

INDIVIDUAL DECISION MAKING

5.1 INTRODUCTION

The normal course of our lives can be viewed as a series of decisions in which we choose among the various job offers, insurance policies, and evening entertainments that we encounter. Some decisions are easy and painless; others are difficult and troublesome. What makes decisions difficult is the existence of doubt, conflict, or uncertainty.

The uncertainty may stem from incomplete knowledge about the world, as when the outcomes of the decision depend on some future state or event. The uncertainty may also stem from lack of knowledge about oneself, as when one is not sure which of several possible outcomes would be most satisfying. The difficulty in deciding whether or not to carry an umbrella, for instance, results from uncertainty about the weather. The difficulty in deciding which color car to buy, on the other hand, results from uncertainty as to which color would be most pleasing.

This chapter deals with two types of decision problems: decisions

with incomplete knowledge and decisions with unsure preferences. Both types of decision problems are accompanied by uncertainty. In the former type, however, the uncertainty concerns the future state of the world, whereas in the latter it concerns the decision maker's own state of mind.

These two paradigms do not exhaust the wide variety of decisions faced by individuals, nor does the discussion cover the range of problems that have been tackled successfully by students in the field. The purpose of this chapter is not to describe the current state of the art but rather to introduce the reader to some of the basic concepts of psychological decision theory.

Decision theory is the study of how decisions are or ought to be made. Thus it has two faces: descriptive and normative. Descriptive decision theory attempts to describe and explain how actual choices are made. It is concerned with the study of variables that determine choice behavior in various contexts. As such, it is a proper branch of psychology. Normative decision theory is concerned with optimal rather than actual choices. Its main function is to prescribe which decision should be made, given the goals of the decision maker and the information available to him. Its results have a prescriptive nature. They assert that if an individual wishes to maximize his expected gain, for example, then he should follow a specified course of action. As such, normative decision theory is a purely deductive discipline.

Despite their different natures, descriptive and normative theories are deeply interrelated in most applications. In the first place there are a variety of situations, such as most economic investments, in which people try very hard to behave optimally. Moreover, when faced with obvious errors of judgment or calculation people often admit them and reverse their choices. Hence, there is an inevitable normative component in any adequate descriptive theory that reflects people's desire to do the best they can. Second, in most interesting decision problems, optimality is not easily defined. The main goal of a company, for example, may be to maximize its profit, yet its reputation and the morale of its employees are also important for their own sake and are often incommensurable with monetary considerations. In such instances a descriptive analysis of the goals is a prerequisite for the application of the normative analysis. Thus, although descriptive and normative analyses differ markedly in goals and orientations, most of their interesting applications contain both normative and descriptive aspects.

This chapter deals exclusively with the static analysis of individual decision making. A comprehensive analysis of the various decision problems, with an emphasis on their mathematical structure, can be found in Luce and Raiffa's *Games and Decisions* (1957). An illuminating discussion of the applications of normative decision theory to managerial and economic decisions can be found in Raiffa (1968). Fishburn's book (1964) also discusses the applications of decision theory from a normative viewpoint. For review of the literature, from a psychological standpoint, the reader is referred to

Edwards (1954c, 1961) and to the more recent surveys by Luce and Suppes (1965) and Becker and McClintock (1967). Some of the pertinent literature has been collected in a volume of readings edited by Edwards and Tversky (1967).

5.2 DECISION WITH INCOMPLETE KNOWLEDGE: THEORIES OF RISKY CHOICE

Decisions with incomplete knowledge, where one does not know for sure which state of the world will, in fact, obtain, are typically represented in the form of a payoff matrix. A payoff matrix is simply a rectangular array whose rows, denoted $a_1, \dots, a_i, \dots, a_n$, correspond to the alternatives that are available to the decision maker and whose columns, denoted $s_1, \dots, s_j, \dots, s_m$, correspond to the possible states of nature. The entries of the payoff matrix are the outcomes, or the consequences, resulting from a selection of a given row and column. Thus the o_{ij} entry of the matrix represents the outcome obtained when the individual chooses alternative a_i and nature, so to speak, chooses state s_j .

To illustrate, imagine an individual deciding whether or not to carry an umbrella to work. Naturally, the decision would depend on the relative discomfort of getting wet as opposed to the inconvenience of carrying an umbrella. Another major consideration would be the relative likelihoods of the two relevant states of nature, rain and no rain. The choice situation can thus be summarized by the payoff matrix displayed in Fig. 5.1.

It must be pointed out, however, that the representation of an actual decision problem by a payoff matrix is an abstraction in several respects. In the first place the alternative courses of action and the relevant states of the world cannot always be clearly delineated. This does not mean that it is impossible to reconstruct an appropriate representation in these cases but only that a clearly formulated representation may be difficult to obtain. Moreover, this representation is clearly not unique. There are many ways of

		States of nature	
		s_1 — Rain	s_2 — No rain
Alternatives	a_1 — Carry the umbrella	o_{11} — Stay dry carrying the umbrella	o_{12} — Stay dry carrying the umbrella
	a_2 — Do not carry the umbrella	o_{21} — Get wet without carrying the umbrella	o_{22} — Stay dry without carrying the umbrella

Fig. 5.1

structuring or representing a decision problem in a payoff matrix form. The art of finding the "right way" of structuring a decision problem contributes a great deal to its successful solution. An appropriate representation of a decision problem in a payoff matrix form, therefore, is the result of an adequate formulation rather than a substitute for one. It is nevertheless a useful analytic tool.

The representation of decisions with incomplete knowledge by payoff matrices suggests distinguishing among three states of knowledge or forms of information under which such decisions are made: certainty, ignorance, and risk.

In decision making under certainty, the decision maker knows exactly which outcome results from each choice. After the choice is made, this known outcome obtains with certainty regardless of nature's choice. This occurs when all the consequences in each row of the payoff matrix are identical, or, equivalently, when the obtained state of the world is known for sure. The choices between roast beef and steak for dinner, or between a trip to Florida or to California are examples of decision making under certainty. Because no random or chance process is involved, such choices are also referred to as riskless.

