CHAPTER 10

Decision Theory

[I]n our modern chess you must constantly be thinking of your opponent. "If I go here and then here, he will counter with that and I'm busted." . . . You don't, or at least you shouldn't make a move without considering what the reply will be.

—REUBEN FINE, *Chess the Easy Way*

"Because that's where the money is."

—BANK ROBBER WILLIE SUTTON, ASKED WHY HE ROBBED BANKS

In the last chapter, we talked about how goal-directed actions are implemented once a goal has been decided upon. Now let us ask: How do we decide what action to take, or what goal to seek?

Any action requires a decision (Fig. 10-1), if only between doing something and doing nothing. In this chapter, we will examine the *theory of decision-making.*

If reinforcement theory is the heart of modern behaviorism, decision theory is the mainstay of the cognitive approach to motivation. It is fundamental to economic theorizing, too, so that it provides a point of convergence between psychology and economics (pp. 369–71). It first arose from the mathematical analysis of gambling games, so it has close historical ties to statistics and probability theory. And, as we shall see, it can be applied to a surprisingly wide range of psychological phenomena.

A rich and important theory! Let us see how it works.

OVERVIEW OF DECISION THEORY

Consider how chess is played.* We examine the board and, for each move we consider, we imagine the resulting position and our opponent's possible replies. "If I go here and then here, he will counter with that and I'm busted"—that is, I lose. That's no good. However: If I check *there* with the knight, his only legal

*If you don't play chess, think of checkers or tic-tac-toe; the principle is the same.

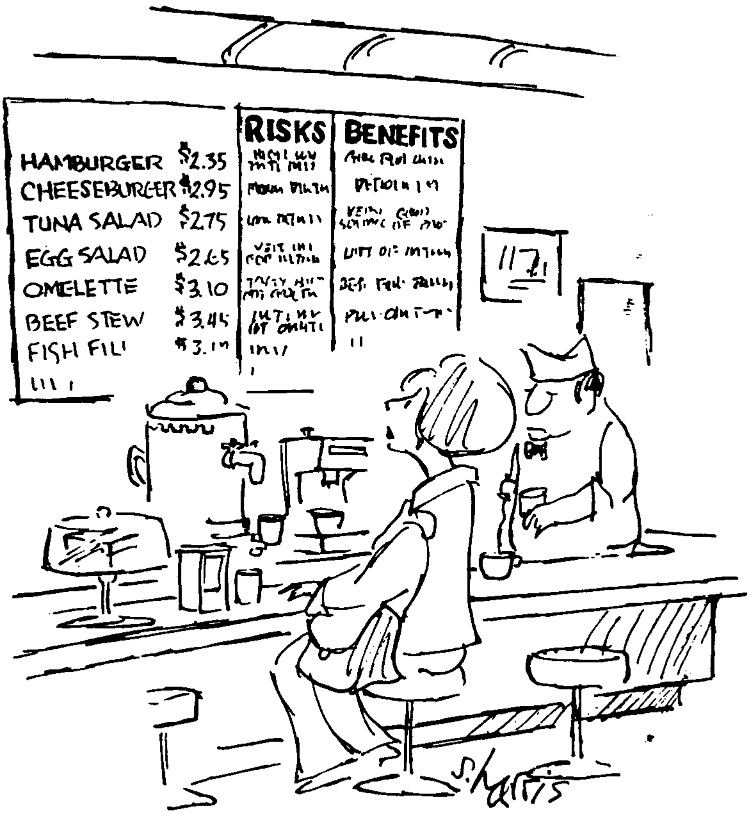

Figure 10-1 Any action involves a decision.

move is to put the king in the corner. Then my rook slithers down to the last rank, and it's checkmate. Isn't that nice! Do it.

Here we have decision theory in a nutshell. We imagine each possible action, consider its consequences, and perform the action whose expected or imagined consequences we most prefer.

We can do this on any of several levels. Given a goal, we can ask: Which of the possible actions is the best way to attain that goal? Here is a rat in a maze, and its goal is finding food: Should it turn left or right? Here we are in a strange town and our goal is to find our hotel: Should we turn left or right? My goal is to win this chess game: Shall I take the knight he offers or will it give him too strong an attack? In each case the question is: Which of the available actions is the best route to the goal?

Or, at a higher level, we can face the problem of comparing goals with each other. Stay home or go to a movie? Accept John's proposal of marriage, or George's, or neither? Become a psychologist, or go to law school?

But let's look at some specific examples.

Case 1: Outcomes Certain

Suppose we are in the presence of a very rich experimenter, whom we trust not to have any tricks up his sleeve. He says, "Would you like me to give you a dollar? Yes or no?" We will probably say "yes," for we get something of higher value for saying "yes" than for saying "no."

Similarly, if he says, "I'll give you either $10 if you shrug your shoulders, or $20 if you scratch your nose," we would of course scratch our nose. We choose the action whose outcome we most prefer. Simple.

Of course, in real life, choices can be very difficult even if outcomes are certain. Shall we take this apartment, with its pros and cons, or that apartment with *its* pluses and minuses? The problem is that each apartment is a mixed bag of the good and the bad. Indeed, it might be a good idea to "decompose" these alternatives, listing the features of each apartment, assigning positive and negative values to each feature, and adding them up to *calculate* which is the better choice.[1]—Or it might *not* be such a good idea, as we will see (pp. 394–95)!

But back to the rich experimenter.

Case 2. Outcomes Uncertain

Suppose he says this: "If you say 'yes,' I'll give you a dollar. If you say 'no,' either I'll give you $10 or I won't." Now we hesitate. We would rather have $10 than $1, so saying 'no' has higher possible value. But how likely is it that we will get that higher value? We would ask for more information; we need to know the *probabilities* of the outcomes, not just their values.

So the experimenter says, "Okay. If you say 'yes,' you get the dollar. If you say 'no,' I'll toss a coin: Heads I give you $10, tails I give you nothing." Do we take the dollar or go with the gamble? Probably we'd gamble. It might occur to us that a fifty-fifty chance at $10 would give us an *average* payoff, if the gamble were repeated many times, of $5 a game. That is the *expected value* of the gamble. Notice that it is not a matter of what we "expect" to happen, in the cognitive sense. It is a mathematical "expectation," and it means the *average* value that would result from a given choice if it were made many times in identical situations.

As exercises, let's consider some more things the experimenter might say. For each one, you should stop to calculate what your choice should be, before looking at the bottom of the next page for the correct choices.

1. "If you say 'yes,' you get $2. If 'no,' the odds are 1 in 10 that you'll get $5, and 9 in 10 that you'll get nothing."
2. In the previous example, make it $100 instead of $5.
3. In example 1, suppose the odds of getting $5 were fifty/fifty rather than 1 in 10.
4. "You can have $10 with probability 1/4, or $100 with probability 1/100. Which gamble would you prefer?"

1. See for example Dawes, 1988.

5. "If you like, you can play this game: Toss an honest die, and if it comes up 1 you get $30, otherwise nothing. Will you pay me $10 to play the game?"*

Rational Choices and Good Outcomes—Don't Confuse Them!

One more point before we move on. Suppose our friend Chris did accept the foolish gamble described in Choice 5. She paid her $10, rolled the die, and it did come up 1—so she received $30 on the deal. We're tempted to say, "Well, it wasn't such a dumb choice after all, right?"

Wrong. It was still a dumb thing to do. Look at it this way: If it were a smart thing to do, Chris should go on doing it, no? But if she did, she would become a money tree, paying on average $50 for every $30 she took in.

This may seem obvious, but in fact we do tend to judge decisions by their outcomes rather than by the rationality of the thought processes that entered into them. There may be important psychological reasons for this (compare the "representativeness heuristic," pp. 408–09), but it still is a mistake, and many injustices occur because of it. In an uncertain world, even a sound decision can turn out badly, while remaining sound.**

A Case Study: Of Crime and Punishment

Now let's leave our experimenter and look at the real world. A few years ago, the New York legislature was outraged, as many legislatures are today, at the traffic in illegal drugs that was sweeping the cities. The lawmakers responded by greatly increasing the penalties for such trafficking, so that long, long prison sentences awaited those convicted of drug dealing.

What was the effect on the drug trade? Virtually none.

Why not? Probably because the penalties, harsh as they were, had little chance—low *probability*—of actually being imposed. Only a few dealers were arrested; of these, only a few were convicted and sentenced. And even a very high penalty, or *cost,* becomes a very low *expected cost* if it has a very low probability of occurring.

It is for reasons like this that many criminologists suggest that we try to increase, not the severity, but the certainty—the *probability*—of punishment for crimes. That would entail, not longer sentences, but such things as more police,

*Here is what you should have done:

1. A return of $5 with probability 1/10 gives an expected return of $5 × 1/10 = $0.50. That is less than $2, so take the sure thing; say "yes."
2. A return of $100 with probability 1/10 gives us an expected value of $10, much greater than $2. Take the gamble and say "no."
3. The expected values are $5 × 1/2 = $2.50, as against $2. Again we should take the gamble.
4. $100 × 1/100 = 1, whereas $10 × 1/4 = $2.50. Take the $10 gamble.
5. Since the odds of rolling a 1 are 1 in 6, the expected return from the game is $30 × 1/6 = $5. Since that is much less than $10, you should refuse to pay $10 to play.

**For discussion see Dawes, 1988; Redelmeier, Rozin, and Kahneman, 1993.

more courts (to reduce the clogging of the court system which encourages plea bargains), and the like. To do that, it is argued, would increase the *expected cost* of crime much more than would an increase in penalties that probably won't be imposed.

The Maximization Principle

Now let us generalize these ideas.

Decision theory assumes that we perform the action whose imagined outcome is of highest value to us, or, as we say, highest *utility*. Saying it another way, we choose among actions so as to *maximize utility*.*

Now for many decision problems, since the outcomes haven't happened yet, we cannot be certain what they will be. Then we have to consider how *likely* we think the various possible outcomes are. What we do, the theory says, is maximize the mathematically *expected utility*. For each action we might perform, we consider all the possible outcomes; for each outcome, weight (that is, multiply) its utility by its probability; and add the resulting products to get the *expected utility* of that action. Then we take the action whose expected utility is highest.

In many real-life situations, it is not clear that we do the kinds of explicit *calculations* that our earlier exercises permitted. Even then, however, decisions do depend on two factors that we feel could be represented by numbers: the utilities—how much do we like or dislike each possible outcome?—and the probabilities—how likely is it that a given outcome will follow a given action? If these probabilities are known, then numbers can represent them. When they are not known, still one usually has some idea of the likelihood of various events, and so we can speak of estimated probabilities—or, as they are called, *subjective probabilities*—as entering into our calculations. Multiplying the utility of each outcome by its subjective probability gives us its *subjective expected utility,* and we can choose so as to maximize that.

A Look Backward

The important thing to see is how our decisions are affected jointly by the *likelihoods* of outcomes (probabilities or subjective probabilities, a cognitive factor), and their costs or values to us (their utilities, a motivational factor). Increasing one has the same effect as increasing the other (when two numbers are multiplied, increasing either one will necessarily increase the product). There can also be trade-offs between the two; increasing one can compensate

*The term *utility* is borrowed from economics. We need such a term to describe the value to us, the "niceness," of various outcomes. Words like "niceness" or "preferredness" are awkward. We don't want to call it "value," for that connotes monetary value, which may not be what we mean. Nor do we want to call it "pleasantness," for if we must decide between raking the leaves and cleaning the sink, pleasure may have very little to do with it. The term "utility" sounds strange, but at least it is neutral.

for decreasing the other. In the above set of problems, we made different decisions in Cases 1 and 2, because the utilities were different, and in Cases 1 and 3, because the probabilities were different.

