

FOUNDATIONS OF TARGET-BASED DECISION THEORY

Robert F. Bordley; MC 482-D20-B24; Renaissance Center; P.O. Box 100; Detroit, Michigan 48265-1000, U.S.A.

Key Words: Utility Theory, Targets, Customer Satisfaction, Reliability Theory, Utilitarianism

Contents

1. Bentham and Utility-Based Decision Analysis
2. Hobbes and Decision Analysis
3. Target-Based Decision Analysis
4. Bounded Rationality and Target-Based Decision Analysis
5. Pedagogical Advantages
6. Improved Modeling of Individual Choice
7. Better Linkages with Finance
8. State-Dependent Utilities
9. Better Linkages with Practice
10. More Consistent with Psychological Evidence
 - 10.1. Cognitive Psychology
 - 10.2. Depth Psychology
11. Conclusions

Glossary:

Behavioral Decision Theory: A field of psychology focused on understanding how individuals actually do make decisions, in light of an understanding of how a fully rational person would make decisions.

Bounded Rationality: A theory which presumes that individuals either lack the intellectual capacities or time required to explicitly optimize some objective function. Hence it presumes that individuals actually focus on finding decisions that satisfy certain target requirements.

Decision Analysis: A profession concerned with helping organizations make decisions by having them explicitly quantify their preferences, their beliefs and the alternatives available to them, using various principles of rational choice.

The Rational Man Hypothesis: The assumption that individuals act as if they were implicitly maximizing some objective function.

Utilitarianism: A philosophic doctrine arguing that individuals and society should explicitly quantify their beliefs in the form of an objective function and make decisions which maximize this objective function

Utility: A function used to quantifying an individual's degree of preference for some consequence or decisions

Summary

Decision theoretic axioms are commonly interpreted in terms of utility functions, reflecting the philosophy of utilitarianism. But an alternate philosophy, suggested by Thomas Hobbes, leads to a target-based interpretation of the decision theoretic axioms. As we show, this target-based interpretation is more consistent with how people actually decide and hence may be much easier to apply than the standard utility-focused decision analysis.

1. Bentham and Utility-Based Decision Analysis

Von Neumann and Morgenstern (1947) enunciated various axioms of rationality which implied that

- (1) For any rational individual, it was always possible to define the utility of a consequence as that probability p making the individual indifferent between receiving that consequence and receiving a lottery with a probability p chance of leading to the best possible consequence and a $(1-p)$ chance of leading to the worst possible consequence.
- (2) The rational individual, when choosing among several possible decisions, would always choose that decision whose possible consequences have the maximum expected utility.

How does one interpret utility? The word, "utility", as used by Von Neumann and Morgenstern is the same term which 19th century utilitarianism used to describe a measure of pleasure. Hence by using the term "utility", Von Neumann and Morgenstern were suggesting that there was some parallelism between their utility theory and utilitarianism.

The utilitarians argued that individuals all sought pleasure and that pleasure should be quantified in the form of a utility function. They developed a calculus of utility which allowed an individual to systematically calculate the utility or pleasure associated with any course of action and argued that individuals should choose those actions which achieve the maximum overall utility. Indeed a careful review of the elements of utilitarianism shows that it is remarkably analogous to modern utility theory. Thus the leading proponent of utilitarianism, Jeremy Bentham (1781), wrote:

"By the principle of utility is meant that principle which approves or disapproves of every action whatsoever according to the tendency which it appears to have to augment or diminish the happiness of the party whose interest is in question...Sum up all the values of all the pleasures on one side, and those of all the pains on the other. The balance, if it be on the side of pleasure, will give the good tendency of the act upon the whole."

As Sabine (1961) noted,

"Bentham...assumed, as had commonly been done by the hedonist moralists, that pleasure and pain are commensurable, a given amount of the one offsetting a like amount of the other, and also that they can be added so that a sum of pleasures may be calculated, which will define the greatest happiness...Usually Bentham spoke as if he believed that human beings really do act in accordance with some such mental parallelogram of forces as this, though occasionally he acknowledged that the notion of adding pleasures, and especially the pleasures of different individuals, is fictitious. What is certainly true was that he considered the fiction to be "a postulation without the allowance of which all political reasoning is at a stand".

One could view much of economics, with its assumption of the insatiable utility-maximizing individual, as similarly reflecting this utilitarian vision.

But it's become widely recognized that Von Neumann and Morgenstern's utility measure isn't simply a measure of desirability; instead it's both a measure of desirability and a measure of the individual's attitude toward risk. Because it confounds these two effects, it's common to break up utility assessment into assessing two different measures:

- (1) a value function v which only measures the desirability of a consequence and

- (2) a utility function, expressed as a function of v , which describes how that value function is distorted by an individual's attitude toward risk.

But given this decomposition, Von Neumann and Morgenstern's notion of utility no longer corresponds to the utility measure which the utilitarians used to quantify degrees of desirability. Instead the value function becomes analogous to the utilitarian measure of utility. And what Von Neumann and Morgenstern call "utility" corresponds to something different than what the utilitarians meant by utility. (Given this semantic confusion, it's tempting to follow Howard (1992) in referring to u as a "preference probability." and not as "utility".)

If u is no longer a measure of desirability, then it's less clear how maximizing expected utility in the Von Neumann/Morgenstern sense relates to utilitarianism. The next section explores this issue in more detail.

