The Psychology of Preferences

When people make risky choices, they often do not do so objectively. Experimental surveys indicate that such departures from objectivity tend to follow regular patterns that can be described mathematically

by Daniel Kahneman and Amos Tversky

Scientific American (January 1982): 160-173

Imagine yourself on your way to a Broadway play with a pair of tickets for which you have paid $40. On entering the theater you discover that you have lost the tickets. Would you pay $40 for another pair of tickets? Now imagine you are on your way to the same play without having bought tickets. On entering the theater you realize that you have lost $40 in cash. Would you now buy tickets to the play?

In objective terms the two situations are identical: in both cases you are poorer by $40 than you were earlier, and you face the decision of whether or not to pay $40 to see the play. Nevertheless, most people presented with such a situation say they would be more likely to buy new tickets if they had lost money than if they had lost tickets. One interpretation of the difference between the two situations is that the same loss can be assigned to different "mental accounts." The loss of $40 in cash is entered in an account distinct from that of the play. The loss therefore has little effect on the decision to buy new tickets. In contrast, the cost of the lost tickets is posted to the account of the play. The unexpected doubling of the cost of the play is difficult to accept.

Recent investigations of the psychology of preferences have demonstrated several intriguing discrepancies between subjective and objective conceptions of decisions. For example, the threat of a loss has a greater impact on a decision than the possibility of an equivalent gain. Most people are also very sensitive to the difference between certainty and high probability and relatively insensitive to intermediate gradations of probability. The regret associated with a loss that was incurred by an action tends to be more intense than the regret associated with inaction or a missed opportunity. These observations and others of a similar character contribute to the understanding of how people make decisions and to the elucidation of some puzzles of rational choice.

The origin of the psychology of preferences can be traced to an essay published in 1738 by the mathematician Daniel Bernoulli. In the essay Bernoulli discussed a widespread characteristic of human preferences: risk aversion. To understand risk aversion, imagine that you are given a choice between two options. The first is a sure gain of $80. The second is a risky prospect that offers an 85 percent chance of winning $100 and a 15 percent chance of winning nothing.

Most people who are presented with this choice prefer the certain gain to the gamble, in spite of the fact that the gamble has a higher "monetary expectation" than the certain outcome. The monetary expectation of a gamble is the sum of its outcomes weighted by their probabilities. The monetary expectation of the gamble offered in this case is $100 multiplied by its probability (.85), added to $0 multiplied by its probability (.15), or a total of $85. The monetary expectation reflects the average monetary value of the gamble; if one were to play this gamble many times, the average gain would be about $85 per play. The monetary expectation of the certain gain is $80 multiplied by 1 (certainty), or $80. A choice is risk-averse if a certain outcome is preferred to a gamble with an equal or greater monetary expectation. A choice is risk-seeking, on the other hand, if a certain outcome is rejected in favor of a gamble with an equal or lower monetary expectation.

The hypothesis that people generally make risk-averse choices has been widely accepted by economists who normally assume that a consumer or an entrepreneur will choose a risky venture over a sure thing only when the monetary expectation of the venture is sufficiently high to compensate the decision maker for taking the risk. Psychological studies, however, indicate that risk-seeking preferences are common when people must choose between a sure loss and a substantial probability of a larger loss.

To get a sense of risk seeking imagine you are forced to choose between a sure loss of $80 and a risk that involves an 85 percent chance of losing $100 and a 15 percent chance of losing nothing. Faced with this choice a large majority prefer the gamble to the loss. The monetary expectation of the gamble (−$85), however, is worse than that of the certain loss (−$80). The majority preference is thus an instance of risk seeking.

The contrast between the majority choices observed in these two problems suggests that preferences between gains are risk-averse and that preferences between losses are risk-seeking. This pattern has been observed both in answers to hypothetical questions and in real decision problems where people were paid according to their choices. The same pattern has been confirmed for hypothetical problems that involve outcomes other than money, such as duration of pain and the number of lives that may be lost in an epidemic or saved by medical intervention.

When people make risk-averse or risk-seeking choices they forgo the option that offers the highest monetary expectation. In order to explain such choices Bernoulli replaced the objective criterion of monetary expectation with the subjective criterion of expected utility. According to expected utility theory, each outcome gives rise to a particular degree of pleasure or utility. The utility of a risky prospect is the weighted sum of the utilities of its outcomes, each multiplied by its probability. The central idea of the theory is that utility is not a linear function of money: a gain of $2,000 contributes less than twice as much to utility as a gain of $1,000. As a consequence the prospect with the highest monetary expectation does not necessarily have the highest expected utility. The decision maker is assumed to select the option with the highest utility, whether or not that option also has the greatest monetary expectation.