Finally, I wonder if the reader is experiencing a tingle of recognition here. Our theory bears a distant family resemblance to Hull's algebraic theory of behavior (pp. 202–08)! As with Hull, our tendency to take a given action is seen as dependent on the two factors. One is motivational—*D* (Drive) for Hull, utility here. The other is non-motivational, but based on learning or cognition—*H* (Habit) for Hull, subjective probability here. And the theories even agree in treating the two factors as multiplying to determine the final outcome! The difference is that Hull's reaction potential (*E*) is seen as driving behavior directly. In contrast, the subjective expected utility of an action is a quantity that is *considered*—compared with those of other possible actions—before the choice among actions is made.

A Look Forward: The Problem of Risk Perception

Perhaps this is the place for a word of caution. The fact is that human beings do not always behave in the crisply logical way that describes our experimenter's hypothetical subject—or the real-life drug dealer! For example, we may worry very seriously about risks that are actually minor—refined sugar in our diet, for instance—while ignoring such risks as driving without seat belts, a practice whose effects on our well-being are much more likely (probability) to be very severe (utility).[2] This looks forward to the problem of *availability* in decision-making—a problem for any logical, rational theory of human behavior.

EXTENSIONS AND APPLICATIONS OF DECISION THEORY

Decision theory and its offshoots have provided a wealth of ideas to explore. In what follows, we will look at how it can be applied in a number of different contexts. That can give some feeling for its richness.

Utility Theory

First, let's look more closely at the notion of value or utility.

THE UTILITY OF MONEY In many economic situations, decisions involve the exchange of money or of commodities with monetary value. Can we use the monetary value of a commodity as a measure of its utility? Not

2. Rozin, 1989.

really. It is easy to show that the *value* of money, its utility, is not related to the *amount* of money in any simple way.

Suppose we were in the presence of a *really* rich experimenter, who said to us: "I will give you $5,000 outright, as a gift. Enjoy." Let us agree that that would make us happy.

Now imagine that he said, "No, better still, I'll give you $10,000." That would please us even more—of course! But would it make us twice as happy as the $5,000 gift? Happier, yes; but *twice* as happy? Most people think not.

In fact, when this thought experiment is actually done, a typical reply is that the gift would have to be *quadrupled* in order to double the value of the happiness it would bring.[3] In other words, utility increases as the square root of the dollar value. People differ to some extent in these judgments, of course, but this much is consistent: *The utility of money increases less rapidly than the dollar amount does.*

To show this idea and its consequences, Figure 10-2 presents a graph of the utility of money—its value to us—against the amount of money in dollars. Such a plot is called the *utility function* of money. Notice that it is concave downward, so that as the dollar amount rises, the increase in utility gets less and less. If the increase in utility drops off rapidly, then $10,000 will be less than twice as valuable to us as $5,000 is, as shown.

This bent utility function shows the well-known economic principle of *diminishing marginal utility*. Presumably it applies to any commodity; more of anything brings us less added happiness or utility, when we have more of it to begin with. This simple idea illuminates some important human actions.

UTILITY FUNCTIONS AND EXCHANGES Suppose we buy a loaf of bread from a grocer for a dollar. What makes this exchange possible is that it produces a gain in utility for both participants. We would rather have the bread than the dollar. The grocer would rather have the dollar than the bread. So we swap, and everyone is happier.

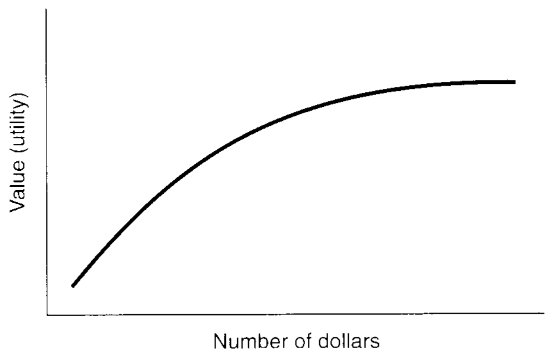

Figure 10-2 A utility function for money. The value or utility of money increases more slowly than its dollar amount does, and the increase in value gets less and less as the dollar amount increases.

3. Galanter, 1962.

But why? A loaf of bread is not *always* more valuable to us than a dollar is. What determines whether it is or not? How much of each commodity—bread and money—we already have. That determines where, on the utility function for each commodity, we and the grocer find ourselves.

Suppose we have plenty of money in our pocket, but little or no bread in the house. Then our position is high on the utility function for money, where the slope is small and a dollar gained or lost doesn't mean much. But we are low on the utility function for bread, so a loaf gained means a lot. In that case, we would buy the bread, accepting a small loss in utility (the dollar) for a big gain in utility (the bread).

But now suppose the converse case. We have plenty of bread in the house, but little money. Now we are low on the utility function for money, where the curve is steep; a dollar lost is a substantial loss in utility. We are high on the utility function for bread, where the curve is shallow; one more loaf wouldn't give us much satisfaction. Under those conditions we would not pay a dollar for a loaf of bread. We might even sell one of our loaves ourselves.

The grocer of course is in a different position. He, we assume, always has more bread than he wants for himself; that's what makes him a grocer. For that reason, a sale is always a gain in utility for him.

The important point to see is that the bent utility function—the principle of diminishing marginal utility—underlies many cases of sale or barter. If we consider how much of human affairs depends on such exchanges, we begin to see just how important a principle it is. It makes the world go round—or so it has been said.

Signal-Detection Theory

To this point, we have considered outcomes that follow actions with probabilities that are fixed. But suppose that the outcome of an action depends on something else—some X that may or may not be true, and we are not sure whether it is true or not. Such decisions are the province of *signal-detection theory*.

Examples abound. Consider:

1. The nation on our border is mobilizing its armies. Does it intend to launch an attack on us? Shall we mobilize our own forces? If we do, and if our neighbor does intend to attack, we will be ready to meet the attack (good). But if it doesn't, we will have gone to needless expense and been needlessly provocative (bad).
2. In dim light, a robber takes your wallet at gunpoint. At the police station, you are shown a photo of a known holdup person, and the police say to you, "Is this the robber?" If you say yes, and you are right, you have helped to catch a robber and you might even get your wallet back (good). But if you are wrong, you will have fingered an innocent person (bad).

3. Suppose the accused person is brought to trial, and each juror, having heard all the evidence there is, must now vote "guilty" or "not guilty." Same dilemma: A vote of guilty, if the accused *is* guilty, is a vote to punish the guilty (good); but if he is not, then it is a vote to condemn the innocent (bad).
4. We are chatting with a fellow student whom we do not know. Shall we invite him or her out for coffee? If we do, then if the stranger is receptive, we will have a new friend, perhaps a new relationship (good); but if not, we will be rejected (bad).
5. Suppose we are sitting in an experimental chamber with headphones on, and a constant white noise (a kind of shooshing sound) is being fed into our ears. On each trial, we get a cue that tells us to listen carefully; and a "signal"—perhaps a very, very faint tone—either is, or is not, added to the noise. On each such trial, we must decide whether the signal was presented or not. If we say "Yes, there was a signal on that trial," and there was, we are correct (good); but if there wasn't, then we are incorrect (bad).

It is from that kind of experiment that signal-detection theory gets its name. So let us look at that experiment more closely.

On each trial, there are four possible outcomes, because there are two possible states of affairs (a signal was presented or it wasn't), and two possible responses: "Yes, I think there was a signal," or "No, I don't think there was a signal" (Fig. 10-3). We assume, remember, that the signal is so faint that the subject *cannot be sure* whether it is there or not, and so she will make mistakes. And there are two kinds of mistakes she can make: She can say no when the answer is yes (a "miss"), or yes when the answer is no (a "false alarm").

Now, presumably the subject wants to perform accurately, and would rather not make mistakes. She places some positive value on right responses and some negative value on wrong ones. Suppose we take experimental control over these values, by explicitly paying her for right responses and fining her for mistakes. We set up what is called a *payoff matrix* (Fig. 10-4). In the first of

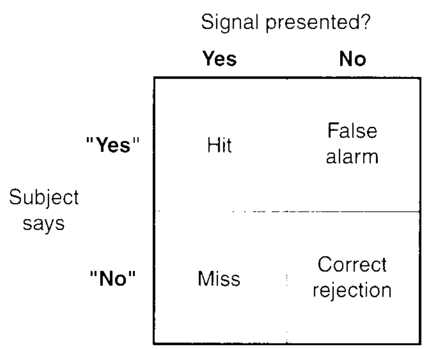

Figure 10-3 Possible outcomes for a signal-detection trial.

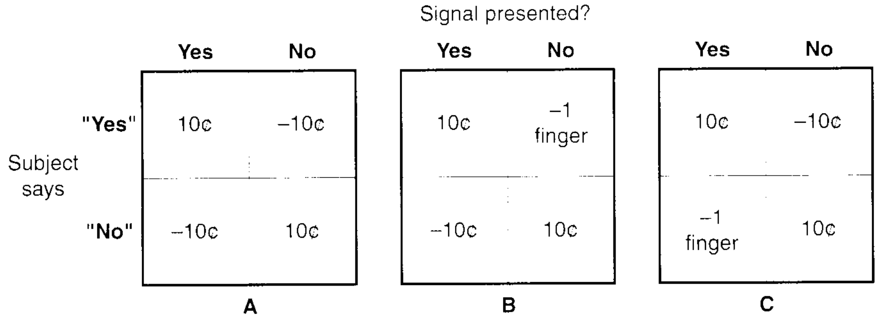

Figure 10-4 Three possible payoff matrices in a signal-detection experiment. As compared with the symmetrical payoff matrix A, matrix B strongly encourages a "no" response, and matrix C strongly encourages a "yes" response.

these (Fig. 10-4A), the payoffs are symmetrical: a mistake of either kind carries the same cost.

But we could just as well set up the payoffs as in Fig. 10-4B). A miss carries only a trivial cost of 10 cents. But if the subject ever makes a false alarm—saying "yes" when she should say "no"—we chop off a finger, then and there! (Note to human subjects committees: This experiment is imaginary.) What will happen? Obviously, the subject will almost never say "yes." The cost of being mistaken in doing so is simply too great. So, over a series of trials, she will have very few false alarms (good), but also very few hits (bad).

Now, after a series of trials like that, suppose we change the payoffs—and tell the subject we are doing so. Suppose we set them as in Figure 10-4C. Now it is a false alarm that carries only a small penalty, but a miss—failing to say "yes" when she should—now costs a finger each time it happens. Obviously, the subject will now say "yes" on nearly every trial. She will have many hits (good), but many false alarms too (bad).

So, by changing nothing but the payoff conditions, we can induce the subject to say "yes" on every trial, or on none of the trials. It would seem likely that less extreme variations in the payoff conditions could produce less extreme response strategies, so that we could place the likelihood of saying "yes" at intermediate values. And that is exactly what happens when such experiments are done (Fig. 10-5).