2. Hobbes and Decision Analysis

Bentham's "mathematics of hedonism" presumed that all individuals are, or should be, understood as seeking pleasure in some form. But Bentham's philosophy built on an earlier philosophic tradition started by Thomas Hobbes. As Macpherson (1962) writes in his review of political philosophy,

"the foundations of utilitarianism starts at least as far back as Hobbes...the utilitarian doctrine...is at bottom only a restatement of the individualist traditions which were worked out in the 17th century: Bentham built on Hobbes"

But even though Bentham built on Hobbes, he also appears to have significantly altered Hobbes's original argument. Thus Hobbes (1691) originally argued

I put for a general inclination of all mankind, a perpetual and restless desire for power after power, that ceaseth only on death. And the cause of this is not always that a man hopes for a more intensive delight, than he has already attained to; or that he cannot be content with a more moderate power; but because he cannot assure the power and means to live well which he hath present, without the acquisition of more.

Hobbes (1690) also wrote:

There be some, that taking pleasure in contemplating their own power in acts of conquest, which they pursue further than their security requires...others that otherwise would be glad to be at ease within modest bounds...those men who are moderate, and look for no more but equality of nature, shall be obnoxious to the force of others, that will attempt to subdue them. And from hence shall proceed a general diffidence in mankind and mutual fear of one another .

MacPherson (1962) summed up this Hobbesian perspective by saying that

Some, not all, men innately desire ever more power and delight, while the rest desire only to continue at their present level....Every man's innate desires are indeed incessant, but not every man's are for an increased level of satisfaction or power. All men in society... do seek ever more power, but not because they all have an innate desire for it. The innately moderate man in society must seek more power simply to protect his present level

In other words, Hobbes argued that there were two kinds of individuals:

- (1) An insatiable minority desiring ever more power and delight.
- (2) A majority of moderate but insecure individuals seeking enough to meet various needs. Because of uncertainty about is required to secure those needs, these moderate individuals constantly try to acquire as much as possible.

Bentham's formulation seems to have emphasized the insatiable minority. Hobbes's formulation indicates that both types of individuals exist. We will refer to the insatiable minority as Benthamite individuals and the moderate majority as Hobbesian individuals. Hobbes argued that most individuals---as opposed to being maximizers of pleasure---are actually oriented toward maximizing their chances of meeting certain well-defined needs. Thus Hobbes' s philosophical perspective is fundamentally different than utilitarianism.

Is the Hobbesian perspective inconsistent with utility theory?

3. Target-Based Decision Theory

We first use work by Castagnoli and LiCalzi (1996), and extended by Bordley and LiCalzi (2000), to show that any Hobbesian individual can be viewed as maximizing expected utility in the sense of Von Neumann Morgenstern. Note first that an individual who is only interested in meeting certain well-defined needs can be viewed as assigning a utility of one to an outcome which always meets all the needs and a utility of zero to an outcome that never meets all the needs. But Von Neumann and Morgenstern defined the utility of any intermediate outcome x as that probability u such that an individual is indifferent between getting x and getting a lottery offering a probability u of meeting the needs and a probability $(1-u)$ of never meet the needs. If the individual is only interested in meeting the needs, then the individual can only be indifferent between x and a u chance of meeting the needs if x offers the individual a u chance of meeting the needs. Since the individual prefers the option with the highest chance of meeting the needs, the individual will prefer the option with highest utility.

Now suppose the individual is concerned with a decision which has various probabilities of yielding different outcomes. Since the utility of each outcome represents the probability of meeting the needs given the outcome, the expected utility of the decision will be the expected probability of meeting the needs given the decision. Since the individual likewise chooses that decision with maximum probability of meeting the needs, the individual will choose that decision with maximum expected utility. Thus an individual with a Hobbesian psychology will maximize expected utility where the utility of a consequence is the probability of meeting the needs given that consequence.

We now wish to prove the reverse, i.e., to prove that any individual who makes decisions which maximize expected utility in the sense of Von Neumann and Morgenstern can be thought of as acting like a Hobbesian individual. To make this argument, recall that the utility function prescribes how the individual ranks various consequences x as well as how the individual ranks all possible gambles. But as previously noted, it's common to try to separate out the individual's preference over consequences from his attitude toward risk by defining:

- (1) A value function $v(x)$ which ranks deterministic consequences x but does not describe the individual's attitude toward gambles

- (2) A utility function, u^* , defined over levels of the value function (and not over consequences) which describes the individual's preference over various gambles whose outcomes are various levels of value (with $u(x)=u^*(v(x))$).

Since the value function $v(x)$ ranks consequences, x , in the same way as the original utility function $u(x)$, there also exists a monotonic transformation f (i.e. the inverse of u^*) such that v is a monotonic function of $u(x)$, i.e. $v(x)$ equals $f(u(x))$.

Now suppose U is a uniformly distributed random variable. If q is any number between zero and one, we can always write q as the probability of q exceeding U . Since utility functions are bounded from above and below, we can always scale them to range between zero and one. We can then similarly write $u(x)$ as equal to the probability of $u(x)$ exceeding U . If we define the random variable T as equal to $f(U)$, then $u(x)$ equals the probability of $v(x)$ exceeding T . Hence any utility function $u(x)$ can always be written as the probability of a value function $v(x)$ exceeding T .

