APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville
Massachusetts Institute of Technology

McGraw-Hill Publishing Company
New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

CHAPTER

18

VALUE AND UTILITY FUNCTIONS

18.1 THE ISSUE

As indicated in the introduction to evaluation, Chapter 10, people's valuation of benefits and costs are often distinctly nonlinear. When this is the case, the evaluation of alternatives cannot properly use the methods of Chapters 11–13. Indeed, if the value of units of a product of a system varies according to its total quantity, we cannot legitimately add them to calculate total benefits or costs. (For example: $X^a + X^a = (2X)^a$ only if the expression is linear, $a = 1$.) When values are nonlinear, we need new procedures of evaluation.

This chapter begins the sequence presenting the ways to evaluate systems when values are nonlinear (Chapters 18–21). This chapter lays the theoretical foundation. The following two chapters indicate how nonlinear values can be measured over first one and then many dimensions. Chapter 21, finally, extends the analysis to the most complex situation of all, collective decisionmaking among groups with nonlinear values over many dimensions.

In practice, nonlinear values are routinely integrated with decision analysis, Chapter 16. This is because significant projects involving risk also normally require this treatment. Most texts on decision analysis in fact combine these two

elements. They are separated here to make the presentation clearer. Section 19.8 illustrates their combination.

18.2 OBJECTIVE

This chapter describes the theoretic basis for representing the relative intensity of the nonlinear values that a person or organization places on consequences. These are collectively referred to as preference functions, $PF(\cdot)$, which assign a value to the quantity $(\cdot)$. These mathematical functions permit the analyst to estimate the value of possible choices, and then to advise his client, a single decision-maker or some group, as to the best choice.

There are two basic means of representing nonlinear preferences for possible benefits or losses: the value and the utility functions. These closely related concepts differ principally in the generality of their application. Value functions are based on just a few basic assumptions and are valid for almost any circumstances; they are so general, however, as to be useless for practical analyses (see Section 18.4). Utility functions result from a further set of assumptions and are thus less general; they are in fact particularly applied to situations involving uncertainty (see Section 18.5). Utility functions are most useful in practice and routinely used in decision analysis. Each concept is presented in detail, starting with the set of axioms that constitute its basis. These sets of axioms have practical consequences for the use of the function and its measurements which are described.

Semantic caution: The definition of utility differs between economics and the field of decision analysis within operations research. The economics literature generally refers to all preference functions as utility functions, whereas in decision analysis this term has a specific, limited meaning. As the distinction between value and utility functions is useful, it is maintained here.

18.3 NONLINEARITY OF PREFERENCES

People in general do not attach the same value to each unit of benefit they receive or of cost they pay. For example, when you are hungry you are most interested in the first plate of food, less in the second, and even less in the third or fourth, when you are getting full. The relationship between a set of consequences, $\mathbf{X}$, and the *measure of preference, PM*, is in general nonlinear and possibly highly complex. Mathematically, we can think of

$$PM = PF(\mathbf{X})$$

where $PF(\mathbf{X})$ is the preference function.

A common form of nonlinearity of preference is that referred to as *diminishing marginal "utility"* by economists (see preceding Semantic Caution). It reflects the fact that people commonly attach less and less value to each additional unit of a benefit they might receive. The way you probably react when hungry illustrates the phenomenon: you will place a high value on your first plate of food and then, as the need for food decreases, your value per plate of food decreases. At some level of benefit, X_s (pronounced "ex-ess"), saturation may occur. The marginal increase in value then decreases toward zero. The total measure of preference correspondingly increases at a decreasing rate. Figure 18.1 suggests the relationship.

A parallel form of nonlinearity is suggested by a person's typical response to uncertain situations. In choosing between alternatives with different degrees of risk, people often select the option with a lesser benefit on average, but which is more certain. For example, imagine yourself evaluating two possibilities for investing \$10,000 of your savings for a year: you may either put it in a bank and receive \$900 interest, or put it in the stock market, a possibility that you estimate is equally likely (probability, $P = 0.5$) to earn you \$0 or \$2000. On "average," the riskier investment is more profitable: the expected value of its gain is \$1000 ($= \$0(0.5) + \$2000(0.5)$). But when you make the choice only once you of course do not see the average, only the particular outcomes. In this situation most people usually prefer the more certain choice, although less profitable on average (see box). This behavior is generally described as "*risk aversion.*"

Exploring Nonlinear Preferences

The concept of nonlinear preferences can be made vivid by some simple experiments. These can be organized either in class or as exercises. Any number of persons can be involved. One person should be the organizer of each game.

The object of the exercises is to show the diversity in preferences that routinely occurs among individuals, and some of the factors that change the preference functions. This is done in two complementary steps.

At the start of the procedure, everyone should have two pieces of paper and the organizer some cash. The participants will submit bids for money that the organizer offers to give away. The results consist of histograms of the frequency of offers at different levels, compared with the expected value of the money offered.

In the first part, the organizer offers to give away a small amount—say \$1—on the toss of a coin: the organizer gives the dollar to a bidder if the coin shows a head, but keeps it on tails. The organizer then asks each participant to submit a bid for the opportunity to receive the money. If the organizer accepts the bid, the participant pays the bid and hopes for the best. The participant is equally likely to net \$(1 − Bid) or \$(−Bid). The expected value of the net monetary gain is \$(0.50 − Bid).