Notice that all this has nothing to do with how *accurately* our subject is performing. It has to do only with how inclined she is to say "yes" as opposed to "no" when in doubt. But we can get from the experimental data a measure of the subject's accuracy, too. That will be shown by *how much the hit rate exceeds the false-alarm rate* (Fig. 10-6). To the extent that the subject more often says "yes" when she should than when she shouldn't, she is performing accurately. Notice that the frequency of "hits," by itself, does not give us that information. A high hit rate—even 100 percent—might reflect accurate performance; or it might mean only that the person is saying "yes, yes, yes" a lot, whether the signal is there or not. False-alarm data will allow us to distinguish these.

It is for reasons like this that in the robbery scenario, the police ought not show us a *single* photo and ask, "Was that the criminal?" Our answer depends too much on our inclination to say "yes" when in doubt, and thus too little on

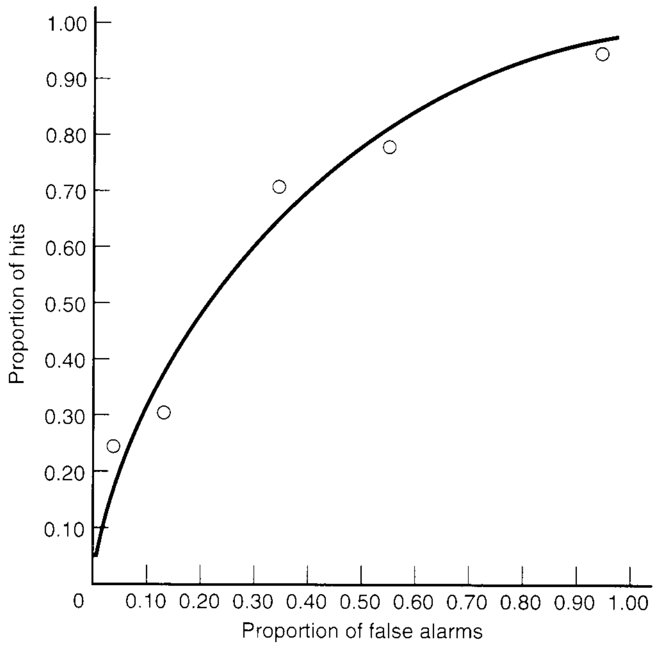

Figure 10-5 Hit rate plotted against false-alarm rate. Each circle represents the outcome of a series of trials under a given payoff matrix. Payoffs that favor a "yes" response—high reward for hits, low cost for false alarms, or both—will lead to high rates of both hits and false alarms, as in the circle high to the right. Payoffs that discourage a "yes" response—low reward for hits, high cost for false alarms, or both—will lead to low rates for both, as in the circle low to the left. Less extreme payoff conditions will lead to less extreme performance, as for the remaining circles. (Data from Green and Swets, 1966.)

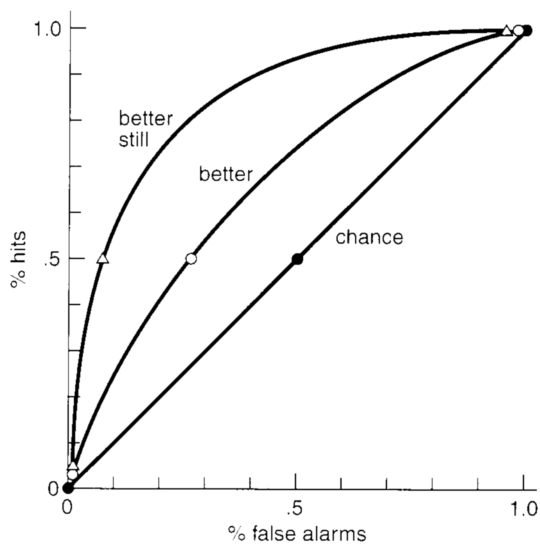

Figure 10-6 Data for hypothetical signal-detection experiment. Notice that the hit rate may be the same (here, .50) for data points on all three curves, but differences in the false-alarm rate allow us to distinguish between very poor, chance-level performance (filled circles) and a high level of performance (triangles).

whether the photo really does or does not show the criminal. The purpose of police lineups and the like is to give the viewer a chance to make mistakes (false alarms), and thus allow the police to distinguish an accurate witness from a yes-person.

And that example allows us to take these ideas back to the real-world scenarios with which we began this section. We see that in every one of them, our decision problem comes down to this: What are the payoffs for various outcomes? More specifically: *Of the two kinds of mistakes we could make, which is worse, and how much worse?* Is it worse to mobilize our armies when we don't have to, or to fail to mobilize when we're about to be attacked? To punish the innocent, or to let the guilty go free? To miss an opportunity to make a friend, or to be put down and rejected?

As an exercise, the reader should apply this kind of analysis to the various scenarios with which this section began. Let's look more carefully at Scenario 4, for example. Which would be worse, to miss an opportunity for a friendship, or to be rejected? If we are very much hurt by rejections, then to ask and be rejected (a "false alarm") might be the worse kind of mistake, carrying a high cost (Fig. 10-7A). That would make it likely that we'll refrain from asking. Our decision would be different if we're not bothered by a turndown (Fig. 10-7B) *or* if the other person is the person of our *dreams*, so that a missed opportunity would be even worse than a rejection (Fig. 10-7C). In either case we would be likely to go ahead and make the offer, with high hopes.

(A question for the interested student to ponder: How much of what we call "personality characteristics"—shyness, for example, or introversion, or sensitivity to rejection—can be seen as depending on the costs and values we assign to various outcomes of our social decisions?)

Finally, in this example we have focused on the *motivational* component of the theory—the payoffs for various outcomes—for that, after all, is our topic. But the *cognitive* component will be just as important—how *likely* do we think we are to be turned down? This will depend on our beliefs about other people: do we think most people are friendly, or cold or hostile? It will also depend on our beliefs about ourselves: Do we think that we are reasonably attractive and acceptable, or that we are unattractive and likely to trigger rejection? A lot of ideas come together in this "simple" analysis!

Before going on, readers should pause to convince themselves that all this does make sense. The numbers in Figure 10-7 should not be taken literally—we made them up, after all—but our decisions do consider "quantities" that represent probabilities and the costs and values of various possible outcomes. Which outcomes are likely or unlikely? What outcomes are good, and how good? Bad, and how bad? That is true of our social decisions as of any others (Fig. 10-8).

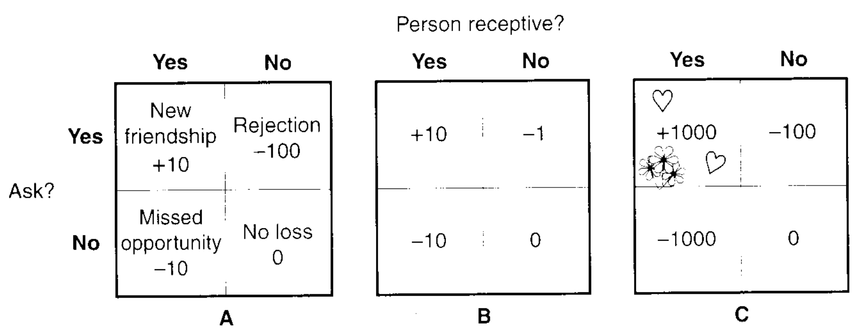

Figure 10-7 To ask or not to ask? Three possible payoff matrices.

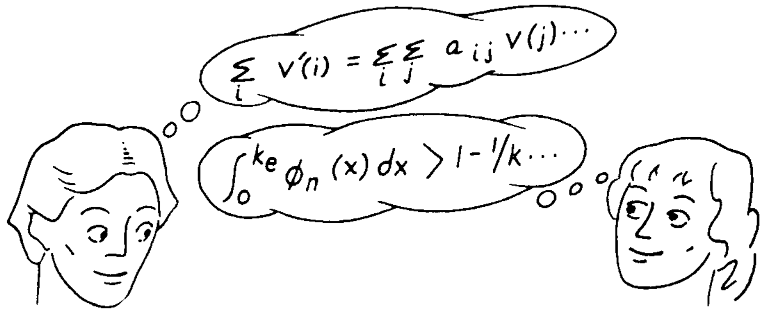

Figure 10-8 Social decisions require some quick intuitive calculations.

These examples can only give you the flavor of an extensive literature on signal-detection theory. Whole books have been written just on the psychology of eyewitness testimony, as in Scenario 2 above,[4] and in them, the concepts of signal-detection theory figure very prominently indeed.

The Costs of Rationality

We mentioned earlier that a rational, utility-maximizing decision is not to be confused with a good outcome. So it is only right to note that there are cases in which strict, utility-maximizing rationality can take us places we would not wish to go. It can, in a word, lock us into decisions whose outcomes are not good for us or for anyone else.

THE TRAGEDY OF THE COMMONS In 1968, economist Garrett Hardin published a thoughtful article entitled "The Tragedy of the Commons." Imagine a pasture that everyone in a community can use as a place where his cattle can graze. That pasture is a *commons,* shared by all the herdsmen.

Now it is only rational for each herdsman to try to maximize profits. How many animals should he have? Well, he may ask at any point: What is the utility to me of adding one more animal to my herd?

Adding another animal involves both gain and loss. The gain is the profit from the animal that is added. We assume that the herdsman gets all of that. There is also a long-term cost, which results from the effects of overgrazing. But that cost is *shared by all the herdsmen there are.* The cost to each individual herdsman is divided by the total number of herdsmen; and so, for each individual, it is low.

Thus it may be to the advantage of *every* herdsman to add another cow, and another, ... Each herdsman gets all the profit, and only a fraction of the penalty, by doing so. "Each man is locked into a system that compels him to increase his herd without limit—in a world that is limited."[5] And this happens, not out of wickedness or even shortsightedness, but simply because each herdsman rationally seeks to maximize utility.

4. Loftus, 1979.
5. Hardin, 1968, p. 1244.

As another example, consider the problem of pollution. "The rational man finds that *his share* of the cost of the wastes he discharges into the commons [for example, the rivers and lakes] is less than the costs [which fall on him alone] of purifying his wastes before releasing them."[6] If so, it is rational for him to pollute rather than purify.

So, by behaving as a rational decision maker, each person perpetuates a system in which everyone loses in the long run. There ought to be a law against such decision situations! And, in fact, there are such laws: legislation about pollution or about land use is an attempt to break up just such traps as these. Stiff fines for pollution, for example, are costs imposed on *individual* manufacturers rather than spread over all of them. They are intended to make it *rational* for each individual to refrain from polluting.

THE PRISONERS' DILEMMA GAME A "game," in decision-theory parlance, is any situation in which the outcomes depend jointly on what two or more actors, or "players," decide to do. And there's a particularly nasty kind of game known as the "prisoners' dilemma." It gets its name from the following scenario:

Two players, Smith and Jones, are arrested and charged with (say) bank robbery. At the police station, they are separated from each other, and each is told the following:

> "Look, Smith [or Jones], we don't have enough evidence to convict you of the robbery. We need a confession. If you each confess, you'll get ten years in prison. If neither of you confesses, we'll convict you of a minor charge and you'll each get one year. But if one of you confesses and the other doesn't—then we'll grant immunity to the one who confesses, and he'll get off scot free. But we'll throw the book at the other—thirty years in prison. And we're telling your partner the same thing, in another room, even as we speak."