Now consider a gamble d which yields consequence x with probability $p(x)$. We can define a random variable V whose probability $p(v)$ of yielding a value v is the probability of the gamble yielding a consequence x with $v(x)=v$. Suppose that V is independent of T . Then the expected utility of the gamble is the product of $p(x)$ and $u(x)$, summed over all consequences x . Since $u(x)$ can be written as a function, $u^*(v)$, of the value function v , we can rewrite this expected utility as the product of $p(v)$ and $u^*(v)$ summed over all values v . Since $u^*(v)$ is the probability that v exceeds T and since V and T are independent, this equals the probability that V exceeds T . Hence

$$\sum_x p(x)u(x) = \sum_v p(v)u^*(v) = \sum_v p(v)Pr(v>T) = Pr(V>T)$$

In summary, Von Neumann and Morgenstern's rationality axioms are universally interpreted as implying the existence of a utility function such that the individual always prefers that gamble with the higher expected utility. But they can also be equivalently interpreted as implying the existence of a value function, v , and a random variable T such that the individual always prefers that gamble which has the highest probability of yielding a value exceeding T . Both interpretations are equally consistent with the axioms.

Since the utility of a consequence which always falls below T is zero, it's natural to interpret T as representing the minimum value required to meet the individual's needs. Since T is uncertain, the individual acts as though there were some uncertainty about this minimal level of requirement. This interpretation closely parallels Hobbes's model of the individual. Hence Von Neumann and Morgenstern's rationality axioms can be interpreted as implying that a rational individual acts as if he were a Hobbesian, i.e., as someone striving to maximize the probability of having enough value to exceed the minimum requirements for contentment.

Since value seems to correspond to what the utilitarians called utility, this says the individual tries to maximize the chances of having an adequate utility in the sense of utilitarianism. Note that in this interpretation, individuals, contrary to the assumptions of modern economics, are not insatiable. As Hobbes indicated, their needs are moderate. They only appear to be insatiable because of uncertainty about what is required to meet those needs.

Thus the assumptions of rationality are consistent with two different interpretations of human psychology:

- (1) The insatiable, utility-maximizing individual of utilitarian philosophy
- (2) The moderate, perpetually insecure, individual of Hobbesian philosophy

Hobbes, of course, argued that the majority of people were moderate but insecure and that only a few were insatiable. If Hobbes is indeed correct, then some (perhaps most) potential clients of decision theory would find the target-based interpretation of decision more appealing than the utility-based interpretation.

Since historical accident has led us to focus on the Benthamite utility-based approach, the remainder of this paper focuses on articulating the Hobbesian target-based approach.

4. Bounded Rationality and Target-Based Decision Analysis

Hobbes had argued that the bulk of mankind were moderate but insecure. Hence the insatiable utility-maximizing individual of Benthamite psychology only describes a small, but powerful, portion of humankind. To provide some evidence for this Hobbesian perspective, we turn to Simon (1978)'s work on bounded rationality.

As is well-known, Simon noted that there were serious costs associated with the memory and computations required to calculate the utility of various outcomes and choose the outcome of highest utility. Most decisions don't seem to be worth the time and effort required to make such computations. Even for decisions which are worth such time and effort, few, if any, individuals seem oriented toward making such computations. For this reason, Simon enunciated his famous theory of bounded rationality. In this theory, an individual has certain pre-specified requirements. If those requirements are met, the individual continues with his current decisions. When those requirements are not met, the individual actively searches for alternative decisions. Instead of looking for the "optimum" decision, the individual adopts the first alternative he discovers which satisfies the requirements.

Since Simon's original challenge, there has been decades of experimental work testing and applying utility theory. In his Nobel Prize Lecture, Simon (1978) indicated that this work had soundly refuted the idea that individuals optimize a utility function:

"there can be no doubt that...the assumptions of perfect rationality... do not even remotely describe the processes that human beings use for making decisions in complex situations."

Instead he argued that individuals clearly seem to be oriented toward achieving targets. As he writes,

"...Research in information processing psychology provides conclusive evidence that the decision-making process in problem situations conforms closely to the models of bounded rationality...Today we have a large mass of descriptive data, from both laboratory and field, that show how human problem solving and decision making actually take place in a wide variety of circumstances. Theories to account for the data incorporate the notions of bounded rationality: the need to search for decision alternatives, the replacement of optimization by targets and satisficing and mechanisms of learning and adaptation. If our interest lies in descriptive theory(or even normative decision theory), it is now entirely clear that the classical and neoclassical theories have been replaced by a superior alternative that provides us with a much closer approximation to what is going on.."

Simon's work establishes that decision-makers consciously make decisions by satisficing targets and not by optimizing utility functions. Hence his work rejects the utilitarian language in which utility theory has been cast.