The organizer plots the distribution of the bids before selecting the highest bidder (this increases the interest). The results are typically as on the top of Figure 18.2. Points to be noted are:

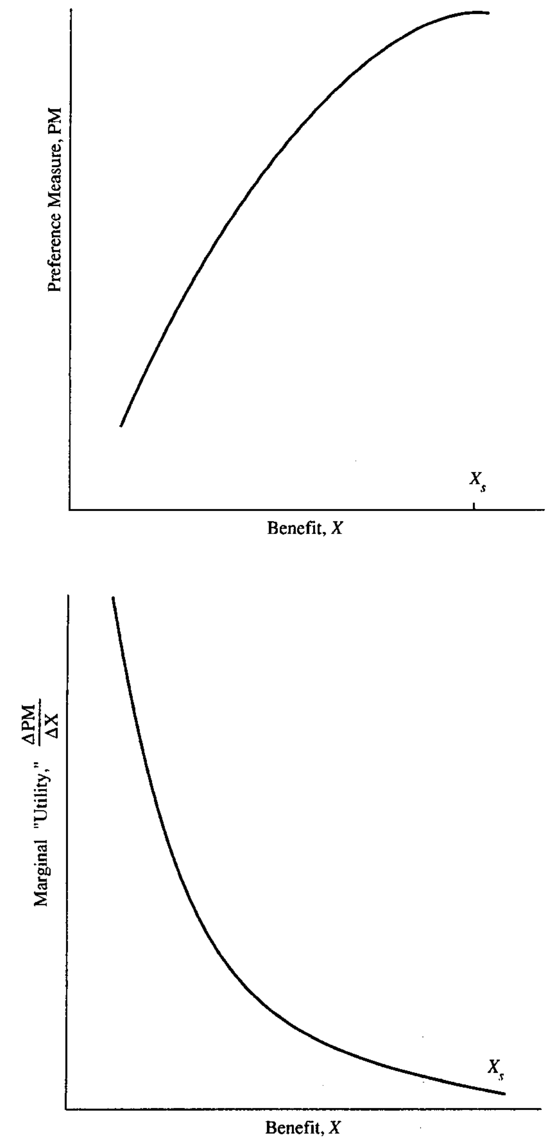

FIGURE 18.1
A nonlinear preference function with diminishing marginal "utility."

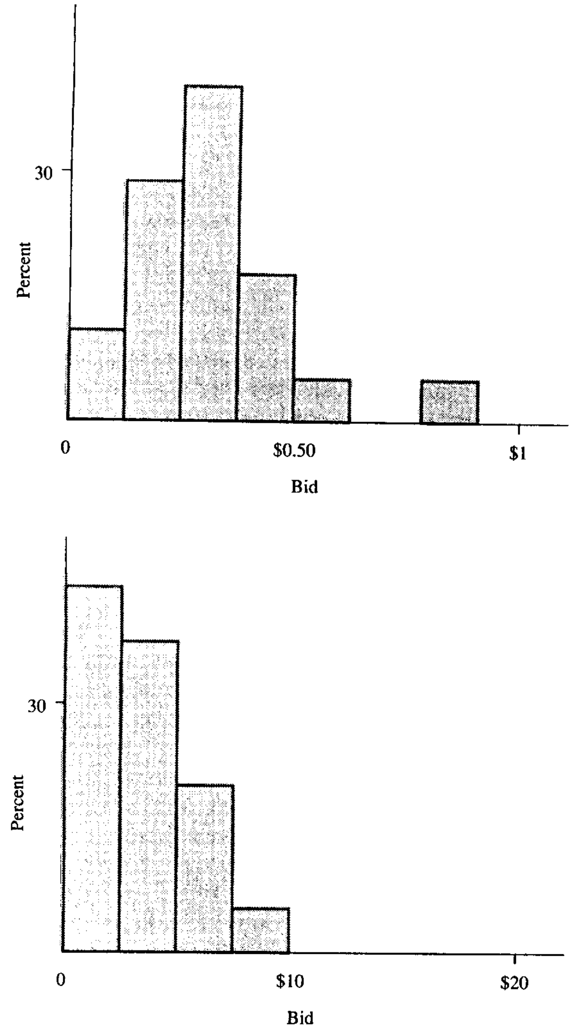

FIGURE 18.2
Typical distribution of bids for 50:50 chance on $1 (top) and $20 (bottom).

1. There is great diversity among the individuals.
2. Most offers reflect "risk aversion" in that the individuals bid less than the average of the offer.
3. Some individuals may show "risk preference" in that they bid more than the average.

The second exercise differs from the first by the size of the prize. A much larger amount, say $20, is offered. The results are as at bottom in Figure 18.2: people appear much more "risk averse." This and subsequent discussion among the participants brings out the fact that risk aversion and preferences generally depend on the amounts at stake, the particular circumstances of each individual (cash on hand, when they last received money, and so on), as well as personal preferences.

Note: No bid is more correct than any other. People differ, and legitimately and appropriately make different choices. The analyst must respect these choices.

The phenomena of diminishing marginal "utility" and "risk aversion" are closely related. This is illustrated with the help of Figure 18.3. We suppose that a person's preference function exhibits diminishing marginal "utility," similar to that of Figure 18.1, and show that it implies "risk averse" behavior. For each of the consequences of the choice with the uncertain outcomes, given in the preceding example ($0 and $2000), we can identify the measures of preference (PM_0 and PM_{2000}) and calculate the average measure of preference of the choice, $EV(PM)$. This is simply the expectation:

$$EV(PM) = PM_0(0.5) + PM_{2000}(0.5)$$

We can see from the shape of the preference function that the benefit that is equivalent to this average measure of preference,

$$X_e = PF^{-1}[EV(PM)]$$

is less than the average consequence, $1000. For most people, it is typically also less than $900. In this case $900 has a higher measure of preference than X_e and is preferred to it. The "risk averse" behavior, being the choice of the lesser benefit that is certain ($900) over the higher expectation of the risk situation, is thus associated with the nonlinearity of the preference function similar to that represented by diminishing marginal "utility."

Semantic caution: The term "risk aversion" is often used in a most misleading way. The decision analysis literature measures "risk aversion" by the extent to which the preference function is nonlinear (See Section 19.7). The implication is that all nonlinearity is due to the presence of uncertainty. This is false because, as the previous discussion shows, "risk aversion" is compatible with and partially explainable by the decreasing marginal "utility" associated with saturation. The same remark applies to "risk preference" discussed in what follows.

