Chapter 2

Problem Solving and Decision Making

MAJOR TOPICS:

ACTIVITIES INCLUDED IN SOLVING PROBLEMS

In everyday usage, "problem," "decision," and "choice" tend to be used more or less interchangeably. As you approach graduation, you become increasingly aware that you must "decide" on your future activities—whether or not to apply to graduate school, look for a job, or take some time off. You have to "choose" among various alternatives. You may feel the whole business becoming a "problem" that you feel you have to "solve." Perhaps "problem" has some negative overtones, and "choice" seems more pleasant, but, for most of us, the three terms mean pretty much the same thing.

Decision theorists can also be careless in using these three words. If pressed, however, most would admit to the value of keeping the three somewhat distinct, with "problem solving" as the broadest term, and "choice making" the narrowest. In the broadest sense, a "problem" is said to exist when an individual becomes aware of a significant difference between what *actually is* and what is *desired*. I examine my checking account and find myself almost out of money; I want a new stereo, which costs $800: I have a problem. Or, you have been getting Cs in a course, and want an A: you have a problem. Or, the air is getting more and more polluted and we would like it clean: we all have a problem.

If a "problem" is "a significant difference between actual and desired," then a "solution" is clearly something that reduces this difference to the point at which it loses significance. A problem is "solved" when the difference becomes too small for us to care about, not necessarily when it reduces to zero. Most of us can live with reality being short of perfection without being aware of a problem! The various activities we call "problem solving," then, are the things we do to reduce the difference between actual and desired states to an acceptable level.

Problem solving embraces a number of component activities, all tightly interwoven with one another. Idealizing the process, we can distinguish the following five components: problem identification, alternative generation, choice, implementation, and monitoring.

Problem Identification

In life, problems do not come to us prepackaged, the way they do on final examinations. Generally we start with a feeling that something is wrong and have to devote a good deal of effort to get a clear understanding of the problem. Think, for example, of how much time and effort a medical doctor devotes to diagnosis—that is, to turning our hazy complaint about not feeling well into a precise statement of what the real problem is. Similarly, an engineer might invest weeks or months of effort refining a client's vague problem ("excessive energy costs") into a precise understanding of exactly what energy sources are currently being used, for what purposes, at what costs, and so on. In many cases, once the problem has been correctly identified the solution is trivially easy.

Most of us are strongly oriented to solving problems and tend to spend too little time formulating them carefully. Some authors have suggested that, in addition to weighing the familiar Type 1 and Type 2 errors, we should beware of "Type 3 errors"—solving the wrong problem. Changing the problem formulation just a little often has dramatic effects on how we go about solving it. For example, if we formulate the problem of air pollution as "I'm coughing too much," a bottle of cough suppressant is a plausible solution. If, however, we formulate the problem in terms of air quality, we are more likely to think of solutions such as moving away or campaigning for reduced auto and smokestack emissions. An even more suggestive formulation sees the pollution problem as "valuable chemicals in the wrong places," which starts us thinking about possible profits to be made, not just costs to be borne, in keeping the materials in their proper places (and out of people's lungs).

A key part of the problem identification or formulation phase is clarifying exactly what will count as a satisfactory solution. What criteria will it have to meet? What constraints do we want to apply? Take the pollution problem again. The air currently carries a burden of various concentrations of several gases and particles, a state we find unacceptable. But what levels *are* acceptable? Do they all have to reach zero before we will regard the problem as solved? And are we prepared to pay *any* price for getting our solution? Closing every factory? Taking every car off the road? Probably few of us want to go this far. But getting the problem properly formulated requires us to say just how far we do want to go, what our goals or objectives or criteria are for an acceptable solution. The first cluster of problem-solving activities, then, are those concerned with moving from our initial, generally vague sense of what the problem is to a clear specification of it: in what ways, exactly, does the actual situation differ from what we wish it to be? And what criteria must a possible solution meet before we will regard it as acceptable?

Alternative Generation

A second cluster of problem-solving activities involves efforts to think up one or more actions that might help the situation. This may require considerable creativity or merely routine searching around. It may turn up hundreds of likely-looking candidates in five minutes, or it may take months of hard work to find even one. We may be able to come up with several broad alternatives in an instant—for example, "ferry," "bridge," or "diversion" as alternatives to a valley-crossing problem—but still have a lot of work to do before we can really start to evaluate each. Which are feasible? What will each cost? What environmental impact will each have?

As with the identification phase, we often short-circuit the alternative generation phase in our eagerness to get to the solution. Once we see a problem in a certain way, we tend to think of only one class of possible solutions, and often pass up much neater, cheaper, or better solutions. We have all had the infuriating experience of pounding away on one method of solution to a mathematical problem, only to find later that a different approach cracks it in three lines. Many of what we think of as "elegant" problem solutions share this trick of seeing a way of solving it quite different from the candidates we were considering. Sifting through a list of alternative solutions often blinds us to the fact that there exists a clear winner that never made it to the list. We carefully examine every mousetrap and poison on the market, and never examine the possibility of simply getting a cat.

Choice

Once the problem is clearly identified and the alternative possible solutions are clearly laid out, we are ready to make our choice. The problem formulation tells us what we are looking for: for example, an amplifier with the following frequency response, at the following power output, with not more than this percentage of distortion, and at less than this price. Our alternative generation has provided us with information about the characteristics of the amplifiers available in the shops. We now sift through the list, weighing each against our criteria, until we find the one that meets our needs most closely. Then we choose that one.

Anyone who has managed to reduce any important problem to such a trivial choice procedure deserves a medal. In reality, choice is vastly more complex than this simplified sketch suggests—and, for most of us, generates a lot more anxiety and difficulty than the simple, mechanical checklist approach would imply. In our amplifier purchase decision, for example, we are likely to encounter difficulties, such as:

- What if more than one alternative meets the criteria? Both the Yagimoto X (at $720) and the Boomer (at $750) exceed my performance specifications, and both are under $800. The Boomer looks nicer, but the X has more power. What is my tradeoff between price, power, and appearance?

- What if none of the alternatives meets the criteria? All the amplifiers meeting my specifications cost over $800. Should I spend more? Or should I relax my criteria? Again, the tradeoff problem.
- What if I cannot be sure? Other things being equal, an amplifier that performs well for ten years is more attractive to me than one that needs fixing in five years. But even well-made instruments sometimes break within a week, and even poor ones sometimes last for years. How can I estimate the chances of a breakdown? And how do I crank this information into my calculations?