Of course, it's common to ignore bounded rationality on the grounds that it's not a model of how individuals ought to behave. The fact that individuals happen to make decisions using targets doesn't imply that it's best to make decisions using targets. Nonetheless getting an individual to actually use decision analysis requires an individual to understand it. As Howard (1992), one of the leading proponents of decision analysis, notes, decision analysis is hard to apply without considerable training:

"One could ask whether the difficulty of performing good decision analysis is not itself a criticism of decision analysis. The analogy I like here is brain surgery. Because effective brain surgery is difficult does not mean that there is anything wrong with it. I would no more expect a person with little training to complete an effective decision analysis than I would expect him to perform a successful brain operation....An important and fundamental difficulty of decision analysis is that it is expensive. The level of skills required assures that it will always be relatively expensive even in a world where computation is cheap. I find that decision analysts require 3-4 years of graduate education and at least 2 years of practical experience before they can be considered fully trained. Only a tiny fraction of decisions will ever be aided by professional decision analysis; we can only hope that among them will be the most important decisions."

Hence it's very hard for untrained individuals to apply decision analysis. For this reason, Simon's alternative theory of bounded rationality is appealing. Unfortunately Simon's theory allows for behavior which is irrational from the perspective of utility theory.

Is there any way of retaining the normative validity of utility theory while incorporating the practicality of the theory of bounded rationality? The target-based interpretation of decision analysis might be a possible solution. Like bounded rationality, target-based decision theory treats individuals as trying to achieve targets or, more precisely, as trying to maximize the probability of achieving targets. But there is one important difference. Simon focused on whether or not the value function exceeded a pre-specified target; target-based decision analysis focuses on whether the value function exceeds a random variable, T .

In order to reconcile the notion of a random variable T with Simon's notion of a fixed pre-specified target, recall that any mathematical model is necessarily an abstraction from the real world and only focuses on a small number of variables. As the physicist, Brody (1994), writes

"When doing scientific research (and in everyday life), we never deal with the universe as a whole, we select from it a segment designed to contain everything relevant to our purpose and nothing irrelevant; following the physicist's custom, we shall talk about the physical system. This subset of reality is called the system. Everything outside the system is the environment."

Savage (1964) referred to these abstractions as "small worlds" and, in the seminal book that launched his formulation of decision theory, acknowledges that it is not possible to subsume the "real world" within a model:

"Any claim to realism made by this book...is predicated on the idea that some of the individual decision situations[small worlds] into which actual people tend to subdivide the single grand decision do recapitulate in microcosm the mechanism of the idealized grand decision.(pg.83)...I am unable to formulate criteria for selecting these small worlds and indeed believe their selection may be a matter of judgement and experience about which it is impossible to enunciate completely and sharply defined principles.(pg.16)"

We could view the value function as an assessment of the value, v , of a consequence based on only those factors considered in the small world. We could model the impact of the neglected factors by a random variable, e . Hence the actual value of the consequence in light of all factors is some function of the "small world" value function, v , and e . Requiring that this function of v and e exceed some specified target, t , is --- given some mild separability conditions----equivalent to requiring that v exceed some function of t and e . If we let the random variable T be this function of t and e , then requiring that the small world value function exceed the random variable T is equivalent to requiring that the actual value exceed t .

Now Simon's theory mandates that the individual focus only on whether the observed value, v , exceeds the target, t . But Simon's theory clearly recognizes the limitations of human cognition and, in fact, is commonly viewed as replacing the "*global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by the organization*". Hence his theory, at least implicitly, recognizes that the decision-maker's assessment of v doesn't reflect all relevant factors. But since his theory is heuristic, he ignores the fact that the estimated "small world" value function probably deviates from the "actual" value function.

Target-based decision analysis proposes that the decision-maker explicitly recognize this inaccuracy by introducing an error term, e . From the target-based perspective, the individual is interested in whether the product of v and e exceeds the target. Hence we could view target-based decision analysis as accepting the satisficing model but as attempting to explicitly model the errors arising from bounded rationality.

5. Pedagogical Advantages

Utility-based decision theory has evolved several important concepts, e.g. risk-aversion and multi-attribute utility theory. It is critical to retain these concepts in any decision theory, whether utility-based or target-based. But as we now show, all of the important properties of utility function correspond, in the language of target probability density, to well-established properties of probability densities. Hence the critical concepts associated with utility-based decision theory are preserved in target-based decision theory.

The index of absolute risk-aversion, i.e., the derivative of the logarithm of marginal utility, is a central feature of utility analysis and measures the degree to which an individual will sacrifice expected value in order to avoid risk. Since the utility function corresponds to the cumulative probability distribution of a random variable T , this classical absolute risk-aversion index corresponds to the derivative of the logarithm of the probability density of T . But this corresponds to the Pearson index which is already commonly used for characterizing probability distributions. Just as the index of absolute risk-aversion is widely used to classify utility functions so the Pearson index is widely used to classify probability densities.

Multi-attribute utility theory extends standard utility analysis to the case in which the decision-maker values each consequence on the basis of various attributes of that consequence (e.g. quality, cost, performance, etc.). Given the assumptions of multi-attribute utility theory, one can define attribute-specific utilities measuring the utility of a product based on its rating on a given attribute. Multi-attribute utility theory then writes the utility of the consequence as a multi-linear¹ function of the utility of various attributes of that consequence. When we rewrite this in target-based language, the attribute-specific utilities correspond to the probability of the consequence satisfying the requirements on a specific attribute. Multi-attribute utility theory then writes the probability of the consequence satisfying the overall requirements as a multi-linear function of the probability of the consequence satisfying the requirements on various specific attributes. But this corresponds precisely to the formalisms of reliability theory where the probability of the system surviving is written as a multi-linear function of the probability of various components of that system surviving. The widely used multiplicative special case of multi-attribute utility theory corresponds, in the target-based language, to reliability theory applied to a system with all components either in series or in parallel.