What should the prisoners do? Figure 10-9 presents the payoff matrix for the game. As we look it over, we recognize a horrible truth: *No matter what Jones does, Smith is better off confessing!* To confess, therefore, is clearly his optimal strategy. And the same is true for Jones. So if both play rationally and confess, they'll each get 10 years in prison, where they could have held it to 1 year for each by remaining silent.

Imagine Smith's agony, even if he is an honorable thief:

> "It would be best for us both to keep silent; a year in prison isn't so bad. But what a reward I might get for confessing, if Jones doesn't! Worse still: What a risk I'd be taking if I didn't confess! If I don't confess and Jones does, I'll spend my best years in durance vile [Smith is a literary thief as well]. Moreover, Jones has the same pressures to confess as I do; and I know that, which adds to my fear that he'll confess; and he knows *that*, which adds to his fear that *I'll* confess, and so he has all the

6. Ibid., p. 1245, my italics.

Jones \ Smith	Confesses	doesn't confess
Confesses	Smith 10 years Jones 10 years	Smith 30 years Jones free
doesn't confess	Smith free Jones 30 years	Smith 1 year Jones 1 year

Figure 10-9 A prisoners' dilemma payoff matrix.

more reason to protect himself by confessing. No, I'd better play it safe and confess."

Jones, we may be sure, is saying the same things to himself.

In short, logic inexorably leads to a joint decision whose consequences—10 years in prison for each—is far from the best possible outcome for either player.

We might suppose that it would be a good thing if the two prisoners could agree, beforehand, not to confess. But that would only bounce the problem back a step, for then each must decide: Shall I keep the bargain or break it? By the same logic, to break the agreement is necessarily more rational than to keep it.

Real-world examples of the prisoners' dilemma are very much out there. Watching a ball game, we would all be more comfortable sitting than standing. But if other people stand, why, we ourselves must stand as well in order to see at all! Everyone ends up standing in self-defense. Hardly the most comfortable outcome—but even worse, for us, would be to remain sitting while others stand. (President Jimmy Carter once pointed out that exactly the same considerations can drive inflation. No one likes inflation. We'd prefer it if prices and wages were stable. But if other people get higher wages and we don't, then prices will rise and we'll end up behind; so we'd better press for higher wages too. And if others reason the same way, ...)

Or consider arms races. Two nations are wasting enormous sums of money on armaments, each to protect itself from the other. Each would prefer to disarm and spend its wealth on something better. But either nation could be in deep trouble if it disarms and the other does not.

A LOOK BACKWARD—AND FORWARD Like the tragedy of the commons, prisoners' dilemma–type payoff matrices can make the rational choice a very bad choice—even for the actor himself. The important thing to see is that these unfortunate choices do not result from malice or shortsightedness. They result from the simple logic of utility maximization.

From a theoretical point of view, it is significant that people do not, in fact, always do the coldly rational thing in such situations. We refrain from despoiling nature, out of respect for posterity—even though, as the saying goes, "What did posterity ever do for us?" People do, under some conditions, cooperate (with the other player, not with the police!) in prisoners' dilemma games.

In short, there may be times when (1) we do not in fact behave as purely rational maximizers, and (2) it is just as well that we do not! Keep that idea in mind; we'll return to it.

DECISION THEORY AND EVOLUTION

Animals, like people, face choices. A stickleback, faced with a rival, must fight or run away. A ram must court this mate—or not. A bird can look for food right *here* or right *there*. Can decision theory help us understand these cases too? Yes, in a way. It can help us understand how motivational mechanisms evolved—how they came to be the way they are.

Evolution and Maximization: The Progeny Payoff

The basic idea is that how a creature spends its time and energy will affect its *reproductive success*—the number of descendants it has (pp. 34–37). What evolution is all about is: Who gets to be an ancestor? If, within a population, those with long necks are ancestors to more descendants than those with short necks, then there will be more long-necked descendants in later generations. Presently living animals bear the genes of those, and only those, who succeeded in becoming ancestors.

Thus, the number of descendants a creature has, its reproductive success, represents a kind of utility—a progeny payoff—for allocating its resources in the way it does. A present organism's repertoire of action patterns, releasing stimuli, unconditioned reinforcers, and other inherited behavior mechanisms can therefore be seen as ones that led to high reproductive success on the part of its ancestors. We can ask: What payoffs in progeny led our ancestors to evolve *these* mechanisms rather than others?

Before we get into examples, however, we must pause to get one thing clear. To this point, we have been using decision theory as a *cognitive* theory of choice. We've assumed, in other words, that the actor is in fact doing the relevant calculations (utility × probability) and basing choices on the results—or, at least, doing something very much like that.

In this section, we will use the theory in a different way, a purely descriptive one. If we say that Actor A is maximizing utility, we will mean by it only that Actor A is doing something that *has the effect* of maximizing utility—whether the actor knows anything about it or not. In other words, we will be describing the *outcomes* of actions, not their causes—and often the "actions" will be those of ancestors long dead, not those of any actors who are onstage now.

But even when used in this descriptive, behavioristic way, decision theory can give us a handy way of describing what goes on. Let us look.

Foraging and Food Selection

What should an animal do about eating? Where is it most likely to find food? How much of each prey type should a predator eat? These are decisions an animal must make, or behave as if it made.

Consider the shore crab. These creatures feed on mussels, whose shells they must break open to get at the morsel of meat inside. The meat, when eaten, provides energy. But breaking open the shell *costs* energy. Actual calculations (by humans, not by crabs!) show that mussels of intermediate size will maximize the benefit/cost ratio for the crab. Smaller mussels yield too little energy; larger ones cost too much. The benefit/cost ratio is shown by the solid curve in Figure 10-10. And the bar graph below shows that the food selections actually made by crabs, when offered a choice, fit that curve quite well.[7]

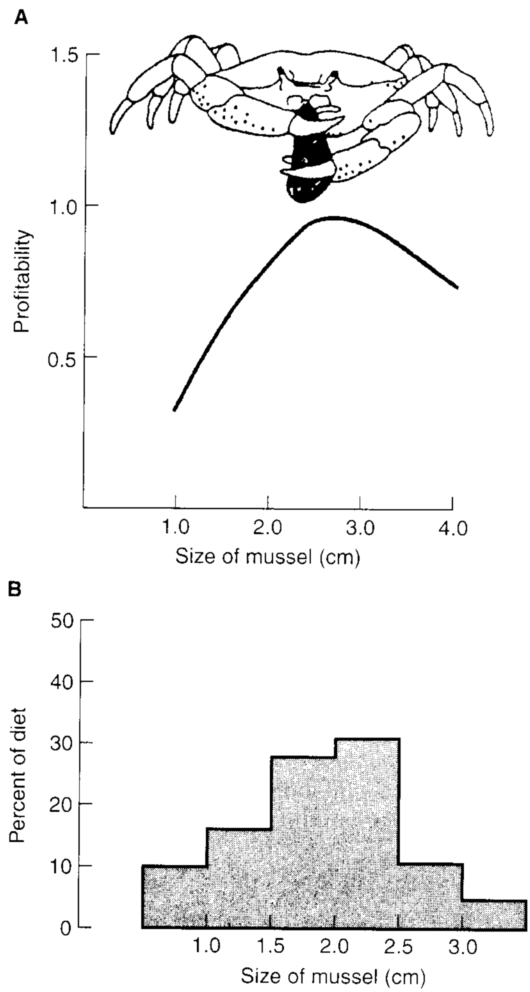

Figure 10-10 An example of optimization. Shore crabs *(Carcinus maenas)* prefer to eat the size of mussel that gives the highest rate of energy return. *(A)* The curve shows the calorie yield per second of time used by the crab in breaking open the shell. *(B)* The histogram shows the sizes eaten by crabs when offered a choice of equal numbers of each size in an aquarium.

7. Adapted from Rozin and Schull, 1988.

But we see at once that the *cognitive* theory of decision-making does not apply at all to this case. No one is suggesting that the crabs do the relevant calculations in their heads—crabs are lousy at arithmetic. Rather, what presumably happened is this: Over the generations, ancestral crabs left more descendants if their nervous systems were "calibrated" so that prey of near-optimal size were the most potent releasing stimuli for feeding, because crabs so calibrated wasted as little energy as possible in feeding themselves. And today's crabs are so calibrated because they *are* those descendants.*

Mate Selection

Back in Chapter 5, when we discussed the Coolidge Effect (pp. 125–31), we pointed out that the way for a male to have the most descendants—to maximize his progeny payoff—is to mate with as many females as possible.

Not so the female. She has only so many possible children, whether she stays or strays. And in some species such as most mammals and birds, she has a substantial investment in each one; she eats for it before it is born and feeds it after it is born. The way for the female to maximize reproductive success, with the limited number of offspring she can have, is to see that as many as possible of her few children survive to adulthood after they are born.[8]

All this is not so bad if the female can raise infants by herself. She can then have about as many surviving children if the male deserts her as if he stays around to participate with her in child care. But what if raising children really requires both adults? In many birds, for instance, keeping a nestful of babies fed is a full-time job for both parents. Unless both parents participate, the young will die.** And in many such species, monogamous male-female pairs are formed. All the babies would die if the male did not do his share—and neither male nor female would leave any descendants.

But if there is *any* reasonable chance that the female can do the job of child-rearing by herself, then the male might still do better to desert! Even if most of the offspring die, still, a male can father so very many children that his mathematically *expected* number of descendants—number of babies times survival rate—could be higher if he mates with many females than if he stays to raise the young with one.

Now the female is not helpless in the face of this vicious payoff structure that favors male fickleness. What she can do—and in many species does do—is *refuse to mate* unless the male shows that he is likely to stay around to raise the young with her. She might, for instance, require him first to build an elaborate house for her, a nest or bower, before she will accept him as a mate. Or she

*For simplicity, we assume that selection is instinctive in this case, though in fact, recent experience affects it as well. For discussion and further examples see Rozin and Schull, 1988.

8. Trivers, 1971.

**Notice that successful ancestry requires not only our having children, but also having them survive to give us grandchidren and thus keep us going as ancestors.

might insist that he first establish and defend a territory, driving other males away (pp. 112–13).

Notice that setting such tasks for the male is doubly to the female's reproductive advantage. First, they make it more likely that the male she mates with will be healthy and strong. Second, they *reduce the expected utility of desertion for the male*. If all the females of the species behave like that, then the male cannot just fertilize when he feels like it. If he leaves his mate to seek a new one, he will have to begin the whole business of courtship all over again with her. This being so, it may now maximize his progeny payoff to have fewer babies, with one mate, and work hard to make sure that those few survive.

The female's finickiness about mates, in such species, must have evolved because finicky females had more descendants than non-finicky ones. And, in a population of finicky females, faithful males had more descendants than the non-faithful ones, who were forced to waste too much time and effort in multiple courtships. We can say that finickiness and fidelity, for females and males respectively, maximized reproductive success among such birds' ancestors.

But notice that again we use *maximize* in a strictly descriptive sense. No one is suggesting that the birds reason all this out. Rather, these courtship rituals evolved simply because the successful female ancestors were those who did require them as releasing stimuli for their own sexual receptiveness. Thus they acquired faithful mates, and a higher grandchild tally as a result; and the multi-great-grandchildren are those that are with us now.

