P.P. Wakker

VALUE THEORY: AN ANALYSIS OF CHOICES UNDER RISK*

Daniel Kahneman and Amos Tversky

The Hebrew University, Jerusalem

The idea that men's choices under uncertainty maximize expected utility has long dominated the analysis of risky decisions. Expected utility theory was first formulated three centuries ago by Daniel Bernoulli (1738), and first axiomatized three decades ago by von Neumann and Morgenstern (1944). It was further developed by Savage (1954) who integrated the notion of subjective probability into utility theory. The theory has been applied both normatively and descriptively. It has been applied as a descriptive theory in economics, to explain various phenomena such as the purchase of insurance and the relation between spending and saving (see Arrow, 1971). It has been applied as a normative theory in decision analysis to determine optimal decisions and policies (see Raiffa, 1968). Indeed, most students of the field regard the axioms of utility theory as canons of rational behavior in the face of uncertainty, and they also regard them as a reasonable approximation to observed economic behavior. Thus, it is assumed that all reasonable people would wish to obey the axioms, and that most people actually do, most of the time.

* Paper presented at a conference on Public Economics, Jerusalem, June, 1975. This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under contract N00014-73-C-0438 to the Oregon Research Institute, Eugene. Additional support for this research was provided by the Research and Development Authority of the Hebrew University, Jerusalem, Israel.

In the present paper we show that utility theory, as it is commonly interpreted and applied, is not an adequate descriptive theory of individual choice under risk, and we propose an alternative model, called value theory.

Decision making under uncertainty can be viewed as a choice between gambles or lotteries. For simplicity, we shall restrict the discussion to gambles with monetary outcomes and (so called) objective probabilities, although many of our conclusions apply to more complicated gambles as well. In particular, we will be primarily concerned with gambles of the form (x, p, y) where one receives an amount of money x with probability p and an amount y with probability $1-p$. The option of receiving x with certainty is denoted by (x).

The application of expected utility theory to choices between gambles is based on the following three tenets:

(i) $U(x, p, y) = pu(x) + (1-p)u(y)$.

That is, the overall utility of a gamble, denoted by U, equals the expected utility of its outcomes. It is assumed that the gamble with the higher utility is chosen in any comparison. However, preferences between gambles depend on one's asset position.

(ii) (x, p, y) is acceptable at asset position w if and only if $U(w + x, p, w + y) > u(w)$.

That is, the gamble (x, p, y) is acceptable at a given asset position if the utility resulting from adding that gamble to one's assets exceeds the utility of those assets alone. Thus, the domain of the utility function u is final consequences (which includes one's asset position) rather than gains or losses.

(iii) u is concave ($u'' < 0$).

That is, the utility function for money is negatively accelerated, or equivalently the decision maker is risk averse. Formally, a person is risk averse if and only if he prefers any sure thing x over any gamble with expected value x. The presence of risk aversion is perhaps the best known generalization about preferences between gambles. For example, most people prefer 450 for sure over the gamble (1000, 1/2, 0) although its expected actuarial value is 500. Similarly, most people are not willing to accept the gamble (1001, 1/2, -1000) although its expected value is positive. Indeed, the presence of risk aversion led the early decision theorists of the 18th century to propose that utility is a concave function of money. In the modern axiomatic approach to utility theory, the concavity of the utility function is not assumed, but inferred from preferences. For modern treatment of risk aversion, see Pratt (1964), and Arrow (1971). Although utility functions with convex regions have been considered to accomodate gambling (Friedman and Savage, 1948; Markowitz, 1952), most applications of the theory employ concave utility functions. Thus, risk aversion is explained by the shape of the utility function for money, with no reference to risk per se.

We shall argue that these three tenets of utility theory are incorrect as a description of choice behavior. Specifically, we shall demonstrate three phenomena, the reflection effect, the certainty effect and the isolation effect, which respectively violate assumptions (iii), (i) and (ii) above. These effects were observed in the responses of several hundred students to problems involving choices between pairs of hypothetical options.

The reliance on hypothetical questions to test the descriptive adequacy of utility theory will certainly raise some doubts about the validity of the method, and the generalizability of the results. We are keenly aware of this difficulty. However, it should be pointed out that the various methods by which utility theory has been tested also suffer from severe drawbacks. Real choices can be investigated either in the field, by naturalistic or statistical observations of economic behavior, or in the laboratory. Field studies can only provide for rather crude tests of qualitative predictions, because probabilities and utilities cannot be adequately measured in such contexts. Laboratory experiments have been designed to obtain precise measures of utility and probability from actual choices, but these experiments typically involve contrived gambles for small stakes, and a large number of repetitions of very similar problems. Consequently, the results tend to reflect demand characteristics and other extraneous effects which depend on display and mode of presentation. Indeed, very few consistent findings have emerged from the experiments designed to shed light on the descriptive adequacy of utility theory.

By default, the method of hypothetical questions emerges as the simplest procedure by which a large number of theoretical questions can be investigated. The use of the method relies on the assumption that people have some idea of how they would behave in actual situations of choice, and on the further assumption that the subjects have no special reason to disguise their true preferences in responding to hypothetical questions. The relation between people's responses to hypothetical questions and their actual choices is an important empirical problem that deserves investigation. Note, however, that the applications of decision analysis to real problems also rely on hypothetical question to measure the probabilities and utilities of outcomes. In all the hypothetical problems discussed in this paper, the modal answer was chosen by more than 2/3 of respondents. At the very least, the pattern of these answers reveals that most people believe they would make choices which can be shown to violate utility theory. If subjects are reasonably accurate in the prediction of their behavior, then the prevalence of responses that violate utility theory provides some presumptive evidence that the theory indeed fails to describe actual choices.

The Reflection Effect

Table 1 displays the modal preferences obtained for ten problems, each involving a choice between a semi-positive or semi-negative gamble, (i.e., gambles in which one of the outcomes is zero) and a sure gain or loss equal to the expected value of that gamble. The preference relation is denoted by >. The two choice problems in each row are mirror images, that is, the signs of all outcomes are reversed.

The risk-aversion hypothesis of utility theory implies that the certain option should be preferred to the gamble in all ten comparisons. The preferences summarized in Table 1 clearly violate this prediction.