Hence the concepts which are central in using utility theory do translate to analogous concepts in target-based decision theory. The index of risk-aversion becomes the Pearson index. Multi-attribute utility theory becomes reliability theory. Indeed this suggests that multi-attribute utility theory can be viewed as an extension of reliability theory to describing preferences.

This may help us solve a major obstacle to disseminating decision theory in statistics courses. Most first year courses in statistics, including engineering statistics, ignore decision theory. In those few cases where it is treated, textbooks generally focus on maximization of expected value and ignore the utility function altogether. (Indeed, even in introductory books devoted to decision theory, utility is often introduced in the next to last chapter.) Since the statistics curriculum is already overcrowded with topics like probability theory, reliability theory and experimental design, there seems to be little hope of arguing that time be taken from other topics in order to treat utility.

But if decision theory can be repositioned as an extension of reliability theory, we may be able to provide a course which can cover both reliability theory and decision theory in less time than would be needed to cover both topics separately. Since reliability theory is already an established topic in the statistics curriculum, this suggests a possible way of introducing decision theory into statistics education. A recent paper in the *Journal of Statistical Education* (Bordley, 2001) championed a proposed course based on such an integrated treatment of reliability theory, decision theory and other tools of statistics. As that paper noted, this also allows the class to be focused around the solution of a high-level business problem. This is important to many statistics teachers who are currently extremely concerned about making statistics education relevant.

By integrating reliability theory and decision theory, this approach also greatly increases the chance that students will leave the class with a working knowledge of both topics.

6. Improved Modeling of Individual Choice

Economic modeling commonly measures a customer's preference for a good or service as a function of the value provided by that good or service. However the literature on customer satisfaction emphasizes

¹ A function $f(a,b,c,\dots)$ is multi-linear if it is linear in each argument. Thus the function $f(x,y)=x+y+xy$ is multi-linear while the function $f(x,y)=x+y+xy+x^2$ is not.

that customers, especially in service industries, value products by looking at the disparity between how those products performed and their expectations or requirements for how the product should perform. This is especially important since customer expectations dynamically change over time. For example, customer expectations for product quality are now generally much higher than they were in the past.) The standard static utility which doesn't explicitly incorporate customer expectations is not readily adapted to modeling these effects. As a result, marketing scientists often complement the standard utility analysis approaches of economics with "gap analysis" which focuses on how far products fall short of customer expectations.

Target-based decision analysis provides a way of reconciling utility analysis with gap analysis. The target-based approach writes the utility of consequence c as the probability that $v(c)$ exceeds a random threshold, T . If there is some uncertainty in v , then we define a random variable V which reflects the uncertainty in the value assigned to consequence c . As a result, the expected utility of consequence c is the probability that V exceeds T .

We now interpret T as reflecting customer expectations or requirements and $(V-T)$ as the gap between product performance and customer expectations. Since T is a random variable, T will change as the customer gets new information. This is consistent with the fact that customer expectations do dynamically change over time.

We can use this formalism to develop a predictive model of customer choice. We note first that the target-based approach assumes that V and T are independent, which is consistent with Savage's condition on how the uncertainty assessed within the small world (the uncertainty about V) should relate to uncertainties outside that small world². Suppose that both V and T are normally distributed with means EV and ET respectively and standard deviations $s(V)$ and $s(T)$ respectively. Then the expected utility of c is the probability that V exceeds T or the probability that $(V-T)$ exceeds zero. Note that $(V-T)$ has mean $E(V-T)$ and standard deviation $s^2(V)+s^2(T)$. If we let U^{**} be a standardized normal random variable and define R by

$$R=[EV-ET]/[s^2(V) + s^2(T)]^{1/2}$$

then the probability that $V-T$ exceeds 0 is the probability that R exceeds U^{**} .

Thus the expected utility of c equals the probability that R exceeds U^{**} . If we let F be the cumulative normal distribution, then the expected utility of c is $F(R)$, i.e., the expected utility of c is a monotonically increasing function of R . Hence ranking outcomes according to the value of R gives us the same ranking as ranking outcomes according to expected utility. Thus gap analysis---if we define the gap using R --- is fully consistent with utility analysis.

Now suppose that $s^2(T)$ is much larger than $s^2(V)$. (This corresponds to assuming that the probability density of T is fairly flat, i.e. that the individual is very uncertain about T .) Then the ranking of outcomes given by R is approximately the same as the ranking of outcomes given by $EV-ET$ (and thus by EV). In other words, ranking outcomes in order of expected utility is approximately the same as ranking outcomes in order of expected value. An alternative extreme case involves setting $s^2(T)=0$. In this

² Savage suggests "That the events that represent the outcome of the small world lottery are all statistically independent of the grand-world acts, or functions, that typically enter as prizes in a lottery" (pg. 90-91).

case, the individual is only interested in choosing that alternative with the maximum chance of having a payoff exceeding ET .

As a paper in the *Journal of Service Research* (Bordley, 2001) noted, this model also makes predictions about customer behavior. Suppose the individual's beliefs about a product's performance, prior to trying it, are described by a random variable. Then after his first experience, he will update his beliefs about V . This will cause EV to increase if the experience was favorable and decrease if it is not. The variance of V , $s^2(V)$, will generally decrease. Hence adding information generally favors a product as long as you expect the product to outperform expectations ($EV > ET$) and generally hurts a product as long as you expect the product to fall short of expectations ($EV < ET$).

