

Decision Making: Rational, Nonrational, and Irrational

Herbert A. Simon

In a style easily accessible to readers without a background in psychology, the author describes the current state of knowledge about human decision-making and problem-solving processes. He also explains how progress has occurred in this field and what some of the central implications for management and management training are. In a final section of this article, the author provides responses to questions from the audience that significantly extend the scope of the article.

PSYCHOLOGY AND EDUCATION

Both psychology and education are disciplines concerned with learning. In this article, I have much to say about learning, and management, as well. Between psychology and education, there needs to be a much broader two-way street. This should not be simply a street through which all of the great ideas that psychologists develop and research in their laboratories can be turned over to education for implementation. It should be a genuine two-way street in which contact with the process of education and educational institutions can provide psychologists with basic research sites, research materials, and insights into the process that can flow back to psychology.

I spent half of my career in professional schools—various schools of business and engineering—and it has always seemed to me that contact with what a humorist would call the real world (as if our universities weren't part of that world) is usually the best source of research problems. Simply reading the journals and finding out what your colleagues did last month may keep

Guest Editors' Note: This article is a modestly edited version of an invited speech delivered at the 1991 annual meeting of the University Council for Educational Administration held in Baltimore.

© 1993 The University Council for Educational Administration

you in touch with the current fads, but the topics we are dealing with, learning and management, refer to phenomena that are going on all the time in the real world. Professional schools have an opportunity for deep contact with that real world. Through such contact, they can become the source of the problems that we should be trying to understand in our research. It is terribly important to build and maintain that two-way street between psychology and education. I don't need to tell you how narrow that street was and how little traffic was on it during most of the history of behaviorism.

We are talking about a psychology (primarily cognitive psychology, but other areas of psychology also) that is more prepared than it has been for several generations to deal with problems at the level of complexity encountered in real educational institutions. Hence there is something worth communicating back and forth now. Psychology, as a discipline today, is prepared to look at real-life situations, applied situations, as a source of ideas and research problems in a way that psychology a generation ago was not.

RATIONALITY

Decision making is a process of considerable interest to both psychologists and educators who look to real-life situations for the source of their research problems. A central question that arises immediately, in trying to understand decision making, is where rationality comes into the picture. We talk about the *rational*, the *nonrational*, and the *irrational*. I don't mean anything very complicated by rational. Behavior is rational, and the decisions leading up to behavior are rational if it turns out that the behavior prescribed is well adapted to its goals—whatever those goals might be. Rationality is the set of skills or aptitudes we use to see if we can get from here to there—to find courses of action that will lead to the accomplishment of our goals. Action is rational to the degree that it is well adapted to those goals. Decisions are rational to the extent that they lead to such action.

Irrational? Well, that is easy then. Irrational means poorly adapted to goals. As far as I can see, there is no such thing as excessive rationality; we use that phrase glibly but doubtfully. When you accuse somebody of being excessively rational, you must have something else in mind. You must have a quarrel with that person's goals and not with his or her thinking.

Nonrational reminds us of something else again. It reminds us that the goals themselves have to be postulated somehow in the decision-making process, except insofar as certain goals may themselves be instrumental to other goals: We go to school *in order* to get an education—or so we tell students they ought to. But somewhere there has to be a fulcrum for the whole

business. As Archimedes said, "If someone will give me a foundation for my fulcrum I can move the whole world." But he needed a foundation.

Those final goals, the things that somehow or other we regard as the ends in themselves (except insofar as they have side consequences and except insofar as they really are thought of as implementing other goals) have nothing to do with rationality. They must come from somewhere else. I will use the term nonrational to refer to those aspects of the decision process that relate to these very final goals.

DECISION MAKING

Now what is meant by the term *decision making*? I use that term very broadly to encompass three classes of things. The first class is finding problems that need attention and attending to them. Our characteristics as human beings living in the kind of world we live in (probably living in the kind of world that human beings have lived in since the species was invented) have always provided us with a very large reservoir of potential problems that might be attended to. One of the important skills we have to acquire in life, and one that has to be in view of our institutions if they are to operate properly, is the skill of finding and attending to problems. You might say, "Well, who has to find problems? I already have some." When I speak about finding and attending to problems, I mean finding and attending to problems that are so important that they ought to be taking up the very small budget of attention that we really have. Those hours between getting up in the morning and going to bed at night are all the hours we can give, no matter how many problems are in our world. So, by *finding and attending to problems*, I am referring to setting priorities and setting those priorities appropriately. Deciding what we, as individuals, and what our organizations will be attending to is a crucial part of the decision-making process.