It is important to realize that this notion of certainty depends vitally on the definition, or the level of analysis, of the consequences. If receiving a particular dinner in a restaurant is regarded as a consequence, then the choice between entrees is a riskless one, because one certainly gets the dinner one orders. Yet, if enjoyment of the meal is viewed as the proper consequence, the choice is no longer riskless, because there is a great deal of uncertainty associated with the outcomes that depends on the quality of the restaurant and the competence of the chef and that might not be completely known in advance. Similarly, enjoyment of a trip depends on the weather in California and in Florida, which is not at all certain. What seems to be a decision under certainty relative to one level of analysis, therefore, may turn into a decision under uncertainty on further analysis of the consequences. For a penetrating analysis of this issue, see Savage (1954, pp. 82-91).

If decision making under certainty is one extreme case in which one knows exactly which state of the world will obtain, then decision making under ignorance¹ is the other extreme in which one knows exactly nothing (or better yet, nothing exactly) about which state of the world will obtain. A decision whether to carry an umbrella, in the absence of any information whatsoever about the likelihoods of the weather conditions, is an example of decision under ignorance. Decision problems of this kind are rare because in most situations one has some information about the likelihood of the relevant states of nature. In fact, some of the approaches to decision under ignorance

¹Decisions under ignorance are commonly referred to as decisions under uncertainty. Because the latter term is also used in a broader sense, however, the former term is preferred.

attempt to reduce them to decisions under risk where the relevant information about the states of the world is utilized in the analysis.

In decision making under risk, it is assumed that the individual can evaluate the likelihoods of the various states of nature. More specifically, his beliefs about the likelihoods of the relevant states can be expressed by some (possibly subjective) probability distribution. Risky choices are essentially gambles whose outcomes are determined jointly by the choice of the individual and the result of some specified random process. The decision maker cannot know, therefore, which state of the world will obtain, but he knows (approximately, at least) the probabilities of occurrences of the various states. In some instances, such as in gambles based on unbiased dice, the objective probabilities are known exactly. In other gambles, such as business investments or insurance policies, only rough subjective estimates of the probabilities are available.

The main part of this section is devoted to the major theory of decision making with incomplete knowledge: the expected utility theory. In the next three subsections the basic model is introduced, the axiomatic structure is discussed, and some empirical tests of the theory are described. Finally, alternative theories that are not based on the expectation principle are studied in the subsequent two subsections.

The Expected Utility Principle

The study of decision making under risk dates back to the 18th century when French noblemen asked their court mathematicians to advise them how to gamble. Although utility theory has changed a great deal since those days, its basic problem remains essentially unchanged. To illustrate, imagine being offered a simple two-outcome gamble at a fixed cost c , where you may win either $\$x$ or nothing depending on a toss of a fair coin. For which values of x would you accept the bet? If you reject the offer, no money changes hands. If you accept it, you receive either $x - c$ (your win minus the price you paid) if a head occurs, or $-c$ (the price you paid) if a tail occurs. The situation is summarized by the payoff matrix displayed in Fig. 5.2. Given a fixed price,

		States	
		Head	Tail
Alternatives	Accept	$x - c$	$-c$
	Reject	0	0

Fig. 5.2

the problem is to formulate a decision rule that would determine the values of x for which the gamble would (or should) be accepted.

One simple decision rule is to compute the expected value of each alternative and to choose the one with the higher expected value. The expected value of an alternative or a gamble is the sum of its outcomes, each weighted by its probability of occurrence. More formally, the expected value of a gamble, with outcomes $x_1, \dots, x_n$ obtained with probabilities $p_1, \dots, p_n$, respectively, equals $\sum_{i=1}^n p_i x_i$. (For further discussion of the expected value notion, see the appendix.)

The expected value of rejecting the offered gamble, denoted $EV(R)$, is clearly zero because $EV(R) = (\frac{1}{2})0 + (\frac{1}{2})0 = 0$. The expected value of accepting the gamble, denote $EV(A)$, is given by the equation $EV(A) = \frac{1}{2}(x-c) + \frac{1}{2}(-c) = (x/2) - c$. Thus $EV(A) > EV(R)$ whenever $x/2 > c$. According to the expected value principle, therefore, one should accept the offered bet if and only if the cost (c) is less than one half the prize (x). The expected value of a bet can be interpreted as the average outcome resulting from playing it an indefinite (or very large) number of times. Thus, if the gamble is played a very large number of times, one can practically assure himself of an average gain of $(x/2) - c$. We speak of gambles as being favorable, unfavorable, or fair, according to whether their expected values are positive, negative, or zero.

If one follows the expected value rule, therefore, one should accept all favorable bets and reject all unfavorable bets. This is not, however, what people actually do, nor is it what they feel they ought to do. In the first place people gamble by accepting bets whose expected values are negative. For otherwise gambling casinos would be out of business and they are not. Second, people buy insurance and in so doing they are paying the insurance company in order to get rid of an unfavorable gamble. Many people pay a monthly fire insurance premium, for example, to avoid the relatively small chance of losing the value of their property as a consequence of fire. In most cases, however, the price paid for the insurance is higher than the expected value of the undesirable gamble. Otherwise, insurance companies would be out of business and they are not. Moreover, people feel that buying insurance (and in some cases even gambling) is a rational form of behavior that can be defended on normative grounds. Hence, the expected value model is inadequate on both descriptive and normative accounts. In particular, it implies that one should be indifferent with respect to all fair bets. Hence, one should not object to tossing a fair coin to decide whether he wins or loses \$1,000, for example. Most reasonable people would probably reject the gamble, however, because the potential gain of \$1,000 does not quite compensate for the potential loss of \$1,000.