7. Better Linkages with Finance

Because of its access to volumes of extensive empirical data on prices, the theory of finance has become a well-established and empirically based science. But many of the key concepts of finance appear to have only a very loose relationship to utility theory.

Two specific examples are the Sharpe ratio and risk-adjusted present value. But as we now show, the target-based perspective clarifies the relationship between finance and decision theory.

We first focus on the Sharpe ratio. Let V be the return on a stock-market portfolio and let T be the return from a benchmark portfolio. One of the key measures of portfolio performance is the Sharpe ratio, defined by

$$G = [EV - ET] / [s^2(V) + s^2(T)]^{1/2}$$

As the previous section noted, this formula can be derived by assuming that utility is a cumulative normal function. Assuming such a functional form for utility seems very artificial. But in target-based decision theory, this corresponds to assuming that the random variable T is normally distributed. Assuming that T is normally distributed is a very natural assumption. For example, section 4 indicated that the randomness in T arose from deviations between the value as estimated in the 'small world' and the value that would be estimated if all factors were taken into account. Since it's likely that the error, e , in estimating the value function could be viewed as the accumulation of many small errors, assuming that T is normally distributed is not unreasonable. Thus target-based decision theory provides a natural way to motivate the Sharpe ratio.

The certainty equivalent of a gamble V , i.e. that known payoff c such that the utility of c equals the expected utility of V , is defined by $u(c) = Eu(V)$. In the target-based formulation, this implies $\Pr(c > T) = \Pr(V > T)$. If we define U^{**} to be a standardized normally distributed random variable as before, this implies $\Pr((c - ET) / s(T) > U^{**}) = \Pr(R > U^{**})$. Hence $(c - ET) / s(T)$ must equal R . This implies

$$c = ET + s(T) \{EV - ET\} / [s^2(V) + s^2(T)]^{1/2}$$

If we define

$$r = [1 + s^2(V) / s^2(T)]^{1/2} - 1$$

then this formula becomes

$$c = ET + [EV - ET] / [1 + r]$$

where we could interpret r as a risk-adjusted discount factor. Hence target-based decision theory is likewise consistent with the notion of risk-adjusted discount factors which are widely used in finance.

8. State-Dependent Utility Functions

In medical applications, the utility of getting a certain payoff can depend very dramatically upon the state in which you get that payoff. This has led to some efforts to modify utility functions to make them state-dependent. Surprisingly enough, target-based decision theory lends itself very naturally to describing state-dependent preferences. As we will show, it also provides very simple ways of modeling state-dependence.

Recall that the equivalence between expected utility theory and maximizing the probability of exceeding T presumed that V and T were independent. Relaxing the assumption of independence means that target-based decision analysis is no longer consistent with standard expected utility theory. Instead it becomes identical to expected utility theory with state-dependent utilities.

In our derivation of R , we found that $R = E[V - T] / s(V - T)$ when V and T were normally distributed. If V and T were independent, then $s^2(V - T) = s^2(V) + s^2(T)$ which gave us the Sharpe Ratio. Allowing utility functions to be state-dependent implies that V and T can be correlated. If ρ denotes the correlation between V and T , then $s^2(V - T) = s^2(V) - 2\rho s(V)s(T) + s^2(T)$. Using this expression for $s(V - T)$ in our formula for the Sharpe ratio gives the revised Sharpe ratio

$$\{EV - ET\} / [s^2(V) - 2\rho s(V)s(T) + s^2(T)]^{1/2}$$

Using this expression for $s(V - T)$ changes our formula for certainty equivalent to

$$c = ET + [EV - ET] / [1 + r]$$

where r is now defined by

$$r = [1 - 2\rho s(V)s(T) + (s^2(V)/s^2(T))]^{1/2} - 1$$

Hence adding a single parameter incorporates state-dependent utility effects into target-based decision analysis. High values of ρ imply that when the payoff is good, expectations are also good and, conversely, when the payoff is poor, expectations are also poor. Hence increasing ρ reduces the risk in the gamble and implies that the certainty equivalent is closer to EV . Conversely reducing ρ increases the risk in the gamble and implies that the certainty equivalent is closer to ET .

Thus target-based decision theory lends itself very naturally to modeling state-dependent preferences.

9. Better Linkages with Practice

Good decision-making involves:

- (1) Careful identification of the problem
- (2) The generation and evaluation of proposed solutions in order to select a solution

(3) Effective implementation of the chosen solution

Traditional decision analysis focuses on the second step. But there are many examples of technically excellent decision analyses that fail to get implemented. As a result, considerable attention is now focused on implementation.

Implementing a solution disrupts the status quo and whatever current plans the organization was implementing. (It's common to refer to these current plans as momentum.) This is often quite costly. If members of the organization feel that momentum was preferable to the proposed new solution, they may either passively resist or actively sabotage implementation. If a decision-maker ignores their resistance and successfully implements the proposed solution, that decision-maker could then be held accountable if the proposed solution ends up being worse than the status quo.