There is a second and more familiar (but not as familiar to economists as it is to psychologists) part of the decision-making process. Once we know what problem we have on our hands, we have to start thinking about what alternatives, what kinds of solutions, might deal with that problem. Solutions aren't handed to us. We are not given an inventory or a list of solutions. "Oh, you have a problem? Well, consider one of the following." Alternatives themselves have to be invented. The whole activity that we call design, the things that engineers and architects and, as a matter of fact, all the rest of us spend the bulk of our serious time doing is dreaming up, elaborating, and crafting possible solutions to the problems that we have decided are our priorities.

Third and finally, there is the matter of evaluating those solutions and choosing among them. But if the first two jobs have been done well—namely, deciding what to attend to and doing a good job of designing possible courses of action, then in many ways this process of evaluating and choosing—I won't say is trivial—but is really not the major part of the job. I have left Point 4 out, of course. I have left out the implementation, for that is not going to be the major burden of this piece; but I will say something about it before I am done.

The decision-making process and the rationality that underlies it, or should underlie it, has been a central topic of consideration in the social sciences. All of the social sciences, in one way or another, have to deal with what is rational behavior and how it comes about. My whole career has depended on that fact. I have been accused of flitting from one social science to another, but the secret really is that I have been preoccupied with one topic, decision making and rationality, all of my life. It just happens that that topic cuts squarely across all of the human sciences because it is what we human beings are doing a great deal of the time. You don't really have to change very much, except a bit of your vocabulary, to move from one of these fields to another, if you stick to the topic of decision making.

ECONOMIC RATIONALITY

Rationality and decision making are necessarily central topics of economics, management, and, of course, psychology. Economics has a very explicit theory about what is rational behavior, whereas psychologists talk less often about rationality and more often about thinking or maybe "good" thinking. And there are very considerable differences between the psychologist's and the economist's approach to rationality, which is the principal approach one finds in the management literature. In economics, rationality is looked at as something substantive. You decide what is rational by looking at the action taken. Does it in fact achieve the goals? We can call that substantive rationality. In psychology, the concern with rationality is procedural. How does one go about it? What processes have to take place? What mental processes occur? (The mind is again in the head in postbehaviorist psychology today.) What mental processes have to take place in order that decisions will be reached that are rational and well adapted to the goals? And so economics and psychology have gone rather separate ways.

I will say just a few words about economic theory before concentrating on psychology. You can see where my loyalties lie. The theory of economic

rationality focuses only on the substance of choice, to the relative neglect of the process, and it really deals with only a third of the three parts of the choice process. Economic theory does not deal with focus of attention or emphasis; economic theory does not deal with where alternatives of choice come from; economic theory deals solely with the question of how, given a menu of possible actions, you choose among them. You all know what the answer is. You choose the one with the greatest utility; you maximize utility. Sometimes utility, by imperceptible steps in the economics textbooks, creeps over from utility to monetary profit. Or, if you are dealing with consumers, it may be wealth. But that is not the essence of the theory, that is just mismanagement of it. At the essence are two things, a focus on the choice among alternatives using some vague criteria of utility (anything can be poured into the definition of utility) and, second, people reaching for the best. Nothing is good enough unless it is optimal.

Of course, that is where my concept of bounded rationality comes in. Maximizing utility bears no resemblance whatsoever to what we human beings actually do. The idea that we even have a conception of what would be optimal behavior in the complex situations of life is unbelievable from the beginning. In our more modest moments, I think we are willing to believe that the world is far more complicated in its structure than our minds are able to grasp in its totality. Therefore, we get along in this world by having some priorities, by dealing with the things that really have to be dealt with, and by finding courses of action that are not disastrous. That is putting it a little negatively; putting it more positively, we look for courses of action that are satisfactory. Well, what is satisfactory?

We have built into us a mechanism called aspiration levels. We usually make assessments of what the world is likely to provide for us if we work at it a little bit: what it is reasonable to expect to achieve, if we make good decisions and if we follow them out. So we form aspirations. Many, many years ago (it must be 40 or 50 now), *Fortune* magazine took a poll of the American public, asking people how much *more* income they would have to have each year so that they would stop worrying about money. They got a very clear answer, and the interesting thing was that they got the same answer independently of the income level of the people they polled—millionaires down to the rest of us. The answer was "I need 10% more than I have right now." That is a very characteristic property of our aspiration levels. They adjust to realities; they adjust to our expectations, based on experience (our experience and the experience of others), about what is attainable. And that becomes a major standard by which we judge whether a solution to a problem is a satisfactory solution.

Of course, we do this because we are simply incapable of doing anything very different. Because of the limits of our rationality, we are simply not in a position to say: "Here are all the possible alternative courses of action; I have evaluated them; I know this one is going to work out best, so I will take it as my option." There is a limit on our knowledge. Who knows what the weather is going to be in Baltimore tomorrow morning? Or in Pittsburgh? There are limits on our knowledge, limits on our ability to compute, to work out the consequences of what we do know. And, of course, there is a further problem, because final values themselves are not subject to a rational calculus, they are simply posited. We also have the problem, if we are talking about social action, of agreement on values. So for all of these reasons, we have to find a way of reaching good decisions, reasonable decisions—satisficing decisions, I like to call them—a way that is consistent with and compatible with the limitations on our mental capabilities.