Choice theorists tend to approach these questions by building a formal model, in this case, a model of the purchaser's objectives. For example, the theorist might suggest that the purchaser evaluates different amplifiers by a simple weighted average of price, power, appearance, and reliability:

$$\text{Value} = a_1(\text{Power}) - a_2(\text{Price}) + a_3(\text{Appearance}) + a_4(\text{Reliability})$$

The purchaser indicates how important each of the features is by setting the four weights $a_1 - a_4$. Each possible amplifier is then simply "scored" on each of the four dimensions, and a value calculated for each. The assumption is that the purchaser will be best pleased with the amplifier that achieves the highest "value" by this method.

Such an approach to making complex choices can often be useful, but it needs to be examined carefully. The model is so clear and straightforward that it is easy to forget that it *is* merely a model, an attempt to represent our desires in mathematical form. The equation above, for example, implies that the purchaser is prepared to make tradeoffs between each of the features of an amplifier; a little more of one feature will make up for a little less of another. This might be true up to a point, but would any purchaser really want an amplifier that was handsome, powerful, and cheap, but totally unreliable?

These objections can, of course, be taken care of by making the model more complex, for example, by setting lower bounds on each feature. The point here is simply that choice is a complex process, and one should look carefully at any procedure that offers to make it easy. It is all too easy to be misled by elegant formal devices like the model presented above and forget to ask the crucial question: "Is this what I really want?"

Implementation

We have defined problem solving as closing the gap between real and desired. A problem is not solved by analysis, creativity, or choice: we need to *do* something to bring about the change. Solutions need to be implementable, and they need to be implemented. Will Rogers once proposed that the way to eliminate U-boats in the North Atlantic was to boil the entire ocean, so that they would turn pink and pop to the surface like shrimp. When asked how he planned to do this, he protested that he was the grand strategist, and left all the details to his subordinates. Many of us take a similar, if less extreme, position in our first work assignments after graduating. We are trained in the analytical skills of problem solving, and tend to stop after we have chosen our recommended solution and written a report showing why the solution is so good. Few real problems actually get solved this way. We need to build in the implementation phase of the work, to find the money, time, effort, and resources needed, to make sure that the other people involved understand what is happening, and that they are motivated to play their parts.

All this is perfectly obvious, but still worth saying. One of the biggest gaps between academic training and actual practice is that the former deals only with the analytic phases of problem solving, while the latter depends on getting action. Operations research (OR) analysts, for example, are notorious for their fascination with the quantitative, modeling parts of their work and neglect of actual implementation of their solutions. Too few have the real-world skills, interest, and just plain savvy to realize that there is more to problem solving than merely finding the optimal solution. A truck-routing schedule that involves scrapping the current fleet of trucks, having drivers work two-hour shifts with three-hour breaks, or that requires the dispatcher to start work at 3 A.M. is not likely to be implemented, no matter how elegantly it solves the abstract problem of routing.

Monitoring

An astute problem solver retains a little humility. No matter how thorough he or she has been in moving through the problem formulation, alternative generation, choice, and implementation phases, there is still a chance that things will go wrong. Perhaps the problem was formulated with an important criterion left out (for example, the dispatcher's working hours). Perhaps the "best" solution was not implemented properly, or the assumptions about how it would work are not borne out in practice. Or, most critically, perhaps the problem situation has changed, and the problem we were trying to solve is not the problem we are currently facing.

For all these reasons, good problem solvers monitor the results of the implementation carefully, to check whether or not it is working as anticipated. This critical feedback loop may provide nothing more than the warm glow of seeing that one's plans worked out as expected. Much more often, however, there will be smaller or larger deviations from these expectations and further problem-solving efforts will be required. Indeed, as we shall see later, there are plenty of situations where attempting solutions amounts to little more than trial and error. In these cases, good monitoring and feedback become especially critical. Trial and error only works if you can see where you have made an error!

PROBLEM-SOLVING SEQUENCES

It is wrong to suppose that these five activities fit neatly together into a sequence over time:

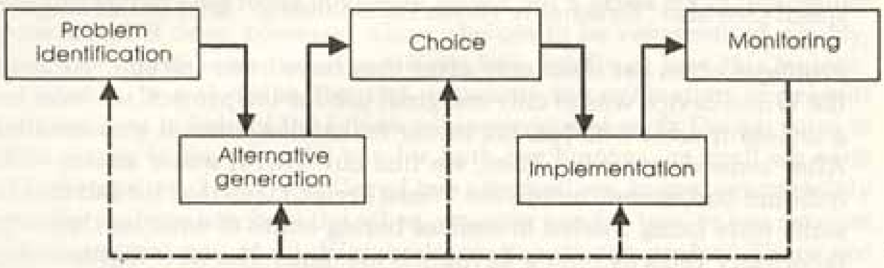

The primary virtue of such a diagram is that it identifies the different activities that go into problem solving. It is, however, downright misleading to expect that these activities flow (or should flow) from left to right in a fixed sequence. As we have noted, people jump to and fro between activities in anything but an orderly way; and work at one phase is extensively influenced, or even distorted, by the others. The illustration may be a useful way of reconstructing what one has done, in a report to one's teacher or one's boss, but it is not likely that you will ever see any real problem get solved this way.

Solutions often chase problems. Most of us carry around ideas that might be politely called "solutions looking for problems." Perhaps we have always wanted to get a Wizzo ten-channel high-precision solidstate digitizer for our lab. Or, more sinister, we would love to find a way to get rid of Joe, that obnoxious senior technician who is the only person who knows how to fix the old centrifuge. When a well-funded new research project comes along, our thoughts naturally run to a design that absolutely requires a Wizzo; and, when a budget cutback threatens, the centrifuge (and thus Joe) seems like the only sensible place to make economies. This sounds cynical, and may be, but it may be entirely unconscious. Our pet "solution" just shapes the way we see the problem.

Criteria often chase solutions. Suppose we got our Wizzo digitizer. Looking back a year later, it turns out that it did not work especially well on the project we bought it for. But we find ourselves justifying it on other grounds. It is important for laboratory morale that we have the most modern equipment; it kept us in touch with the Wizzo Company, so we get first news of other advanced equipment. Neither of these aims occurred to us at the time we made the purchase, but, looking back, they seem entirely sensible objectives. We are endlessly inventive in discovering ways of making our past decisions look sensible, not just to our bosses and our colleagues, but to ourselves. Read any autobiography to see how rarely people made mistakes—in retrospect! (See also "Rationality versus Rationalizing," later in this chapter.)