How do people identify the outcomes of a decision? It has generally been assumed, following Bernoulli, that utilities are assigned to states of wealth. We depart from this tradition and analyze choices in terms of changes of wealth rather than states of wealth. The classical analysis implies that preferences reflect a comprehensive view of the options. In contrast, we propose that people commonly adopt a limited view of the outcomes of decisions: they identify consequences as gains or losses relative to a neutral point. This form of mental accounting can lead to inconsistent choices, because the same objective consequence can be evaluated in more than one way.

To explain choices we introduce a function that associates a subjective value to any amount that may be gained or lost. We call it a value function. As in classical utility theory, risk seeking and risk aversion are explained by the curvature of the function that relates subjective values to objective outcomes. Consider first the value function for gains. We assume that gains have positive value and that a zero gain, which merely retains the status quo, has a subjective value of zero. We also propose that the value function for gains is concave downward, so that each extra dollar gained adds less to value than the preceding one. The concavity of the value function is compatible with the common impression that the difference between a gain of $100 and one of $200 is more significant than the difference between a gain of $1,100 and one of $1,200. This shape of the value function favors risk aversion.

Consider the choice between a sure gain and a gamble that offers some probability of a larger gain and some probability of no gain at all. The advantage of the large gain over the sure thing is evaluated in a shallow region of the value function, where increments of money produce relatively small increments of value. The advantage of the sure outcome over no gain is evaluated in the steepest region of the function, where each dollar makes a greater difference. Because value is nonlinear the sure outcome is closer to the large gain in terms of value than in terms of money. The shape of the value function thus favors a risk-averse preference for sure outcomes over risky prospects [*see illustration below*].