Table 1

(1000, 1/2, 0)	<	(500)	(-1000, 1/2, 0)	>	(-500)
(200, 1/4, 0)	<	(50)	(-200, 1/4, 0)	>	(-50)
(5000, 1/4, 0)	<	(1250)	(-5000, 1/4, 0)	>	(-1250)
(200, 1/100, 0)	>	(2)	(-200, 1/100, 0)	<	(-2)
(5000, 1/1000, 0)	>	(5)	(-5000, 1/1000, 0)	<	(-5)

took a [illegible] on ~~[illegible]~~ Budget Constraint [illegible] the effect on the [illegible] [illegible]

The striking fact about Table 1 is that preferences in the two columns are reversed. In the first three rows of the Table, where the probabilities of gains or losses are moderate or high, preferences are risk-averse in the positive domain (left column) and risk-seeking in the negative domain (right column). In the last two rows of the Table, where probabilities and expected values are low, preferences are risk-seeking in the positive domain and risk-averse in the negative domain. This pattern of preferences is summarized in Table 2.

Table 2

	Positive Domain	Negative Domain
Moderate or high Probability	Risk-aversion	Risk-seeking
Low Probability	Risk-seeking	Risk-aversion

The Table suggests that risk-aversion is limited to moderate-probability gambles in the positive domain, and to low-probability gambles in the negative domain. It also suggests that risk-seeking is not limited to the purchase of lottery tickets which involve small probabilities

Suppose we have the Follow's Gambles

$$\left.\begin{array}{l} (-20, \frac{1}{4}, 0) \gtrless -5 \\ (-8, \frac{1}{4}, 0) \gtrless -2 \end{array}\right\}$$

[illegible]

of gain, " , but is also manifest in choices involving moderate or high probabilities of loss. An account of this pattern of preferences will be presented later. For the moment, we wish to underscore the observation that, for a large class of comparisons between gambles and their expected values, the reflection of options around zero reverses the preference relation. We refer to this phenomenon as the <u>reflection effect</u>.

<u>The Certainty Effect</u>

Table 3 displays modal preferences for a set of four choice problems. In this Table, the problems on the right are obtained from those on the left by reversing the signs of outcomes. The problems in the bottom row are

Table 3

(1000, 1/2, 0)	<	(450)	(-1000, 1/2, 0)	>	(-450)
(1000, 1/20, 0)	>	(450, 1/10, 0)	(-1000, 1/20, 0)	<	(-450, 1/10, 0)

obtained from those in the top row by reducing the probabilities of gains or losses by a factor of 10. In particular, each option on the bottom row can be expressed as a mixture of the corresponding option on the top row and zero, with a probability of 1/10. For example, (1000, 1/20, 0) = ((1000, 1/2, 0), 1/10, 0) and (450, 1/10,0) = ((450), 1/10, 0). The substitution axiom of utility theory asserts that the preference order is invariant over probability mixtures. That is, if an option A is chosen over another option B, then (A, p, 0) must be preferred over (B, p, 0), for any probability p.

The preferences in Table 3 violate this axiom, in both the positive and the negative domains. Since $(450) > (1000, 1/2, 0)$, it follows from utility theory, with $u(0) = 0$, that $2u(450) > u(1000)$. However, the fact that $(450, 1/10, 0) < (1000, 1/20, 0)$ implies the reverse inequality. Thus, the estimated utility of a gain of 450 is higher when it is inferred from a choice in which this outcome is certain, than when it is inferred from a choice in which the same outcome is not obtained with certainty. Precisely the same analysis applies to the negative options. Since $(-1000, 1/2, 0) > (-450)$, it follows from utility theory that $u(-1000) > 2u(-450)$. However, the fact that $(-450, 1/10, 0) > (-1000, 1/20, 0)$ again implies the reverse inequality. Thus, the disutility of a loss of 450 appears greater when it is certain than when it is not. This type of violations of utility theory, which arises when an outcome appears in both risky and riskless options, is called the <u>certainty effect</u>.

The certainty effect is not restricted to choices between gambles with monetary outcomes. To illustrate, we presented subjects with the following choices:

A1: A one-week tour of England

A2: A 50% chance to win a 3-week tour of England, France and Italy.

The same subjects were asked to choose between:

B1: A 10% chance to win a one-week tour of England

B2: A 5% chance to win a three-week tour of England, France and Italy.

The majority of subjects chose A1 over A2, and B2 over B1, in defiance of the substitution axiom of utility theory. Thus, the one-week tour of England appeared relatively more attractive in the first problem, where it enjoyed a certainty advantage over the alternative, than in the second problem, where both outcomes were uncertain.

The isolation effect

According to utility theory, choices are made in terms of their final concequences. Thus, (x, p, y) is as desirable as (z) in asset position w whenever

$$u(w + z) = U(w + x, p, w + y) = pu(w + x) + (1 - p)u(w + y).$$

In contrast, we propose that people usually evaluate a gamble in isolation, that is, in terms of the gains and losses associated with that particular gamble, rather than in terms of the final asset positions that may result from playing it. Consequently, different formulations of a choice problem that are identical in terms of final assets can elicit different preferences. [illegible]

To demonstrate, we presented subjects with the following choice problems:

Problem A. In addition to whatever you own, you have been given 1000. You are now asked to choose between the following options:

$A_1 = (1000, 1/2, 0)$ $\qquad$ $A_2 = (500)$

Problem B. In addition to whatever you own, you have been given 2000. You are now asked to choose between the following options.

$B_1 = (\cancel{-1000}\ 0, 1/2, \cancel{0}\ -1000)$ $\qquad$ $B_2 = (-500)$

The overwhelming majority of subjects chose A_2 in problem A, and B_1 in problem B. These preferences are inconsistent with utility theory, because, in terms of final assets:

$$A_1 = (w + 2000, 1/2, w + 1000) = B_1 \text{ and } A_2 = (w + 1500) = B_2,$$

where w denotes the subject's initial wealth. In fact, problem B is obtained from problem A by adding 1000 to the initial bonus, and subtracting 1000 from the outcomes of both options. Evidently, people respond differently to the two formulations. They evaluate the options in isolation, without combining them with the bonus that was added to their wealth. This behavior is called the <u>isolation effect</u>.