We can sometimes "soup up" these capabilities a little bit. Writing was a great invention because among our terrible limitations is the limitation of short-term memory. You know the difficulty we have of keeping more than one phone number in mind as we go from the telephone book to the phone. But writing came along, and we can write the telephone numbers down, if we are clever enough to think of it. Today, we also have other kinds of help, many of them due to the computer, many of them due to mathematics. We have had now, for 50 years, the disciplines called management science and operations research.

How do operations research models or management science models help us, with or without computers? Most of these models are built around the economist's view of rationality. They try to optimize something. In a so-called linear programming problem, you define something you are trying to maximize and indicate some limitations of resources, some constraints, and some boundaries on the actions you ought to take. Then you give this information to your friendly computer, and it reports out the optimum. For example, if you give it information on the number of calories, vitamins, and so on in each of the foods you can buy in your supermarket and the price tags on those and the number of calories you need for a healthy diet, a linear programming problem will determine the lowest-cost diet that will meet all of those needs. It will optimize. You had better like pork liver and cabbage, however (it is not necessarily as bad as that, because of course you could put another constraint on the problem and say no liver). When you look at how we use these management tools (and they have shown themselves to be extremely valuable, especially at middle management levels) to overcome our bounded rationality—and even the bounded rationality of computers—

we do several things. First, we cut the problem down to size; we horribly oversimplify the problem. We replace the original problem by this simplified problem, shaping it carefully so that it is within the bounds of whatever computational resources and information we have available. Having done that, we can optimize this approximate problem and thereby produce a satisficing solution (maybe) to the real-world problem. We shouldn't suppose that computers or numerical techniques or any things of this sort have exempted us from the million-year restriction human beings have had of satisficing instead of optimizing. It simply expands a little bit our capabilities of considering a wider range of solutions and evaluating them a little better. But we are still, in our thinking, employing severely limited approximations to the real world.

PSYCHOLOGICAL RATIONALITY

Let me stop economist-bashing now and move on to psychological theory and what has been learned about the actual processes that human beings use in decision making. We have made a very great advance, over the past 30 years, in our understanding of the processes that humans use for arriving at decisions. Attention and motivation (how things get on to the agenda of priorities) is still in an embryonic state. People are beginning to investigate it now, but I can't report dramatic progress yet in that area; give it another generation or so. The real progress has been made on the second and the third aspects of decision making, as I defined it earlier: theories about how we generate schemes of action, how we solve problems and how we arrive at choices as a result of solving those problems. A good deal can be said today about these matters that is applicable to real-world situations in education and elsewhere.

How was such progress made? As you have probably discovered, there are still a few rats in some psychological laboratories. But, by and large, the rat is out and human beings are in. Not only are we willing to talk about minds again, but we have found some ways of studying the human mind. Now it should be made clear that rats do have some advantages, particularly because you are allowed to do a lot of things with rats that you are not allowed to do with human beings (especially if you have a local "human subjects" committee). A major problem in any science is how to find instruments that will allow you to look at the relevant phenomena. We have gained a couple of very powerful instruments in the last 35 years—not as powerful as we would like, but very much more powerful than we had before.

COMPUTER SIMULATION OF THINKING

The first of these powerful instruments is computer simulation of human thinking. If you think of a computer as some kind of device that crunches numbers, then you need to think again. Computers, of course, can be made to crunch numbers if you give them the appropriate programs. But they can be made to do lots of other things. I suspect that among those of you who use computers many more use them as word processors than use them to do simultaneous equations (unless you've got your family accounts on a spreadsheet). So you already know that computers do things besides crunching numbers. You know that they deal with patterns or symbols. You don't care how those symbols are represented inside the machine, whether they are electrical or magnetic or made of water or whatever it is (maybe, hidden inside the machine, there are those old IBM punch cards). What you care about is that computers are general-purpose, symbol-processing, or pattern-processing devices. For that reason, and because when human beings think they are also manipulating patterns inside their heads, computers can be used to simulate human thinking.

Let me be clear what I *don't* mean by computers "simulating" human thinking. First, I don't mean that every computer program that is inverting a matrix is following processes like those that are used by human beings. Most of us don't even know how to invert a matrix, and those of us who do know would take far too long to do it to make it worth our while. So when computers are inverting a matrix (or when computer programs play chess), they are usually not doing anything that resembles very closely what a human being does when engaged in the same activity. On the other hand, there are some computer programs that can be demonstrated to be playing chess in a very humanoid fashion indeed.