Hence it is important that the proposed solution appear to be clearly better than momentum. This implies that the proposed solution should be relatively easy to communicate and implement. It also implies that the proposed solution needs to have a good chance, should it be implemented, of outperforming momentum. This leads to an informal decision criterion: choose the decision with the maximum chance of outperforming momentum.

Now standard decision analysis treats all decision alternatives as equal and evaluates them using the same metrics. It does not give the momentum alternative a preferred status. (In a few cases, analysts attach a cost of change to each of the proposed new solutions in order to give momentum special treatment. But this can be awkward.) Hence there is a discrepancy between the formal criteria which we often use in a decision analysis and the informal criteria which are clients often use. As Bordley (2001) noted in the *Journal of the Operational Research Society*, the target-based interpretation of utility offers an easy way to resolve this discrepancy.

Consider the possible future outcomes that might result if the organization implements momentum (e.g., if it does nothing). Assign each of these outcomes a possible value. Then we can think of momentum as a random variable, T , which leads to various possible values. Suppose we define the utility of a consequence as the probability that it outperforms momentum and the expected utility of a decision as the probability that the decision outperforms momentum. Choosing the alternative which maximizes utility involves choosing the alternative with the maximum chance of outperforming momentum. (Since the payoffs from momentum and the payoffs from any proposed solution are likely to be correlated, this utility function is implicitly state-dependent.) Assigning a utility of u to momentum implies that we only implement an alternative different from momentum if that alternative has at least a u chance of being better than momentum.

There's an interesting connection between this formulation and the theory of hypothesis-testing. Suppose the null hypothesis is the hypothesis that the proposed alternative is not better than momentum. Rejecting the null hypothesis corresponds to assuming that the proposed alternative is significantly better than the momentum. Since the utility of the proposed alternative corresponds to the probability of it outperforming momentum, the higher the utility, the more that the proposed alternative will be significantly better than momentum and the greater the confidence with which we can reject the null hypothesis.

10. More Consistent with Psychological Evidence

10.1. Cognitive Psychology

Kahneman and Tversky (1979) postulated that individuals make choices to optimize a value function which is S-shaped with the inflection point corresponding to the reference point. They note that the reference point is readily manipulated and showed how prospect theory with an S-shaped value function seems to explain a wide variety of empirical phenomenon. Heath, Larrick and Wu (1999) showed that one could further validate prospect theory by interpreting the reference point as an individual's goal. In this case, prospect theory can explain a wide variety of known empirical phenomenon in motivational psychology.

Kahneman and Tversky's value function is somewhat related to Friedman and Savage's S-shaped utility function. Friedman and Savage (1948) showed that this S-shaped utility function can explain why individuals simultaneously buy insurance and lottery tickets. At lower levels of income, individuals were risk-prone; at higher levels of income, individuals were risk-averse. Their formalism differs from Kahneman and Tversky's in assuming that the inflection point isn't manipulable but is a stable aspect of the individual's preferences.

Target-based decision theory provides some motivation for this S-shaped utility function. From the target-based perspective, an individual is only interested in getting an income which satisfies his needs. Unfortunately he is not certain about how much income is required. Let t be the individual's best estimate of the level of income required. But the individual recognizes that the required level of income could be either less than t or greater than t . His uncertainty about this required level of income can be described by defining a random variable T with a single mode at t .

For this distribution, the probability of T equaling some value x initially increases, as x increases, until $x = t$, and then decreases. This implies that the cumulative probability of T is convex for x less than t and concave for x greater than t . Hence the individual is inclined to take risks when x is less than t and inclined to avoid risks when x is greater than t . This reproduces the Friedman and Savage S-shaped utility function where the mode, t , corresponds to the inflection point of the Friedman-Savage utility and the reference point of Kahneman and Tversky's value function. The target-based model also explains why the individual is initially risk-prone and then risk-averse. When x is less than t , the individual knows he probably doesn't have enough income to meet needs and will therefore be inclined to take risks---since he is only concerned with meeting needs. When x exceeds t , the individual knows that he probably does have enough income to meet needs and will therefore be inclined to avoid risks.

10.2. Depth Psychology

Hence target-based decision analysis is already more consistent with empirical evidence in cognitive psychology than the concave utility function commonly used in utility-based decision analysis. However target-based decision analysis may also be more consistent with empirical evidence from areas of psychology outside cognitive psychology.

In the standard utility function, asking an individual for his preference between various gambles allows us to define utility over certain consequences as well as over gambles involving those consequences. In this formalism, the only uncertainty associated with choice arises from uncertainty about consequences. If the consequence was completely articulated and completely known, there would be no uncertainty.

But depth psychology, including the psychoanalytic movement, postulates that individual behavior is motivated by various needs of which the individual is only partially aware. Even if an individual knew all the potential consequences of some action, that individual could not be fully certain about whether this consequence would lead to contentment. Thus depth psychology suggests that in addition to there being uncertainty about external consequences, there is also an internal uncertainty about whether or not the individual's needs would be satisfied by those consequences.

Standard utility-based decision theory could attempt to incorporate these effects by modeling the individual's internal state as part of the external consequences. But this can be awkward. In contrast, we could directly interpret target-based decision analysis as modeling the individual as if he were uncertain about what would be required to satisfy his needs. In this case, the uncertainty about V reflects "external" uncertainty about the external world and the uncertainty about T reflects "internal" uncertainty about whether a given consequence will, in fact, satisfy those needs.