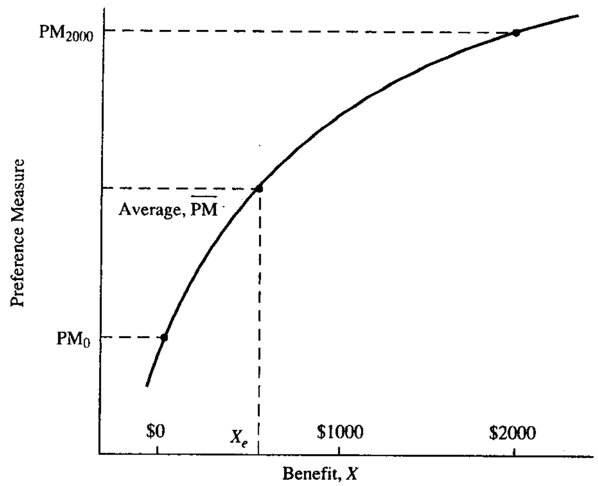

FIGURE 18.3
Representation of risk averse behavior: The measure of preference for the example risky situation, X_e, is less than the $1000 expectation of the situation.

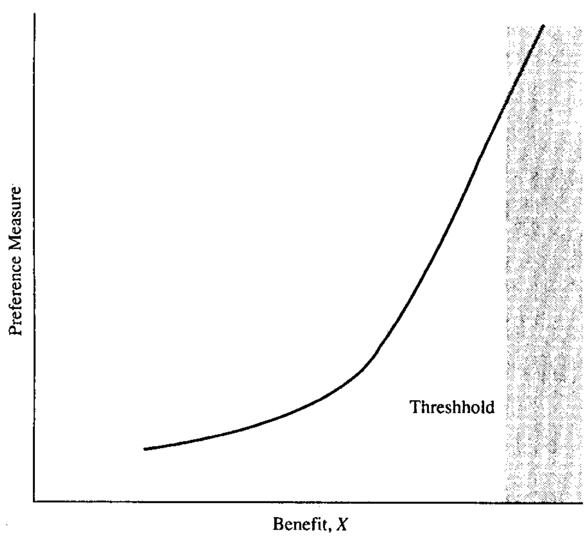

FIGURE 18.4
Preference function compatible with a threshold effect or risk preference.

People's preferences are not always compatible with diminishing marginal "utility" or "risk aversion." The inverse behavior, although less common, occurs frequently enough to be ordinary. Graphically, the preference function would be as in Figure 18.4. It would be likely to occur when there is some threshold to the consequences, above which they really matter and below which they are much less significant. For example, one might have such a preference for the speed of response to an emergency: if the police arrive fast enough, for example, they may prevent a crime—worth a great deal—if not, they can only record the event—worth relatively little.

Similarly, people faced with risky choices may show "*risk preference*." The contrary of "risk aversion," it means that they choose options that on average are less profitable, but that have the possibility of obtaining significant rewards. Entrepreneurs drilling for oil have been documented to behave in this way. Individuals buying lottery tickets, where the average total prize money paid is about 30% of the money collected, also exhibit "risk preference." The interpretation is that the people believe that winning the big prize would lift them above the threshold of their usual standard of living to a totally new life—for instance, one in which they might not have to work—and attach great importance to this possibility.

Preference functions need not be all one way or the other. They can easily be mixed (Figure 18.5). The inflection in the curve is associated with the thresholds.

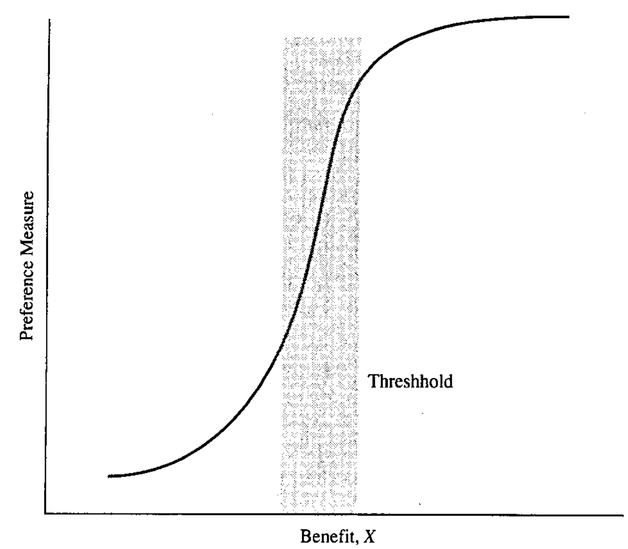

FIGURE 18.5
Preference function reflecting two different patterns of choice above and below a threshold.

Taking up the example of the speed of the police response to an emergency, one might attach only decreasing importance to the number of minutes they arrive in advance; what really counts is that they arrive in time. People confronted with risky choices may likewise exhibit "risk preference" for small amounts and "risk aversion" for larger ones. This behavior has been particularly documented for contractors who normally operate with a budget for unforeseen contingencies: so long as their choices fall within this budget, they may accept risks that they would avoid if the sums were larger.

Preference functions are usually assumed to be monotonically increasing with benefits. Strictly speaking, this need not be the case. One's valuation for the quantity of salt in one's food is an example of a preference function that first increases and then decreases. Some salt is appreciated, but beyond the right amount, more salt becomes increasingly distasteful (Figure 18.6). The justification of monotonic increases in practice lies in the idea that a person can avoid excess amounts of a consequence that would decrease his or her preference. One can, for example, simply not use any more salt than one wants, even if one has it available.

The problem for evaluation is to determine the form of the preference function, and to do so as analytically as possible. This is where value and utility functions are useful.