What is the difference? The difference is that computers, as we know, are very fast. We talk about strange numbers like microseconds, gigaseconds, and megabits. We talk of millionths of this and billionths of that; computers are very fast. The human mind, alas, is a very slow device. It takes approximately 1/1000th of a second to get a message across a single synapse from one neuron to the next. And, to do anything interesting, you must get quite a few messages across quite a few synapses. No human being can recognize his or her mother coming down the street (or father for that matter, or brother) in less than approximately one second. Who would walk half a block to buy a computer that could not do the simplest act of recognition in less than a fraction of a second?

Clearly we humans cannot do the gear spinning that computers do when they are crunching numbers or doing some of the other things they can do. We must get by in some other way. Well we do, and we know today pretty well how we get by: what the other way is. Human beings are sometimes confronted with a situation in which there are zillions of possibilities. It is said that there are least 10^{40} possible games of chess— 10^{40} branches in the tree of "If I do this, and he does that, etc." The number 10^{40} is 10 with 40 zeros after it, and I have been told (I don't know who told me this) that that is more than the number of molecules in the world. Some computers have been programmed to recognize 10^8 possibilities. Even with 10^8 you are hardly starting to crawl up the hill. When such a computer reviews 10^8 possibilities before it makes its move, it does not play chess, it plays an approximate game of chess. It simplifies the situation until it can make the calculations. That means it pretends that there are only 10^8 positions instead of 10^{40} , and it almost gets away with it, probably not with a Fisher or a Kasparov, but with most of the rest of us.

What does a human chess player do? We have very good evidence about this. Chess has been the *Drosophila* of artificial intelligence research. It is a standard organism for a great deal of research, starting back a long way, but particularly with the Dutch psychologist, Adriaan DeGroot. There is strong evidence that a chess grand master in a difficult position probably does not look more than 100 positions ahead, branching through 100 possibilities before selecting the correct move. There is also good evidence that that grand master does not look at more positions than an ordinary amateur player does before making a move. Well, what is the difference, then, between the grand master and the ordinary player? The difference is that the grand master looks at the *right* moves. The way we humans get by with a very limited computing capacity is through using a tremendous amount of selectivity in what we look at. We solve problems by searching, just as computers do, and just as cats do when you put them in a Thorndike puzzle box. But we search extremely selectively. We depend on a vast body of knowledge about the structure of the problem situation to search efficiently. And when we don't have knowledge of that structure, we don't do very well. There are lots of problems that we don't solve; lots we don't know the answer to. People have been trying to prove or disprove Fermat's last theorem for over 300 years, and they still don't have it.

We have used computers to study human thinking by simulating it. But we have done this by writing programs for the computer in such a way that, instead of spinning its wheels very fast and trying everything, it looks at things very selectively and tries to make use of the same kind of sly cunning (we call it heuristic rules of thumb) that every human being uses in solving

problems. To give you an example, one area we have been studying in recent years (besides chess) is the process of scientific discovery. That area was picked because nobody can accuse a human being or a computer that makes a serious scientific discovery of doing something trivial. Because this is something we admire in human beings, if we can simulate it we have learned something about an important human activity.

Kepler found some important laws of astronomy, one of which is his third law. It is a very simple law to state. The planets are going around the sun with a certain period of revolution, and they are at various distances from the sun. Copernicus was probably the first who arrived at reasonable estimates of the distances; the periods have been known since ancient times, and the problem was first addressed by Aristotle. Kepler took as his problem the explanation of the relation between these periods and the distances. It is well known that the outer planets, those farther from the sun, take longer to get around the sun than the nearer ones. After a lot of work, Kepler arrived at what we call Kepler's third law: The period varies with the $3/2$ power of the distance. In our laboratory, we have a little computer program; we call it Bacon, after Sir Francis, because like Sir Francis it is very fond of induction. You can give Bacon data; you can, in fact, give it exactly the same data that Kepler had about the periods of the planets and their distances from the sun. If you don't have his data available, you can get them from the *World Almanac*. If you give Bacon that data, after a little while Bacon will find (Bacon's job in life is to find regularity in any data you give it) a law connecting the periods and the distances, and the law will be Kepler's law.

In the case of the computer, we can find out what happens in the "discovery" of Kepler's law step by step: It's a little too late to ask Kepler. We ask the computer to pick out a pattern. It gets the answer on the third try. Now there are an infinite number of mathematical functions to test with those data. Why is Bacon successful on the third try? It is successful because it doesn't search randomly or by brute force; it searches very selectively, trying to make use of the result of each step it takes to extract information that will suggest what it should try next. It first tries a linear function, then it tries a quadratic, and from the misfit of these functions to the data it figures out that the proper function is the $3/2$ power. Interestingly enough, Kepler also made at least several tries, because he published an earlier and incorrect solution with the same quadratic law that Bacon gets on its second try. Bacon kept going, and so, ultimately, did Kepler, and they both got it on the third try.