Problems often get clear only after they have been solved. Although the Wizzo device was of only marginal use for the project, we have had a lot less trouble with Joe, the senior technician, since it was installed. After some discreet inquiries, we find out why. A major source of his habitual bad temper before the Wizzo arrived was that Joe felt that his skills were being wasted in endless boring hours of entering data from laboratory records onto a keypunch machine. The new digitizer does all this automatically, and he is now able to do much more of the interesting, skilled work on which he prides himself. His mood is, understandably, much improved. The Wizzo solved a problem—Joe's dissatisfaction with his work assignments—that you didn't even realize existed. Once again, the orderly problem-solving sequence seems to have run backward.

Summary

We have defined the general process of problem solving as a series of activities aimed at reducing or eliminating a perceived difference between the way a situation actually is and the way we want it to be. Within this general process, we have discussed five kinds of activities or "phases." The first three—problem identification, alternative generation, and choice—are included in what is normally referred to as "decision making." Two further phases—implementation and monitoring of the chosen solution—emphasize the action-taking side of problem solving, the requirement that we do not simply solve the problem on paper, but put our thoughts into action and watch how they turn out.

Logically, we can think of these five phases as running in an orderly way from problem identification to solution monitoring. However, in practical settings, we often find the phases much less clear: we move to and fro through the different activities in all sorts of orders, rarely completing one phase before we move on to another, and with our activities at one phase shaping what we do at another.

SOME SIMPLE CHOICE MODELS

As we noted earlier, decision theorists have tended to focus on the "choice" part of the problem-solving process, ignoring the other parts. That is, they generally assume that the problem is well formulated, that all the alternatives are known, that the consequences of each alternative are understood, and that whichever alternative is chosen will be implemented. This focus obviously leaves out a great many interesting processes! It does, however, allow choices to be represented clearly, and it will help later discussion if we look briefly at how this is done.

Choice, to a decision theorist, concerns the evaluation of present alternatives in light of the future consequences of each. Do you want to toss a coin to see who will pay for both our lunches, or shall we each pay for our own? You are offered two alternatives. In one, you certainly pay for just one lunch; in the other, you may pay for two, or you may get lucky and not pay at all. Which alternative do you prefer? There are two useful ways to represent such choices: decision trees and payoff matrices. A decision tree is simply an orderly network linking alternatives to consequences, as shown in Figure 2.1. It is usual to use small squares to show decision points (here, "Toss" or "Don't toss"), and small circles to show chance events (here, the result of the toss).

An alternative way of representing the same problem is shown in Figure 2.2. Here the alternatives are represented by rows, the result of the coin-toss by columns, and the cell entries show the outcomes.

Clearly, neither of these representations is much help for a problem as simple as this. Their real value is in helping to clarify more complex situations. Decision trees are particularly useful when there are several chains of choices all linked together. For example, in trying to estimate the risk of a major accident at a nuclear power plant, one needs

Figure 2.1. A Simple Decision Tree for the Lunch Problem

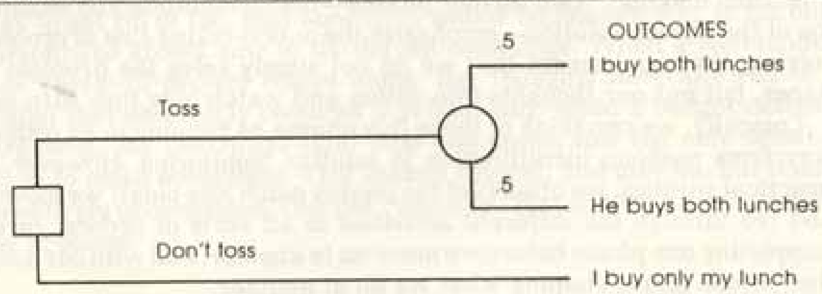

Figure 2.2. A Simple Payoff Matrix for the Lunch Problem

ALTERNATIVE	COIN COMES UP	
	Heads	Tails
Agree to toss	Buy both lunches	Buy neither lunch
Don't agree to toss	Buy my lunch only	Buy my lunch only

to consider many complex sequences of events: IF the main coolant pump fails AND the backup system is out for maintenance AND the operator opens the wrong valve accidentally OR this instrument malfunctions . . . A decision tree is of great value in tracing out the thousands of such strings of events, and estimating the chances of each occurring.

Payoff matrices are especially useful for representing situations in which the consequences of your decisions are influenced by the decisions of others. For example, your company's decision whether or not to market a particular product needs to consider whether or not your competitors introduce a similar product. As the number of possible alternatives open to each side grows, it is useful to construct a payoff matrix for each pair of alternatives, to be sure that all the possibilities are properly considered.

DECISION RULES AND UNCERTAIN CONSEQUENCES: THE EXPECTANCY MODEL

Decision trees and payoff matrices both focus our attention on selecting an action now in light of its likely future consequences. Looked at in this way, a decision involves two sorts of guesses about the future:

1. What will happen if I choose a particular action?
2. Will I like that set of consequences or not?

Decision theorists have generally considered the second guess as being relatively straightforward. They assume that decision makers know what they want, or have "stable preference orderings" over different outcomes. (There are, however, situations in which this assumption seems clearly wrong. March [1978] discusses several situations, such as eating unfamiliar foods or listening to unfamiliar music, in which one's actions are intended to *change* one's tastes.) Generally, though, tastes or preferences for outcomes are treated as "given."