The fact that (500) > (1000, 1/2, 0) while (-500) < ($\cancel{-1000}$ 0, 1/2, -1000) demonstrates the sensitivity of the preference order to a translation of the entire choice problem by 1000. In contrast, when subjects were asked to assume different asset positions when choosing between gambles, it was found that a difference of 1000 in initial wealth had no effect on preferences. This analysis suggests that the carriers of utility are not final asset positions, but rather the gains and losses associated with a gamble.

THEORY

Gambles of the form $(x; p, y)$ can be classified according to the signs of their outcomes. A gamble is called mixed if one outcome is positive and the other outcome is negative. It is called semi-positive or semi-negative, respectively, if one outcome is zero and the other is positive or negative. It is called purely positive or negative, respectively, if the outcomes are either both positive or both negative.

The present theory is formulated in terms of two basic equations. The first applies to mixed, semi-positive and semi-negative gambles, and the second applies to pure gambles. As in utility theory, the present theory associates a value V with each option so that the preference ordering between options coincides with the ordering of their V-values. That is, $A > B$ if and only if $V(A) > V(B)$. The overall value of a gamble is expressed in terms of two functions π and v. π associates with each probability p an uncertainty weight $\pi(p)$. v assigns to each outcome x a number $v(x)$ that reflects the value of that outcome. Note that V is defined over gambles, whereas v is defined on monetary outcomes. The two scales coincide for degenerate gambles where $V(x, 1, 0) = V(x) = v(x)$.

what does this mean

The first basic equation of the theory is

(1) If $x \geq 0 \geq y$ or $x \leq 0 \leq y$, then

$$V(x, p, y) = \pi(p)v(x) + \pi(1-p)v(y), \quad \text{where } v(0) = 0.$$

Equation (1) is analogous to the expected utility formulation, except that the classical utilities and probabilities are replaced by values and uncertainty weights. Note that v is defined for monetary changes (i.e., gains and losses) from a given asset position. In principle, therefore, for each asset position there is a different value function. However, we have already observed that the preference order between gambles is essentially unaffected by moderate changes in asset position. Hence, it appears that over a reasonably wide range of assets the value function is approximately the same.

From Equation (1), we can evaluate the cash-equivalent of a gamble. An amount $C(x, p, y)$ is the cash-equivalent of the gamble (x, p, y) if the subject is indifferent between playing the gamble and receiving, or paying, $C(x, p, y)$ for sure. It follows, therefore, that $V(x, p, y) = V(C(x, p, y)) = v(C(x, p, y))$, and hence, by Equation (1),

$$C(x, p, y) = v^{-1}(V(x, p, y)) = v^{-1}(\pi(p)v(x) + \pi(1-p)v(y))$$

whenever $x \geq 0 \geq y$ or $x \leq 0 \leq y$.

The second basic equation applies to purely positive and purely negative gambles. We propose that such gambles are analyzed into two components. (i) The riskless component, i.e., the minimum gain or loss which is certain to be obtained or paid. (ii) The risky component, i.e., the additional gain or loss which is actually at stake in the gamble. Equation (2) describes the manner in which the riskless and risky components are combined to determine the cash-equivalent of a pure gamble.

(2) If $x > y > 0$ or $x < y < 0$ then

$$C(x, p, y) = y + C(x-y, p, 0).$$

Equation (2) asserts that the cash equivalent of a pure gamble equals the riskless component of the gamble, i.e., y, plus the cash-equivalent of the residual semi-positive (or semi-negative) gamble, i.e., $C(x-y, p, 0)$. For example, $C(400, 1/2, 100) = 100 + C(300, 1/2, 0)$, and $C(-500, 1/10, -200) = -200 + C(-300, 1/10, 0)$. Note that the risky component $(x-y, p, 0)$ of the weight reflects a person's readiness to gamble on an event, rather than his degree of belief in its occurrence.

The uncertainty weight associated with an event is a function of one's subjective probability for that event. In the problems analyzed in the present paper, we assume that the subject adopts the stated values as his subjective probabilities, and we may therefore express the uncertainty weights as a function of stated probabilities. We turn now to discuss the properties of uncertainty weights.

First, we assume that π is a strictly increasing function of p, with $\pi(0) = 0$ and $\pi(1) = 1$.

Second, we assume that $\pi(p) > p$ for small values of p and $\pi(p) < p$ for moderate and large values of p. That is, the uncertainty weight exceeds the stated probability for events of low probability, while for all other events the uncertainty weight is smaller than the corresponding probability. This property of uncertainty weights should not be confused with the well-known observation that subjects overestimate low probabilities and underestimate high ones. Errors in the estimation of probability represent a discrepancy between judged probabilities and so-called objective probabilities. The uncertainty weight represents a further distortion that is applied to the subject's probability.

called subcertainty

Third, we assume that $\pi(p) + \pi(1-p) < 1$, for all $0 < p < 1$. Since $\pi(1) = 1$, this assumption implies that sure things are overweighted relative to uncertain options. This is the essence of the certainty effect.

Figure 1 presents an example of a relation between uncertainty weight and probability which illustrates these assumptions. Note that, over most of