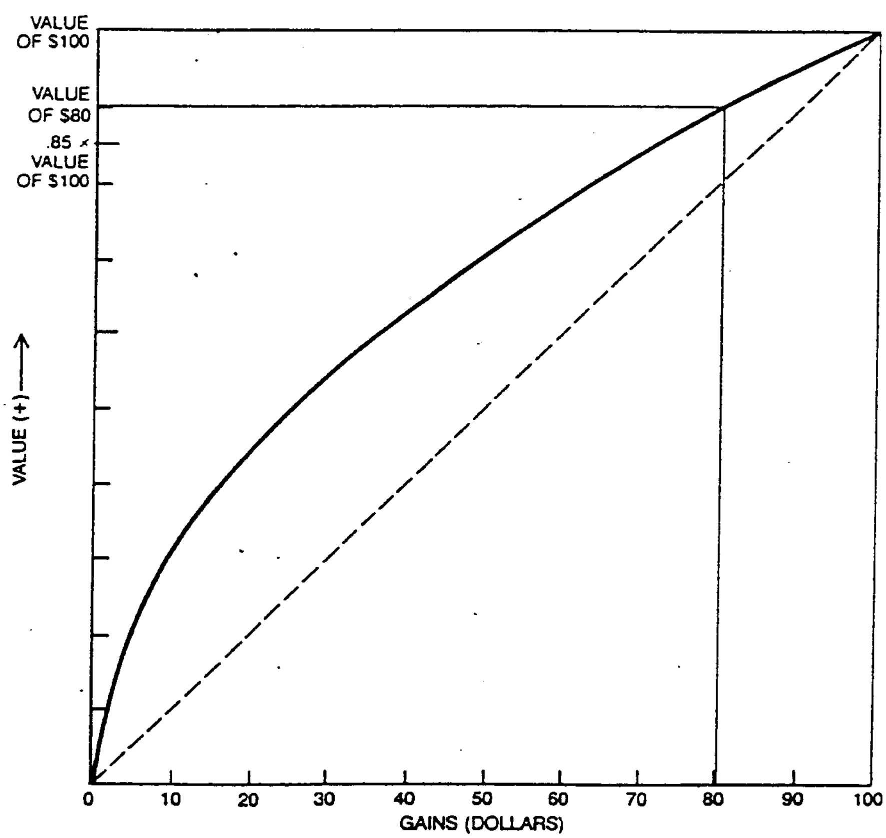

CONCAVE VALUE FUNCTION helps to explain preferences where gains are concerned. A value represents the subjective attractiveness of a gain (or loss). The function assigns a value to each possible gain. Gains are indicated on the horizontal axis, values on the vertical axis. The status quo is assigned the value of zero. (The units of value are arbitrary.) The concave function becomes progressively flatter as the amount of the gain is increased. Risk-averse preferences are common when people make choices among gains. In a risk-averse choice a sure gain is preferred to a gamble with a higher monetary expectation. The monetary expectation of a gamble is the sum of its outcomes multiplied by their probabilities. Hence an 85 percent chance of winning $100 has a monetary expectation of .85 multiplied by $100, or $85. This is greater than the monetary expectation of a sure gain of $80 ($80 multiplied by certainty, or $80). The value of the sure gain, however, is greater than that of the gamble, which is .85 multiplied by the value of $100. The value function thus predicts that the sure gain will be preferred to the gamble.

To explain the common observation of risk-seeking preferences in choices that involve losses, we assume that the value function for losses is convex, so that each extra dollar lost causes a smaller change in value than the preceding one. This proposal is compatible with the common impression that the difference between a loss of $100 and one of $200 appears more significant than the difference between a loss of $1,100 and one of $1,200. The convexity of the value function for losses favors a risk-seeking preference for a gamble over a sure loss. The advantage of the best outcome of the gamble (no loss) over the sure loss is evaluated in the steepest region of the value function. On the other hand, the advantage of the sure loss over the worst outcome of the gamble is evaluated in a flatter region. As a result the sure loss is relatively closer to the worst outcome on the value scale than it is on the money scale [*see illustration on page 164*].

The properties of the value function can be inferred from observed choices. To undertand how this procedure works imagine you have been offered a choice between winning a cash prize and having a 50 percent chance of winning $100. What amount would make the sure prize just as attractive as the bet? A choice of $35 as the matching prize is common. If you find $35 as attractive as the gamble, your value for this amount must equal your value for the gamble. On the assumptions that we introduced above the value of the gamble is 1/2 the value of $100, since the probability of winning $100 is 1/2 and the value of the other outcome (winning nothing) was assumed to be zero. Consequently the choice of $35 to match the gamble indicates that the value of $35 is half the value of $100.

Now find the prize amounts that match a 50 percent chance of winning $200, $500, $1,000 and $2,000. If you try this thought experiment, you will naturally find that the amount of the prize increases with the size of the bet. You may also observe that the prize is roughly proportional to the stake. A person who matches a prize of $35 to an even chance of winning $100 is likely to match a prize of $300 or $350 to an even chance of winning $1,000. Thus as the amount of the bet has increased by a factor of 10 the prize has grown by almost the same factor.

It can be proved mathematically that the value function of an individual who sets prizes exactly proportional to stakes can be described by a power function. In such a function the value of an amount of money equals that amount, raised to some power. The exponent of the value function can be estimated from choices. If the prize that matches a 50 percent chance of winning an amount of money is always 35 percent of that amount, the exponent of the value function for gains is about 2/3. Actually the proportionality of prizes to stakes is only approximate. The prize generally increases more slowly than the stake. Because the deviation from proportionality is small, however, a power function provides a convenient description of the values of gains.

A similar analysis can be applied to losses. An even chance of losing \$100 or losing nothing is often matched by a sure loss of about \$40, and an even chance of losing \$1,000 or nothing by a sure loss of about \$400. If the sure loss that matches a 50 percent chance to lose an amount is always 40 percent of that amount, the exponent of the value function for losses is about 3/4.

In order to piece together the value functions for both gains and losses it is necessary to study risky prospects that involve both positive and negative outcomes. Consider a gamble in which you face a 50 percent chance of losing \$100 and a 50 percent chance of winning a cash prize. What is the smallest prize that would make this bet acceptable to you? It is a commonplace that the pleasure of winning a sum of money is much less intense than the pain of losing the same sum. Accordingly people accept a gamble at even odds only when the possible gain from the gamble is substantially larger than the possible loss. For example, you may find that a 50 percent chance to lose \$100 is unacceptable unless it is combined with an equal chance to win \$200 or more.

By observing the choices people make among gains and losses we have specified some features of the value function. The approximate proportionality of prizes and bets is reflected in the power relation between gains and their values. Risk aversion and risk seeking are described by the concave and the convex segments of the value function. The asymmetry in the response to gains and to losses is expressed in the greater steepness of the value function for losses. It is important to emphasize, however, that the value function is merely a convenient summary of a common pattern of choices and not a universal law.