On the first guess—what consequences will flow from different actions—decision theorists distinguish three situations:

1. **Certainty:** The consequences of each action alternative are known for sure. Action 1 will lead to Outcomes A and B; Action 2 will lead to Outcomes C and D. In this situation, choice reduces to simply evaluating each package of consequences against one's known preferences, and selecting the one which, on balance, offers the most of what one wants.
2. **Risk:** One does not know for sure the consequences of each alternative, but can assess a probability of each consequence resulting from each alternative. Action 1 will lead to Outcome A with probability 0.5, and to Outcome B with probability 0.1. Action 2 will lead to Outcome C with probability 0.7, and to Outcome D with probability 0.3. In this situation, it is often reasonable to apply the decision rule: "Discount the value of each outcome by the probability of its occurrence"—the so-called "expected value rule." For example, one might evaluate a simple gamble such as a coin toss with a one-dollar stake as being worth 50 cents. Over a long series of such bets, one would collect the dollar about half the time, so a 50 percent discount would be reasonable. (Again, there are interesting exceptions to this rule. For example, would you rather have $1 million for sure, or toss a coin for $5 million? The latter has a higher "expected value" ($2.5 million), but most people would hesitate to gamble the opportunity to be rich for life, and would choose the certain $1 million. However, the payoff-discounted-by-probability rule seems generally sensible for small stakes and repeated plays.)
3. **Uncertainty:** In this situation, one knows what consequences might follow from each alternative, but cannot assess the relevant probabilities. Action 1 might lead to Outcomes A and B, Action 2 might lead to Outcomes C and D, but the relevant probabilities cannot be estimated. A variety of interesting decision rules have been suggested for this situation (see, e.g., Raiffa [1968] for an excellent introduction). A cautious decision maker might select a "minimax"

strategy, choosing the alternative that offers the best set of outcomes if everything goes as badly as it could. A bold optimist might go with a "maximax" strategy, choosing the alternative that will yield the highest value if everything goes as well as it could. (Note, by the way, that it is easy to think of situations in which one is much more uncertain than this formal treatment of "uncertainty" allows. For example, one might have no idea of what alternatives are available, what the consequences of each might be, or how desirable each is. Decision theorists have had little to say about such extremely uncertain situations, though they are clearly common in the real world.)

Organizational researchers have made great use of decision models derived from the "expected-value" decision rule. In its simplest form, this "expectancy model" suggests that people's actual choices are, in fact, guided by the "payoff-discounted-by-probability" rule, though the outcomes considered, their desirability, and the probabilities of each flowing from a given alternative are all assessed subjectively. (Note that there is a long step here from the decision theorist's suggestion that people *should* choose in this way to the researcher's interest in describing how people *actually* choose.) The basic form of the expectancy model can be shown as a decision tree:

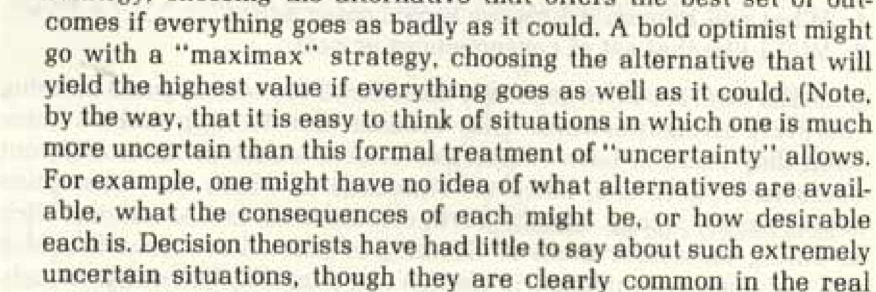

Decision rule: Select action alternative with highest net value, where

$$\text{Net value of action}_i = \sum_{\text{All } j \text{ Outcomes}} P_{ij} V_j$$

In this form, the model is merely a generalization to several outcomes of the decision rule discussed earlier under "Risk." Each action is linked probabilistically to several outcomes; the action's "net value" is simply the sum of the value (positive or negative) associated with each outcome, discounted by the probability of the action leading to that outcome.

For descriptiv
terms in the ab
about what con
assumption here
need to know wh
will lead to a pa
bility of that con
forms the heart
motivation, a the

GOOD DECISION

In the "game" c
one of the six ch
to his head and
empty chamber,
As it happens
The weight of th
empty chamber
odds considerab
even so, most o
insanity or extr
"play" is not a
There is an i
judge how good
sions can have
have good outco
you will lose in
decide badly an
If we cannot
there any way t
sion theorists w
which the decisi
dural rationality
the best chance
drawing on a l
criteria by whic
decision maker:

1. Reviews a w
2. Reviews the

For descriptive purposes, researchers relabel the "probability" terms in the above model with "expectancies"—subjective beliefs about what consequences will follow from each alternative. The assumption here is that, if we wish to understand people's choices, we need to know what *they think* the probability is that a particular action will lead to a particular consequence, not what the *objective* probability of that consequence might be. A descriptive model of this sort forms the heart of what is called the "expectancy theory" of work motivation, a theory we will examine in more detail in Chapter 4.

GOOD DECISIONS AND GOOD DECISION PROCESSES

In the "game" called Russian Roulette, the "player" puts a bullet in one of the six chambers of a revolver, spins the magazine, puts the gun to his head and pulls the trigger. If he is lucky, the hammer hits an empty chamber, and he collects the bet. If not, he is dead.

As it happens, the odds are said to be much better than one in six. The weight of the single bullet tends to make the cylinder stop with an empty chamber at the top, or firing position, improving the player's odds considerably. So, in most cases, the player will win the bet. But, even so, most of us would regard "playing" this game as an act of insanity or extreme desperation. Even if one wins, the decision to "play" is not a sensible one.

There is an important general point to be made here: one cannot judge how good a decision was by its outcomes. Even very good decisions can have disappointing outcomes, and very bad decisions can have good outcomes. Of course, if you bet against the odds repeatedly you will lose in the long run. But in any particular decision, you may decide badly and still get lucky (or vice versa).

If we cannot judge the quality of a decision from its outcomes, is there any way to distinguish good and bad decision making? Most decision theorists would argue that there is: examining the process by which the decision was arrived at. We look for what is called "procedural rationality," for a decision-making procedure that seems to offer the best chance of making a sound decision. Janis and Mann (1977), drawing on a large body of previous analysis, have identified seven criteria by which such "procedural rationality" may be judged. The decision maker:

1. Reviews a wide range of possible alternatives
2. Reviews the full range of objectives and values

3. Weighs both costs and risks, positive and negative consequences, that could flow from each alternative
4. Searches intensively for relevant new information on each alternative
5. Assimilates this new information carefully
6. Reexamines all alternatives before making a final decision
7. Makes detailed implementation and contingency plans, covering both expected consequences and anticipated risks. (Janis and Mann, 1967: p. 11)

These two authors describe such a process as "vigilant decision making." It is worth noting that they do not expect "vigilance" as the normal pattern. In fact, much of their work is concerned with the circumstances under which people will depart from "vigilance," and thus adopt less successful decision processes. (We shall have more to say about this in Chapter 9, on group processes.) We should emphasize, again, that no procedure guarantees satisfactory outcomes. The claim is only that a rational procedure, such as that sketched above, gives us the best chance of making good decisions and thus, on the average, enjoying satisfactory outcomes. We still need good luck and a benevolent environment to make any given decision a winner!