These brief glimpses of an active research area should further emphasize the broad domain of the maximization concept and the theory of decision-making. As a way of looking at the problems an animal faces, and as a way of describing how it solves them, the idea has applications far beyond the cognitive domain in which we first encountered it. In the next section, we will see still more.

CHOICE AND REINFORCEMENT

In the last section, we saw the convergence between decision theory and the study of instinctive repertoires by ethologists. But what about learned repertoires? In this section, we will again see convergence, this time between decision theory and the laboratory study of operant behavior.

Concurrent Schedules

Consider a pigeon in a Skinner box, confronted with a pecking key (pp. 295–96). Suppose that at irregular intervals, food reinforcement is set up on a given key so that the next peck on that key is reinforced with food. This is a *variable-interval* (VI) reinforcement schedule. Nothing indicates to the bird when reinforcement is available; any peck could be reinforced as far as the bird can tell, and so the pigeon pecks away at a steady rate.

Let us add a complication: another pecking key. Now there are two response keys, and reinforcement can be programmed at variable intervals independently on each response key. The availability of reinforcement on one key has nothing to do with whether or not it is available on the other. And there is still no signal to tell the bird when either key is set up to deliver reinforcement. Now we have *concurrent variable-interval schedules* of reinforcement.

The Matching Law

Now let us ask: What is the effect of different rates of reinforcement on the distribution of pecking behavior? Richard Herrnstein ran a series of sessions using concurrent VI schedules with pigeons.[9] The total number of reinforcements available in an hour-long session was held constant at 40; but Herrnstein varied the proportion of reinforcements that could be obtained by pecking each key. For example, on some sessions each key might be programmed to yield 20 reinforcements per hour. In other sessions, one key might be programmed to offer 10 reinforcements per hour and the other key 30 reinforcements per hour; in yet other sessions, the two keys might offer 5 and 35 reinforcements per hour. And so on.

Figure 10-11 shows some striking findings. The bird's allocation of pecks exactly matches the ratio of reinforcements available on the two keys.

It doesn't have to be that way. A pigeon could obtain all the reinforcements available, with *any* distribution of pecking behavior over the two keys, as long as it pecked each key at least occasionally. In fact, however, the bird distributes its pecks to match the relative frequencies of reinforcement, and with

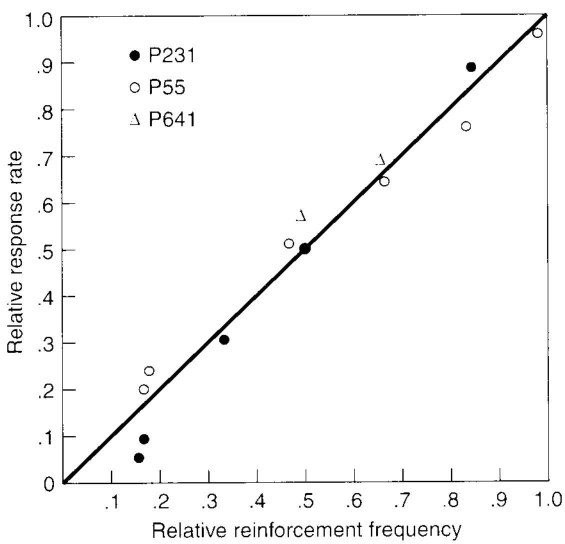

Figure 10-11 Pigeons distribute pecks between two keys to match the relative rates of reinforcement on the two keys. Each symbol represents an individual pigeon. (From Herrnstein, 1970.)

9. Herrnstein, 1970.

remarkable precision. If one key pays off three times as often as the other, the bird pecks that key three times as often as the other; if twice as often, twice as often; and so on. That is the *matching law*.

The matching law turns out to have considerable generality. It holds for key-pecking in pigeons under a variety of conditions. It is also seen in bar-pressing experiments in rats, and with button presses or even eye movements in human beings.[10]

The Law of Relative Effect

What the matching law says is this: The effect of a reinforcer on response rate is not absolute. Rather, its effect is relative to the effects of other reinforcers available in the situation. A given reinforcer, on a given schedule, will have strong effects if few other reinforcers are available. If many are available, it will have weaker effects. For that reason, the matching law is sometimes called the *law of relative effect*. It turns out to have some intriguing implications.

For example, it follows from the law that we could reduce responding to Key A in either of two ways: (1) by reducing the rate at which we reinforce it, or (2) by *increasing* the rate at which we reinforce some alternative, such as pecks at Key B. This is exactly what happens.

The same idea has been extended to behavior therapy in humans (pp. 297–99). For example, there is the case of a mildly retarded male who showed outbursts of assaultive behavior. Extinction of the behavior would have been too slow and dangerous; and the man's parents were reluctant to use punishment. Instead, the man was reinforced by tokens, exchangeable for money, for various activities having to do with personal hygiene, chores, and the like. Reinforcing these actions, *which had nothing to do with the problem behavior*, resulted in a prompt and substantial reduction in that behavior.[11]

Examples of the principle in everyday affairs are easy to find. Suppose we live in a community where there are few or no good eating places. We try a new restaurant, and—surprise!—we get an excellent meal. Our likelihood of going back to that place will become very high, so the good meal is a potent reinforcer. But if there already were many good places to eat in the neighborhood, we could try the *same* new restaurant and get the *same* excellent meal—but our likelihood of going back would not increase as much. The reinforcer is the same, but it has less effect if other reinforcers are available.

Or: Consider the horrible change that comes over the most pleasant, civilized people when they get behind the wheel of a car. They bully; they jockey; they cut each other off; they risk disaster to get one car length ahead of the other guy. Why? One possibility is: the matching law:

> In these contexts, people have little to reward them besides the minor victories of the bus or the road. The relative reward is therefore large, so it provokes a rich flow of action.[12]

10. Schroeder and Holland, 1969.
11. McDowell, 1982.
12. Brown and Herrnstein, 1975, p. 84.

In other words, the small reinforcement of getting one car length ahead of another driver can become a potent reinforcer if no other reinforcers are available. And when we're stuck in a car in rush-hour traffic, few reinforcers *are* available. Getting a bit ahead is about the only reward we can obtain. So we may work hard, and risk much, to obtain it.

Melioration

Okay, but how does it work? How does distribution of responses come to match relative rates of reinforcement?

For concurrent reinforcement situations, there are different theories about just what the pigeon is doing. Many of them assume that in matching response allocation to reinforcement frequency, the bird's behavior is maximizing something, e.g., the momentary probability of reinforcement[13] or the rate of reinforcement over short periods.*

Herrnstein's view is different. As he sees it, the process that produces matching, called *melioration,* is just a restatement of the reinforcement principle: Other things equal, a response that is reinforced more often than another will come to occur more often than that other. And if the pigeon comes to peck either key more than the matching relationship says it should, then *rate* of reinforcement (reinforcements per response) will be greater for the "neglected" key when it does get pecked, and the bird will therefore drift toward pecking it more. Thus deviations from matching are corrected by differences in average rate of reinforcement.**

Notice a few things about this idea. First, the process is assumed to be continuous over time. Decision theory, as we developed it earlier, suggests that a decision is a distinct *event:* The actor calls "Time out!" and considers her

13. See for example Shimp, 1975; Platt, 1979.

*In fact, a great deal of attention has been paid to the relation between matching and maximizing. This literature has become very technical; the reader who is interested, and comfortable with mathematics, must be referred to original sources (for example, Herrnstein and Vaughan, 1980; Herrnstein, 1993; Staddon, 1993).

**To see this, let's work through an example. Suppose a pigeon can emit 1000 responses in an hour's session. It can distribute these in any way it likes between two alternatives, say Key A and Key B. Suppose there are a total of 90 reinforcements available in that hour, 60 from Key A and 30 from Key B, but that reinforcements on both are available at unpredictable intervals.

Suppose that at first, the bird responds equally at the two keys. Over the session, its total return from Key A would be 60 reinforcements per 500 responses, or on average, 0.12 reinforcements per response. Its return from Key B would be 30/500, or 0.06 reinforcements per response. And, because the rewards are delivered at variable intervals, those two *average rates* of return would be present from the outset of the session. But wait! The bird is getting more reinforcements per response from Key A. If rate of response increases with rate of reinforcent, the bird ought to begin directing more of its responses toward Key A.

Suppose that it does, and that it overshoots: it begins pecking 90 percent of the time at Key A. But then its rate of return from Key A would be 60/900, or 0.07, whereas from Key B it receives 30/100, or 0.3. So now it should drift back toward spending more pecks on Key B. Let the process continue, and response distribution should settle down to *match* reinforcement distribution over the two keys. That, we recall, is exactly what happens.

options, calculates her expected utilities, chooses the best course of action, and then steps back into life and implements it. But melioration is an ongoing process in which an actor keeps a kind of running track of relative payoffs, drifting from Action A toward Action B whenever B's *rate* of return (payoff per unit time or effort) exceeds A's.

Second, this account does not speak of the subject as *maximizing* anything. The bird simply responds to rates of reinforcement. Indeed, Herrnstein argues that melioration is a more important psychological process than maximization, in that it better describes how animals (including humans) actually behave most of the time. In many cases, it is true, melioration and maximization predict the same actions, or close to them. But there are cases where they make different predictions; and in such cases, there is indeed evidence that animals and humans are more likely to match than maximize.[14]

Third, this theory does not try to specify the cognitive *mechanisms* that produce the actions we see. Thus this kind of theorizing—like theories about progeny payoffs—is descriptive, not cognitive.

It does, however, put the behaviorist-cognitivist debate in a new light. Let us see how.

A Look Backward

With the concept of melioration, we come to the cutting edge of modern behaviorism. And an interesting thing has happened in this research area: Behaviorism and cognitive psychology, once fierce rivals in debate, have become complementary.

As I see the current situation, few writers would deny that cognitive mechanisms operate even in animal behavior. The experimental evidence for cognitive maps, search images, response-reinforcer expectancies, and the like, is simply too strong (Chap. 9). The fact remains that we know little about how these processes work, or just how they take hold of overt behavior in any given case. And, for some purposes, we do not have to know.

Consider melioration. The animal's behavior, we say, tracks the average rate of return for each of the response alternatives. But surely the pigeon does not *literally* count reinforcers, count responses, and then whip out its calculator to divide one sum by the other! We can say, descriptively, that it behaves *as if* it did *something* like that; but what that "something" is, further research must unravel. We don't know.[15]

Meanwhile, however, we can try to specify more precisely the relation between behavior and environment, e.g., whether the animal is maximizing or meliorating or what. In doing so, we can leave the underlying cognitive processes unspecified, just as the cognitivist can leave underlying physiological processes unspecified. We can, in other words, try to state more clearly *what*

14. See Mazur, 1994; Herrnstein, 1990.
15. Gallistel, 1990.

the system is doing. Questions about *how* it does it can wait (compare pp. 23–24).

And, recognizing that, cognitive and behaviorist theorists find themselves with little left to argue about. Perhaps we can all live in the same house after all.

A LOOK BACKWARD: THE THREE FACES OF DECISION THEORY

In applying decision theory to a number of contexts—eyewitnesses, coffee dates, grocery stores, and even foraging and key-pecking—we see that the theory has a number of quite different uses. It really is not one theory, but three. It is time to look over these three faces, and distinguish them.