18.4 VALUE FUNCTION

A *value function*, $V(\mathbf{X})$, is a means of ranking the order relative preference between sets of consequences, of benefits and costs, $\mathbf{X}$. It assigns a number to every $\mathbf{X}$ such that for any two sets, $\mathbf{X}_1$ and $\mathbf{X}_2$, one is preferred to the other ($>$),

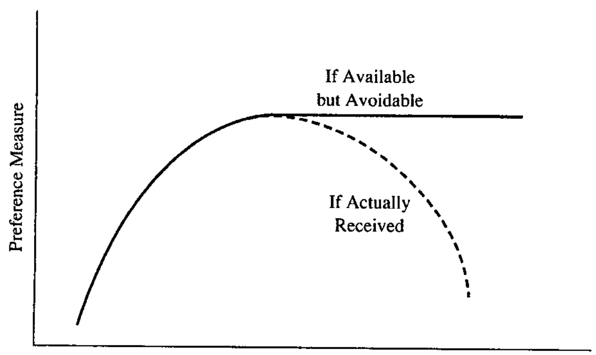

FIGURE 18.6
Preference functions for a benefit that is eventually undesirable.

only if its value is greater than the other's:

$$\mathbf{X}_1 > \mathbf{X}_2 \text{ only if } V(\mathbf{X}_1) > V(\mathbf{X}_2)$$

Similarly, $\mathbf{X}_1$ will be indifferent to or preferred by $\mathbf{X}_2(\mathbf{X}_1 \sim \mathbf{X}_2$ or $\mathbf{X}_1 < \mathbf{X}_2)$ depending on their relative values.

Axiomatic basis. The existence of a value function depends on three basic axioms, that is, assumptions about the situation. Mathematically two other technical axioms are also required, but these have little significance for the analyst. Each important axiom is described in turn, together with the several names by which it is called in economics and operations research.

The most basic axiom concerns the existence of relative preferences for all outcomes. It asserts that

> For every possible pair of consequences, $\mathbf{X}_1$ and $\mathbf{X}_2$, in the domain of interest, a person will either prefer one to the other or be indifferent between them:
>
> $$\mathbf{X}_1 > \mathbf{X}_2, \quad \mathbf{X}_1 \sim \mathbf{X}_2, \quad \text{or} \quad \mathbf{X}_1 < \mathbf{X}_2$$

It is referred to as the axiom of completeness or complete preorder. It represents a reasonable assumption. Although in general a person may confront many situations involving choices which have quite different attributes and may, thus, find it impossible at the moment to state his or her relative preferences, that does not mean that the person does not ultimately have relative preferences. The axiom is equivalent to the assumption that people do, in the end, make choices and can thus express their preferences.

The second axiom is that of transitivity. It assumes that

> For any three possible sets of consequences, $\mathbf{X}_1$, $\mathbf{X}_2$, and $\mathbf{X}_3$, if $\mathbf{X}_1 > \mathbf{X}_2$ and $\mathbf{X}_2 > \mathbf{X}_3$ then the preference is transitive such that $\mathbf{X}_1 > \mathbf{X}_3$.

This is reasonable for any individual or group with a common set of preferences. It may, however, not be true for groups in general. This is because different individuals may rank choices in different orders and this fact, combined with the way the group organizes or votes upon its collective choice, may lead to intransitive choice by groups. This phenomenon is discussed in detail in Chapter 21.2.

The third important axiom is that of monotonicity. It simply assumes that more of a good thing is better (or more of a bad thing is worse). For reasons indicated in the previous section, when the desirability of salt in food was discussed (Figure 18.6), this is a reasonable assumption in most practical situations.

The monotonicity axiom is equivalent to the Archimedean Principle that the value of any item in a series can be represented as a weighted average of the value of the extremes. That is, expressing the greatest and least $\mathbf{X}_i$ as $\mathbf{X}^*$ and $\mathbf{X}_*$, respectively, For all $\mathbf{X}_i, \mathbf{X}_j$ within the range of interest, $\mathbf{X}^* \succcurlyeq \mathbf{X}_i$, $\mathbf{X}_j \succcurlyeq \mathbf{X}_*$, there is a number between zero and one, $0 \le w \le 1$, such that some other $\mathbf{X}_K$ lying between $\mathbf{X}_i$ and $\mathbf{X}_j$ can be expressed as

$$V(\mathbf{X}_K) = wV(\mathbf{X}_i) + (1 - w)V(\mathbf{X}_j)$$

This can only be true if preferences are monotonic.

Consequences. If the assumptions implied by the preceding axioms are acceptable, then it can be demonstrated that a value function exists. The nature of a value function should be understood carefully. The value function defines which possibilities are preferred to which others: it defines the *order* of preferences. Unless additional assumptions are made, it says nothing about the intensity of the preferences.

The units of the value function have no intrinsic meaning. Any value function can be transformed by a monotonic function, $MF(\bullet)$, and the result will represent exactly the same preferences as before because:

$$MF(V(\mathbf{X}_0)) > MF(V(\mathbf{X}_1)) \text{ only if } V(\mathbf{X}_0) > V(\mathbf{X}_1)$$

For example, the value function over two dimensions:

$$V_1(X_1, X_2) = X_1^2 X_2$$

will give exactly the same ordering over combinations of X_1 and X_2 as

$$V_2(X_1, X_2) = 2\log(X_1) + \log(X_2)$$

even though the units would be quite different. Value functions such as V_1 and V_2 which are monotonic transforms of each other and thus lead to the same ordering over $\mathbf{X}$ are said to be *strategically equivalent*.

Because the units of the value function are meaningless, little attention is paid to describing its functional form. Attention is focused instead on the one aspect that can be defined unambiguously: indifference between different sets of consequences, $\mathbf{X}$. Zero difference in value function is not altered by any monotonic transformation and is something definite. The effort is thus to identify sets of consequences which are equally preferred or indifferent to each other.

The structure of the value function is commonly portrayed by contours of equal preference between sets of consequences. These contours may be called *isovalue lines* in general (the economics literature calls them *isoutility lines*). Strategically equivalent value functions have identical isovalue lines. Figure 18.7 illustrates the situation.

