; see Barlow and Proschan (1975). If we interpret the survival of a component (respectively, of the system) as the event that a single-attribute (respectively, the overall) target is met, the two formulations are mathematically equivalent.

This implies that the problem of maximizing a multiattribute utility function can be mapped to the problem of minimizing the probability of failure in a fault tree (more generally, an event tree). This suggests a tool to circumvent the well-known limitations of assuming utility independence in multiattribute utility assessment.

Consider, for instance, an individual trying to assess how the multiattribute utility of a house depends upon its price, its location and its facilities. The assumption of utility independence is necessary to quantify how much of the overall utility of a house depends on the single-attribute utilities of price, location and facilities. It postulates that each single attribute must contribute to the overall utility independently from the other attributes. This strong assumption is often implausible.

In contrast, reliability theory employs event trees which describe how the achievement of a higher-level objective is influenced by various combinations of lower-level objectives. For example, an event tree might specify that a house is worthwhile by considering whether it has a reasonable price *or* if it has both a satisfactory location and adequate facilities. Taking expectations over this tree defines the probability of achieving the overall target as a function of various combinations of price, location and facilities. Equation (8) reduces to $U(\mathbf{x}) = P(\mathbf{x} \in A)$, where A is not necessarily a cartesian product and therefore cross-dependencies may be explicitly modelled.

Further work is needed to assess the possibility of framing multiattribute utility assessment as a reliability problem on an event tree, but this approach holds some promise. For example, we can develop an alternate way of deriving the multiattribute utility function by identifying the combinations of lower-level targets whose attainment is necessary for the achievement of the overall target.

Define first a (possibly uncertain) acceptable target for each of the higher-level and lower-level objectives. Thinking of a pyramid with the overall target T at the top, there are second-level targets T_i supporting T ; and then third-level targets supporting the second-level targets T_i , etc. For each higher-level target, specify how its achievement is related to the attainment of its (possibly uncertain) lower-level targets: whether we need to meet all of them; or if it suffices to meet at least one; or if we have to meet at least a majority of the supporting lower-level targets; etc. Finally, we would compute the probability of achieving the higher-level targets as the probability of attaining the right combination of lower-level targets.

Since this enables us to translate the fundamental objectives hierarchy into the standard event tree of reliability analysis, we are able to exploit the hierarchical structure of the event tree to compute these probabilities. This makes it possible to use a “bottom-up” approach to the computation of the probabilities of meeting the higher-level target.

5.2. Multiattribute group decisions

Some important applications of decision analysis concern multiattribute group decisions. Due to limitations of space, we focus on a comparison between the utility-based and the target-based approaches which only scratches the surface of the problem. For concreteness, suppose that we are to consider a major corporate decision involving the chief officers. Both approaches would begin by trying to develop collectively an understanding of which lower-level attributes support the higher-level attributes.

Successively, the utility-based approach generally interviews each agent separately and tries to summarize his assessments in a set of numbers. For example, each of the executive officers might be asked to assess his willingness to pay for market share, profits and innovation. The assessments from each of the executive officers would then be presented to the entire board of executives and the utility-based decision analysis would creatively search for a decision that each officer deems superior to the *status quo*.

The target-based approach can usefully complement this approach, by suggesting that the previously constructed hierarchy of attributes be viewed as an event tree. This should help to focus the group discussion on the likelihood of an acceptable performance in a higher-level objective conditional

on the attainment of acceptable (or unacceptable) levels of performance on the lower-level objectives. For example, the financial officer may believe that an acceptable corporate performance is solely determined by whether the profit is sufficiently high; the operations officer may think that the only key component is an acceptable level of innovation; and the sales officer may hold that it is uniquely determined by a reasonably high level of the market share. At this point, the dialogue should focus on why there are such pronounced differences between the different perspectives. Possibly, this dialogue would lead the group to unearth some higher-level goal like long-run prosperity as the ultimate criterion, with different members of the group disagreeing over whether or not a target of long-run prosperity requires an acceptable performance in all three dimensions. By confronting the officers' views about higher and lower-level targets, this process should make it easier to foster convergence of opinions than a utility-based perspective would.

The critical contribution of the target-based approach is that it may help to translate much of the disagreement between individual utility functions into a disparity of opinions about which (and how) some targets impact other targets and over the probabilistic assessments. We believe that there are situations in which the group can discuss how various targets impact other targets much more easily than they can settle disagreements over values. It is likely that some of the existing methodologies for handling uncertainties (e.g., influence diagrams or tornadoes) might also be used to help modelling uncertainties about targets. Hence, we feel that a target-based approach to multiattribute group decision making may turn out to be easier to use than utility-based decision theory.