Similar difficulties, arising from a gambling puzzle known as the St. Petersburg paradox, led Daniel Bernoulli, as early as 1738, to formulate the expected utility principle. The expected utility of a gamble, with out-

comes $x_1, \dots, x_n$ obtained with probabilities $p_1, \dots, p_n$, respectively, equals $\sum_{i=1}^n p_i u(x_i)$, where $u(x_i)$ is the utility of the i th outcome. The expected utility principle asserts that the gamble with the highest expected utility is to be chosen. The decision rule proposed by Bernoulli, therefore, is also based on the expectation principle, but it replaces the objective scale of value by a subjective scale of utility. The introduction of a subjective scale results in a more general and plausible model that seems to resolve the difficulties arising from the expected value model. (More than 100 years later, the Bernoullian notion of a subjective scale became the cornerstone of quantitative psychophysics founded by G. T. Fechner.)

The advantages of expected utility theory over expected value theory are numerous. In the first place it allows individuals to have different utilities for money and hence different preferences among gambles. This is essential for any interesting descriptive or normative theory, as individuals' preferences are not (and for this matter need not be) independent of their attitude toward risk. Second, it has been assumed that the more money one has, the less he values each additional increment, or, equivalently, that the utility of any additional dollar diminishes with an increase in capital. These considerations, as well as the insurance phenomena, have led to the decreasing marginal utility hypothesis according to which the utility function is concave, or negatively accelerated. (The similarity to Weber's law is not accidental. In fact, the logarithmic utility function proposed by Bernoulli was the one put forth by Fechner as the form of the general psychophysical law.)

To demonstrate the explanatory power of the expected utility principle, let us suppose, for illustrative purposes, that the utility for money has the following form:

$$u(x) = \begin{cases} \sqrt{x} & \text{if } x \geq 0 \\ -(x^2) & \text{if } x \leq 0. \end{cases}$$

Thus the subjective value of gain grows as the square root of the actual value, whereas the (negative) subjective value of loss grows as the square of the actual value. The graph of the function is portrayed in Fig. 5.3. Although the choice of this particular utility function is arbitrary, it was proposed as early as the 18th century as a prototype for Everyman's utility function, and Stevens (1959) has defended it on the basis of some experimental evidence. If we examine the proposed utility function, it is easy to see that it is strictly concave for any two points that are not in the $[-1, 1]$ interval. That is, a straight line connecting the utilities of any two points (outside that interval) lies entirely below the curve. Stated algebraically,

$$pu(x_1) + (1-p)u(x_2) < u[px_1 + (1-p)x_2]$$

for all x_1, x_2 outside the $[-1, 1]$ interval and for any $0 < p < 1$.




Fig. 5.3 Geometric illustration of the proposed utility function.

The left side of the above inequality, however, is the expected utility of a gamble where one receives x_1 with probability p and x_2 with probability $1-p$. Its right side is the utility of the expected (monetary) value of the same gamble. The above inequality asserts that, given a concave utility function and the expectation principle, one should always prefer receiving a given sum of money over taking any gamble whose expected value equals that sum. Thus the decreasing marginal utility hypothesis, embodied in the concavity of the utility function, can account for the common risk-averse tendency to reject fair bets and to buy insurance policies.

For concreteness, consider an individual who wishes to insure his \$3,000 car against theft. Suppose insurance costs \$30 a year and that there is 1 chance in 1,000 that the car would be stolen during this time. If he insures his car, he will lose \$30 a year, irrespective of whether the car is stolen, as the insurance company will reimburse him in case of theft. If he does not insure his car, he will lose \$3,000 if his car is stolen and nothing otherwise. This situation is summarized in the payoff matrix displayed in Fig. 5.4. It is easy to see that the expected value of buying insurance, denoted

$EV(B)$, is less than the expected value of not buying insurance, denoted $EV(NB)$, because $EV(B) = -30$ whereas $EV(NB) = .001(-3,000) + .999(0) = -3$. If monetary values, however, are replaced by utilities according to the proposed utility function, the following computation shows that the expected utility of buying insurance, denoted $EU(B)$, is greater than the expected utility of not buying insurance, denoted $EU(NB)$. For: $EU(B) = u(-30) = -(30^2) = -900$, whereas $EU(NB) = .001 u(-3,000) + .999 u(0) = -.001(3,000^2) + .999(0) = -9,000$. Thus, although $EV(NB)$ exceeds $EV(B)$, $EU(B)$ exceeds $EU(NB)$. The expected utility principle, in conjunction with a concave utility function, therefore, can account for the purchase of insurance despite its smaller expected value.

The utility function suggested can explain not only the rejection of favorable gambles but also the acceptance of some unfavorable ones. Consider, for example, a simple dice game where you pay \$.50 as a participation fee and you have one chance out of six to win \$2.75. It is not difficult to see that this gamble, denoted G , has a negative expected value, yet its expected utility is positive. To verify this assertion note that

$$EV(G) = \frac{1}{6}(2.25) + \frac{5}{6}(-.50) = \frac{9}{24} - \frac{10}{24} = -\frac{1}{24} < 0,$$

whereas

$$EU(G) = \frac{1}{6}u(2.25) + \frac{5}{6}u(-.50) = \frac{1}{6}(1.50) + \frac{5}{6}(-.25) = \frac{6}{24} - \frac{5}{24} = \frac{1}{24} > 0.$$

Hence, by an appropriate choice of a utility function, the purchase of insurance, as well as gambling behavior, can be rationalized. Indeed, a similar rationalization has been proposed by Friedman and Savage (1948).

It is important to realize that we have not shown that people insure their property or gamble on horses in order to maximize some utility function.

		States of the world	
		Car stolen ($p = .001$)	Car is not stolen ($p = .999$)
Alternatives	Buy insurance	-30	-30
	Do not buy insurance	-3000	0

Fig. 5.4

All we have shown is that although these phenomena are incompatible with the expected value principle, they can be accounted for, if an appropriate utility function is introduced.