RATIONALITY

Most of us pride ourselves on being rational, especially in matters to do with our work. It is, however, often difficult to pin down exactly what we mean when we claim to have acted "rationally." Is it rational for you to be reading this book? Was it rational for me to write it? I hope neither decision is obviously crazy, but there is a long gap between "clearly rational" and "obviously crazy." What is changing as we move from one to the other?

"Rational" is generally used to describe decisions or actions based on thinking and reasoning. It is perfectly sensible to snatch your hand away from a hot stove, but we generally think of that as a reflex action rather than a rational one. It is perfectly sensible to eat plums (if you happen to like them), but we generally see that as a matter of taste, not a matter of rationality.

We can define rational behavior in terms of the general problem-solving model discussed earlier. Strictly, we would be choosing rationally only if:

1. We could specify exactly what we were aiming for
2. We examined every possible alternative action open to us
3. We could predict all the consequences of each alternative
4. We chose the alternative that generated most of what we wanted (and least of what we did not want).

This set of requirements will be familiar to anyone who has studied economics. Economists build much of their theory on the assumption that people, at least in the aggregate, act rationally in this sense.

Practically, of course, we cannot hope to be fully rational by these standards. Our values and aims are at least somewhat unclear and conflicting. We can never be sure that we have examined every possible alternative. Even for the ones we do examine, we can only anticipate some of the more immediate consequences, and have great difficulty judging how likely each is. And, as we try to do more and more of each of these things, we get more and more overwhelmed with information, and it becomes harder and harder to find the time, energy, and sheer mental ability to digest it all and pick the best alternative.

(If you find yourself agreeing that this is a fair description of other people, but not of yourself, think for a moment about some trivial purchase you have made recently—say, a toothbrush. What would it take to make this a strictly rational purchase decision? How worn was your old one? How serious is using an old toothbrush, in terms of later dental decay? How good was the new one? Did you check other shops for a better price? Within what area? What does it cost, in time, gasoline, and wear on the car to drive to another store? And so on, with no limit in sight. Obviously, it is crazy to devote a full-scale effort at being rational to trivial decisions like buying toothbrushes; but it should be just as obvious that we could not achieve complete rationality, even if we were to try.)

SATISFICING

We should not assume that, because people are not completely rational in the strict sense, they are irrational. It is clear that people often, especially in matters connected with their jobs, *intend* to be as rational as possible. Herbert Simon, who won a Nobel prize for his studies of organizational decision making, suggests that we should think in terms of "bounded" rather than strict rationality. These "bounds" are imposed both by our weaknesses as decision makers and by the difficulties of the situations we have to cope with. We get by reasonably well in most cases by simplifying the decision process. That is:

1. We consider only a few major value dimensions
2. We examine only a few possible alternatives
3. We consider only some of the main consequences of each alternative
4. We choose the first alternative that reaches some level we think of as "good enough."

To emphasize the distinction between this model and the strictly rational process, Simon uses the term "satisficing" for the boundedly rational process, in contrast to "optimizing" for the strictly rational process. He suggests that people are better described as "satisficers" than as "optimizers." Satisficing certainly fits better with the psychological evidence we shall look at in Chapter 3 than does optimizing.

The idea of satisficing as a decision process has several important implications for understanding how decisions get made in organizations. One is that, while there is only one optimal solution to most decisions, there may be many equally good satisficing solutions. For example, I may review alternatives in a different order than you do. If we each stop when we find an alternative that is "good enough," we can easily find ourselves disagreeing about what should be done. This will often mean that the final choice has to be made on some ground other than the logic of the problem itself, for example, which of us is senior, which of us is more persuasive, which of us gets in first. Such factors do, in fact, make a great deal of difference in most organizational decisions—as they could not if decisions were made on strictly rational grounds.

A second interesting implication of "satisficing" is that, if we couple it to a mechanism for resetting our ideas of what is "good enough," we will often edge in on an optimal solution. Researchers use the term "level of aspiration" for the standard that an individual currently regards as "good enough." For most of us, if we do better than our "level of aspiration" on one try, we tend to set it higher the next time. Conversely, if we fall short of our "level of aspiration," we tend to revise it downwards. This resetting can get us quite close to optimality, without having to go through all the difficult steps noted earlier. For example:

Suppose you are in charge of a small manufacturing plant, and have to make a decision as to how much the plant will produce. If you set production targets too low, the plant will have idle capacity, sales will be lost, and profitability will suffer. On the other hand, if you try to produce too much, machine maintenance may be neglected, people will work too hard (and perhaps get sick or injured), and so on. Finding the optimal balance is, clearly, a very tricky decision. As a satisficer, however, suppose you set a target at 105 percent of last month's production. If you achieve this level without too much strain for a few months, you may reset your aim to 110 percent. Perhaps this level shows a number of danger signals, and cannot be reached without stress. You reset your aim down a little, perhaps to 108 percent, and find that this works satisfactorily on a regular basis. You have found an acceptable balance between overproduction and underproduction —without ever having to solve the highly complex problem of finding the analytic optimum.

There are, of course, traps to trying to solve problems in this way. As an analogy, consider the problem of getting to the top of a mountain (where the top represents the true "optimal solution"). If you have an accurate map of the entire mountain, you can locate the top and plan the best route to get there—that is, you can use the "analytic" approach. Suppose, however, that you do not have a map. You may do quite well by merely checking which way is uphill from where you are at any moment, and climbing that way. This will work fine if the mountain is a single peak, like a cone. However, if the mountain has several minor peaks, as well as the real one, you may need to go *downhill* at some point on your way to the top. The satisficing-and-search method will trap you near one of the minor peaks. In other words, it is a potentially "conservative" approach. It will help you make minor improvements from where you are now, and with luck these improvements may add up to a very large gain. But it may trap you in dead ends, while the really large payoffs are somewhere else.