The Cognitive Theory: Maximizing as Mechanism

First, decision theory can be a cognitive theory about *how* decisions get made. That theory is the one we explored first: The actor anticipates the outcomes of possible actions, and chooses the action leading to the highest expected utility. This implies that the actor is actually performing the cognitive operations that calculate expected utilities, or, at least, doing something very much like that.

The Descriptive Theory: Maximizing as Outcome

Second is the descriptive, or behavioristic, use of decision theory. We might say that Subject Y is maximizing X, and mean by it simply that what Subject Y does *has the effect* of maximizing X. This need not imply anything whatever about what, if anything, the subject is thinking or calculating or anticipating. It says nothing about the mechanisms that caused the actions we see.

For example, in Darwinian theory (pp. 34–37), evolution is a matter of maximizing fitness, where fitness is defined simply as reproductive success—leaving descendants. But a given creature, making its living in the world, is not calculating its reproductive success. With some human exceptions, it is not even trying to leave descendants—not all humans know what makes babies, and it is unlikely that any other animal does. The behaving animal is only trying to find food or a mate, avoid dangers, and the like.

This means that fitness and reproductive success, though important concepts, are not *motivational* concepts in Darwinian theory. Reproductive success is not part of what causes a living creature to do what it does. It is a term that we, the observers, use to describe the *outcomes* of its actions—or the actions of its ancestors, long dead.

The Normative Theory: Maximizing as Best Choice

Finally, decision theory can be a *normative* theory. That means that it describes, not what we do, but what we ought to do. It gives us a way of calculating the *best* selection from the options we are confronted with.

For example, suppose we are playing poker, blackjack, or any game of chance in which various events occur with fixed probabilities and in which fixed gains or losses—say, monetary payoffs—attend the various outcomes. Now it can be shown that if we actually do the decision-theory calculations, and play the game according to those calculations, then in the long run we will win more money (or lose less) than if we play it any other way.

If we do play that way, then we might say that we are playing *rationally*. We are logically calculating what is the best thing to do, and then doing it.

The normative theory is not the same as the cognitive one. If we calculate wrong, or if we don't know how to do the calculations, then we cannot play the game rationally in the strict normative sense. We can still play it cognitively. We may decide to stay in the hand because we believe that our hand is strong, and so we expect to win. But if we are operating on hunch, not on calculation, then we will not—except by accident—consistently make the normatively correct decision. And we will lose out, over the long run, to a player who does.

The Three Faces Compared

It is easy to confuse these uses of the term "maximizing." They can shade into each other to such an extent that the theorists themselves may have trouble distinguishing them on occasion.[16]

What makes things so confusing is that, very often, all three apply. Consider again the expert poker player, who calculates that her best action is to throw in her hand. She is doing the necessary calculations in her head and basing her decision on them; the cognitive usage applies. She is also doing what is in fact the best thing to do if she wants to make money; the normative usage applies. And we can describe her behavior by saying that it maximizes expected utility, even if we don't know whether she is actually doing the calculations or not; the descriptive usage applies.

But the three uses do not have to coincide. Where they don't—as in mate selection or foraging, for example—it is important not to confuse them (compare pp. 379–82).

SOME COMPLICATIONS IN DECISION-MAKING

In introducing decision theory, we considered the decision maker as weighing alternatives, calculating expected utilities, and choosing an action on the basis of those calculations. But that *is* a theory; it may or may not describe how decisions are made by real-world actors. When we look at what actually happens, we find that matters are not so simple as we have supposed.

16. See for example Mazur, 1983a, b; Staddon and Hinton, 1983.

The Problem of Capacity

Thinking and planning, like courting or eating, take time and effort. And certain constraints will be imposed by time limitations—we have only so much time and energy to devote to any activity.

Consider this: We encounter a tiger. The tiger has no cage around it, and somehow looks as if it has not eaten recently. What do we do? Do we stop to consider all possible actions—we might try to pet the tiger, we might sing a song to sooth its savage breast, we might try to hypnotize it—and carefully calculate expected utilities? Of course not. There isn't *time*. We must do something *now*, even if it is not a carefully calculated maximizing action.

Indeed, in extreme cases, time constraints make a truly maximizing strategy quite impossible. Take the problem of selecting a marriage partner. Obviously, it is out of the question to interview *all possible* marriage partners before selecting one. So that decision—an important one—must be made in some other way.

Conversely, suppose we are deciding where to go for lunch. Do we sit down to consider all possible places, calculating quality, price, cost of getting there, . . . ? Probably not. It simply *isn't worth it*.

Fully rational decisions, then, could take too much effort or too much time—more than we have, or more than they are worth. So we turn to shortcuts—ways of simplifying the decision process—so as to reduce the time and effort it takes.

Alternatives to Maximizing

Suppose we were deciding where to go for lunch. We could go through a careful, maximizing decision process; but is there time? Is it worth it? No? Well then, how else might we decide?

SATISFICING We could begin by setting up a standard: What outcome would be *good enough*? Then we consider various possible actions one by one and pick the *first* one whose expected outcome is good enough.[17] That strategy is called *satisficing*.

Satisficing differs from maximizing in several important ways. First, satisficing may require much less time and effort, since we need not consider all the alternatives; we stop when we find one that is good enough. That is the good news. The second is not such good news: when we satisfice, we pick the first alternative that is good enough, so we may well miss another that is better still. Finally, in maximizing we can consider the possible actions in any order, since we're going to consider them all anyway. But if we satisfice, the order in which we think of alternatives makes a great difference. If we consider restaurants in alphabetical order, we might decide that Archie's Place is good enough, and go

17. Simon, 1976.

there. If we considered them in reverse alphabetical order, maybe Zorba's Place would be good enough and we'd go there instead.

Satisficing may be far more commonly used than maximizing. Shoppers often satisfice. Rather than insisting on the very best necktie or hamburger in town for the money, one may choose the first that is *satisfactory* in quality and price. Obviously the time and effort saved is considerable.

BETTERING Another kind of decision strategy goes by various names, such as *incrementalism* or *muddling through*.[18] *Melioration* (pp. 385–86) is a special case of it. Let us call it *bettering*. The decision rule is as follows: Whatever you're doing now, if an option comes along that has higher expected utility than that, switch to it. Otherwise go on with what you're doing.

This strategy differs from the other two in that it is a continuous process. Maximizing or satisficing requires that one take time out, step back from life to make a decision, and then step back into life and implement it. Bettering, on the other hand, describes an actor who is going on with the routine business of living, but always alert to the possibility of making life better for herself, and changing from one action to another if it has that expected outcome.

Examples abound. We might put our savings in a bank, leave them there, and forget about them, until we see that another bank offers higher interest rates. Then we switch banks.

Or think again about where to eat lunch. Suppose our group goes to Gary's Grill for lunch every day. We don't really *decide* to go there, we just go; it's routine. But today someone says: "Let's all go to Francoise's Fine Filet and Frogs' Legs Place. My treat." We all agree—of course! The food is much better there than at Gary's, and for us it's free. Our lunching strategy is a bettering one: We just do the usual thing, unless a better option is put before us. If one is, we take it.

Bettering is a fine strategy for reducing the time and effort allocated to decision-making. No decision is required at all unless a better alternative comes along. It has its dangers, however. If the actor never steps back to decide on long-term goals, she will be alert only to immediate short-term ones. It will be difficult to take one step back to go three steps forward, to take a loss now for a bigger gain later.

Then there is the converse danger: A series of small, bettering changes can lead step by step to a destination that would not have been chosen had one foreseen it during a moment of thoughtful decision. "A stepwise increase in commitment can end up locking the person into a career or marriage without his ever having made a definite decision about it."[19]

Thus, some writers picture the human actor as operating mostly on habit and inertia—they even refer to him as a "reluctant decision maker." He goes about his routine by rote, unless he sees risks or costs in doing so, including the

18. Herrnstein and Vaughan, 1980; Janis and Mann, 1977
19. Janis and Mann, 1977, p. 35.

"cost" of passing up something better. That placid existence may be punctuated with minor decisions—where to go for lunch, which necktie to buy, and the like. These are probably decided by a *satisficing* strategy.[20] Indeed, it is just such trivial decisions that the satisficing concept was designed to account for.

AH, BUT WHICH STRATEGY? The recognition of alternative strategies does make decision theory more flexible and, perhaps, more realistic. But if we wish to develop it into a theory of purposive behavior, these options pose a problem. We can make decisions in more than one way, we say. But then which strategy shall we use? Must we first decide *that?* And if so, on what basis?

Thus it has been suggested that satisficing, as a way of choosing a necktie or a lunch, really is an optimizing strategy if one figures in the costs of acquiring all the information we would need in order to use a maximizing strategy. That sounds fine, but in practice, it could trap us in an infinite regress. Must we *decide* whether to consider these costs? Must we first consider just what information we would need in order to consider them? What information do we need in order to know what information we need? We seem to have a formula for paralysis here (Fig. 10-12).

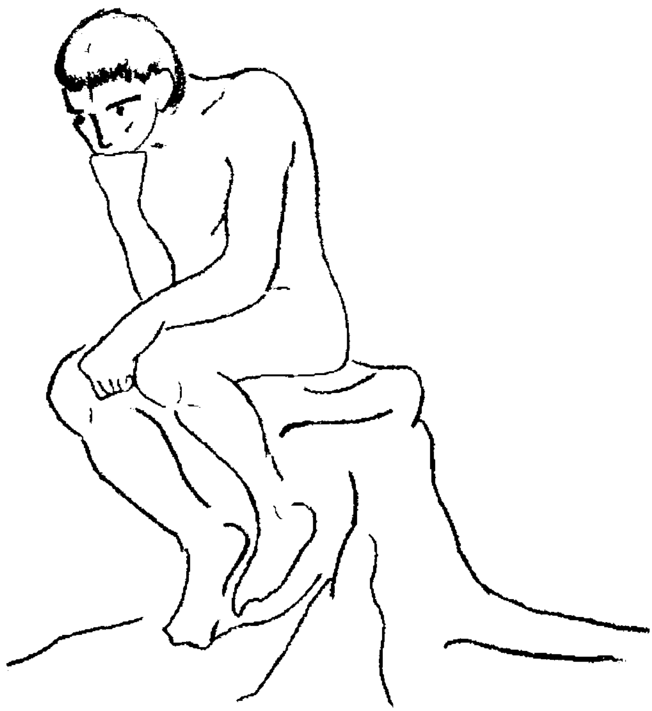

Figure 10-12 *"I'm deciding whether to decide whether or not it would be worthwhile to gather all the information I would need in order to decide whether. . . . Don't bother me!"*

20. Simon, 1976.

Of course, research can be, and is, directed toward the question of strategy selection: What determines when we use one strategy and when another? The problem is that we don't have a normative theory that tells us when we *should* use one strategy or the other. And that could land us in big trouble because, lacking that, we have no solid criterion for determining when a decision is *rational* and when it is not!

In short, we see theoretical storm clouds on the horizon. They are about to get darker.