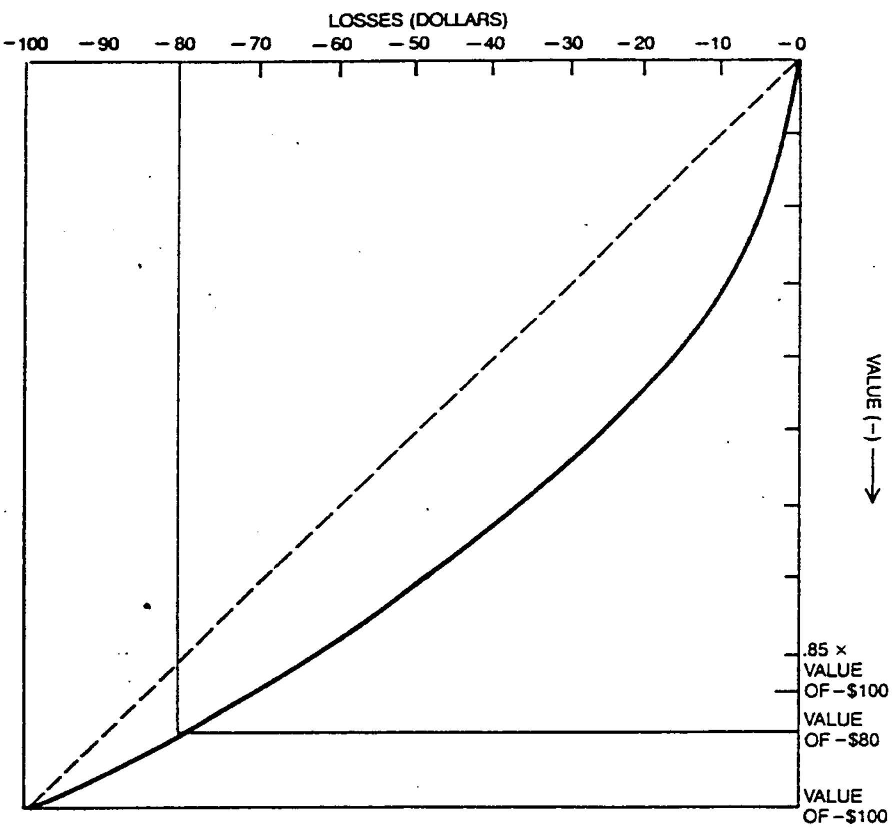

CONVEX VALUE FUNCTION **explains preferences where losses are concerned. The function assigns a unique negative value to each monetary loss. Losses are indicated on the horizontal axis to the left of the origin; values are shown on the vertical axis downward. The status quo is assigned the value of zero. The convex function becomes progressively flatter for larger losses. Risk-seeking preferences are common when people make choices among possible losses. In a risk-seeking choice a gamble is preferred to a sure loss that has a greater monetary expectation. The value function predicts that a gamble offering an 85 percent chance of losing \$100 and a 15 percent chance of losing nothing will be preferred to a sure loss of \$80. The monetary expectation of the gamble (−\$85) is less than that of the sure loss (−\$80). The value of the gamble, however (which is .85 multiplied by the value of −\$100), is greater (has a smaller negative magnitude) than that of the sure loss. Because the gamble has a greater value it will be preferred. The convex value function therefore helps to explain risk-seeking preferences.**

Naturally individuals differ in their attitudes toward risk and toward money, and no single value function can describe the preferences of all individuals. Choices also can vary substantially with the method used to elicit preferences. For example, the acceptability of a risky prospect may depend on whether the respondent matches a gain to a set loss, matches a loss to a set gain or adjusts the odds when the amounts at stake are fixed. In addition the near proportionality of prizes to stakes breaks down beyond the range of moderate gains and losses. For extremely large gains the value function becomes practically flat as individuals can no longer distinguish one huge gain from another. In the negative domain, however, the function can be very steep in the vicinity of monetary losses that would force a substantial change in a person's way of life.

The discussion to this point has assumed that the contribution of each outcome to the value of a risky prospect is weighted by its probability. There is evidence that this assumption is often violated and that the "decision weights" that multiply the values of outcomes do not coincide with the probabilities.

Some properties of decision weights can be explored by the following informal experiment. Imagine that you can improve your chances to win a very desirable prize. Would you pay as much to raise your chance of winning from 30 to 40 percent as you would to raise your chance from 90 percent to certainty? It is generally agreed that the former offer is less valuable than the latter. It is also agreed that an increase from impossibility to a probability of 10 percent is more significant than an increase from 30 to 40 percent. Thus the difference between certainty and possibility and the difference between possibility and impossibility loom larger than comparable differences in the intermediate range of probability [*see illustration on page 140*].