LIMITATIONS ON RATIONALITY

As we have seen in several of the examples discussed thus far, rationality is very difficult to achieve. Indeed, it is probably no exaggeration to say that it is impossible to achieve, in any strict sense, once one gets beyond highly simplified situations such as gambling games. Even when we try our best to be rational, we cannot achieve it fully. It will be convenient to list here the major constraints that limit our approximations to rational choice.

Human Information-Processing Limits

We are able to think of only a few—a very few—bits of information at a time. Without trying to specify exactly how many, a reasonable estimate is somewhere in the region of seven distinct, meaningful "chunks" of information (Miller, 1956), a humblingly small number. These "chunks" can vary in information content from single digits to complex symbols or images, but we seem to be able to process only about seven of them in our short-term memories. Clearly, we are better off using rich "chunks" than otherwise, hence the critical role of powerful notational systems such as matrices: they allow us to think about a complex set of relationships as a single symbol. Several fascinating studies of chess players suggest the same point. Grand masters are no better than novices at remembering random arrangements of chess pieces on a board. They are, however, hugely superior at remembering positions from real games. It appears that they have the ability to store board positions as powerful "chunks," where a large region of the board is stored as a single "chunk": a "Sicilian corner," for example. Thus, learning dense, informationally rich chunks is a powerful way of using what little short-term memory we have, but about seven bits remains the upper limit. This is clearly a critical constraint on our ability to weigh up mentally the many factors that enter into any complex decision.

Situational Pressures

The organizational world is often a busy, even a hectic place, especially as one moves into the ranks of management. Studies of how typical managers spend their time (for example, Mintzberg, 1973) have found what experienced managers already know: the day is broken up into many small activities, one thing is constantly interrupted by another, and attention must be rapidly redirected from one matter to another. In fact, Mintzberg found that a typical managerial activity took less than ten minutes! Obviously, it is sometimes possible to allocate longer chunks of time to an important issue, to break up the problem into smaller pieces that can be dealt with singly, or to work on it outside of normal business hours. But, equally obviously, such long blocks of uninterrupted time are the exception, not the rule. Time and attention are scarce managerial resources. Most decisions are made with a minimum of either, suggesting a very different way of deciding than the cool, reflective, careful analysis implied in our earlier discussion of rational decision processes.

Informational Limits

Managers appear to live in an informationally rich environment. Their world is full of memos, reports, computer printouts, production figures, personnel reports, journals, laboratory reports, and so on. But does this make them well informed in terms of the decisions they have to make? In many cases, it does not. As Ackoff (1967) suggests, it may be more accurate to think of managers as suffering from an overabundance of irrelevant information. True, the particular number you are

the critical role of
ey allow us to think
ymbol. Several fasci-
oint. Grand masters
om arrangements of
superior at remem-
hey have the ability
ere a large region of
cilian corner," for
chunks is a power-
we have, but about
critical constraint
that enter into any

ic place, especially
ies of how typical
, 1973) have found
is broken up into
upted by another,
atter to another. In
vity took less than
o allocate longer
the problem into
rk on it outside of
ng blocks of unin-
and attention are
ade with a mini-
leciding than the
ier discussion of

vironment. Their
oduction figures,
so on. But does
ns they have to
ggests, it may be
m an overabun-
number you are

looking for may be right on your desk. But if it is buried in the middle of a 200-page printout, it might as well not be there at all.

An English organizational researcher, Anthony Downs, suggests three important features of a manager's information supply. (1) When a manager starts to think about a problem, available relevant information is likely to be scarce. (2) Additional information can frequently be found, but at a significant cost in money, effort, and delay. (3) Some of the most critical information is never available at decision time, that is, information about the future. This last, of course, is inherent in any decision process, and the risk or uncertainty of future events must be considered in making the decision. In the next chapter, we shall look at some of the errors that we are all prone to make in predicting events.

As a practical matter, we should not expect decisions made by real people in real organizations to meet the criteria for strict rationality. Put bluntly, in any reasonably complex problem we are generally not smart enough, and we do not have (and cannot get) enough information to choose rationally. Thus such simplifying decision processes as satisficing are not simply a matter of taking a lazy, second-rate approach. In most settings they are all we are capable of.

RATIONALITY VERSUS RATIONALIZING

As we have seen, choosing rationally is a very difficult business, especially when the decision is an important one. However, we have strong cultural values in favor of rational choice, and these values are reinforced in the context of organizational life. We are insulted if a friend says of a decision we have made, "That's irrational!" If the comment comes from our boss, it is worse than an insult: it is a threat to our job.

The central difficulty about choosing rationally is that it turns on our guesses about the future. *What if* we choose Alternative A? *What if* the stock market falls twenty points? *What if* I ask that person for a date and he/she turns me down? Questions like this are immensely easier if we are looking back in time, rather than forward. (See the next chapter for a discussion of prediction errors.) Once we know what has happened, it is usually easy to reconstruct our choices so that we look rational in retrospect—this is the process we call "rationalizing." The sorts of remarks that tip you off that someone (yourself, maybe?) is rationalizing include:

"I realized it was risky at the time." (Something has gone wrong.)

"I had a hunch that it would work out." (Something has gone right.)

"It got messed up down the line." (I made a good decision but someone else implemented it badly.)

"Looking back, what I was really trying to do . . ." (The choice did not meet the original objectives, but here is another set of objectives it *did* meet.)

As usual, it is a lot easier to spot these tricks when someone else is playing them, but it is worth bearing in mind that we all spend a lot of time rationalizing—making sense of the past, rather than of the future. This has costs. For one thing, it takes up time and energy we could be better using for something else. Rationalizing also prevents us learning from our mistakes: if we see no mistake, there is nothing to learn from! And it sets false expectations, for us and for others. If everyone else is making all these superrational decisions, how come I find them so difficult? Maybe I should be cautious and not confess my confusions. This leads to defensive decision making. It becomes less important to make the right choices than it is to have a defensible way of making them.

Rationalizing is both common and dangerous. As we shall see later in this book, we all do it, very often without being aware that we are. We are taught from an early age to "act sensibly," so that rationalizing becomes second nature to us. Organizations reemphasize this value. As a result, much time and energy are absorbed in making sense of the past instead of making sense of the future. And choices are often made so as to be defensible, instead of being right.