You should note the resemblance between the value functions in Figure 18.7 and demand functions, representing the quantity of an item that will be purchased as a function of its price, typically as illustrated in Figure 18.8. There is a close connection between them. Formally, the demand function represents the quantity of goods a person or group will purchase at various prices. The demand

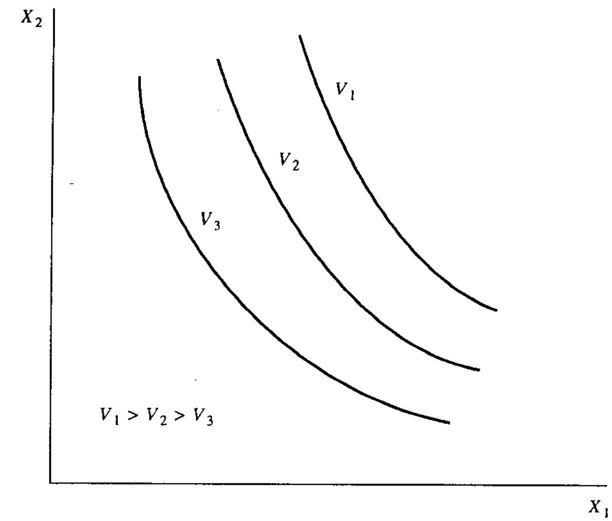

FIGURE 18.7
Isovalue contours representing the structure of a value function for two variables.

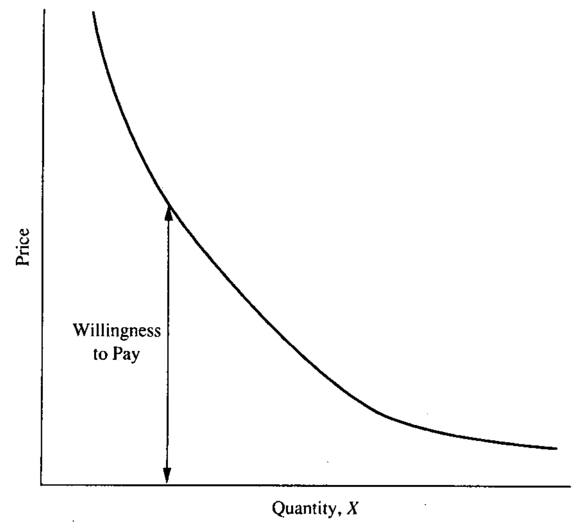

FIGURE 18.8
A demand function can be viewed as a representation of diminishing marginal "utility."

function represents the amount that a person will pay for each incremental unit of an item. This amount, the *willingness to pay*, can be viewed as the change in value associated with each extra unit. The typical demand function thus represents diminishing marginal "utility."

Measurement. Value functions are not measured directly, as a practical matter. This is only attempted in special circumstances, generally in research programs, and when additional assumptions are believed to hold. This is a consequence of the observation that quite different functions may be strategically equivalent and that units of value have no intrinsic meaning. There is thus much controversy about the significance of any attempt to measure value functions.

For the purpose of evaluation, the information contained in a value function can be obtained indirectly. This is done by estimating the demand function using a statistical analysis of choices people actually make. This analysis portrays revealed preferences, revealed in the sense that the choices a person makes define to some extent that person's underlying preferences. As a practical matter, value functions are never estimated for use in a decision analysis or an evaluation.

18.5 UTILITY FUNCTION

The *utility function*, $U(\mathbf{X})$, is a special kind of value function with a noticeable advantage: its units do have meaning relative to each other, unlike value functions. The utility function exists on a particular cardinal scale, on which values can be calculated meaningfully. It is then possible, within the limits defined by the axioms of utility, to measure the utility function.

This advantage has a most important practical consequence. It means that, provided always that the axiomatic assumptions are valid, it is possible to evaluate choices analytically, even when people have nonlinear preferences. The procedures for doing this are those of decision analysis, Chapter 16, and the way they can be combined with utility functions is illustrated in Section 19.8.

Lotteries. The discussion of utility relies heavily on the concept of a lottery. This is a simple risky situation, which can be visualized as a one-stage chance node as presented in the presentation of decision trees (Section 16.4). The different kinds of lotteries are defined as follows.

A *lottery* is characterized by the set of possible outcomes, $\mathbf{X}_i$, which will occur with probability P_i. Graphically, it is usual to show a lottery as a set of branches, as in Figure 18.9. The lottery can also be written as $(X_1, P_1; X_2, P_2; \ldots)$, being the sequence of pairs of outcomes and associated probabilities. The order in which these pairs are given has no particular significance.

When there are only two possible branches to the lottery, and where P_2 must be complementary to P_1, a more compact notation is used: $(X_1, P; X_2)$. This is called a *binary lottery*.

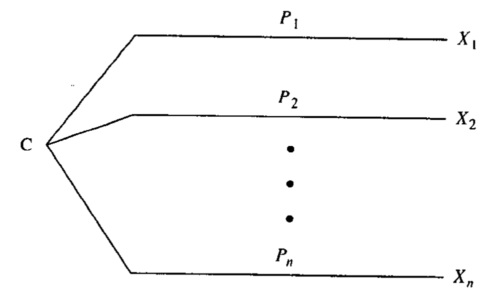

FIGURE 18.9
A lottery represented as a set of branches.

When one of the outcomes of a binary lottery is the status quo, and thus equals zero, we have the *elementary lottery*: (X, P).

Axiomatic basis. The assumptions underlying the utility function are all those required for value functions, plus some others. These consist of two assumptions concerning probability, and a key assumption about the nature of a person's preferences for probability. This last axiom is central; it is also quite controversial and the subject of much research. It represents the main limitation on the applicability of utility functions.