Investigations of risky choices tend to confirm these hypotheses. Low probabilities are commonly overweighted but intermediate and high probabilities are usually underweighted relative to certainty. The impossible event is naturally assigned a weight of zero and the certain event is assigned a weight of one. The overweighting of small probabilities can give rise to risk seeking in the positive domain and risk aversion in the negative domain, in contrast to the prevalent pattern of preferences described above.

The inflated effect of small probabilities contributes to the appeal of lottery tickets and accident insurance, which are concerned with events that are highly significant and relatively improbable. The underweighting of intermediate and high probabilities, on the other hand, reduces the attractiveness of possible gains relative to sure ones. It also reduces the threat of possible losses relative to sure ones. The contribution of decision weights was neglected in the construction of the value function we described above. In that function all observed risk aversion and risk seeking were assumed to be the result of curvature of the value function. If the probability of .5 that was applied in the construction of the function were replaced by a slightly smaller decision weight, the value function would become less curved, although it would retain its general *S* shape.

We have so far been concerned with rules that govern the evaluation of risky options. Another important aspect of the psychology of preferences is how people define the consequences of their choices. The same decision can be framed in several different ways; different frames can lead to different decisions. For example, consider the following problem:

Imagine that in addition to whatever else you have, you have been given $200. You are now asked to choose between (*A*) a sure gain of $50 and (*B*) a 25 percent chance of winning $200 and a 75 percent chance of winning nothing.

Most people given these options make a risk-averse choice, preferring the sure gain (*A*) over the gamble (*B*), which has the same expected value. Now consider the following problem:

Imagine that in addition to whatever you have, you have been given a cash gift of $400. You are now asked to choose between (*C*) a sure loss of $150 and (*D*) a 75 percent chance of losing $200 and a 25 percent chance of losing nothing.

Most people have a risk-seeking preference for the gamble (*D*) over the sure loss (*C*). The options presented in the two problems are identical, however, in objective terms. There is no valid reason to prefer the gamble in one version and the sure outcome in the other. Choosing the sure gain in the first problem yields a total gain of $200 plus $50, or $250. Choosing the sure loss in the second version yields the same result through the deduction of $150 from $400. The choice of the gamble in either problem yields a 75 percent chance of winning $200 and a 25 percent chance of winning $400.

If our respondents took a comprehensive view of the consequences, as is assumed by theories of rational decision, they would combine the bonus with the available options and evaluate the composite outcome. Instead they ignore the bonus and evaluate the first problem as a choice between gains and the second as a choice between losses. The reversal of preferences is induced by altering the description of outcomes. We call such reversals framing effects.

Framing effects arise when the same objective alternatives are evaluated in relation to different points of reference. We asked a large number of physicians to consider the following problem:

Imagine that the U.S. is preparing for the outbreak of a rare Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows: If Program *A* is adopted, 200 people will be saved. If Program *B* is adopted, there is a 1/3 probability that 600 people will be saved and a 2/3 probability that no people will be saved. Which of the two programs would you favor?

The majority response to this problem is a risk-averse preference for Program *A* over Program *B*.

Other respondents were presented with the same problem but a different formulation of the programs: If Program *C* is adopted, 400 people will die. If Program *D* is adopted, there is a 1/3 probability that nobody will die and a