I don't want to explain Bacon at length. I do want to point out simply that, if you are able to introduce into a computer the same kinds of knowledge and the selectivity that can be derived from that knowledge that human beings possess when we know something about an area, then you can frequently get

the computer not only to solve problems that are difficult for human beings to solve (we knew computers could do that all along), but you can also get it to solve the problem in a very humanoid way. You can show that it is following a path very similar to the human path, and in much more detail than in the Kepler example. That gets me to the other tool that has revolutionized psychological research: the so-called thinking-aloud protocol.

THE NATURE OF EXPERTISE

A thinking-aloud protocol is obtained by giving somebody a problem and asking him or her to talk aloud while solving the problem. With some types of problems, this causes difficulty, but, for most problems, it in fact affects neither the solution path nor the time required to solve the problem (there are lots of footnotes to put on that assertion, but I will not get into them now). For many problems, you can get people to think aloud. This shouldn't surprise us because many of us sometimes talk to ourselves while we are solving problems.

Through this technique, we have learned an enormous amount about the steps in problem solving. Notice we are not asking people to give us a theory about problem-solving processes. That would be like going to a meeting of the physics society and having a Geiger counter on the podium giving a lecture on theoretical physics. Our subjects are not giving us a theory. What they are doing is producing data. They are telling us what things enter into their awareness, step by step, as they are solving the problem. Their minds are solving the problem, not giving us the theory of what they are doing.

The thinking-aloud technique has given us a wealth of information about the human problem-solving process. From this has emerged a picture of what an expert is, an expert in any particular domain. A person will be an expert in a particular domain, but the characteristics of expertise seem to hold up as we move from one domain to another. By now, there have been extensive studies made on several dozen domains. The key to the expert's skill (this should not surprise educators) is possession of a large body of knowledge—but not just any kind of knowledge, and not knowledge just organized in any way.

We even have a measure of the size of the body of the expert's knowledge or at least a measure of minimum size. I am not going to insist upon how accurate this measure is. In psychology today, we have a unit that we call a *chunk*. That is not just a slang word; it is a technical unit of measurement of mental storage. By chunk, we mean any unit that has become familiarized to the person who has it. For example, most educated people will have a

vocabulary of about 100,000 words. That means, if you see those words on a page or hear them from someone, you will have at least an approximate notion, or maybe a very exact one, of what those words mean. They will be familiar terms on the page. We spend an awful lot of time with children teaching them to acquire the beginnings of that skill. By the time they get to college, or certainly by the time they get their Ph.D., they have around 100,000 chunks, or familiar words, give or take factor of two. So first of all, we know that, in any area of expertise, the expert in that area has 50,000 familiar chunks or more—something of that magnitude.

What do these chunks do for you? They act like an index to your knowledge. You have all sorts of things stashed away in memory, but they do no good unless you can get at them. In an encyclopedia, the last volume contains an index. If you lose the index volume, you might as well chuck the rest of it out, because you could never find anything. So we have to have an index for what is stored in our memories. Now how is that index used? To recognize things. Any one of us in our area of expertise can recognize cues in the situations that we encounter in our daily work, and those cues immediately access, in our memory, all the things we knew about them. Medical diagnosis is an example. You walk into the doctor's office with some peculiar spots on you, and it does not take long for the doctor to say, "chicken pox," if that's what it happens to be. (There may be trivial reasons why he says that—he knows that there is an epidemic of chicken pox running around, and sometimes the best way to predict the doctor's diagnosis is to find out what has just been published in the *Journal of the American Medical Association*.)

The fact of the matter is that any expert can recognize the symptoms, the clues, to the bulk of the situations that are encountered in his or her everyday experience. The day would simply not be long enough to accomplish anything if cues didn't do a large part of the work for the expert. Then again, back to our *Drosophila*, the chess players; in the case of chess, the grand master in a tournament game will take, on the average, 3 to 5 minutes to make a move because that is all the time allowed. You might think that a chess game is a very slow business, but time is of the essence for someone who is playing. Many games are lost because of the official time limits on the game; given a short time, you do stupid things. So time is of the essence. Now that same grand master, having collected 50 people who will pay \$5.00 each to play with him or her, can go around the table with those 50 players, spending just a couple of seconds at each place making a move, going to the next one, another few seconds. At the end of the evening, the grand master will have won 47 of the 50 games, drawn 2, and maybe one little kid beat him or her. Now how can that be? How can such instantaneous play take place? It is very simple, if you are a grand master. When you go to each board, either you see

a cue or you don't. If you don't see a cue, you just make a developing move; that is, you make a move that gets your pieces out into play where they are mobile. But sooner or later, because the opponent isn't as good a player as the master, he or she will make a move that creates a situation on the board that the grand master can recognize as a weakness in the opponent's position. And then the master does what memory tells him or her to do in those positions. There is no problem to solve. The experts see the cues staring in their faces and respond accordingly, and at the end of the evening they have won 47 of the 50 games. That is the way experts get through the day.