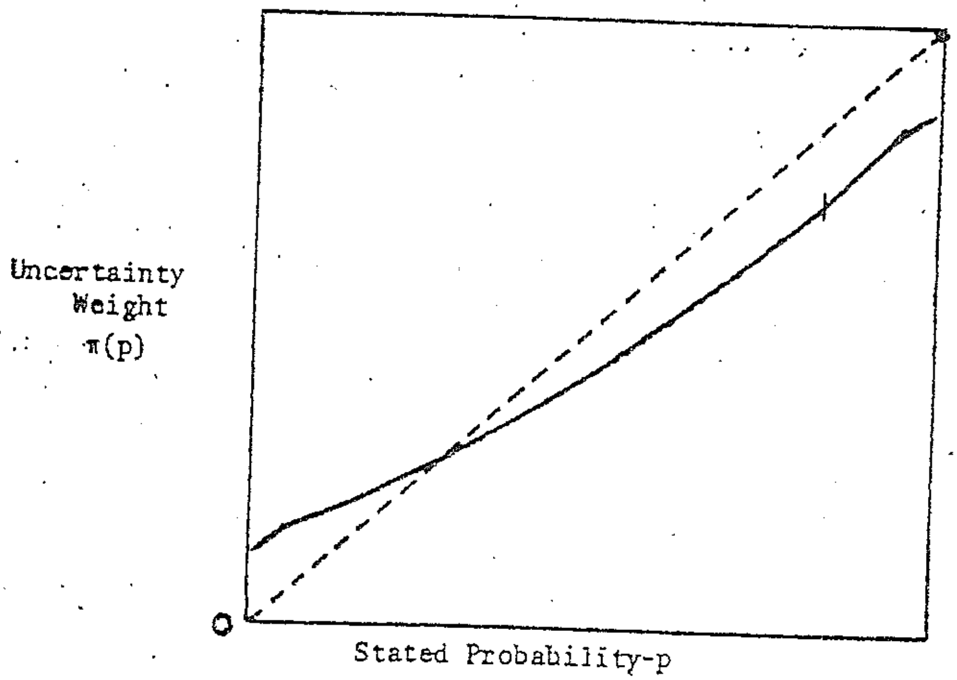

Fig. 1. Illustrative uncertainty weight function

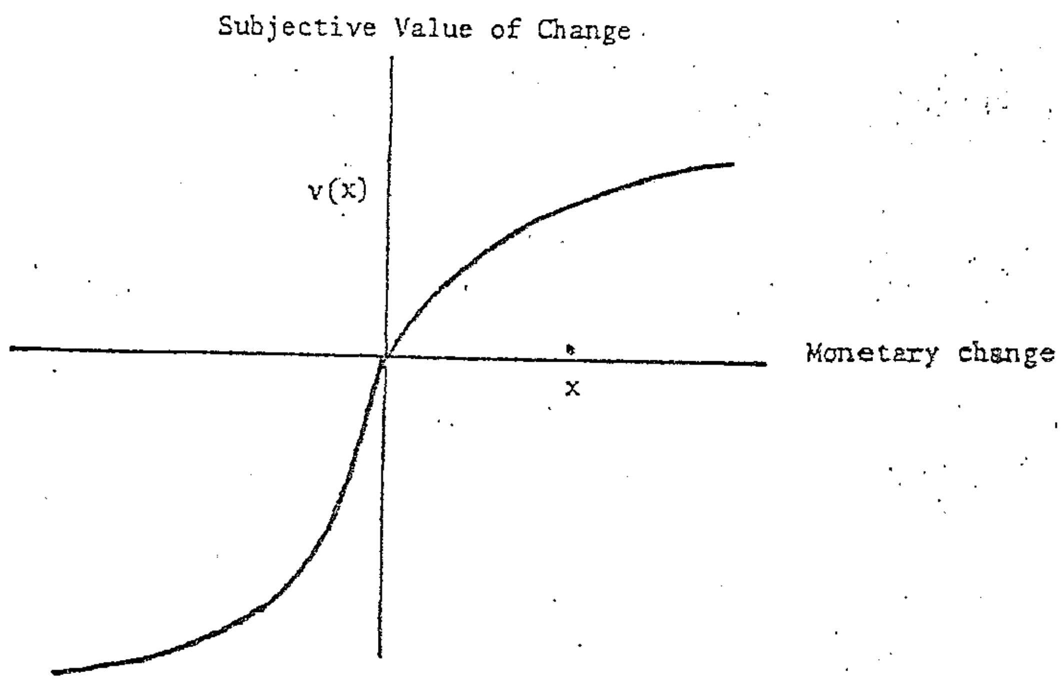

Fig. 2. Illustrative value function

convexity in middle is not mentioned

the range, π is regressive, i.e., the slope of π is less than one. The function has been drawn effectively discontinuous at the end points.

The present treatment is not restricted to gambles with stated probabilities that are accepted by the subject. In the absence of stated probabilities, we assume that the subject establishes his probabilities for the outcomes, and applies the π-function to these probabilities. In these cases, the regressiveness of uncertainty weights with respect to objective probabilities will be further enhanced by the tendency to overestimate low probabilities and underestimate high ones.

This discussion has distinguished subjective probability, which is a measure of degree of belief, from the uncertainty weight that is inferred from gambling decisions. Our analysis suggests that attempts to infer subjective probabilities from risky choices actually recover uncertainty weights, which reflect attitudes to risk as well as degree of belief.

The Value Function

The present theory introduces a value function v whose domain is changes in wealth, i.e., gains and losses. Naturally, v is strictly increasing and v(0) = 0. We do not assume that choices between gambles are independent of assets, since a drastic change in wealth can surely change one's preferences. Strictly speaking, therefore, v is a function in two arguments: current wealth and magnitude of change. It seems, however, that the preference relation between gambles is relatively insensitive to wealth and highly sensitive to changes. Furthermore, the value function for changes appears to have the same general shape at different levels of wealth. Con-

May be hard to measure on the basis of gambles in isolation. The definition of gamble we drawn 17 consider sequently, it is possible to explain most choices without reference to initial wealth.

We turn now to describe the properties of the assumed value function. The first property refers to the relative impact of gains and losses. For most people, losses loom larger than gains. For example, the great majority of subjects exhibited the following preference order:

$$(250, 1/2, -250) > (1000, 1/2, -1000) > (2000, 1/2, -2000).$$

Although people occasionally engage in fair symmetric bets, it appears that in general the preference order for gambles of the form $(x, 1/2, -x)$ is inversely related to the magnitude of x. Therefore, by Equation (1), if $x > y$

$$\pi(1/2)[v(x) + v(-x)] < \pi(1/2)[v(y) + v(-y)]$$

Hence

$$v(x) - v(y) < v(-y) - v(-x).$$ let y=0

Consequently, $v(x) < -v(-x)$, and $v'(x) < v'(-x)$ for all x, provided the first derivative of v exists. That is, the value function is steeper for any given loss than for the corresponding gain.