The Problem of Availability

There was once a king who died of a fever, brought on by sitting too close to a hot brazier. He continued to sit there, literally roasting to death, because the servant whose job it was to move the brazier couldn't be found. Apparently it never occurred to His Majesty to move the brazier himself. It wasn't the sort of thing that kings do for themselves—and so he simply did not *consider* doing it. It was not, as we say, *cognitively available* to him.[21]

We too may fail to solve problems because certain ideas are not cognitively available to us. Consider the following riddle:

> A boy and his father were in an automobile accident. The father was killed instantly. The boy was alive but badly injured, and was rushed to the hospital and prepared for emergency surgery. But the surgeon, on seeing the boy, said, "I can't operate on that boy. He's my son."

How can that be? (No, it has nothing to do with adoption or reincarnation.) If you see the solution right away, try the riddle on your friends and see how long it takes them to find it. If you don't see it right off, think about it until you do. Then think about the implications of the fact that the solution—a perfectly obvious one—has a way of being cognitively unavailable to many people in this society.

Then again, if some actions don't occur to us, this is not necessarily a bad thing. Consider this example. A popular view of moral behavior—dating back to Thomas Hobbes if not earlier—is that we refrain from stealing or lying because of fear of disapproval or punishment. Sometimes, no doubt, self-control does depend on these expected outcomes; but it can also depend on what options we consider at all. In one study, children of Mennonite and of neighboring non-Mennonite families were interviewed as to what behaviors should or should not be performed. The Mennonite children seldom even mentioned the moral values of their community that prohibited stealing, lying, and the like. It seems that those well-internalized values placed these actions outside the range of options to be considered.[22] Such actions were not cognitively available; they simply did not *occur* to the children.

21. Tuchman, 1984.
22. Radke-Yarrow et al., 1983.

IMPLICATIONS OF AVAILABILITY No strategy of decision-making—maximizing or any other—will lead us to choose a response that is never considered. Nor does decision theory tell us what responses we *should* consider. It only specifies ways of choosing among the options we do consider. So our strategies for decision-making, such as the maximization of utility, come in only after we "list" in our minds the options that we will consider. They provide no guidance as to how to make the list.

And what *that* means is that decision theory is incomplete as a cognitive, descriptive, or normative theory of action. To its logic about how to make rational decisions, we need to add considerations that are not *logical* but *psychological:* What alternatives are, or are not, even considered in a decision situation?

WHAT DETERMINES AVAILABILITY? What exactly does it mean to say that an option is "available" or that it isn't? In some cases, it may have to do with how readily a response—or a goal or a plan—can be called into play by a higher-order system. This might reflect untaught properties of a creature's cognitive apparatus.

For instance, Bolles's concept of species-specific defense reactions (pp. 255–56) could be restated this way: Certain actions, but not others, are readily available as ways of coping with danger. Conversely, consider again the fact that pigeons have trouble *refraining from pecking* as a food-getting response (pp. 318–19). The utility-maximizing response is simply not readily available to the actor, for whatever reason. In humans too, there is some evidence that certain cognitive actions are readily available for some problems, but not for others that are logically, but not psychologically, equivalent.[23]

Second, one's learning history may make certain options readily available, others not. The culture in which one is raised may have a decisive influence here: To a Mennonite child, stealing is not cognitively available as an option, whereas to an inner-city American youth it may well be. Indeed, society's learning curriculum may try to short-circuit the decision-making process altogether, by making only *one* alternative cognitively available as the only right or possible thing to do. If this is successful, an actor may "just do" the socially approved thing, "mindlessly" if you wish, without really making a decision at all.[24]

Finally, what we might call "local," situational events may promote or discourage the availability of certain thoughts over the short term. We speak of "hints" that jog certain ideas into availability, perhaps by association (shades of Thomas Hobbes!).

For example, think again about the riddle of the surgeon's son (p. 392). Suppose you tried the riddle out on a friend, and she was unable to solve it. Now, if she had just come from a discussion of the problems faced by profes-

23. Barkow, Cosmides, and Tooby, 1992.
24. Quinn, 1992.

sional women in this society, then the solution might have come to her mind at once.

And, Gentle Reader: If *you* failed to solve the riddle, does that reference to "professional women" enable you to solve it now? If so, then you see how a situational influence—the text you are reading, right this minute—can jog your mind to increase the *availability* of an idea! (Do not feel bad. An incredibly high proportion of college students, women as well as men, fail to solve that riddle. Think about that.)

The Availability of Reasons

We may make poor decisions, then, because too little information is cognitively available to us. We may also do so because too much is! Irrelevant considerations may throw us off the track of sound decision-making.

That at least is one way of interpreting a disturbing series of studies by Timothy Wilson and his colleagues.[25] In one experiment, college students were allowed to make selections from a number of posters. They could take home the ones they selected and put them up on their walls. Some were asked to make careful decisions, "decomposing" the various prints into their pluses and minuses and weighing these against each other.[26] Later on, these subjects were significantly *less* satisfied with the choices they had made than subjects who had simply gone with their overall impressions. Similar results were found for a variety of other products, and even for how much dating couples liked each other!

Why should this be so? Wilson and his colleagues suggest the following: When we sit down to list pros and cons, we do it verbally, in a kind of internal conversation with ourselves. (This is all the more true, of course, if we list the pros and cons in written words.) The problem is that some of the pros and cons will be easy to put into words. Others—our intuitions, emotional reactions, and gut feels—will not be. So, when we make our lists, they will likely include those pros and cons that are easiest to put into words. Those that are hard to put into words will likely be omitted or downplayed.

But in fact, these hard-to-describe reactions may actually be among the most important determinants of our true preferences. If so, then to omit or downplay them is to omit or downplay the most important influences on what we like and how much we like them. In that case, a careful listing of *verbalizable* pros and cons could lead us away from, not toward, a clear appreciation of what our own likes and dislikes really are. It's the sort of thing we mean when we say, "I talked myself into it," especially when we add or imply, "And I wish I hadn't!"

Wilson does not of course claim that this always happens, and that a careful decision will necessarily be a bad one. But his studies show that it can be.

25. Wilson et al. 1993; Wilson and Kraft, 1993.
26. Dawes, 1988.

Therefore, as so often, the question is: When does this happen, and when not? And why?

The Framing of Decisions

As a normative theory of rational decision-making, decision theory assumes that the possible actions, and the expected utilities of their outcomes, are established. We know what we want, what our options are, and what is likely to happen if we do each.

But it just is not that simple. Amos Tversky and Daniel Kahneman have investigated this matter in an ingenious series of thought experiments. And they have shown how the *same* decision problem can have very different rational solutions, depending on how one looks at it—where more than one way of looking at it may be quite legitimate.

In one of their studies, all subjects read the following "stem":

> Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:

> Then half the subjects read the following continuation:

> 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 other half saw the following:

> 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 2/3 probability that 600 people will die.
> Which of the two programs would you favor?[27]

Looking over the whole business, we see that the two possible decisions are identical in the two cases. Programs A and C describe the same outcome; Programs B and D describe the same outcome. If A is preferred to B, then C should be preferred to D. But in fact, of the subjects who had the A-B choice, 72 percent preferred A; whereas of those given the C-D choice, 78 percent preferred D!

There is no way we can resolve this wild inconsistency without elaborating our theory of decision-making. The choices change, depending on how subjects are invited to look at the problem.

27. From Tversky and Kahnemann, 1981.

Why? Because A and B treat the problem in terms of *lives saved*. C and D treat it in terms of *lives lost*. Let's examine the utility functions that characterize these scenarios (Fig. 10-13). We assume that the diminishing-returns principle applies to lives saved and lost, as well as to money. Then we see that saving 200 lives for sure has higher expected utility than a 1/3 chance of saving 600 lives (Fig. 10-13A). If we look at the problem in terms of lives saved, we should take the sure thing—as most subjects did.

But what if we look at it in terms of *lives lost*? Then a 2/3 chance that 600 people will die has less *negative* expected utility than 400 deaths for certain (Fig. 10-13B). If we look at it that way, we should go with the gamble. Most subjects did.

But the only difference is in how one states the problem! To save 200 lives *is* to lose 400 lives. By looking at it one way or the other, we can make either decision the rationally correct one.

Now we see why the notion of "framing" is appropriate here. It is as if each subject were looking at the problem through a window that showed one utility function or the other. Different window, different decision—even for the same outcome.

Well, which frame is the correct one to use? How ought we to look at the problem? *Decision theory does not tell us how we ought to look at it.*

As a normative theory, decision theory tells us what we ought to do if we want to maximize expected utility of lives saved. It tells us what we should do if we want to minimize expected cost of lives lost. But if we can look at the same problem in either of those two ways, the theory is simply no help. It applies only *after* the utilities have been framed. Framing them is up to us.

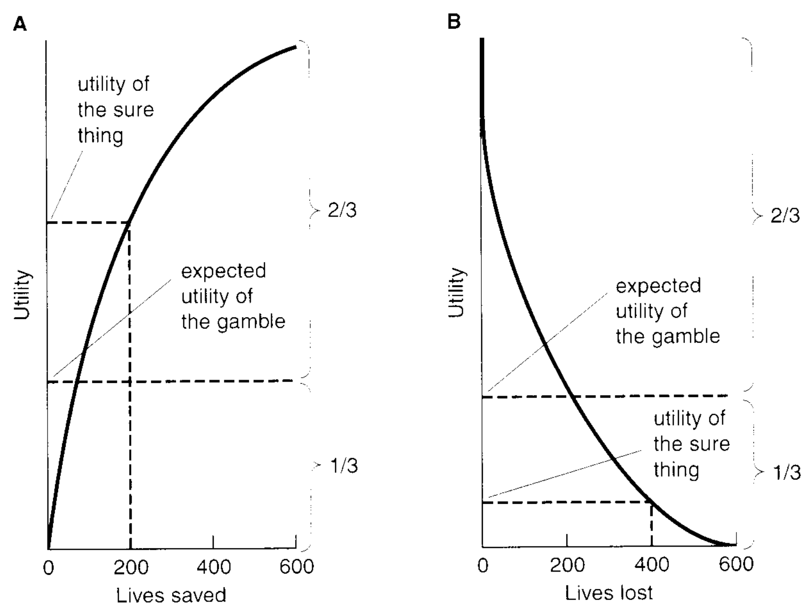

Figure 10-13 The framing problem. (A) If one considers lives saved, then 200 saved for sure has higher positive utility than the expected utility of the gamble (1/3 the utility of 600 lives saved). (B) If one considers lives lost, then 400 lives lost for sure has lower utility—or greater *negative* utility—than the expected utility of the gamble (1/3 the utility of no lives lost).

The framing effect operates in real-life decisions, too. Surgery for lung cancer was seen as less attractive, relative to radiation therapy, when the two were compared using mortality (lives lost) rather than survival (lives saved) figures. The effect was just as large with physicians as with lay people.[28]

Other examples? Consider two merchants, who accept both cash and credit cards. One gives a *discount for cash*. The other charges a *surcharge for credit cards*. These could amount to exactly the same charges, but doesn't one sound like a bargain and the other like a rip-off?

Then there is a personal experience of the author's, and you have my solemn word that this actually happened. In a course I taught, I gave students the option of writing a term paper, to be worth extra points toward the course grade. On an anonymous course-evaluation questionnaire, one student complained that such a policy was unfair. It would be more fair, said she or he, to require a term paper of everybody—and then deduct points for those who didn't write one!