The axioms concerning probability are straightforward. The first concerns existence: that probabilities exist and can be quantified. This is a reasonable proposition, however much difficulty there might be in obtaining good estimates of probability in any specific case (see Section 15.1).

The second axiom on probability can be viewed as the monotonicity axiom for probability. It says, in effect, that a higher probability of a benefit is better. Formally, it states that a person prefers a greater chance of getting a desirable prize to a lesser chance.

> Given two choices, each with the same uncertain outcomes $\mathbf{X}_1$ and $\mathbf{X}_2$ where $\mathbf{X}_1$ is preferred to $\mathbf{X}_2$, a person will prefer the choice with the greater probability of getting $\mathbf{X}_1$ (and the complementary, lesser probability of getting $\mathbf{X}_2$).

As a specific example, this axiom assumes that you would prefer a 10% chance of getting a $1000 to a 5% chance, and that the greater the chance of the prize, the more valuable it is to you.

The key axiom of utility theory concerns the nature of a person's preferences with regard to probability: it assumes, in effect, that a person's preferences are linear in probability. The axiom is usually stated in different terms, however, and is referred to as the *substitution* or *independence axiom*. Either way, the axiom leads to the assertion that a person's preferences for an item should vary linearly with the probability of its occurrence. This is a testable proposition that can be compared to actual behavior, as is done a few paragraphs onwards.

This *substitution* axiom simply states that

> If a person places an equal value on two possible outcomes, A and B, then these can be substituted for each other in any choice involving uncertain outcomes without changing that choice. Thus, if **A** ~ **B**, then $P(A) + (1 - P)C \sim P(B) + (1 - P)C$.

It asserts, in effect, that equals can be substituted for each other. Stated this way, the axiom seems absolutely obvious. There is a subtlety, however, which is at the heart of the controversy concerning this axiom.

The subtlety of the substitution axiom lies in the idea that the outcomes that can be substituted are not necessarily identical kinds of items but are items that can be quite different in character. One may be a certain outcome, and the other may be a lottery. The substitutions supposed by the axiom may thus radically change the risks faced by a person, for example from certainty to uncertainty or in general from one probability distribution to another. The axiom implies that the substitutions can occur regardless of the other opportunities in front of a person and, thus, regardless of how these substitutions alter the probabilistic distribution of the consequences. For this reason the substitution axiom is also called the *independence axiom*.

The conditions called for by the substitution/independence axiom will only be met if the person's preferences are linear in probability. If they are not, changing the probability distribution could change the preference over choices involving uncertain situations.

The controversy about this axiom arises because repeated experiments demonstrate that people often act as if their preferences were nonlinear in probability. A classic instance of this is the "certainty effect" evidenced by individuals prepared to pay disproportionate amounts to achieve certainty in an uncertain situation. The experiment, simply stated, confronts the individual with the idea that he or she holds a certain number of lottery tickets to win a prize and asks how much the person would be willing to pay to obtain another ticket. The results are typically that the respondent would pay some average sum, except if there were the opportunity to buy the extra tickets to complete the set and thus guarantee the prize. A famous example of behavior contrary to the substitution/independence axiom is the Allais "paradox"—which is only a paradox if your premise is that the axiom is correct (see box).

The essential practical questions raised by the controversy are: do we assume that the substitution axiom is so obviously rational that any behavior contrary to it is a mistake, however well informed and intelligent the person may be? If so we should somehow disregard some of the person's stated preferences, and substitute ones we deem correct. (This can be taken as being presumptuous, if not arrogant.) Or do we suppose that the axiom is sometimes deficient for some subtle reason and reject it—and consequently the utility function which depends upon this axiom—when people express contrary preferences? Logic by itself does not resolve this question.

The Allais "Paradox"

This is a classic example of a situation in which many people choose as if their preferences were nonlinear in probability. This is contrary to the premise of the substitution/independence axiom. For those who believe this axiom must hold, the example is a "paradox."

For others, the author included, this case is simply an example of the complexity of human behavior. The author's own position is that there is nothing unusual about acting contrary to any norm or axiom which, when seen by itself, appears obvious: there are numerous situations when different norms conflict and it is impossible to satisfy all of them simultaneously. For example, we may agree that two important norms for behavior are truthfulness and kindness. Yet, when someone confronts us with the question "How do I look?" do we always manage to be both absolutely truthful and kind? Probably not: If we are absolutely truthful we may be unkind and vice-versa. It simply may not be possible to conform to all norms simultaneously.

The example can easily be run as an individual or classroom demonstration. It consists of two sets of choices. The organizer presents each set and asks the participants to express their preference on a piece of paper. The observation is that many participants, typically over half, indicate combinations of choices which are inconsistent with the substitution axiom.

The first choice is between a definite fortune and a gamble for a greater or equal amount but with a small but perceptible probability of getting nothing:

$$A = \$1 \text{ million}$$

or

$$B = (\$5 \text{ million}, 0.10; \quad \$1 \text{ million}, 0.89; \quad \$0, 0.01)$$

In practice, almost everyone responds: $A > B$.

The second choice is between two similar gambles:

$$C = (\$1 \text{ million}, 0.11; \quad \$0, 0.89)$$

or

$$D = (\$5 \text{ million}, 0.10; \quad \$0, 0.90)$$

People normally divide in their preferences between C and D. For a majority, though, $D > C$. Their reasoning apparently is that the great difference in prizes outweighs the small difference in probability.

The number of participants making any choice can be presented in a simple matrix as in Figure 18.10. Only the pairs of choices AC and BD are consistent with the substitution/independence axiom.

That the predominant pair of preferences, AD, is inconsistent with the substitution axiom is demonstrated as follows. By the monotonicity/Archimedean axiom we know that there is some w such that

$$\$1 \text{ million} \sim (\$5 \text{ million}, w; \quad \$0)$$

	A	B
C	30	5
D	55	10

FIGURE 18.10
Typical distribution of participants by pairs of choices. Shaded combinations are consistent with utility axioms, others illustrate Allais "paradox."