Still more worrying, some recent studies (e.g., Staw, 1980; Staw and Fox, 1979) suggest that our tendency to rationalize the past may even shape our future decisions. For example, if we have been responsible for funding a research project that seems to be failing, we are prone to pouring more money into it, in hopes of bailing it out and thus making our original decision look good, instead of cutting our losses and going on to something else. As well as wasting time and effort "reinventing" the past, rationalizing may distort our future decisions, too.

PROGRAMMED VERSUS UNPROGRAMMED DECISIONS

Given the time and energy that are used to solve a problem, even if we are not especially rational about it, we all develop ways to economize on this effort. Some of these ways seem to be "wired in" as physical reflexes. We do not *decide* to snatch our hands from a hot flame, or to blink when an object comes close to our eyes. At a second level, most of us develop habitual ways of doing things, so that they take virtually no mental effort or attention. For example, I have a fairly automatic sequence of activities for getting up in the morning. From the time the alarm rings to the time I get into my car, I have gone through 45 minutes of tooth-brushing, showering, coffee-drinking, and so on,

without thinking abou
being a morning pers
example, by a telepho

Between elaborate
large category of "p
habits, in the sense of
are they full-scale dec
in computer programm
calling up different br
vated, a predictable se

For example, once I
same way. If a certain
seems to be heavy on t
regular search routine
you recorded my route
dozens of different alte
this is not to say that I r
one will work best that
really generated by a fe
by each.

Organizations develo
many of them highly el
computer programming,
get a lot done without thi
for example, has "prog
buildings cleaned, and
without anyone thinking
ment has to solve a larg
grammed" way before
student should learn, ho
will live, and so on. But o
almost effortless. A few
ment? specialty? financ
unfolds automatically.

This marvel of efficien
make the work very bor
motor is complex and fu
formance specifications i
problem is that programs
they have to treat people
beings experience this as
categories, we are forced
dealt with as a special cas
program turns on having a

without thinking about it or even being aware I am doing it. In fact, not being a morning person, I am rather upset if I am forced to think, for example, by a telephone call in the middle of my routine.

Between elaborate problem solving and such habitual routines is a large category of "programmed decisions." These are not really habits, in the sense of unfolding in the same way every time. But nor are they full-scale decisions. They are more like packaged subroutines in computer programming. They may have a variety of "IF" statements calling up different branches, but, once a particular branch is activated, a predictable sequence unfolds within it.

For example, once I am in my car, I do not always drive to school the same way. If a certain traffic light is red, I turn earlier. If the traffic seems to be heavy on the freeway, I get off onto a side street. I have a regular search routine for finding a parking place on the campus. If you recorded my route from home to office for a year, you might find dozens of different alternatives, all of which I used at least once. But this is not to say that I review them all each morning and choose which one will work best that day. My complex pattern of switching routes is really generated by a few choices, with large chunks of routine evoked by each.

Organizations develop hundreds of decision "programs" of this sort, many of them highly elaborate. Just as with packaged subroutines in computer programming, these decision programs let the organization get a lot done without thinking very much about each case. A university, for example, has "programs" for admitting students, for getting the buildings cleaned, and so on. Once they are set up, they run almost without anyone thinking about them. A newly created university department has to solve a large number of difficult problems in an "unprogrammed" way before the first student arrives: what materials the student should learn, how he or she is to be examined, where he or she will live, and so on. But once you have the "program" established, it is almost effortless. A few quick decisions—regular or advanced placement? specialty? financial aid?—and the whole complex sequence unfolds automatically.

This marvel of efficiency has its costs, of course, One is that it may make the work very boring for the employees. Designing an electric motor is complex and fun, but not if it is reduced to plugging the performance specifications into a design package. A second, and familiar, problem is that programs require standardized rules to evoke them, so they have to treat people and problems in preset categories. Human beings experience this as depersonalizing. If we happen not to fit the categories, we are forced into one that does not quite fit, or we are dealt with as a special case—often painfully, since the efficiency of the program turns on having as few special cases as possible. Anyone who has tried to take a nonstandard mix of courses will know the problems one can encounter.

An important aspect of understanding how organizations work is to be able to trace out these decision programs. What matters do they cover? How elaborate are they? How are exceptions dealt with? Organizations differ a great deal in these areas. We tend to think of government agencies as having very elaborate decision programs: for example, what welfare benefits is a particular client entitled to? (These rules and procedures, as well as being relatively efficient, help to offset personal biases on the part of the welfare caseworker.) On the other hand, a small research-based company might be highly unprogrammed, with very few standard procedures and most issues decided on an *ad hoc* basis. These two organizations will operate very differently—and they will be very different as places to work in!

DECISION EVENTS AND DIFFUSE DECISIONS

Thus far most of the discussion has considered a decision as essentially an event. A single individual, at some point in time, weighs up a set of alternatives against a set of criteria and choses one alternative to implement. In short, decisions are identifiable events; decision makers are identifiable individuals.

However, the stress on the single decision maker, and the single decision event, is misleading as a description of what actually happens in organizations.

Some years ago, when I was working on my dissertation research, I became interested in NASA, the American space agency. At that time NASA was spending hundreds of millions of dollars annually on research projects. The decisions as to what got funded and what did not seemed interesting, and my dissertation was intended to look into how these decisions got made. Obviously, the first thing to do was to find the decision maker and go talk to him or her.

Finding the decision maker turned out to be a frustrating business. No one would admit to making such an important decision. The research scientists on the bench denied making it: they merely submitted outlines of projects they thought would be useful or interesting and waited to see what got funded. The branch chiefs in the laboratories denied making it: they simply reviewed the proposals they received, and sent the most promising ones up the line. Senior laboratory managers denied making it: they merely reviewed projects on the basis of promise, balance across different laboratory areas, and so on. And so it went. At every level of management, everyone I talked to felt that the "real"

decision was made som
recommendations, prop
on—everything except
senior people in the age
As one of them told me,
little actual choice left-
down the line."

It took me a long time
confusing answers was
the equivalent of walki
makes the cars around
is what makes cars—ar
tional decisions. Decisi
and involve many peopl
ent geographical locati
cesses," to emphasize
tional level, and geogra

It is useful to think i

Time-frame
Number of participants
Organizational locations
Geographical dispersion

I have found this idea o
framework in organizi
book. Most of the pres
with focused decision
single point in time. Th
in more depth at how w
ing our decisions. Cha
people make decisions
relatively focused, indi

Later chapters move
processes. For exampl
can be seen in terms

decision was made somewhere else. Of course, each gave advice, made recommendations, proposed, reviewed, coordinated, evaluated, and so on—everything except actually making a decision! Even the most senior people in the agency's management denied making the decision. As one of them told me, "By the time suggestions reach me, there's very little actual choice left—it's already been made by the people further down the line."