The utility analysis is not limited to monetary outcomes and it can also be applied to gambles whose consequences are nonmonetary such as the enjoyment of a play, a loss in an election, or a satisfaction from an accomplishment. In fact, the general notion that people act to maximize their utility has formed the basis for the economist's concept of Economic Man, which originated with Jeremy Bentham and James Mill. The expected utility principle combines this general utility maximization notion with the assumption that the utility of a gamble equals the expected utility of its outcomes. If the available gambles, therefore, are to be played over and over again, then the gamble with the highest expected utility would yield the highest utility in the long run. In spite of the intuitive appeal of this rationale, the application of the expected utility principle to essentially unique choice situations requires an independent justification. Modern utility theory provides such a justification in the form of an axiomatic foundation of the expected utility principle.

Modern Utility Theory

Modern utility theory was first developed by von Neumann and Morgenstern as an appendix to their famous *Theory of Games and Economic Behavior* (1947). The theory consists of a set of axioms about preferences among gambles. The basic result of the theory is summarized by a theorem stating that if an individual's preferences satisfy the specified axioms then his behavior can be described, or rationalized, as the maximization of his expected utility. Because the axioms can be regarded as maxims of rational behavior, they provide a normative justification for the expected utility principle. Although several axiomatizations of utility theory have been developed in the last two decades, we present the original formulation of von Neumann and Morgenstern with a few inessential modifications.

The axioms are formulated in terms of a preference-or-indifference relation, denoted $\succsim$, defined on a set of outcomes, denoted A . Later, this set is enriched to include gambles, or probability mixtures, of the form (x, p, y) , where outcome x is obtained with probability p and outcome y is obtained with probability $1-p$. Given the primitives $\succsim$ and A , the following axioms are assumed to hold for all outcomes x, y, z in A and for all probabilities p, q that are different from zero or 1.

A1

(x, p, y) is in A .

A2

$\succsim$ is a weak ordering of A , where $>$ denotes strict preference and $\sim$ denotes indifference.

A3

 $[(x, p, y), q, y] \sim (x, pq, y).$

A4

If $x \sim y$, then $(x, p, z) \sim (y, p, z).$

A5

If $x \succ y$, then $x \succ (x, p, y) \succ y.$

A6

If $x \succ y \succ z$, then there exists a probability p such that $y \sim (x, p, z).$

It is important to realize that utility theory was developed as a prescriptive theory, justified on the basis of normative considerations alone. The close interrelationships between normative and descriptive considerations, however, suggest that utility theory may also be used as a psychological theory of decision making under risk. In discussing the axioms, therefore, we examine them from both normative and descriptive viewpoints.

The first axiom is what is technically called a closure property. It asserts that if x and y are available alternatives, so are all the gambles of the form (x, p, y) that can be formed with x and y as outcomes. Because gambles are defined in terms of their outcomes and their probabilities, it is assumed implicitly that $(x, p, y) = (y, 1-p, x)$. The second axiom requires the observed preference-or-indifference relation to be reflexive, connective, and transitive. That is, for all gambles x, y, z the following conditions are satisfied:

1. Reflexivity: $x \succeq x$.
2. Connectivity: Either $x \succeq y$ or $y \succeq x$ or both.
3. Transitivity: $x \succeq y$ and $y \succeq z$ imply $x \succeq z$.

A detailed discussion of these properties is given in the appendix. Reflexivity is empirically trivial because any gamble is obviously equivalent to itself. Connectivity is also innocuous because any two gambles can be compared with respect to preference. Although transitivity might be violated in certain contexts, it is, nevertheless, a very compelling principle. It is certainly imperative on normative grounds, and it is a plausible descriptive hypothesis.

Axiom 3 is a reducibility condition. It requires that the gamble (x, pq, y) , in which x is obtained with probability pq and y with probability $1-pq$, be equivalent, with respect to the preference order, to the compound gamble $[(x, p, y), q, y]$, in which (x, p, y) is obtained with probability q and y with probability $1-q$. Compound gambles differ from simple ones in that their outcomes are themselves gambles rather than pure outcomes, such as monetary values that can be won or lost. Note that the final outcomes of both the simple and the compound gambles are x and y . Furthermore, the probabilities with which x and y are obtained are the same in both gambles.

This follows from the fact that the probability of obtaining x in the compound gambles is the probability of obtaining (x, p, y) in the first stage (i.e., q) multiplied by the probability of obtaining x in the second stage (i.e., p), which equals pq . (Assuming that the probabilities of the two stages are independent.) Consequently, the probability of obtaining y in the compound gamble is $1 - pq$, and hence the two gambles eventually yield the same outcomes with the same probabilities. Thus axiom 3 asserts, in effect, that the preferences depend only on the final outcomes and their probabilities and not on the process by which they are obtained. Normatively, it makes perfect sense to suppose that the choices are invariant with respect to rearrangements of the gambling procedure, as long as the outcomes and their probabilities remain unchanged. If, on the other hand, people have aversions or attractions associated with the actual gambling process, they may not be indifferent between the compound and the corresponding simple gamble.

In general, the psychological interpretation of the axioms raises intricate problems. If all gambles are presented to the individual in terms of their final outcomes and their associated probabilities, then A3 is trivially substantiated. If the gambles are displayed in terms of their immediate rather than final outcomes, the relationship between the compound and the simple gambles may very well escape the subject and A3 can be easily violated. Moreover, two gambles that are formally identical may elicit different responses from the subject because of differences in display, context, and other situational variables. An individual may reject a bet offered to him by a friend, for example, though he may gladly accept a formally identical bet in a gambling casino. The interpretation of utility theory as a behavioral model, therefore, has to be supplemented by a psychological theory that accounts for situational variables that affect risky choices. In the absence of such a theory, the applicability of utility theory is limited to specific contexts and its explanatory power is substantially reduced.

The fourth axiom is a substitutability condition. It states that if x and y are equivalent, then they are substitutable for each other in any gamble, in the sense that $(x, p, z) \sim (y, p, z)$ for any p and z . This axiom excludes the possibility of interacting outcomes in the sense that the probability mixture of x and z can be preferred to the probability mixture of y and z , although x and y , taken alone, are equivalent.