It is generally assumed that utility is a concave function of money, i.e., that the utility function is negatively accelerated, or that $u'' < 0$. We hypothesize that v is concave above zero and convex below zero. The value-difference between a loss of 1100 and a loss of 1,200 appears smaller, to most people, than the value-difference between a loss of 100 and a loss of 200. An illustration of a value function with these properties is displayed in Figure 2. The hypothesis that the value function is convex for $x < 0$ is supported by the observation that, for the great majority of our subjects, Evidence against

(-250, 1/2, 50) > (-200, 1/2, 0). This observation implies that

$$v(50) - v(0) > v(-200) - v(-250).$$

That is, the rise of the value function in the interval (0, 50) is greater than the rise in the interval (-250, -200), and hence v cannot be concave everywhere. Moreover, since $v(x) < -v(-x)$

$$v(0) - v(-50) > v(50) - v(0) > v(-200) - v(-250)$$

and

$$v(-250) > v(-200) + v(-50)$$

Hence, the value function for negative changes behaves in a super-additive fashion which is necessary, though not sufficient, for convexity.

The proposed value function has three essential features. First, it is defined for gains and losses rather than for wealth. Second, it becomes steeper as one approaches the origin from either above or below. Third, it is steeper for negative changes than for comparable positive changes.

These features of the hypothesized value function are compatible with some basic principles of perception and judgment. First, the sensory apparatus is attuned to the detection of changes or differences rather than to the evaluation of absolute magnitudes. In every sensory domain, the past and the present context of experience define an adaptation level, and stimuli are perceived in terms of their relation to the adaptation level (Helson, 1964). Thus, an object at a given temperature may be experienced as hot or cold to the touch, depending on the temperature of objects to which one has adapted.

Second, for many sensory and perceptual continua the sensitivity to changes decreases as one moves away from the adaptation level. This property of the perceptual system is clearly adaptive. It maximizes the sensitivity of the system to small changes that are most commonly encountered.

Third, the greater sensitivity to negative rather than to positive changes is not specific to monetary outcomes. It reflects a general property of the human organism as a pleasure machine. For most people, the happiness involved in receiving a desirable object is smaller than the unhappiness involved in losing the same object. A high sensitivity to losses, pains, and noxious stimuli also has adaptive value. Happy species endowed with infinite appreciation of pleasures and low sensitivity to pain would probably not survive the evolutionary battle.

Considerations of simplicity have led us to limit the present discussion to two-outcome gambles. However, value theory can be readily extended to n-outcome gambles. Let $(x_1, p_1; \ldots; x_n, p_n)$ denote a gamble in which the amount x_i is obtained with probability p_i, $1 \le i \le n$, and $\Sigma p_i = 1$. In this case, Equations (1) and (2) are generalized as follows:

Whenever $\min x_i \le 0 \le \max x_i$, then $V(x_1, p_1; \ldots; x_n, p_n) = \sum_{i=1}^{n} \pi(p_i)v(x_i)$.

If $\min x_i = y > 0$, then $C(x_1, p_1; \ldots; x_n, p_n) = y + C(x_1 - y, p_1; \ldots; x_n - y, p_n)$.

If $\max x_i = z < 0$, then $C(x_1, p_1; \ldots, x_n, p_n) = z + C(x_1 - z, p_1; \ldots; x_n - z, p_n)$.

The empirical question regarding the invariance of values and uncertainty weights over the number of outcomes deserves careful investigation.

Very important

IMPLICATIONS

The certainty effect

Let us consider some of the implications of value theory, starting with the analysis of the certainty effect demonstrated earlier in this paper. According to Equation (1)

$$(1000,1/2,0) < (450) \quad \text{iff} \quad \pi(1/2)v(1000) < \pi(1)v(450)$$

On the other hand,

$$(1000,1/20,0) > (450,1/10,0) \quad \text{iff} \quad \pi(1/20)v(1000) > \pi(1/10)v(450)$$

Consequently, the above preferences (which violate utility theory) are obtained iff

$$\frac{\pi(1/2)}{\pi(1)} < \frac{v(450)}{v(1000)} < \frac{\pi(1/20)}{\pi(1/10)}$$

Recall from Fig. 1 that $\pi(1) = 1.0$ and that π is a regressive function of p. In particular, the uncertainty weights satisfy

$$\frac{\pi(1/2)}{\pi(1)} < 1/2 < \frac{\pi(1/20)}{\pi(1/10)}$$

Hence, substitutability is violated for any pair of outcomes for which the ratio of values falls between the corresponding ratios of uncertainty weights.

The same argument shows that the corresponding preferences among negative gambles, i.e.;

$$(-1000,1/2,0) > (-450) \text{ and } (-1000,1/20,0) < (-450,1/10,0)$$

are obtained iff

$$\frac{\pi(1/2)}{\pi(1)} < \frac{v(-450)}{v(-1000)} < \frac{\pi(1/20)}{\pi(1/10)}$$

Thus, the violation of substitutability for both positive and negative gambles is explained by the nature of the uncertainty weights. Note that the special properties of the value function were not utilized in this analysis. The certainty effect could be obtained, for example, with a linear value function.

Perhaps the best-known counter example of utility theory has been proposed by the French economist Maurice Allais (1953). As we shall show, Allais' example can be regarded as another manifestation of the certainty effect. Consider the following decision problems.

Problem A. Choose between:

A_1: 1,000,000 for sure

A_2: 1,000,000 with probability 0.89
5,000,000 with probability 0.10
0 with probability 0.01

Problem B. Choose between:

B_1: 1,000,000 with probability 0.11
0 with probability 0.89

B_2: 5,000,000 with probability 0.10
0 with probability 0.90

Most people choose A_1 in Problem A and B_2 in Problem B. These preferences violate expected utility theory. To demonstrate, suppose $u(0) = 0$, hence the former choice implies $0.11\ u(1{,}000{,}000) > 0.10\ u(5{,}000{,}000)$, while the latter choice implies the reverse inequality. Specifically, the above choices violate Savage's sure-thing principle, according to which the preference order between options remains invariant upon changes in the value of

Yes . . .

an outcome that is common to both options.