As we look over these last sections, we notice that our concept of rational decision-making is becoming much less clear than we might have thought it would be. We saw, for instance, that two quite different solutions to the same decision problem could be considered normatively optimal, depending on how the problem is framed.

It is important to see just how troublesome this is. What it means is that in cases where utilities can be framed in different ways, we have no criterion for rational behavior! That's because we have no normative theory for the framing process itself. We can try to construct descriptive theories about how people do frame problems, and cognitive theories about the mental operations by which they do so.[29] But no theory tells us how they *should* go about it.

A Look Backward: Some Problems for Decision Theory

Historically, decision theory emerged from the analysis of games of chance; and for that purpose, it did very well. The alternatives, probabilities, and payoffs all are fixed and are known; and the mathematics of decision theory do tell us how best to play poker, blackjack, and the like.*

To extend the theory to choices of all kinds seemed natural enough. And it fit well with common sense: The actor does what he *believes* (subjective probability) will get him what he most *wants* (utility). Simple!—or so it seems at first.

As a general theory of decision-making, however, we have seen that the theory faces some problems.

28. Redelmeier, Rozin, and Kahneman, 1993.
29. See Tversky and Kahneman, 1981, 1983.
*Though even in that domain, non-logical factors may enter in, as when an adult intentionally loses a game to a child.

Rational Decision-making: Do We Do It? There is a wealth of experimental evidence that humans do not, in fact, behave as fully rational decision makers.[30] To give just one example: A collection of professional scientists, mathematicians, and philosophers played a simulated prisoners' dilemma game. Now, as we saw earlier (pp. 377–78), one's rational "move" in the game is necessarily to "confess." But among these professionals, all thoroughly familiar with the logic of the game, the plain fact is that only some chose "rationally." Some did, some did not.

Even where rational calculations do seem to describe our decision processes, they may come in after the decision, not before. As Robert Zajonc points out:

> We sometimes delude ourselves that we proceed in a rational manner and weigh all the pros and cons of the various alternatives. But this is probably seldom the actual case. Quite often "I decided in favor of X" is no more than "I liked X." Most of the time, information collected about alternatives serves us less for making a decision than for justifying it afterward.[31]

Other writers too have suggested that we may use the maximization idea not to guide a decision, but as a *reason* by which we explain the decision, to others and to ourselves.[32]

Could We Do It? If we do not always make coldly rational decisions in practice, could we do so if we tried? It is not clear that we could.

Sometimes a maximizing strategy is impossible in principle, as in selecting a marriage partner. And even a careful maximizing decision about where to go for lunch would obviously be more trouble than it is worth. But then, how do we decide whether that is so? An infinite regress looms: Must we rationally decide whether to rationally decide . . . ? Do we take time to decide whether we have time . . . ?

Finally, even where maximizing is possible in principle, we may lack the necessary capability. A really good poker player may do the mental arithmetic necessary to figure the probability that drawing another card will lead to a winning hand and, knowing that, the size of the pot, and the cost of staying in the hand, calculate that dropping out is the correct choice. But most of us, the author definitely included, usually cannot do that: we don't know how, or we lack the necessary skill and concentration for such mental arithmetic, or both. I might make the same decision—to drop out—but base it not on calculation but on a vague intuitive feeling that it would be best to do so.

Let us stay with that example for a minute, for it may help bring the problem into sharper focus. If good poker players calculate expected utility and maximize it, what do poor players do? Whatever it is, (1) it is *not* identical to the process that decision theory describes, and (2) it is probably the more usual

30. See for example Nisbett and Ross, 1980; Kahneman, Tversky, and Slovic, 1981.
31. Zajonc, 1980, p. 155.
32. See for example Shafir, Simonson, and Tversky, 1994.

mode of human decision-making. There are, after all, more poor poker players than good ones.

Well, what does a poor poker player do? If he throws in his hand, it is because of a vague hunch that his hand isn't very good and the cost of staying in is high and, well, it doesn't seem worth the risk. Most of us make most decisions by habit or hunch or gut feel. In experienced poker players, these hunches may be "calibrated" by the player's history of successes and failures in situations like this one. In other words, they may be calibrated by their *environmental consequences,* somewhat as reinforcement theory would have it (see the previous chapter). Or they might depend on rules of thumb, or *heuristics,* which guide our decisions (see the next chapter!).

Either way, decision theory will have to be expanded before it can become an adequate theory of human behavior. It will have to consider the habits, heuristics, and hunches that do not flow from its arithmetic, but that do describe most of the decisions we make, most of the time.

SHOULD WE DO IT? As a cognitive theory, we see that decision theory has its limits. What about as a normative theory? *Should* we maximize expected utility? It is not clear that we should.

There clearly are cases in which to be *rational* might not be, well, *sensible.* We don't have much time to calculate when faced with an uncaged tiger; and tigers aside, choosing a lunch place rationally could leave us no time to eat! And even if we do have time, careful, thoughtful decision-making does not always work out well in practice, as we have seen.

Finally, there are cases where we do not maximize and are glad of it! In the prisoner's dilemma game, a maximizing strategy leads inexorably to an unwanted outcome. But as we have seen, even knowledgeable professionals did not, in fact, all play the game rationally. And one of the players, a philosopher, gave his reason: "I would rather be the person who buys the Brooklyn Bridge than the person who sells it."[33] This looks forward to the role of emotion in decision-making, and the strategic uses of emotion (Chap. 12).

The Bottom Line

The conclusion we reach is simply this. If the behaviorist approach to purposive behavior—reinforcement theory—is incomplete (Chap. 8), well, so is the prototype of the cognitive approach—decision theory. And if neither one will do, that means that at this point we simply do not have an adequate theory of purposive behavior. Such a theory, when it comes, will probably borrow ideas from both of the other theories; but it may well look quite different from either one.

33. For discussion of this experiment and related issues, see Hofstadter, 1985.

MODELS IN THE MIND

Since the last two chapters have been concerned with the cognitive approach to motivation, it is worth pausing to make clear the relations between the two chapters.

We get the future into the present, cognitive theory says, as an image of a potential state of affairs. Thus such images can represent the *imaginary* consequences of actions we *imagine* taking.

Here is where decision theory comes in. The decision maker is assumed to imagine possible actions and their consequences in deciding what to do. Thus the actor creates an internal model of the world. Then she runs the model faster than the world runs, and, from what happens in the model, she predicts the consequences of actions before they are performed. And she decides what action to take, based on those predictions. Specifically, she calculates the subjective expected utilities of the various actions and chooses the action whose subjective expected utility is highest.

The process reminds us of the use of computer models to simulate the world. Such models allow us to calculate what *would* happen if such and such were done. The advantages of this are very great. We can ask such questions as: If we made this or that design change, would the airplane be likely to crash? And we can ask the question without actually trying the experiment, risking a valuable real-world plane and a valuable real-world pilot. In a similar way, we can consider the costs and benefits of an action without the risk and effort of actually performing it in the real world. Decision theory as a cognitive theory is a simplified, first-pass attempt to describe how we do that "considering."

SUMMARY

As reinforcement theory is basic to modern behaviorism, so decision theory is basic to cognitive theories of action. It treats our decisions as guided by two factors: the utilities, or values to us, of the possible outcomes of an action (a motivational factor), and the probabilities of the various outcomes (a cognitive factor). The utilities, weighted (multiplied) by the probabilities, of the possible outcomes of each action sum to give us the *expected utility* of the action in question. Then we choose the action with the highest expected utility (or, where probabilities are unknown and must be estimated, the highest *subjective expected utility*).

The theory can be elaborated in a variety of ways. In *utility theory*, the principle of diminishing marginal utility shows that the value of a commodity to us is likely to decrease as the amount we already have increases. This simple psychological principle makes possible much of our economic activity, even such a simple matter as buying bread from a grocer. Often we must make a decision in which the outcome depends on whether some condition, X, is true or not when we are not sure whether it is true or not. Our decision then will depend on how likely we think it is that X is true (the cognitive factor) and the costs or benefits of possible outcomes (the motivational factor). *Signal-detection theory* applies decision theory to such cases. But rational decision-making does not always lead to the happiest possible outcomes, as shown by the *tragedy of the commons* and the nasty *prisoners' dilemma.*

Decision theory can also shed new light on the evolution of behavior. Species evolve because some forms leave more descendants than others, and so we can think of number of descendants as an outcome—a "progeny payoff"—for "deciding" to evolve certain anatomical structures *and* tendencies to behave in certain ways. Thus shore crabs are "calibrated" so as to select prey that maximizes the benefit/cost ratio (energy gained/energy lost) in handling the prey. The Coolidge Effect (Chapter 4) can be seen as increasing the reproductive success of the male, though the female may increase her own reproductive success by behaving in ways that oppose it. In all this, however, decision theory is used descriptively, not cognitively. No one suggests that the animals do the calculations. If they behave as if they did, it is because ancestral animals that behaved that way left more descendants than those that did not; so modern animals, who *are* those descendants, have inherited the corresponding instinctive tendencies.

Reinforcement theory also makes contact with decision theory, as in the analysis of concurrent reinforcement schedules. The *matching law*—organisms divide their choices so as to match the relative amount of reinforcement for the various actions—gives rise to the *law of relative effect:* The effect of a reinforcer is relative to what other reinforcers are available. This work also has led to the notion of *melioration* as an alternative to the principle of maximizing expected utility. These ideas are yielding new points of convergence between behavioristic and cognitive theories.

Looking over these developments, we see that decision theory can be used in at least three quite different ways. It can be a *cognitive* theory, describing the calculations an actor performs before making her decision. Or it can be purely *descriptive,* describing the outcomes of a decision while making no assumptions about how the decision was made. Finally, it can be a *normative* theory, showing us what we ought to do in a given situation if we wish to maximize utility. These are quite distinct uses of the theory; in discussing evolution, for example, we use the theory as a descriptive one, but not as a cognitive one.

Decision theory is, after all, a theory. In real life, we may not behave as the theory suggests we do. First, we have only so much time and effort to devote to decision-making (we could not possibly consider *all possible* marriage partners, for example); and a careful, utility-maximizing decision strategy could take too long in emergencies or be too much trouble for trivial decisions. There are alternatives to maximizing, such as *satisficing* or *bettering*. But how do we decide among strategies? Do we decide how to decide? Do we decide how to decide how to decide? . . . An infinite regress threatens, and decision theory does not tell us how to avoid it.

Then too, much depends on what options we consider—which ones are cognitively *available* to us. Again, decision theory is no help; it tells us how to weigh the options we do consider, but not which ones we should consider.

Conversely, we may give more weight to certain considerations—those most easily verbalizable, for example—than these considerations ought to have.

Yet another complication is the problem of *framing*. Do we think in terms of lives saved or lives lost? What is the optimal decision may depend on how we look at a problem, and again the theory may be silent about how we ought to look at it.

All told, decision theory has limitations as a theory of action. Humans do not in fact always behave as rational decision makers. In some cases—perhaps many—we could not do so if we tried. And that is perhaps just as well, for it is not always clear that we *ought* to behave in a strictly "rational," payoff-maximizing way. The crisp logic of decision theory must come to grips with the fuzzy "psychologic" of actual decision-making, which includes considerations of what we do or do not think about (the problems of capacity and availability), of how we think about it (the problem of framing), and of what we really want.