That raises a question that has been debated a great deal. Let me put it in a paradoxical form. The question is whether human reasoning or thinking is "logical" and whether it would be a good thing if it were. We hear conflicting opinions about this. There is, in fact, a good deal of sentiment around that says, "All thinking is logical." When this is said, others object that something besides logic has to enter in. Well, partly we are quarreling about definitions of words, and partly we are quarreling about a phenomenon. The people who question whether thinking is logical will claim that many of the best decisions that people make are intuitive or based on insight or on judgment. What is intuition and what is insight and what is judgment? How do we decide that somebody is exhibiting intuition? Well, we present that person a problem—or the world does; the person comes up with a solution very rapidly (maybe the solution is even right); and when you ask that person how he or she arrived at the solution, he or she says, "Well, I don't know, I just used my judgment" or "I just used my intuition." That is, that person honestly tells you that he or she doesn't know how that decision was formed.

This phenomenon is also familiar under another name. It is exactly the phenomenon of *recognition*. That's back to recognizing your mother again; it takes you one second. You do it with high reliability, but if I ask how you do it and you try to describe the cues, I don't think very many of you could give me a description of your mother that would allow me to pick her out from the crowd coming through the airport gate. We are able to recognize someone or something with high reliability, without being able to report the processes we use to do it. There is a big indexing device in our heads (I won't describe how it gets that way here, although we know a lot about that). It is a discrimination net that sorts out the stimuli we see in front of us, picks out the familiar ones, picks out the chunks, and thereby gives us access to the information we have about it. And that is recognition. I do not see any phenomena like what is usually called intuition or insight that cannot be explained by the process I call recognition.

Another idea concerning how the mind functions that is widely discussed is that there are two different kinds of mental processes, one set of them is

controlled by the left half of our brains (in right-handed people) and the other set by the right half of our brains. It is well known that there are differences between the two halves of the brain. All you have to do to prove that is to get hit on one side or the other of your head. Things are very different if you get hit on the left side of your head than if you get hit on the right side of your head. Very often, you can get hit badly on the right side of the head and really never know the difference. Seldom is that the case for the left. So there are differences. But we shouldn't stop there. We know, for example, that recognition of visual patterns has something special to do with the right side, that almost everything verbal has to do with the left side, and almost everything analytic has to do with the left side.

Some people have concluded from this that there is analytic thinking and intuitive thinking. I know absolutely no evidence to support that division. There is nothing really serious you can do without using both sides of your head, and particularly the left side. Some parts of the right side don't seem to be very heavily used, most of the time, but I don't know of any shred of evidence that there are these two species of thinking.

Of course, thinking is more analytic sometimes and more intuitive at other times. In particular, the thinking of experts dealing with ordinary situations is highly intuitive. It becomes analytic only when the going gets tough, when novelty enters into it, when new problems have to be solved. So all of us, if we are to lead productive lives as experts, have to have good equipment both with respect to our recognition capabilities and with respect to our analytic capabilities. We shouldn't suppose that we can go through life being either an intuitive thinker or an analytic thinker. The analytic thinker never would have time to get 8 hours' work done in 8 hours without a large component of intuitive reactions, and the intuitive thinker would never be able to deal with any kind of novelty or any level of difficulty above the routine, without a good deal of analytic capability.

I will not say anything about the computerization of expertise but, obviously, as we begin to understand what an expert is, we can attempt to get a computer to share some of its power, to become a partner in the expertise. Today there are computer programs of quite high quality for medical diagnosis. As far as I know, none is in general use. I think one of the reasons is the problem of who is going to be legally liable for the diagnoses; doctors are a little bit sensitive about issues of legal liability these days. But making use of the notion that a large part of diagnosis is recognition, there have been built a good many expert systems today. The realm of application of expert systems is growing rapidly and is based in considerable part on what we have been learning about the human decision-making process.

CONCLUSION: DECISION MAKING AND MANAGEMENT

I would like to draw a boundary now around decision making and make the obvious statement that of course decision making is not the whole story of management, because decisions do have to be implemented. Much of the effort of managers is not just devoted to making decisions, including all the aspects I have mentioned, but to implementing them. We could try to write down a theory of that, a theory of actually acting like a manager and lots of people have. You have all read the textbooks (maybe I even wrote one or two of them). There is something awfully disappointing about those books. They sound like common sense. We shouldn't be too disgusted about that; common sense is a good thing, but then why do you have to write a book about it? What would you want a manager to know and do? Make decisions when needed, that's the focus of attention in business. Make sure that the decisions are carried out; build organizational loyalty and understanding; delegate; enlist the whole organization's resources in reaching decisions. I haven't even used up one class hour, and the whole course is over.