Extending Equation (1) in the natural fashion and applying it to Allais' example yields

$$A > B \text{ iff } \pi(1.0)v(1{,}000{,}000) > \pi(0.89)v(1{,}000{,}000) + \pi(0.10)v(5{,}000{,}000)$$

and,

$$D > C \text{ iff } \pi(0.10)v(5{,}000{,}000) > \pi(0.11)v(1{,}000{,}000)$$

It follows, therefore, that both preferences hold iff

$$\frac{\pi(1.0) - \pi(0.89)}{\pi(0.10)} > \frac{v(5{,}000{,}000)}{v(1{,}000{,}000)} > \frac{\pi(0.11) - \pi(0)}{\pi(0.10)} .$$

Thus, the observed preferences in Allais' example are accounted for by the following inequality of differences between uncertainty weights (see Figure 1),

$$\pi(1.0) - \pi(0.89) > \pi(0.11) - \pi(0).$$

This [illegible] strange def. of [illegible]

According to the present analysis, utility theory will be violated for any pair of outcomes whose value ratio falls between the respective ratios of uncertainty weights. In this treatment, A is preferred to B not because $u(1{,}000{,}000)/u(5{,}000{,}000) > 10/11$, but rather because the outcome of receiving 1,000,000 enjoys the certainty advantage in Problem A. Since both options in Problem B are uncertain, the certainty advantage vanishes and the preference ordering is reversed.

If this interpretation of Allais' problem is correct, it should be possible to reverse the preference order in Problem B by a sequential representation that restores the certainty effect.

Problem C

Consider the following two-stage game. In the first stage there is a probability of 0.11 to move into the second stage and a probability of 0.89 to end the game at this stage without winning anything. If you move into the second stage you have a choice between

C_1 = 1,000,000 for sure

C_2 = 5,000,000 with probability 0.91

0 with probability 0.09

You have to make your choice before the game starts, that is before the outcome of the first stage is known.

Note that C_1 and C_2 are formally equivalent to B_1 and B_2, respectively. By calculating the overall probability of winning we find that C_1 = (1,000,000, .11, 0) = B_1 and C_2 = (5,000,000, .10, 0) = B_2. Although most subjects preferred B_2 over B_1, they chose C_1 over C_2. To understand this reversal, observe that in the sequential representaion, the subject's choice does not affect the outcome of the first stage. Thus, subjects focus on the second stage and choose as if this stage will, in fact, be reached. Under this interpretation, the certainty effect is restored, and people behave as they did in Problem A. This result further demonstrates the certainty effect, and the manner in which choices depend on the representation of a decision problem.

Insurance and gambling

People spend billions of dollars to purchase insurance policies and lottery tickets, despite the negative expected values of both forms of investment. These behaviors present a major challenge to any descriptive theory of choice under risk. The best-known explanations of insurance and gambling refer to the shape of the utility function for money. The common assumption that utility is a concave function of wealth explains the risk-averse purchase of insurance, but it fails to explain the risk-seeking purchase of lottery tickets. To explain both behaviors, Friedman and Savage (1948) had to invoke a utility function that is concave in some regions and convex in others.

The value function introduced in the present theory explains neither insurance nor gambling. Since it is concave above zero and convex below zero, it favors risk-aversion in the domain of gains and risk-seeking in the domain of losses, and tends to make both insurance and gambling unattractive. In this theory insurance and gambling occur in spite of the value function, not because of it. They are explained in value theory by the properties of the uncertainty weights.

Not really a theory of insurance

The common feature of lotteries and insurance is that they involve a small probability of a large gain or loss. Overweighting this probability increases the attractiveness of a lottery and the aversiveness of an uninsured risk. Because $\pi(p) > p$ for small probabilities, people may be willing to pay more than the expected value of lottery tickets and insurance policies, although the shape of the value function militates against

I think you will find that in reality $\pi(p) > p$ for positive gains
$\pi(p) < p$ for losses

such choices. For larger probabilities, $\pi(p) < p$, and the theory predicts that neither insurance nor gambling will occur when the probability of loss or gain is substantial.

These predictions were confirmed in the data presented earlier in Table 1. It was shown there that people are willing to purchase insurance at expected value when the probability of loss is small (1/100, or 1/1000), but not when the probability of loss is moderate (1/4 or 1/2). This pattern held for both small and large losses (-200 and -5000). Similarly, people prefer a lottery ticket to its expected value when the probability of gain is low, but not when it is moderate. Again, this pattern held for both small and large gains. It is important to note that these conclusions are based on hypothetical choices. Nevertheless, we challenge anyone to make a profit by selling lottery tickets or insurance policies where there is a substantial probability of a large gain or loss. To find buyers for such propositions, additional inducements will probably be required, such as convenient maintenance service in the case of insurance, and an atmosphere of competition or entertainment in the case of gambling.

In sharp contrast to the present treatment, the risk aversion hypothesis of utility theory does not explain the purchase of lottery tickets at all, nor does it predict that the attractiveness of insurance policies is inversely related to the probability of loss. Utility theory can be viewed as an attempt to eliminate the concept of attitude to risk or uncertainty and to explain risky choices solely in terms of attitudes to money or wealth.

The evidence suggests that risky choices cannot be adquately explained by attitudes to gains, losses or wealth. It is necessary, in addition, to consider attitudes to uncertainty that are expressed in the overweighting of unlikely outcomes and in the underweighting of outcomes which are probable but not certain. When the probabilities of gain or loss are substantial, the uncertainty weights, in conjunction with the properties of the value function, produce risk-seeking in the negative domain and risk-aversion in the positive domain. When the probabilities are small, the uncertainty weights can produce risk-aversion in the negative domain (e.g., insurance) and risk-seeking in the positive domain (e.g., gambling).

Value theory does not purport to account for all forms of risk-seeking and risk-aversion. Many factors not included in this theory (e.g., regret, social pressure, superstition, magical thinking) probably play an important role in risky choices. Value theory is an attempt to modify those assumptions of utility theory that are most severely violated, so as to achieve a more realistic account of choice behavior.