The fifth axiom asserts that if x is preferred to y , then it must be preferred to any probability mixture of x and y , which, in turn, must be preferred to y . It is certainly not objectionable for monetary outcomes. An alleged counterexample to this axiom is Russian roulette; players of this game apparently prefer a probability mixture of living and dying over either one of them alone. For otherwise, one can easily either stay alive or kill oneself, without ever playing the game. A more careful analysis reveals, however, that this situation, perverse as it may be, is not incompatible with axiom 5. The actual outcomes involved in playing Russian roulette are (1) staying alive after playing the game, (2) staying alive without playing the

game, and (3) dying in the course of playing the game. In choosing to play Russian roulette, therefore, one prefers a probability mixture of (1) and (3) over (2), rather than a probability mixture of (1) and (2) over both (1) and (2) as the alleged counterexample suggests. The former preference, however, is not incompatible with A5. This argument also demonstrates the errors that can result from an incomplete analysis of the choice situation or from an inappropriate identification of the outcomes. A careful analysis of the payoff matrix is a prerequisite to any serious application of the theory.

The last axiom embodies a continuity or a solvability property. It asserts that if y is between x and z in the preference order (i.e., $x \succ y \succ z$) then there exists a probability p such that the gamble (x, p, z) is equivalent to y . This axiom excludes the possibility that one alternative is "infinitely better" than another one, in the sense that any probability mixture involving the former is preferable to the latter. For a proposed counterexample, let x be the prospect of receiving one dime, let y be the prospect of receiving one nickel, and let z be the prospect of being shot at sunrise. Because $x \succ y \succ z$, A6 requires that there exists a probability p , such that the gamble (x, p, z) in which one receives a dime with probability p or is shot at sunrise with probability $1-p$ is equivalent to receiving a nickel for sure. Some people find this result unacceptable. Its counterintuitive flavor, however, stems from an inability to comprehend very small probabilities. Thus in the abstract, people feel that there is no positive probability with which they are willing to risk their life for an extra nickel, yet in actual practice a person would cross a street to buy some product for a nickel less, although by doing so he certainly increases the probability of being killed. Hence, the initial intuitions that tend to reject axiom 6 seem inconsistent with everyday behavior.

Axiom 6 captures the relationships between probabilities and values and the form in which they compensate for each other. This form becomes transparent in the following theorem of von Neumann and Morgenstern.

THEOREM 5.1

If axioms A1–A6 are satisfied, then there exists a real-valued utility function u defined on A , such that

1. $x \succeq y$ if and only if $u(x) \geq u(y)$.
2. $u(x, p, y) = pu(x) + (1-p)u(y)$.

Furthermore, u is an interval scale, that is, if v is any other function satisfying 1 and 2, then there exists numbers b , and $a > 0$ such that $v(x) = au(x) + b$.

Thus the theorem guarantees that whenever the axioms hold, there exists a utility function that (1) preserves the preference order and (2) satisfies

the expectation principle as the utility of a gamble equals the expected utility of its outcomes. Moreover, this utility scale is uniquely determined except for an origin and a unit of measurement. The proof of this theorem is quite difficult and therefore omitted. A simplified version of the result can be found in Luce and Raiffa (1957, pp. 23-31).

The main contribution of modern utility theory to the analysis of decision making under risk is in providing sound justification for the Bernoullian expected utility principle. This justification does not depend on long run considerations, hence it is applicable to unique choice situations. Furthermore, the axiomatic structure highlights those aspects of the theory, which are critical for both normative and descriptive applications.

Some people, however, remained unconvinced by the axioms. One of them, Allais (1953), argued that the theory of utility is too restrictive and hence inadequate. To substantiate the claim he constructed the following example of two hypothetical decision situations each involving two gambles, expressed in units of a million dollars.

Situation 1. Choose between

Gamble 1. $\frac{1}{2}$ with probability 1;

Gamble 2. $2\frac{1}{2}$ with probability .10,
 $\frac{1}{2}$ with probability .89,
 0 with probability .01.

Situation 2. Choose between

Gamble 3. $\frac{1}{2}$ with probability .11,
 0 with probability .89;

Gamble 4. $2\frac{1}{2}$ with probability .10,
 0 with probability .90.

Most people prefer gamble 1 to gamble 2, presumably because the small probability of missing the chance of a lifetime to become rich seems very unattractive. At the same time most people prefer gamble 4 to gamble 3, presumably because the large difference between the payoffs dominates the small difference between the chances of winning. However, this seemingly innocent pair of preferences is incompatible with utility theory. To demonstrate this, note that the first preference implies that

$$u(\text{gamble 1}) > u(\text{gamble 2})$$

and hence

$$u(\frac{1}{2}) > .10u(2\frac{1}{2}) + .89u(\frac{1}{2}) + .01u(0)$$

so

$$.11u(\frac{1}{2}) > .10u(2\frac{1}{2}) + .01u(0).$$

Similarly, the second preference implies that

$$u(\text{gamble 4}) > u(\text{gamble 3})$$

and hence

$$.10u(2\frac{1}{2}) + .90u(0) > .11u(\frac{1}{2}) + .89u(0)$$

so

$$.10u(2\frac{1}{2}) + .01u(0) > .11u(\frac{1}{2}),$$

which is clearly inconsistent with the inequality derived from the first preference.

How do people react to such inconsistencies between their intuitions and the theory? Some people, who feel committed to their preferences, would undoubtedly reject the expected utility theory. Or, to use Samuelson's phrase, they prefer to "satisfy their preferences and let the axioms satisfy themselves." Others, who feel committed to the theory, tend to reexamine their preferences in the light of the axioms and to revise their initial choices accordingly. An illuminating introspective discussion of Allais's example, from this viewpoint, has been offered by Savage (1954).

Savage admits that, when first presented with Allais's example, he preferred gamble 1 to gamble 2 and gamble 4 to gamble 3 and that he still feels an intuitive attraction to these choices. Yet, he has adopted another way of looking at the problem. One way in which the gambles can be realized is by a lottery with 100 numbered tickets, one of which is drawn at random to determine the outcome according to the payoff matrix presented in Fig. 5.5.