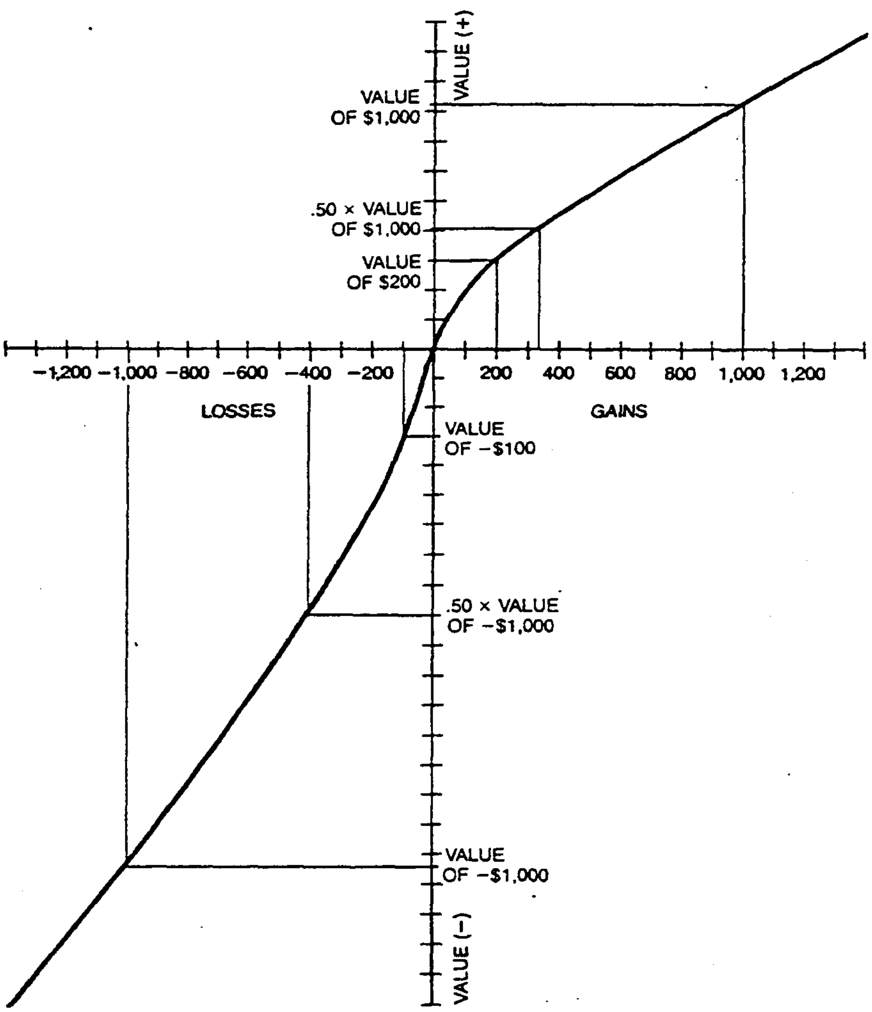

S-SHAPED VALUE FUNCTION combines a concave segment representing the values of gains and a convex segment representing the values of losses. The concave segment to the right of the origin reflects risk aversion in choices between gains; the convex segment to the left reflects risk seeking in choices between losses. The segment for gains represents preferences in which a 50 percent chance to win a given amount is just as acceptable as a sure gain of 35 percent of that amount. For example, a 50 percent chance of winning $1,000 is as acceptable as a sure gain of $350. The value of $350 is therefore half the value of $1,000. Such a pattern of preferences can be expressed by a power function with an exponent of 2/3. (In a power function the value of a gain is equal to the magnitude of the gain raised to a fixed exponent.) The segment for losses represents preferences in which a 50 percent chance of losing a given amount is just as acceptable as a sure loss of 40 percent of that amount: a 50 percent chance to lose $1,000 is as acceptable as a sure loss of about $400. The negative value of −$400 is therefore half the negative value of −$1,000. Preferences where losses are concerned can be expressed by a power function with an exponent of 3/4. The *S*-shaped function reflects the common observation that a loss has a greater subjective effect than an equivalent gain. Because the slope of the function for losses is steeper than that for gains, the negative value of a loss of $100 is equal to the positive value of a gain of $200. The function therefore represents the preference of a person who finds an even chance of winning $200 or of losing $100 just acceptable; this is a typical pattern. The function shown summarizes some common features of preferences observed in tests of many people; individual value functions can vary greatly.

2/3 probability that 600 people will die.

The majority choice in this problem is risk-seeking: the certain death of 400 people is less acceptable than a 2/3 chance that 600 people will die.

It is easy to see that the two versions of the problem describe identical outcomes. The only difference is that in the first version the death of 600 people is the normal reference point and the outcomes are evaluated as gains (lives saved), whereas in the second version no deaths is the normal reference point and the programs are evaluated in terms of lives lost. Because of the *S*-shaped value function and the overweighting of certainty the two frames elicit different preferences.

The decisions we have discussed so far involve a single dimension. Many decisions, however, concern transactions, in which the possible outcomes include connected changes in several dimensions of value. The basic example of a transaction is the purchase of goods after which one has more goods and less money than before. A transaction must be evaluated according to the balance of costs and benefits in a mental account.

The framing of a transaction can alter its attractiveness by controlling the costs and benefits that are assigned to its account, as in our example of the theater tickets. Going to the theater is normally framed as a transaction, in which the cost of the ticket is exchanged for the experience of seeing the play. If the ticket has been lost, buying a new one effectively doubles the cost of the play. In contrast the loss of the cash is not considered a debit to the account of the play. The cash loss affects the decision to buy a new ticket only to the extent that a minor reduction of wealth reduces the tendency to make optional purchases. When we asked a sample of students a version of this question involving a single lost ticket, the majority of those who were asked to imagine they had lost the money said they would buy a ticket, but the majority of those who were asked to imagine they had lost the ticket said they would not.