It took me a long time to realize that the reason I was getting all these confusing answers was that I was asking a stupid question. I was doing the equivalent of walking into an automobile plant and asking, "Who makes the cars around here?" The answer is, of course, that a *process* is what makes cars—and a *process* is what makes important organizational decisions. Decision processes extend over long periods of time, and involve many people at various organizational levels and at different geographical locations. I christened them "diffuse decision processes," to emphasize their dispersion over time, people, organizational level, and geography (Connolly, 1977).

It is useful to think in terms of a spectrum of decision phenomena:

	Focused decision events	Diffuse decision processes
Time-frame	Short	Long
Number of participants	One	Many
Organizational locations	One	Many
Geographical dispersion	None	Large

I have found this idea of a spectrum of decision types to be a very helpful framework in organizing my thinking—and, indeed, for organizing this book. Most of the present chapter, for example, has been concerned with focused decision events: single individuals making choices at a single point in time. The next chapter continues this emphasis, looking in more depth at how we gather and make sense of information in making our decisions. Chapter 4 treats motivation mainly in terms of how people make decisions on such matters as how hard to work—again, a relatively focused, individual-level decision.

Later chapters move into more multiperson, diffuse kinds of decision processes. For example, matters of power, influence, and leadership can be seen in terms of one individual affecting the decisions of another. Chapter 8, on conflict and conflict resolution, clearly has this same interest, since we look at choices shaped not only by one's own interests and desires but by those of others. Communication (Chapter 5) is the process by which information is transmitted and thus the process by which individual decisions are coupled together. The multiperson emphasis emerges still more strongly when we look at decision making in groups (Chapter 9). Finally, in Chapter 10, organizational structures are treated as procedures (both formal and informal) by which diffuse decision processes are made more orderly and structured.

SUMMARY

We have defined a "problem" as a perception that a significant difference exists between actual and desired situations, and a "solution" as an action taken to close this gap. The process between includes several components: problem identification, alternative generation, choice, implementation, and monitoring. These five components do not generally follow this neat sequence, but tend to be interwoven with one another, often in complex ways. No matter what sequence is followed, good outcomes cannot be guaranteed, because the future is always uncertain. However, we did propose several criteria for decision processes that seem to raise the probability that good decisions can, on the average, be made.

It is clear that human beings are not capable of "strict rationality" in the optimizing sense. Our efforts at rationality are constrained by our limited ability to process information, and by scarcity of information, time, and attention in real organizational settings. We should think of behavior in terms of "intended rationality," where we "satisfice" by searching for the first alternative that is "good enough" in terms of some current "level of aspiration." Revising levels of aspiration after successful and unsuccessful search can allow us to make continual incremental improvements, but may trap us into local, rather than global, maxima.

Although rationality is hard to achieve (or even to approximate), we are often expected to pretend that our actions are, in fact, rational. This leads to "rationalizing," attempts to reconstruct the past so as to make decisions that were made appear to have been rational. Rationalizing is dangerous: it makes it difficult to learn from our mistakes, it can lead us to make defensible (rather than good) decisions, and it can trap us into throwing good money after bad, trying to justify the original bad decision.

Most organizational decisions are made in a "programmed" way, with a series of predictable steps unfolding one after the other rather like a computer program subroutine. These "programs" are important to the organization's decision-making efficiency, and free up time and resources for the rare, "unprogrammed" problem. They do, however, standardize procedures, making decisions less flexible, and the work less interesting.

The last section of the chapter suggested a spectrum of decision types, running from the focused decision event (one individual choosing at one point in time) to the diffuse decision process (many individuals involved over a long period of time). Much of the remainder of this book can be seen in terms of this range of decision activities. Individual processes of perception, inference, judgement, and motivation emphasize the decision-event end of the scale. Authority, influence, leadership, and conflict fall in the middle of the scale. Communication, groups, and organizational structure raise issues toward the highly diffuse, decision-process end of the scale.

The processes of problem solving, decision making, and choice clearly rely on the people involved having access to relevant information. In the next chapter, we discuss how people acquire and make sense of information, how the processes can go wrong, and some of the things we can do to avoid these traps.

DISCUSSION QUESTIONS: CHAPTER 2

1. Select a problem that you feel is important, either in a broad social sense (e.g., energy, the environment) or in a personal sense (e.g., how to make more money, what job to choose). For the problem you have chosen:
   a. Write a brief, concise problem formulation, including the criteria for a good solution.
   b. Suggest three alternatives that might solve the problem.
   c. Describe how you would choose among the alternatives. (You may have to invent numbers for costs, probabilities, etc., but try to use reasonable numbers.)
   d. Describe how you would go about implementing your chosen solution.
   e. Suggest how you might monitor the situation so as to be sure that your chosen solution was actually working.
2. After completing the above assignment, reconsider your problem formulation. Do you wish to change or clarify it? Do any new alternatives occur to you? Might your choice now be different?

3. Suppose you are considering two different ways of spending your next day off: going to a concert or going to a picnic. The picnic will be more fun if it's sunny, but the concert will be more fun if it rains.
   a. Draw a decision tree to represent your choice situation.
   b. Draw a payoff matrix to represent your choice situation.
   c. Suppose you value the four possible outcomes as follows:
      Picnic and sunny: + 10
      Picnic and rainy: − 10
      Concert and sunny: + 6
      Concert and rainy: + 8
      Which activity would you choose if you were sure it would be sunny? If you thought the probability of rain were 50 percent? If you had no idea what the probability of rain was?
4. Why can't we evaluate how well a decision was made by just looking at its consequences? How can we rate how well a decision is made?
5. What do we mean when we say a decision maker is "satisficing"? How is this different from "optimizing," and when might the two amount to the same thing? What factors make it likely that satisficing is more common than optimizing in real decisions?
6. What is the difference between "rational" and "rationalizing" activities? Why might the latter be dangerous in organizations?
7. What is the difference between "programmed" and "unprogrammed" decisions? Which is more common in organizations? Why?

Chapter

Percepti
Inferenc

MAJOR TOPICS