Most of us know these things, and even most of our students know them or very quickly learn them after the first week of our teaching. The only thing that keeps us from being great managers, if we aren't great managers, is that we don't have the habit of doing these things. Some of them are unpleasant. Very few managers are sadistic enough to enjoy disciplining their employees. Very few of us get pleasure from dealing with mistakes, particularly our own. "Covering up" happens with great regularity in organizations.

Very few of us find it pleasant to make choices in dilemmas. By *dilemma*, I mean a situation in which all of the choices have really unsatisfactory consequences. We know that human beings behave qualitatively in a different way when faced with dilemmas than when faced with ordinary choices. A child may hesitate between chocolate and strawberry, but not for long, if he or she cannot get both. But between castor oil, and some other distasteful alternative, people hesitate a long time.

Training managers turns out to be much more like training athletes than like educating scientists. But that might be wrong because, come to think of it, proper habits are probably also the key to good science. We may be exaggerating the formal aspect of the training of scientists. Sometimes when I am dealing with my doctoral students, which is a kind of apprenticeship, I really think that I am practicing psychiatry without a license. The problem is not what they know or don't know, the problem is not how smart they are or aren't, the problem is whether they can establish in themselves habits that make good scientists.

We have advanced a very long way in our understanding of the decision-making process. We have something to say about decision making not merely in laboratory situations but also in real-life situations. But I don't want to pretend that we really know as much as we ought to know about the training task: not merely how you give people knowledge but how you develop and mold habits in people, which is a major task in management training. It doesn't usually happen in universities (I won't speak for schools of education, because I have never been in one, but in schools of business it certainly doesn't). If it doesn't happen in the university, then it has to happen after the manager-to-be has left the university.

In conclusion, I am saying, first of all, that we now have a considerable understanding of managerial decision making. The processes that are involved in it are identical with the processes we have been finding in every kind of psychological thinking and decision-making situation. There are no mysteries in it; there is no magic in it; there is no left brain/right brain romance in it. We can teach many aspects of it; we can certainly teach people to be alert to the processes, whether or not we are successful in creating new habits. We can say a lot about *attention directing* and the *intelligence operation*, using that term to mean gathering the information about the environment that is necessary for effective attention directing. We can say a great deal now about the generation of alternatives: how human beings through the chunks that they have in memory, the knowledge they have, are able to search their environments very selectively. We can say a lot about the role of knowledge in expertise.

Earlier I emphasized the 50,000 chunks of information possessed by experts. What I forgot to mention is that nobody, including child prodigies, reaches world-class expertise in any domain (and a lot of them have been studied by now) with less than 10 years of intense application. Mozart was composing music at age 4; Mozart composed no world-class music until at least age 17, as a minimum estimate. Four from 17 is 13. Nobody reaches world-class expertise in any domain with less than 10 years of intense application. We know a lot about the process. Benjamin Bloom, in the School of Education at Chicago, has done some fascinating studies with his students on exceptional talent, as has my colleague, John R. Hayes. So we know an awful lot about this process, and we should be able to put that knowledge to work in our schools in teaching people to acquire expertise. We can increasingly make aspects of expertise automatic, through the classical operations research (OR) tools and also the new artificial intelligence (AI) tools that are based directly on this psychological knowledge.

Finally, and equally important, effective management consists of combining this knowledge and these skills with the habit of applying them. I think

one of our great challenges is to learn how, within institutions of higher learning, we can provide that *skill* training (I will use that dirty word). How can we provide the skill training that is so necessary in conjunction with the knowledge we are also trying to impart to our students?

QUESTIONS FROM THE AUDIENCE

Question: What particular aspect of decision making or problem solving involves a shift of paradigm, a movement from one representation of the problem to another representation of the problem?

Answer: That is a topic now that's receiving a lot of research attention. It is clearly one of the more difficult aspects of problem solving, and so we have put it off for a few years, but now it is a very active area of research. Fortunately, of course, most of the time we don't have to change our paradigm. Most of the time, the expert encounters familiar problems, and one of the things that he or she is at once reminded of is an appropriate representation or paradigm for dealing with that problem. If you present a particular problem to an engineer, the response may be, "Oh, that is that calculus problem" or "I had better use partial differential equations on that" or whatever the case might be. So, fortunately for us, one of the things we learn as part of our education is a collection of paradigms that we bring to bear on problems. Some interesting work has been done on high school algebra, so-called algebra story problems, which shows that a major difference between good and bad problem solvers in algebra is that the good problem solvers are able to classify the problem accurately in such a way that the right paradigm is invoked. Poor problem solvers don't; they look at superficial aspects of the problem rather than the things that really determine the paradigm. They look at a problem and say, "This is about airplanes," rather than saying, "This problem must be the same as those river boat problems, because whether it is an airplane or a boat, it still involves relative motion." Micheline Chi, at the Learning Research and Development Center at the University of Pittsburgh, has done some very interesting work on that particular task.