		Ticket number		
		1	2-11	12-100
Situation 1	Gamble 1	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
	Gamble 2	0	$2\frac{1}{2}$	$\frac{1}{2}$
Situation 2	Gamble 3	$\frac{1}{2}$	$\frac{1}{2}$	0
	Gamble 4	0	$2\frac{1}{2}$	0

Fig. 5.5

An examination of the payoff matrix reveals that if one of the tickets numbered 12–100 is drawn, it does not matter, in either situation, which gamble is chosen. Hence, one should consider only the possibility that one of the tickets numbered 1–11 will be drawn, in which case the two choice situations are identical. Limiting our attention to tickets 1–11, the problem in both situations is whether a 10:1 chance to win $2\frac{1}{2}$ million is preferred to $\frac{1}{2}$ a million with certainty. If one prefers gamble 1 to gamble 2, therefore, he should also prefer gamble 3 to gamble 4, if he wishes to be consistent. In concluding his discussion Savage (1954) writes:

It seems to me that in reversing my preference between gamble 3 and 4 I have corrected an error. There is, of course, an important sense in which preferences, being entirely subjective, cannot be in error; but in a different, more subtle sense they can be. Let me illustrate by a simple example containing no reference to uncertainty. A man buying a car for \$2,134.56 is tempted to order it with a radio installed, which will bring the total price to \$2,228.41, feeling that the difference is trifling. But, when he reflects that, if he already had the car, he certainly would not spend \$93.85 for a radio for it, he realizes that he has made an error (p. 103).

The preceding analysis exemplifies how utility theory can be applied to situations where it seems incompatible with one's intuitions. Here, the theory is viewed as a guideline or a corrective tool for a rational man rather than as an accurate model of his nonreflective choices.

Indeed, MacCrimmon (1967) has presented problems of the kind devised by Allais to upper-middle-level executives in order to study both the descriptive validity and the normative appeal of utility theory. He concluded that his subjects tended to regard most of the deviations from the theory as mistakes and were ready to correct them, if given the opportunity.

Another criticism of expected utility theory revolves around the concept of probability. The theory is formulated in terms of gambles whose numerical probabilities are assumed to be known in advance. Such knowledge, however, is missing in most applications. Can the theory be generalized to situations where no a priori knowledge of (numerical) probabilities is available? The answer is positive, provided some consistency requirements are fulfilled.

Savage (1954) has developed an axiomatic theory leading to simultaneous measurement of utility and subjective probability. We do not wish to present this theory here but we do wish to show how the relation "more probable than" between events can be defined in terms of preferences between gambles. Consider the choice situation, displayed in Fig. 5.6, between two gambles G_1 and G_2 whose outcomes x , y , and z depend on whether E , F , or neither event (denoted $\overline{E \cup F}$) occurs. If neither E nor F occurs, there is no reason to prefer one gamble to another. Thus, assuming x is preferred to y , the only apparent reason for preferring G_1 to G_2 is the fact that E seems more probable than F . Stated formally, E is said to be more probable than F if

		Events		
		E	F	$\overline{E \cup F}$
Gambles	G_1	x	y	z
	G_2	y	x	z

Fig. 5.6

and only if G_1 is preferred to G_2 . Clearly, some assumptions are needed to guarantee that the above relation is well defined. Indeed, Savage has further shown that his axioms are sufficient to establish the existence of a uniquely additive subjective probability function s and an interval scale utility function u such that

1. $x \succsim y$ if and only if $u(x) \geq u(y)$ and
2. $u(x, E, y) = s(E)u(x) + [1 - s(E)]u(y)$,

where (x, E, y) denotes the gamble where x is obtained if E occurs and y otherwise. Because the probabilities, as well as the outcomes, are viewed as subjective, Savage's theory is called the subjective expected utility model, or the SEU model, for short.

The historical development of expectation models reveals a clear trend toward more general and more subjective decision models. In the expected value model, both probability and value are defined objectively. In the expected utility theory, objective values are replaced by utilities, and in the subjective expected utility model, objective probabilities, in addition, are replaced by subjective ones.

The introduction of subjective quantities generalizes the theory in two major respects. They reflect individual differences in the evaluation of outcomes (utilities) and events (subjective probability). At the same time they do not have to be specified in advance, because they can be derived from choices. The common property shared by all objective and subjective expectation models is that the subjective value of a gamble is a composite function of two basic independent factors: the desirability of its outcomes and the likelihood of its events.

Tests of Utility Theory

In recent years there have been several attempts to test the descriptive validity of the subjective expected utility theory. Although the theory may be applicable to many real world problems (such as the selection of a military strategy, a financial investment, or a job offer), it is very difficult to test the theory in

these contexts mostly because of the large number of unknown parameters (utilities and subjective probabilities) and the lack of the appropriate controls. Consequently, students of choice behavior have devised simple experimental paradigms, based on choice between gambles, that enable them to test utility theory under controlled experimental conditions. The price paid for these controls is that the scope of the investigation is necessarily limited to those situations that can be studied in the laboratory. The extrapolation from laboratory experiments to real world behavior is always a risky venture. With this problem in mind we turn now to the discussion of some of the methods used to test the theory and to derive utility and subjective probability scales. More specifically, three such methods are discussed. Reviews of the experimental literature and discussions of the related methodological issues can be found in some of the articles cited in the introduction.

The first experimental study of expected utility theory was conducted at Harvard by Mosteller and Nogee (1951). Their subjects were presented with gambles, constructed from possible hands of poker dice, that they could accept or reject. If a subject rejected the gamble, no money changed hands; if he accepted it, he won x cents if he beat the hand and lost a nickel if he did not. The situation is described in the payoff matrix of Fig. 5.7, where the values are expressed in pennies.