We then substitute the lottery for the $1 million where it occurs. Thus $A > B$ implies:

$$(\$5 \text{ million}, w;\quad \$0) > (\$5 \text{ million}, 0.89w + 0.10;\quad \$0)$$

Since we presume more is better regarding probability, this means that

$$0.11w > 0.10$$

Likewise, $D > C$ implies:

$$0.10 > 0.11w$$

In short, the pairs of choices *AD* and *BC* represent situations for which the substitution axiom does not hold; it leads to contradictory inferences about w. The explanation of the phenomenon appears to be that many individuals are nonlinear in their valuation of probability, especially toward certainty: they weigh the extra 1% chance of success highly when it is the difference between a sure thing and a gamble (*A* vs. *B*) but only lightly when it merely changes the gamble slightly (*C* vs. *D*). This is the essence of the "certainty effect."

As a practical matter, recent research we have conducted indicates that the substitution axiom holds, at least as a first-order approximation, except when the probability of some great consequence is either very small or close to certain. Consequently it would seem reasonable to use utility functions except in those cases. Unfortunately those "low probability, high consequence" situations, such as those associated with nuclear reactors say, are fairly common.

Consequences. The axioms that distinguish utility from value functions lead to two important practical results:

1. The utility of a consequence can be measured on a special cardinal scale (see Section 18.6), which allows calculations and analysis.
2. Measurement of utility is a balancing procedure in which probabilities serve as the weights (see Chapter 19).

These results depend particularly on the substitution/independence axiom and the linearity of value with respect to probability it implies. Indeed this linearity

permits us to state that the utility of any chance at a prize $(\mathbf{X}, P)$ as a multiplication of that probability and utility of the prize:

$$U(\mathbf{X}, P) = PU(\mathbf{X})$$

The value of an entire risky situation, or lottery, is then

$$U(\text{Lottery}) = \sum P_i U(\mathbf{X}_i)$$

This relationship leads both to the cardinality of the utility function and to the possibility of measuring it.

18.6 SCALE FOR UTILITY

The scale on which utility can be measured is a special form of the cardinal scale: the ordered metric. To understand what it implies, consider first the standard cardinal scale, the ratio scale, which is used to calibrate phenomena such as distance and time. The *ratio scale* has two features:

1. Its units are constant, identifiable amounts (such as meters or hours) which can be combined linearly by addition, subtraction, or averaging. For example: 1 mile plus 2 miles is 3 miles; a speed of 30 mph for a day plus 60 mph for 2 days averages to 50 mph.
2. Zero on the scale indicates an absence of the phenomenon being observed. This implies that ratios between measures have meaning. Thus, if an adult is 6 ft (1.8 m) tall and a child 3 ft (0.90 m), we can say that the adult is twice as tall as the child.

This second feature is the reason the standard cardinal scale is known as a ratio scale.

The *ordered metric scale* differs from the ratio scale in that it does not have the second property. Zero on an ordered metric scale has no absolute meaning; it is simply a reference point that could as well be any other number, such as 32 or whatever. The concept of the ordered metric scale is little known although it is widely used in practice, specifically to measure temperature (see box).

Ratios between measures on the ordered metric scale do not have any meaning. Thus if one object has a centigrade temperature of 100°C and another a temperature of 25°C, we cannot validly say that one is four times as hot as the other. Indeed, the Fahrenheit temperatures of these objects are also 212°F and 77°F, and not a 4:1 ratio.

All measurements on an ordered metric scale can be transformed into equivalent measures in an important way, known as a *positive linear transformation*. To understand this, start with the idea that ratio scales can be transformed into equivalents by straight multiplication: we can convert meters into feet and dollars into francs by applying the appropriate positive conversion factor. For example,

$$(\text{distance in feet}) = \left(\frac{10}{3}\right)(\text{distance in meters})$$

Since there is no absolute zero on an ordered metric scale, we can also obtain equivalent scales by adding or subtracting any amount, which simply changes the reference point. We then have the positive linear transformation (*PLT*):

$$PLT(\bullet) = a(\bullet) \pm b \qquad a > 0$$

This is what we use to convert from degrees Fahrenheit to Centigrade and vice-versa (see box).

Measurements of utility, being on an ordered metric scale, can also be transformed into equivalents by a positive linear transformation:

$$U'(\mathbf{X}) = aU(\mathbf{X}) \pm b \qquad a > 0$$

More precisely, $U'(\mathbf{X})$ and $U(\mathbf{X})$ are strategically equivalent, as defined in Section 18.4, in that they would lead to the same results in an evaluation.

Any two utility functions related by a positive linear transformation are said, in the jargon of operations research, to have a *constant shape*. This term is used extensively with the analysis of utility functions over many dimensions (See Section 20.3).

Temperature: An Ordered Metric Scale

Temperature is traditionally measured on an ordered metric scale, either degrees Centigrade or Fahrenheit. (The Kelvin scale, which rests on the relatively modern concept of the absence of heat at absolute zero, is a ratio scale.)

Temperature scales permit proportional averaging, and are thus cardinal. Thus two units of a fluid at 60° plus one unit at 90° equals three units at 70°.

Conventional temperature scales do not imply an absence of heat at 0°; their zero points are simply convenient references. Ratios of temperature therefore have no particular significance; to say that "it's twice as hot today as yesterday" is a meaningless statement—unless one is referring to degrees Kelvin.

Temperatures scales are "constant up to a positive linear transformation." For example,

$$(°F) = \tfrac{9}{5}(°C) + 32$$

or, vice-versa,

$$(°C) = \tfrac{5}{9}(°F) - 17.77$$

Conventional temperature scales were each defined with respect to two absolutely arbitrary reference points. All other points on the scale were determined proportionally to these two. As a matter of convenience these two reference points were labeled 0° and 100° on both scales. For the Centigrade scale, these represent the freezing and boiling points of pure water, at standard atmospheric conditions. For the Fahrenheit scale these represent the freezing point of water saturated with salt and the human body temperature (they did not get the latter quite right in the early 18th century). Although the Fahrenheit scale looks odd, when one is thinking of freezing at 32°F, it is as rational as the Centigrade scale.