Mental accounting is dominated by the tendency to group the costs and benefits associated with an object, as the following example illustrates:

Imagine you are about to buy a jacket for $125 and a calculator for $15. The calculator salesman tells you that the calculator you want to buy is on sale for $10 at the other branch of the store, a 20-minute drive away. Would you make the trip to the other store?

The majority of the respondents who answered this question said they would drive to the other store. Another group answered a similar question in which the cost of the jacket was changed to $15 and the cost of the calculator was changed to $125 in the original store and $120 in the other branch. Of the respondents presented with this version, the majority said they would not make the trip. The total purchase and the consequences are the same in both versions: one has to decide whether to drive 20 minutes to save $5. The contrasting choices in the two versions suggest that the respondents evaluated the saving of $5 in relation to the price of the calculator. In relative terms a reduction from $15 to $10 is more impressive than a reduction from $125 to $120.

Framing effects in consumer behavior may be particularly pronounced in situations that have a single dimension of cost (usually money) and several dimensions of benefit. An elaborate tape deck is a distinctive asset in the purchase of a new car. Its cost, however, is naturally treated as a small increment over the price of the car. The purchase is made easier by judging the value of the tape deck independently and its cost as an increment. Many buyers of homes have similar experiences. Furniture is often bought with little distress at the same time as a house. Purchases that are postponed, perhaps because the desired items were not available, often appear extravagant when contemplated separately: their cost looms larger on its own. The attractiveness of a course of action may thus change if its cost or benefit is placed in a larger account.

If a decision is influenced by the reference point with which possible outcomes are compared, what determines the reference point? The dependence of impressions, judgments and responses on a point of reference is a ubiquitous psychological phenomenon. The same tub of tepid water may be felt as hot to one hand and cold to the other if the hands have been exposed to water of different temperatures. A given income may be considered lavish or inadequate depending on whether one's earnings have recently increased or decreased. In these cases the reference point is the state to which one has become adapted. In many cases, however, the reference