Hence target-based decision theory is inherently more consistent with some of the assumptions of depth psychology.

11. CONCLUSIONS

It's been common to model a rational individual's preferences using utility functions. But there is an alternative, mathematically equivalent, way of modeling preferences using uncertain requirements. Hence a rational individual can be characterized either by utility-based preferences or by target-based preferences. In this paper, we suggest that many individuals will have their preferences better described by the target-based formulation. We then review results in finance, psychology, marketing and organizational science suggesting that this target-based formulation seems much more consistent with empirical work in these areas. We also argue that target-based decision theory may be simpler and easier to teach. For this reason, we believe that the target-based approach to decision analysis deserves further exploration.

Acknowledgements

The author is grateful to Prof. Brachinger for the opportunity to contribute to the EOLSS. I also thank an anonymous referee for helpful suggestions.

Bibliography

Bentham, Jeremy (1781). *Principle of Morals and Legislation*, p.1-4. [This is Bentham's seminal work on the foundations of utilitarianism.]

Bordley, R and M.LiCalzi (2000). "Decision Analysis using Targets instead of Utility Functions." *Decisions in Economics and Finance*. 23,53-74. [This work extended LiCalzi and Castiglione's original work on the equivalence of utility-maximization and maximizing the probability of exceeding stochastic thresholds to non-monetary outcomes. It also discusses practical applications of the idea.]

Bordley, R. (2001) "Integrating Gap Analysis and Utility Theory in Service Research." *Journal of Service Research*. 3,4. [This paper develops a model of customer expectations for service quality using target-based decision analysis.]

Bordley, R. (2001). "Teaching Decision Theory in Applied Statistics Courses." *Jr. of Statistical Education*. 9,2. [This paper develops an integrated applied statistics course including reliability, experimental design and regression centered around the idea of target-based decision analysis.]

Bordley, R. (2001). "Comparing Value-Focused Thinking and Interactive Planning." *Jr. of Operational Research Society*.(in press). [This paper shows that Ackoff's interactive planning model---which is a non-quantitative method for strategic decision making---is closely related to a value-focused thinking--a quantitative approach to strategic decision making.]

Brody, T., de la Pena, L, de la Pena,L., Hodgson,T (1994) *The Philosophy Behind Physics*, Springer Verlag. [This book is a collection of Brody's writings on the stochastic interpretation of quantum mechanics.]

Castagnoli,E. and M.Li Calzi (1996) "Expected Utility without Utility." *Theory and Decision*. 41,281-301. [This seminal paper showed that one could replace the utility function in decision theory by a stochastic threshold.]

Friedman M. and L. J. Savage (1948). "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*. 65,279-304. [This paper argued for replacing the standard concave utility function by a partially convex, partially concave utility function which could explain observed preferences for both gambling and insurance.]

Heath, Ch, R. Larrick and G. Wu (1999). "Goals as Reference Points." *Cognitive Psychology*. 38, 79-109. [This paper applied Tversky's prospect theory to a different field of psychology, human motivation]

Hobbes,T (1691) *Leviathan*,in C.B.Mcpherson (1982). *Leviathan*. Pelican Classics, p.161. [This seminal work models human beings as driven by rational self-interest. This leads to perpetual conflicts which can only be resolved by an authoritarian government.]

Hobbes, T (1640) *Elements*, in F.Woodbridge (1954) *Hobbes Selections*,Charles Scribner, New York, p.54. [This collects Hobbes's thoughts on human psychology, religion and government.]

Howard, R (1992). "In Praise of the Old Time Religion" in W.Edwards(ed.) *Utility Theories: Measurements and Applications*, Kluwer, Boston and Dordrecht,27-55. [This paper critiques some popular alternatives to decision analysis].

Kahneman, D. & A. Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*. 47,263-291. [This seminal paper presents prospect theory, Kahneman and Tversky's theory of the S-shaped value function describing individual reaction to losses and gains.]

Macpherson,C. (1962). *The Political Theory of Possessive Individualism*. Oxford Press, London.p.2. [This book by a noted political philosophy discusses the political philosophy of Hobbes, Locke and other rationalists.]

Sabine, G. (1961). *A History of Political Theory*. Holt, Rinehart & Winston, New York. [This book provides a solid history of political philosophy including such writers as Plato, Hobbes, Aquinas, Locke and Hegel.]

Savage, L (1964). *The Foundations of Statistics*. Wiley Publications, New York. [This seminal book showed that a rational person, whose behavior followed certain axioms, must act so as to maximize the expected value of a utility function.]

Simon, Herbert (1978) "Rationality as Process and as Product of Thought." *American Economic Review*. 68,2. [This article, written upon Simon's receipt of the Nobel Prize, argues that Economics needs to abandon its reliance on the notion of utility-maximization in favor of a notion of satisficing targets.]

Slovic, P. and A. Tversky (1974), "Who Accepts Savage's Axiom?" *Behavioral Science*. 19, 468-373. [This paper suggested that there might be justifiable reasons why human beings don't follow Savage's axioms of rationality.]

Von Neumann, J. and O. Morgenstern. (1947) *Theory of Games & Economic Behavior*. Princeton University Press, Princeton. [This seminal work was the theoretical foundation for individual decision making given known probabilities. It also was the foundation for the theory of games.]