The subjects were shown how to calculate the probabilities of the relevant events and were also given a table with the true odds for all the poker dice hands used in the study. By varying the payoffs, the experimenter finds the value of x for which the subject is indifferent between the two alternatives. After that value is found, its utility may be computed in the following fashion. Because the subject is assumed to be indifferent between the two alternatives, $u(\text{Accept}) = u(\text{Reject})$ and hence, according to expected utility theory,

$$u(0) = pu(x) + (1-p)u(-5),$$

where p denotes the (known) probability of beating the hand. Because

		States	
		Win the hand	Lose the hand
Alternatives	Accept	x	-5
	Reject	0	0

Fig. 5.7




Fig. 5.8 Proportion of times an offered gamble was accepted by one of the subjects in the Mosteller and Nogee experiment.

utility is measured on an interval scale, we can set $u(0) = 0$ and $u(-5) = -1$ to obtain $u(x) = (1-p)/p$.

By varying odds and payoffs, the utility function can be constructed, provided the indifference point can be found. To determine this point, Mosteller and Nogee varied the payoff (i.e., x) systematically and replicated the choices many times over a 4-month period. The proportion of times the subject elected to play each of the gambles was plotted against the corresponding payoff. The point at which the subject accepted the gamble 50 percent of the times was estimated from the graph and taken as the indifference point. An example of such a graph is given in Fig. 5.8. The reader may recall that essentially the same procedure is employed in psychophysics to determine sensation units, or jnd's. Although the proportion of acceptance increased with the payoff, subjects revealed a considerable degree of inconsistency. That is, many gambles were accepted on some replications but rejected on others. This inconsistency is particularly apparent when the graph of Fig. 5.8 is contrasted with the step function (i.e., a function that jumps from 0 to 100 per cent at one point) predicted from any model that does not allow inconsistency.

Fourteen subjects completed the experiment, of whom 9 were Harvard undergraduates and 5 were members of the Massachusetts National Guard. The utility scales derived from the study, in the range between 5 cents and \$5.50, were concave (negatively accelerated) for the students and convex (positively accelerated) for the guardsmen. These utility functions were then used to predict, with moderate success, additional choices made by the same subjects involving more complicated gambles.

Although the Mosteller–Nogee study did not provide a very direct test of utility theory, it lent some indirect support to the expected utility principle, and it demonstrated that it is feasible to measure utility experimentally and to predict future behavior on the basis of these measurements.

One criticism of the experiment of Mosteller and Nogee is that their results are also interpretable in terms of subjective probabilities instead of (or in addition to) utilities. The fact that the objective probabilities were known to the subjects does not imply that the subjects actually used these values in their decisions. Davidson, Suppes, and Siegel (1957) conducted a series of studies at Stanford University designed to meet this objection and to measure both utility and subjective probability. Their approach followed an earlier development of Ramsey (1931), which was based on the idea of finding an event whose subjective probability equals one half. Assuming the subjective expected utility model, this event can be used to construct a utility scale, which can then be used to measure the subjective probabilities of other events.

If an event E with subjective probability of one half exists, then its complement $\bar{E}$ also has a subjective probability of one half because, according to the theory, the (subjective) probabilities of complementary events must sum to unity. Consequently, if one regards E and $\bar{E}$ as equiprobable, he should be indifferent between the gamble G_1 , where he receives x if E occurs and y otherwise, and the gamble G_2 , where he receives y if E occurs and x otherwise. Conversely, if one is indifferent between G_1 and G_2 , for all x and y , then his subjective probability for E is equal to that of $\bar{E}$, which must equal one half.

Davidson, Suppes, and Siegel tried several events, but they were forced to reject coin flips and penny matching because their subjects showed systematic preferences for heads over tails, for example. They finally used a six-sided die, with the nonsense syllable ZEJ printed on three of its sides and ZOJ on the other three. This die came reasonably close to satisfying the subjective equiprobability criterion.

After the desired event E is identified, various gambles can be constructed by varying the payoffs, as indicated in Fig. 5.9.

		States	
		E	$\bar{E}$
Alternatives	G_1	x	y
	G_2	z	w

Fig. 5.9




Fig. 5.10 Bounds for the utility curve of subject 9 in the Davidson, Suppes, and Siegel experiment.

If the subject prefers G_1 to G_2 , for example, then, according to the subjective expected utility model, there exist utility (u) and subjective probability (s) functions such that

$$s(E)u(x) + s(\bar{E})u(y) \geq s(E)u(z) + s(\bar{E})u(w).$$

But because E was selected so that $s(E) = s(\bar{E})$, the subjective probabilities can be cancelled, and we obtain the inequality

$$u(x) + u(y) \geq u(z) + u(w).$$

By an appropriate selection of payoffs, one can obtain upper and lower bounds on the utility function from the inequalities derived from the choices, provided they are consistent with the model.

Out of the 19 subjects who took part in the study, 15 satisfied the theory and the bounds on their utility scales were determined. Because the bounds were, in general, sufficiently close, they permitted an approximate determination of the shapes of the utility functions, which were nonlinear in most cases. An example of the bounds obtained for one of the subjects is shown in Fig. 5.10.

Suppose a new event F is now introduced into the payoff matrix shown in Fig. 5.11. If the payoffs can be selected so that subjects are indifferent between G_1 and G_2 , then according to our theory

$$s(F)u(x) + s(\bar{F})u(y) = s(F)u(z) + s(\bar{F})u(w).$$

		States	
		F	$\bar{F}$
Alternatives	G_1	x	y
	G_2	z	w

Fig. 5.11

Because $s(F) + s(\bar{F}) = 1$, we can solve for the subjective probability of F , which is given by

$$s(F) = \frac{u(w) - u(y)}{u(x) - u(z) + u(w) - u(y)}.$$

Having found the utilities, therefore, by using the special event E , we can now solve for the subjective probabilities of other events as well. Although Davidson, Suppes, and Siegel did not construct subjective probability scales, they applied their procedure to measure the subjective probability of a single event whose objective probability equaled one fourth. The majority of the subjects for whom subjective probability could be calculated tended to underestimate the objective probability.