Alternative Formulations of Choice Problems

One of the major themes of the present analysis is that, in making decisions under risk, people typically isolate a segment of the problem from its broader context. In the sequential version of Allais' problem, for example, the decision becomes operative only if the second stage is reached, and hence the choice is made by considering this stage in isolation. As a result, the sequential and non-sequential versions of the problem elicit different preferences, although the two versions assign identical final probabilities to the various possible gains. Earlier in the paper, we showed that the addition of a constant to all the outcomes of a gamble, which reversed their signs, had a markedly different effect from the addition of the same constant to initial wealth. Thus, two formulations of a choice problem which differ in gains and losses may lead to different preferences, although the two formulations are identical in terms of final assets. Indeed, the assumption that people usually evaluate gains and losses, rather than absolute wealth, is a basic tenet of the present theory. || X Y

What constitutes a gain or a loss, however, depends on the representation of the problem and on the context in which it arises. Consider, for example, a man who has spent an unhappy afternoon at the horse races, has already lost \$100 on the first four races and is now facing a decision on how to place his bet on the fifth and last race of the day. Being aware that he is \$100 out of pocket, he is likely to regard the decision of whether to pay \$10 on a 15:1 longshot as a choice between (40, p, -110) and (-100), rather than as a choice between (140, p, -10) and (0). As will be shown below, the present theory implies that a person is more likely to choose the gamble in the former representation than in the latter. This prediction is confirmed by the well-known observation that betting on longshots increases in the course of the racing day. evidence? [illegible]?

There are many other contexts in which the decision-maker evaluates a gamble in terms of gains and losses that differ from the actual sums of money that will change hands. We propose that the present theory applies to the gains and losses as perceived by the subject. A businessman who is doing less well than his competitors may view the maintenance of the status quo as a loss, although his balance shows a profit. Conversely, an entrepreneur who is weathering a slump with greater success than his competitors may interpret a small loss as a gain, relative to the larger loss that he had reason to expect. The tourist who had budgeted $50 for spending in a Las Vegas casino may regard all available bets as semi-positive, because he feels that the costs have already been paid in advance. These examples involve transformations which modify the perception of the gains and losses associated with the various options. Good

A common class of transformation is translation, where a positive or negative constant is added to all possible outcomes. The race-track example, for instance, was analyzed as a negative translation, where the bettor's losses up to the fifth race are added to all the options that are available to him at that time. S.C.

The present theory entails that, in many situations, negative translation of a choice problem will increase the willingness to take risks. In particular, we now show that if a gamble $(x, p, -y)$ is just acceptable, i.e., indifferent to (0), then a negative translation of the choice problem will make the gamble relatively more attractive, i.e., $(x-z, p, -y-z) > (-z)$, where $x, y, z \geq 0$.

Since $(x, p, -y)$ is indifferent to (0) we obtain, by Equation (1),

$$V(x, p, -y) = \pi(p)v(x) + \pi(1-p)v(-y) = 0 .$$

Letting $f(y) = -v(-y)$, yields

$$\pi(p)v(x) = -v(-y)\pi(1-p) = f(y)\pi(1-p).$$

We wish to show that for any $z \leq x$

$$V(x-z, p, -y-z) = \pi(p)v(x-z) + \pi(1-p)v(-y-z) > v(-z)$$

or equivalently that $f(z) + \pi(p)v(x-z) > \pi(1-p)f(y+z)$.

(By equation (2) the result is readily extended to $z > x$.)

Since f is concave, it suffices to show

$$f(z) + \pi(p)v(x-z) > \pi(1-p)(f(y) + f(z)),$$

or $\pi(p)(f(z) + v(x-z)) > \pi(1-p)f(y)$, since $\pi(p) + \pi(1-p) \leq 1$.

But
$$\begin{aligned} \pi(p)(f(z) + v(x-z)) &> \pi(p)(v(z) + v(x-z)) && \text{(since } f > v\text{)} \\ &> \pi(p)v(x) && \text{(by the concavity of } v\text{)} \\ &= \pi(1-p)f(y) && \text{(by assumption).} \end{aligned}$$

This result shows that negative translation of a choice problem where a gamble is just acceptable makes the translated gamble more attractive than the corresponding sure loss. Thus, the loser who has not made peace with his losses, and the ambitious entrepreneur who is never satisfied with what he can get, will accept risks that others may reject. Positive translation, on the other hand, does not necessarily reduce the acceptance of risks. The present theory allows for situations where (x, p, -y) is indifferent to (0) while (x+z, p, -y+z) > (z), $x, y, z \geq 0$.

We have argued that subjects tend to evaluate risky options in terms of gains and losses. However, people surely can and sometimes will evaluate options in terms of final consequences, as advocated by utility theory. In this case, all gambles will be viewed as positive, and the effective value function will be concave everywhere. Consider, for example, an individual whose current wealth is $60,000, and who is faced with a choice between (-1000, 1/2, 0) and (-500). Our results suggest that this individual would make a risk-seeking choice and prefer the gamble over the sure loss. However, this preference is quite likely to be reversed if a decision-analyst suggests to the individual that he should formulate the alternatives as (59,000, 1/2, 60,000) vs. (59,500). Here, a risk-averse choice appears more likely. Indeed, the individual's experience of the consequences of his choice may be affected by his formulation of the problem. It is conceivable, for example, that the decision-maker may face a sure loss of $500 with greater fortitude if he views the problem in terms of final assets rather than in terms of gain and loss.

The casting of choice problems in terms of final consequences eliminates one major class of risk-seeking preferences, but not another. It eliminates those risk-seeking choices that are due either to the certainty effect or to the convexity of the value function in the negative domain. On the other hand, this formulation does not eliminate a second class of risk seeking preferences which reflect the overweighting of small probabilities, namely, the purchase of lottery tickets.

22

2nd order prob. without KCTH

The Ellsberg paradox

The present treatment has appealed repeatedly to two psychological principles. (i) People commonly do not trace a decision problem to its final consequences. (ii) The intensity of subjective experience is a negatively accelerated function of the magnitude of the stimulus.

These assumptions are applied to explain one of the classic puzzles of utility theory - Ellsberg's paradox. Ellsberg (1961) considered gambles in which one may win 100 or nothing, depending on whether a ball that is to be drawn from an urn is red or black. He demonstrated that people have a strong preference to bet on either color when the urn is known to contain 50% black and 50% red balls than when the urn contains red and black balls in unknown proportions. Further, it was reported that people are willing to pay more for the opportunity to participate in the first gamble than in the second (Raiffa, 1961).

In the choice problems introduced by Ellsberg, one chooses between two gambles that are identical in terms of final outcomes and their probabilities. The only difference between the gambles is that the characteristics of the process that generates the final outcome (e.g., the composition of the urn) are known exactly for one gamble but not for the other. In all cases, people revealed a distinct aversion for the bet in which the composition of the urn was not known, in direct violation of utility theory.