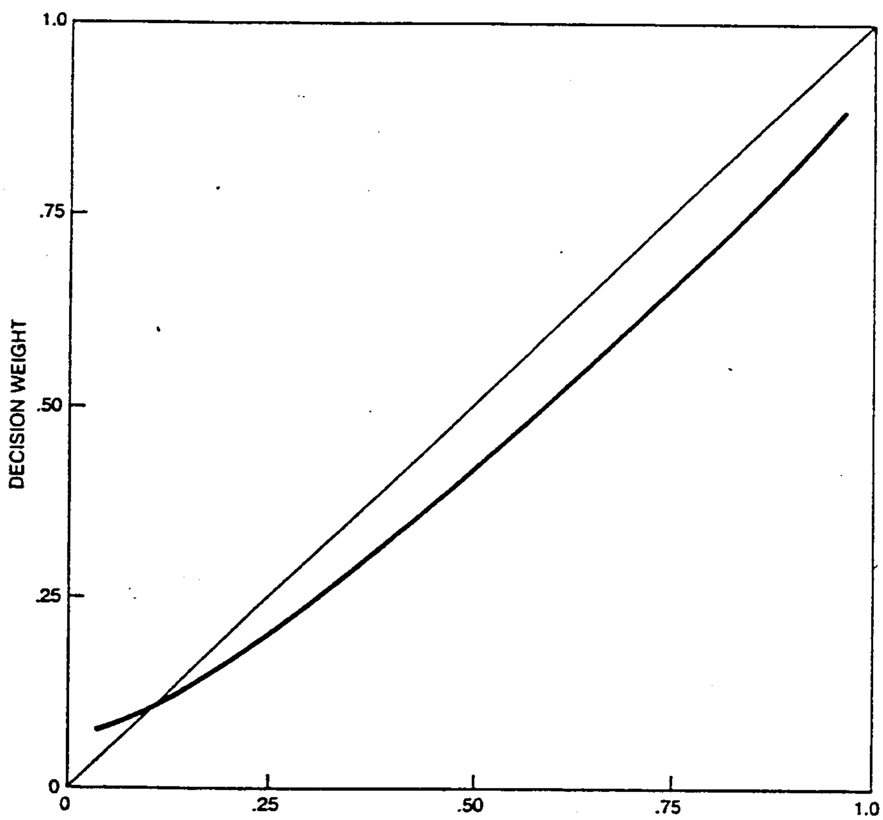

DECISION WEIGHTS express the subjective evaluation of probabilities. The classical analysis of decisions assumes that decision weights (*heavy line*) coincide with probabilities (*light line*). Recent investigations, however, suggest that decision weights do not coincide with probabilities. Because probabilities are used to derive the value of gambles, the subjective response to probabilities complicates the task of formulating a realistic value function. In particular, the difference between certainty and possibility and the difference between possibility and impossibility are given more weight than comparable differences in the intermediate range of probability. The main properties of decision weights are illustrated by the curve shown. Impossibility is assigned a weight of zero; certainty is assigned a weight of one. Small probabilities are overweighted in relation to impossibility; moderate and large probabilities are underweighted in relation to certainty. The overweighting of small probabilities contributes to the attractiveness of both lottery tickets and insurance policies by enhancing the impact of unlikely events.

point is determined by events that are only imagined. Consider the following incident:

Mr. Crane and Mr. Thomas were scheduled to leave the airport on different flights at the same time. They traveled from town in the same limousine, were caught in a traffic jam and arrived at the airport 30 minutes after the scheduled departure of their flights. Mr. Crane is told that his flight was delayed and just left five minutes ago. Who is the more upset?

Almost everyone presented with this problem agrees that Mr. Thomas is more upset than Mr. Crane, although their objective conditions are identical: both have missed their flight. Furthermore, both had expected to miss their flight, so that Mr. Thomas has no reason to be more surprised or disappointed than Mr. Crane. If Mr. Thomas is the more upset, it is presumably because in the act of imagination Mr. Thomas comes closed than Mr. Crane to catching his flight. The frustration experienced in an unsatisfactory situation increases when it is easy to imagine a more desirable alternative. For another example of the same notion, consider the following:

The winning number in a lottery was 865304. Three individuals compare the ticket they hold to the winning number. John holds 361204; mary holds 965304; Peter holds 865305. How upset are they respectively?

There is general agreement that the experience is devastating for Peter, quite severe for May and very mild for John. Here again the ranking corresponds to the degree to which the individuals can be described as having "come close" to winning the prize.

An individual 's experience of pleasure or frustration may therefore depend on an act of imagination that determines the reference level to which reality is compared. It is notable that of imagination by which one creates alternative

Regret is a special form of frustration in which the event one would change is an action one has either taken or failed to take. A natural extension of the hypothesis we have developed for the analysis we have developed for the analysis of frustration is that regret is felt if one can readily imagine having taken an action that would have led to a more desirable outcome. This interpretation explains the close link between the experience of regret and the availability of choice: actions taken under duress generate little regret. The reluctance to violate standard regret because it is easy to imagine doing the conventional thing and more difficult to imagine doing the unconventional one.

A closely related hypothesis is that it is often easier to mentally delete an event from a chain of occurrences than it is to imagine the insertion of an event into the chain. Such a difference in imaginability could help to explain the observation that the regret associated with failures to act is often less intense than the regret associated with the failure of an action. Consider the following:

Paul owns shares in Company A. During the past year he considered switching to stock in Company B, but he decided against it. He now finds that he would have been better off by $1,200 if he had switched to the stock of Company B. George owned shares in Company B. During the past year he switched to stock in Company A. He now finds that he would have been better off by $1,200 if he had kept his stock in Company B. Who feels more regret?

Here again it is generally agreed that George is more upset than Paul, although their objective situations are now identical (both own the stock of Company A), and each reached his situation by deliberate decision.

Apparently it is easier for George to imagine not taking an action (and thereby retaining the more advantageous stock) than it would be for Paul to imagine taking the action. Furthermore, on would expect both me to anticipate the possibility of regret and to act accordingly. In general the anticipation of regret is likely to favor inaction over action and routine behavior over innovative behavior. Our analysis traces these biases of decision to the rule of the cognitive operations by which alternative realities are constructed.

We have considered a variety of examples in which a decision, a preference or an emotional reaction was controlled by factors that my appear irrelevant or inconsequential. Some of these examples illustrate impediments to the achievement of rational decision that were discussed by Herbert A. Simon under the heading of "bounded rationality." The difficulty people have in maintaining a comprehensive view of consequences and their susceptibility to the vagaries of framing are examples of such impediments. The descriptive study of preferences also presents challenges to the theory of rational choice because it is so often far from clear whether the effects of decision weights, reference points, framing and regret should be treated as errors or biases or whether they should be accepted as valid elements of human experience.