A more recent attempt to test utility theory and to measure utility and subjective probability simultaneously was conducted by Tversky (1967). Consider a set of gambles of the form (x, p) in which one wins (or loses) $\$x$ if p occurs and receives nothing if p does not occur. Let $M(x, p)$ be the bid, or the minimal selling price of the gamble (x, p) . That is, $M(x, p)$ is the smallest amount of money for which one would sell his right to play the gamble. According to utility theory, therefore,

$$u[M(x, p)] = u(x)s(p) + u(0)s(\bar{p}),$$

where u and s are the utility and the subjective probability functions and $\bar{p}$ is the complement of p . Since utility is measured on an interval scale on which the zero point is arbitrary, we can set $u(0) = 0$, and after taking logarithms, we obtain,

$$\log u[M(x, p)] = \log u(x) + \log s(p).$$

If both x and p are varied, then the resulting bidding matrix should be additive in the conjoint measurement sense (see Sec. 2.5). That is, the bids can

be rescaled monotonically such that the rescaled bid for the gamble (x, p) equals the sum of the scale values of x and p . The additivity of the bidding matrix provides a method for testing the assumption of independence between utility and subjective probability that lies at the heart of utility theory.

To facilitate the measurement process, an additional assumption about the utility function was explored. Suppose utility is a power function of the form $u(x) = x^\theta$ for some $\theta > 0$. Substituting this form in the last equation yields

$$\log [M(x, p)^\theta] = \log x^\theta + \log s(p);$$

thus

$$\log M(x, p) = \log x + \frac{1}{\theta} \log s(p).$$

Hence, if utility is a power function, then the logarithms of the bids should equal the sum of the functions of the gamble's components. Conversely, it is possible to show that if the utility function is monotonic and if the above equation holds for all x and p , then utility must be a power function.

This prediction was tested in an experiment using 11 male inmates in a state prison in Michigan. They were presented with gambles of the form (x, p) , where the outcomes varied from a gain of \$1.35 to a loss of that amount, and the events varied in objective probability from .1 to .9. Each subject stated his bid, or his minimum selling price, for every gamble. Each gamble was presented once in each of three sessions. Several gambles, chosen randomly, were played at the end of each session.

Analyses of variance applied to the logarithms of the subjects' bids supported additivity in 41 out of the 44 bidding matrices. The data, therefore, can be accounted for by power utility functions, to the accuracy allowed by their own variability.

The derivation of the utility and the subjective probability scales was based on the observation that for any pair of complementary events p and $\bar{p}$

$$M(x, p)^\theta = x^\theta s(p) \quad \text{and} \quad M(x, \bar{p})^\theta = x^\theta [1 - s(p)].$$

Taking logarithms and solving for θ yields

$$\theta = \frac{\log s(p)}{\log M(x, p) - \log x} = \frac{\log [1 - s(p)]}{\log M(x, \bar{p}) - \log x}.$$

Because the denominators can be calculated from the data, estimates of both θ and $s(p)$ can be obtained. The resulting utility and subjective probability functions for a typical subject are shown in Figs. 5.12 and 5.13.




Fig. 5.12 Utility function for a typical subject.




Fig. 5.13 Subjective probability function for a typical subject.

The utility functions tended to be linear for gains and convex for losses. Apparently, the subjects were ready to sell the positive outcome gambles for their expected (monetary) value, but they were willing to pay more than the expected value to get rid of the negative outcome gambles. Such a preference would lead to the purchase of insurance despite an inferior expected monetary value.

Subjective probability scales were linear functions of objective probability for some subjects; most of the subjects, however, overestimated low probabilities and underestimated high ones. This result has also been found in many other studies.

The results obtained in this study have both methodological and substantive implications. From a methodological viewpoint it was shown that the application of additivity analysis to a specified class of gambling experiments yields a simple simultaneous construction of utility and subjective probability scales. From a substantive viewpoint the data support the

multiplicative relation between utility and subjective probability embodied in the subjective expected utility principle. Moreover, the results indicate that (within the range of payoffs investigated) the utility of money can be described as a power function of money with different exponents for positive and negative outcomes.

All in all, the three studies reported in this subsection show that the subjective expected utility model provides an adequate account of most of the data obtained in simple gambling experiments. Such findings are of primary interest to the experimental psychologist who is concerned with the question of how people combine various subjective dimensions (e.g., utility and subjective probability) in order to evaluate, or choose among, alternatives. An economist interested in consumer behavior and insurance purchase or a political scientist interested in presidential decisions, for instance, might feel that although utility theory may be applicable to their field of study the experiments thus far performed are too simple and contrived to provide relevant information. It is our hope, nevertheless, that the mathematical and the experimental methods developed in these studies can be extended to the investigation of complex decisions in real world environments.

Special Preference Models

The subjective expected utility model is not the only theory put forth to account for decisions under risk. More than half a century ago, an economist, Irving Fisher (1906), proposed that people base their choices among gambles on the variances of the gambles as well as on the expectations.

The variance of a gamble G , denoted $V(G)$, is given by the formula $V(G) = E(G^2) - E(G)^2$, where G^2 is the gamble obtained from G by squaring its outcomes. In particular, if G has two outcomes x and y obtained with probabilities p and $1 - p$, respectively, then its variance equals $p(1 - p)(x - y)^2$. The variance is the most common measure of the dispersion of the outcomes, and it seems conceivable that different attitudes toward risk be reflected by preferences for different amounts of variance. Although some types of variance preferences can be accommodated by an appropriate selection of utility functions, Allais (1953) has further argued that, even with the introduction of such utility functions, the expectation principle alone is insufficient to explain risky choices. Allais suggested the investigation of decision models that depend not only on expectation but on other attributes such as variance and skewness as well.

Alternatives to utility theory can be developed by first characterizing each gamble in terms of its objectively defined attributes and then formulating decision rules based on these attributes. Thus, instead of choosing among gambles according to their subjective expected utilities, one can choose among them according to some linear combination of their expectations and variances, for example. The major difficulty with this approach lies in the need to specify in advance the relevant attributes, or dimensions, whose