The two psychological principles stated above, which are incorporated in value theory, jointly provide an explanation of the preferences observed in Ellsberg's problems. By the first principle, subjects fail to trace the problem to its final consequences, and it is therefore natural for them to view the proportion of winning balls in the urn as an outcome. As a consequence, the choice between the two urns becomes a choice between a sure thing (the known proportion) and a symmetric distribution around that proportion. By the second principle, these outcomes are evaluated according to a concave function, and hence the sure thing will always be selected. Thus, the preference for the urn of known composition is a form of risk-aversion. This interpretation is supported by the responses of subjects to the following choice problem:

too many confounding factors

"Imagine that you are to undergo an operation. The two possible consequences of this operation are that it may be fully successful or that it may fail, in which case a new operation will be required. The operation is performed in three clinics in your town. In each clinic there are two surgeons who alternate in operating, and you have no control over the choice of surgeon that will operate on you.

The probability of success in the operation is 50% in each of the clinics, but you have been given the following information:

In <u>Clinic A</u>, the two surgeons each have had a record of 50% successes in this particular operation.

In <u>Clinic B</u>, one of the surgeons has been successful in 70% of cases, and the other in 30% of cases.

In <u>Clinic C</u>, the chances of success are between 50% and 70% for one of the surgeons, between 30% and 50% for the other surgeon.

If you had to choose between these clinics (remembering that you have no control over the choice of surgeon within a clinic) which clinic would you choose? Rank the three clinics according to your preference."

The subjects' ranking indicated a clear preference for Clinic A over Clinic C and for Clinic C over Clinic B. This order of preference is entailed by the assumption that people treat probabilities as outcomes, and evaluate them by a concave value function. In such a function, the difference in value between 50% and 70% is smaller than the difference in value between 50% and 30%. Consequently, Clinic A, which represents a sure thing, should be the most preferred. Clinic B, which corresponds to the gamble (30%, 1/2, 70%) should be the least preferred, and Clinic C, which corresponds to a uniform distribution over the interval [30%, 70%] should be intermediate. Precisely this ordering was obtained. On the other hand, the hypothesis that Ellsberg's paradox stems from an aversion to ambiguity predicts that Clinic C should be the least preferred, because its description is the most ambiguous. Thus, the preferences among clinics can be treated as a second-order certainty effect without any reference to the inherent aversiveness of ambiguity. This is not to deny, howevever, that aversion to ambiguity exists and may play a role in some choice problems.

It is of some interest to note that the phenomena of insurance and gambling, which have been traditionally discussed in terms of the utility function, are explained in the present theory as uncertainty effects. In contrast, Ellsberg's paradox, which has commonly been discussed as an uncertainty effect, is explained in the present theory by reference to the concavity to the value function.

This discussion suggests that the effect described by Ellsberg is not merely an academic curiosity. It reveals a phenomenon that has far-reaching implications for real-world decision problems. Probabilities are naturally viewed as outcomes in many situations, e.g., when one considers the effect of various actions on the probability of a scientific breakthrough, the bankruptcy of a firm, or the survival of a patient. In general, the effect of an action on the critical probability will be harder to predict for some actions than for others. The risk-averse attitude to vague probabilities will induce a bias favoring that course of action for which the relevant probability is least ambiguous. We suspect that this bias plays an important role in many personal and public decisions.

Normative Implications

Utility theory has been used in two ways: as a descriptive theory of human choices under uncertainty, and as a prescriptive theory of how a rational person should behave. The main conclusion of the present paper is that utility theory does not provide an adequate description of the choices people make between risky options. We turn now to examine the prescriptive implications of this conclusion.

The great majority of decision theorists regard the axioms of utility theory as valid normative principles of decision making under uncertainty. They believe that any reasonable man who understands the axioms would wish to satisfy them and would regard their violation as an error. The knowledge that strict adherence to the axioms will protect one from losing propositions (e.g., a Dutch book) increases the normative appeal of the theory. The idea that people wish to obey utility theory, even though they often violate it in practice, is the basis of many prescriptive applications. In decision analysis, for example, the relevant utilities and subjective probabilities of the decision-maker are inferred from his responses to hypothetical choice problems, and these values are then used to select the option whose expected utility is maximal.

To justify this prescriptive application of utility theory, one assumes (i) that the theory is adequate from a normative standpoint, and (ii) that the observed departures from the theory represent error, confusion, and other aberrations.If the violations of the axioms are large and systematic

XX

rather than small and random, it is impossible to infer utilities and probabilities from reported preferences. The application of decision analysis therefore depends on the descriptive validity of utility theory, at least as a first approximation.

We have argued that the observed violations of utility theory are both large and lawful, and that they cannot be treated as random error. Specifically, two sources of violations of expected utility theory have been identified. The first is the tendency to isolate a choice problem from one's assets and evaluate it in terms of gains and losses. The second is the replacement of subjective probabilities by uncertainty weights which reflect attitude to uncertainty and not merely degree of belief. These findings invalidate the attempt to infer utilities and subjective probabilities from preferences. No consistent utility function for wealth can be inferred from the choices of a subject who evaluates gains and losses according to an S-shaped value function. Similarly, one cannot recover a proper subjective probability measure from the preferences of a subject who applies uncertainty weights. The observation that people's preferences vary with the formulation of problems underscores the need for decision aids to help people make more consistent and rational choices. At the same time, these observations put into question the adequacy of the procedures used in decision analysis to elicit utilities and probabilities.

There is a discrepancy between the manner in which the consequences of risky choices are actually perceived and experienced and the manner in which they are commonly interpreted in utility theory. People naturally think of consequences as changes whereas utility theory formulates consequences as states of wealth. If man is constructed in such a way that he is much more sensitive to gains and losses than to absolute wealth, then any attempt to maximize human welfare must recognize this fact. More generally, a normative approach to decision must take into account the nature of man as a pleasure machine. If rational behavior is that which maximizes likely pleasure and minimizes likely pain, then the laws of these experiences should be an essential part of any normative theory of choice. Good

Normative
status of
loss aversion