Semantic caution: Graphically, the notion of utility functions with constant shape can be confusing. When you plot such curves their slopes are generally different in fact. This is because of the a factor in the positive linear transformation. What is really constant is not the shape but the implications for choices: utility functions with constant shape, being strategically equivalent, always indicate the same order of preference among consequences or lotteries.

The above property implies that utility, and ordered metric scales in general, can be defined with respect to any two points one wishes to select. These can furthermore be assigned any convenient value. The ability to define the value of the ends of our scale of utility turns out to simplify the measurement of utility considerably. Conventionally, we thus set the utility for the worst set of consequences, $\mathbf{X}_*$, at zero:

$$U(\mathbf{X}_*) = 0$$

and the utility for the best consequences, $\mathbf{X}^*$, at 1:

$$U(\mathbf{X}^*) = 1$$

Sometimes utility is also scaled from 0 to 100 or, for undesirable outcomes such as fatalities, from −1 to 0. These choices are a matter of personal preference.

At this point we have all the essential theoretical basis for proceeding with actual measurements of utility. Chapter 19 shows how to do this for one dimension, and Chapter 20 shows how we can build on this capability to obtain utility functions over several dimensions, the *multiattribute* utility functions.

REFERENCE

Kahneman, D., Slovic, P., and Tversky, A., eds., (1982). *Judgement and Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, England.

PROBLEMS

18.1. *Diminishing Marginal Value*
Assume that on your next quiz in systems analysis, a grade between 60 and 69 is just passing, D; between 70 and 79 is a C; between 80 and 89 is a B; and over 90 is an A. Consider that at this stage your grade on the quiz would be 60. How many hours would you work to get 70? 75? 80? 85? 90? 95? 100? Plot results and discuss.

18.2. *Cornering the Market*
Suppose that your local organization is having a raffle for a college scholarship worth $5000, and is selling off 1000 tickets. How much might you pay for one ticket? For a 51st ticket, if you already had 50? For the 1000th ticket if you already had the 999 others? Discuss your answers.

18.3. *Assembly Robot (A Third Time)*
See Problem 17.2. Is the criterion of minimizing expected time the right one for the factory manager? Would it be equally good for decisions in other areas (for example, the time required for open heart surgery, for the arrest of an epidemic, or the interception of an enemy missile)?

18.4. *Threshhold*

(*a*) As the founder of a new engineering company you must earn $10,000 a month to pay rent, salaries, and the bank loan. Anything over this is your own salary. Anything under must be added to your debt. Think about and plot the form of your utility for a monthly income of $5000 to $20,000.

(*b*) As an employee in a big company, you know that the average salary increase for your colleagues is 5%. Variations depend on your bosses' assessment of your performance. Plot the form of your utility for a salary increase of 0 to 15%.

18.5. *Strategic Equivalence*

The sales, S, of a new high-tech product are estimated to depend on price, P (measured in $\$ \times 10^3$) and quality C (measured by an index):

$$S = 300P^{-0.5}C$$

(*a*) Plot price versus quantity (the demand function) for $C = 1, 2, 3$. Interpret the meaning of these curves.

(*b*) Suppose we introduce a new index of quality: $C' = 3C^2$. How would we change the formula for sales? How would the new index affect the analysis?

18.6. *Scales*

(*a*) What are the positive linear transformations that give °C as a function of °F? K as a function of °C?

(*b*) Suppose a utility function for X is a linear interpolation between the following points:

X	:	$0	$20	$50	$100
$U(X)$	:	0	0.1	0.8	1.0

Plot $U(X)$. Then plot $U'(X)$ as the positive linear transformation with a range of 0 to 0.4; 0.8 to 1.0; 0.5 to 0.8.

(*c*) The utility for response time, T, to an emergency is

T (min)	:	0	5	10	15	20
$U(T)$	:	1.0	0.8	0.5	0.1	0

Plot $U'(T)$ as the positive linear transformation with a range of 40 to 80.

(*d*) The utility for deaths, D, due to industrial accidents is

D(units)	:	0	1	2	5	10
$U(D)$	:	0	−0.01	−0.05	−0.30	−1.0

Plot $U'(D)$ as the positive linear transformation with a range of 0 to 100.

CHAPTER 19

MEASUREMENT OF UTILITY

19.1 ORGANIZATION

This chapter presents the basic procedures for measuring utility: the Certainty Equivalent (CE) and the Lottery Equivalent/Probability (LEP) methods. The CE approach has been the standard and is still widely used in practice. It suffers from a broad range of practical defects, however. This has led to the development of the LEP, which experts in the measurement of utility now recommend.

Many other procedures are possible. As the next section indicates, we measure utility by solving a simple linear equation with one unknown and a minimum of three independent arguments. These can be permuted in many ways, leading to a variety of measurement procedures, each emphasizing different features. Research into these possibilities is active, and we can expect additional improvements in utility measurement.

The techniques for measuring utility blend two contrasting intellectual paths. First there is the theory of utility, presented in Chapter 18. Then there is *psychometrics*, the science of measuring behavioral responses of people (or animals). Psychometrics has been developed and applied extensively in behavioral research for nearly a century. It demonstrates the techniques that should be used to obtain good measurements of utility, those that are replicable with small margins of error or variance.

The proper presentation of utility measurement is thus inherently complex because we must both integrate psychometrics and cover alternative methods. To simplify matters, we have adopted the following organization. Sections 19.2 and 19.3 provide the theoretical and psychometric basis. Sections 19.4 and 19.5 describe the CE and LEP methods and place them in the overall step-by-step