6. Conclusions

Most of the current discussion in decision theory is based on utility functions. This paper shows that there exists an alternate target-based language for thinking about decision making and that there are contexts in which it may be more useful than the traditional utility-based language. This suggests that there might be a role for target-based models in decision analysis.

This paper is far from having exhausted the range of possible applications of target-based models. However, some potential advantages should be apparent.

- We can replace the somewhat esoteric concept of an NM-utility function by the more widely understood notion of a target. This may facilitate the teaching of decision analysis.
- We can relate the concept of risk aversion to the notion of the agent's uncertainty about his target. This may facilitate the description and the understanding of the causes of risk-aversion.

- We can offer some guidelines on how to model changes in the agent’s future utility function from the current one. This may facilitate practical applications of utility theory.
- We can suggest a new and simple approach to the formulation of a group utility function. This may facilitate both the analysis and the process of aggregating individual preferences.
- We can show that reliability theory and multiattribute utility theory are more intimately related than previously thought. This may facilitate new insights or new models.

For these reasons, we think that the target-based language seems to offer considerable promise for the theory and practice of decision analysis and that this should warrant it further consideration.

Appendix. Savage’s theorem

This appendix reviews Savage’s theorem and shows how to interpret it in a target-based language. We begin with a detailed statement of the theorem in the utility-based language. Let S be the state space, C the set of (sure) consequences and D the set of acts. Assume for simplicity that C is finite. There is a preference relation \succeq' over acts which induces a preference relation \succeq over consequences such that $\delta_x \succeq' \delta_y$ if and only if $x \succeq y$, where δ_x denotes the degenerate lottery yielding x for sure. Hence, C is a completely ordered finite set.

Theorem 1 (Savage, 1954). *Suppose that the agent has a preference relation \succeq' over all acts which satisfies axioms P1–P6. Then the following statements hold.*

- (i) *There exists a unique finitely additive probability measure P on S . For each decision d , this probability measure induces a probability distribution P_d over the consequences associated with d .*
- (ii) *There exists a [utility] function U on C such that $U(x) \geq U(y)$ if and only if $x \succeq y$. The [utility] function U is bounded and unique up to increasing affine transformations.*
- (iii) *The value function v over D defined by $v(d) = \sum_x U(x)P_d(x)$ represents \succeq' . That is, $d_1 \succeq' d_2$ if and only if $v(d_1) \geq v(d_2)$.*

The boundedness of U is proved in Fishburn (1970) for the general case of a set C with an infinite number of consequences, where the additional P7 axiom is required to extend the representation in (iii) to all acts. Here, it is an obvious corollary to the finiteness of C .

In part (ii) of the theorem, we have put in brackets the word “utility” to highlight the fact that the interpretation of U as a utility function is not part

of the statement of the theorem. The meaning of U is open. To provide a reinterpretation, note that (ii) establishes that U is bounded and unique up to increasing affine transformations. Picking the appropriate transformation, we can normalize the range of U to be the interval $[0, 1]$ and change the statement in part (ii) to:

- (ii. a) There exists a unique function $U : C \rightarrow [0, 1]$ which is increasing with respect to \succeq .

This function U can be interpreted as a distribution function on C , defining a probability measure Q on C such that $U(x) = Q\{y : x \succeq y\}$. By a standard probability-theoretic argument, we can pick the probability space (Ω, Π) where $\Omega = [0, 1]$ and Π is the uniform distribution on $[0, 1]$ and construct a random consequence $T : \Omega \rightarrow C$ with distribution function $U : C \rightarrow [0, 1]$; see, for instance, Proposition 3.4 in Fristedt and Gray (1997). Assuming $\Omega \cap S = \emptyset$, we see that the random variable T turns out to be stochastically independent of any P_d .

We interpret the random consequence T as the (subjectively) uncertain target of the agent. The probability distribution of T is given by U ; that is, $P(x \succeq T) = U(x)$. Hence, what previously was the utility function is now the distribution function of the uncertain target: $U(x)$ is the probability that x meets the target. Of course, the more preferable is x , the higher the probability that it meets the target; that is, $x \succeq y$ if and only if $U(x) \geq U(y)$.

Under this interpretation, part (ii) of the theorem becomes:

- (ii. b) There exists a (stochastically independent) random target T with a unique [probability distribution] function $U : C \rightarrow [0, 1]$.

Carrying this interpretation over to part (iii), which in the utility-based language states that $v(d)$ is the expected utility of d , we obtain the target-based formulation of Savage's theorem:

Theorem 2. *Suppose that the agent has a preference relation \succeq' over all acts which satisfies axioms P1–P6. Then the following statements hold.*

- (i) *There exists a unique finitely additive probability measure P on S . For each decision d , this probability measure induces a random consequence X_d with probability distribution P_d .*
- (ii) *There exists a (stochastically independent) random target T with a unique [probability distribution] function $U : C \rightarrow [0, 1]$.*
- (iii) *The value function v over D is the probability that the random consequence X_d meets the uncertain target T :*

$$v(d) = \sum_x U(x) P_d(x) = \sum_x P(x \succeq T) P_d(x) = P(X_d \succeq T).$$

References

1. Bacharach, M. (1975): Group decisions in the face of differences of opinions. *Management Science* **22**, 182–191
2. Barlow, R., Proschan, F. (1975): Statistical theory of reliability and life testing. Holt, Rinehart and Winston, New York
3. Behn, R.D., Vaupel, J.W. (1982): Quick analysis for busy decision makers. Basic Books, New York
4. Berhold, M.H. (1973): The use of distribution functions to represent utility functions. *Management Science* **19**, 825–829
5. Borch, K. (1968): Decision rules depending on the probability of ruin. *Oxford Economic Papers* **20**, 1–10
6. Castagnoli, E., LiCalzi, M. (1996): Expected utility without utility. *Theory and Decision* **41**, 281–301
7. Cyert, R.M., deGroot, M.H. (1975): Adaptive Utility. In: Day, R.H., Groves, T. (eds.). Adaptive economic models. Academic Press, New York, pp. 223–246
8. D'Aspremont, C., Gevers, L. (1977): Quality and the informational basis of collective choice. *The Review of Economic Studies* **44**, 199–209
9. DeGroot, M.H., Mortera, J. (1991): Optimal linear opinion pools. *Management Science* **37**, 546–558
10. Doise, W. (1969): Intergroup relations and polarization in individual and collective judgements. *Journal of Personality and Social Psychology* **12**, 136–143
11. Eliashberg, J., Winkler, R.L. (1981): Risk sharing and group decision making. *Management Science* **27**, 1221–1235
12. Fishburn, P.C. (1970): Utility theory for decision making. Wiley, New York
13. Fristedt, B., Gray, L. (1997): A modern approach to probability theory. Birkhäuser, Boston
14. Genest, C., Zidek, J.V. (1986): Combining probability distributions: a critique and an annotated bibliography. With comments and a rejoinder by the authors. *Statistical Science* **1**, 114–148
15. Harsanyi, J. (1955): Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *Journal of Political Economy* **63**, 309–321
16. Howard, R. (1992): In praise of the old time religion. In: Edwards W. (ed.): Utility theories: measurements and applications. Kluwer, Boston, pp. 27–55
17. Hylland, A., Zeckhauser, R. (1979): The impossibility of Bayesian group decision making with separate aggregation of beliefs and values. *Econometrica* **47**, 1321–1336
18. Kahneman, D., Tversky, A. (1979): Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–291
19. Karni, E. (1985): Decision making under uncertainty: the case of state-dependent preferences. Harvard University Press, Cambridge, MA
20. Keeney, R.L. (1976): A group preference axiomatization with cardinal utility. *Management Science* **23**, 140–145
21. Keeney, R.L. (1992): Value-focused thinking. Harvard University Press, Cambridge, MA
22. Keeney, R.L., Kirkwood, C.W. (1975): Group decision making using cardinal social welfare functions. *Management Science* **22**, 430–437
23. Keeney, R.L., Raiffa, H. (1976): Decisions with multiple objectives: preferences and value tradeoffs. Wiley, New York
24. Manski, C.F. (1988): Ordinal utility models of decision making under uncertainty. *Theory and Decision* **25**, 79–104

25. Moscovici, S., Zavalloni, M. (1969): The group as a polarism of attitudes. *Journal of Personality and Social Psychology* **12**, 125–135
26. Pratt, J.W., Raiffa, H., Schlaifer, R. (1995): Introduction to statistical decision theory. The MIT Press, Cambridge, MA
27. Roberts, K. (1980): Possibility theorems with interpersonally comparable welfare levels. *The Review of Economic Studies* **47**, 409–420
28. Savage, L.J. (1954): The foundations of statistics. New York: Wiley
29. Schervish, M.J., Seidenfeld, T., Kadane, J.B. (1990): State-dependent utilities. *Journal of the American Statistical Association* **85**, 840–847
30. Seidenfeld, T., Kadane, J.B., Schervish, M.J. (1989): On the shared preferences of two Bayesian decision makers. *Journal of Philosophy* **86**, 225–244
31. Shafir, E., Diamond, P., Tversky, A. (1997): Money illusion. *The Quarterly Journal of Economics* **112**, 341–374
32. Simon, H.A. (1955): A Behavioral model of rational choice. *The Quarterly Journal of Economics* **69**, 99–118