There are some problems that are almost impossible until you get the right paradigm, and the question is how that comes about. Some of you may know the problem of the mutilated checkerboard. You have a checkerboard and some dominoes; each domino covers exactly two squares of the board, so you can cover the whole checkerboard with 32 dominoes. No problem, a child could do that. Now some malicious person comes along and snips out the northwest square and the southeast square of your checkerboard, and you only have 62 squares left. Can you cover what's left with 31 dominoes?

When you give that problem to people, they have a representation of it, they have a paradigm. They try covering the board with dominoes, but alas, it never quite works, there are always a couple of squares left over. After about an hour or 2, even if you pay your subjects, they begin to get pretty frustrated. Either because they are exceptionally good at this game, or exceptionally persistent, or more often because you give them some hints, suddenly it occurs to them that the uncovered squares, the ones that they can't cover are always the same color. And shortly after that occurs to them, they notice that each domino covers one red and one black square. They say, "Oh, after the mutilation, I now have two less squares of one color than the other. So, of course, because every domino covers one square of each color, no matter how many dominoes I have I can't cover unequal numbers of the two colors." End of problem.

That is a paradigm shift. Now we know a good deal about it. We know that people don't come equipped with a list of paradigms that they tick off and say, "If this one doesn't work I will try that one." There are a zillion potential paradigms. What subjects do, occasionally, is to focus on things that, under all their actions, remain invariant in the situation. In the midst of the whole confusion, we are always surrounded by, a few things stay constant, and human beings are well tuned to look for constants, to look for the invariant. That is what science is all about, finding invariance in confusing phenomena. Those subjects who are particularly good at focusing attention on invariance sometimes themselves discover the properties that lead to the change of paradigm and the solution. There is now more research going on of this kind, but I wouldn't pretend that we are anywhere near the whole story of where new paradigms come from. We do know they don't come very often. Many of us learned calculus, but only Newton invented it.

Question: I think we need to ask seriously whether we are devoting enough of our attention in the curriculum and in what we do in the classroom to, first, developing these perception skills of acquiring the chunks that enable people to recognize the cues in situations and, second, whether we are doing enough on the habits.

Answer: Let me take up the first topic, of perceptual skills. Consider the textbooks. Every algebra textbook will tell you what rules you can apply to an equation without changing the value of the unknown. You know you can add and subtract, multiply and divide, as long as you do the same thing to both sides. Fine, we teach that in Week 3; in Week 4, we teach solving equations. But, if you look at Week 4 in the textbooks (at least in any textbooks I have examined), you will find precious little information about *when* you should add or subtract or multiply or divide, how to figure that out by looking at the equation, or *how much* you should add or subtract or

multiply or divide. Now there are some simple rules for that. As a matter of fact, they are so simple that we have written a program for a computer, a so-called adaptive production system, allowing that computer to examine a few worked-out examples of solutions of algebraic equations and then to reprogram itself to solve algebra equations. On the basis of what we learned by writing that program, we decided that students could also learn effectively by studying appropriate worked-out examples. And we now have a large-scale experiment going in China. At any rate, the experiment shows that we have very effective algebra instruction, with almost no lectures, with the primary activity of students being the study of worked-out examples plus problem solving. This is not just discovery learning (I can give you reasons why just discovery learning is not the answer) but nearly total reliance on worked-out examples. That experiment came straight out of the theory, and there are a lot of applications for it in algebra. It requires that we teach the perceptual patterns that students need to learn to recognize cues and know what to do and when to do it. By teaching cue recognition, I don't just mean lecturing about it; we have to provide students with experiences with which they train themselves to recognize these chunks.

Your second question is a tough one, and part of the reason it is tough is what I suggested a little earlier—skill is a dirty word around universities. If you put the word *vocational* in front of it, you are really damned. If something is "useful," oh goodness! But we really have to face up to the question of whether a university is capable of training habits as well as imparting knowledge and if that is a reasonable thing to do in a university. If the university is an appropriate setting in which to do it, we'd better find some way of changing our attitudes toward that part of the task. And that problem I will turn over to the political scientists or the sociologists.

Question: How can we use computers to improve teaching effectiveness? What about computer-aided instruction?

Answer: You may have noticed that, although I have spent the last 35 years of my life talking more to computers than to people, I didn't really say very much about computers here. Of course, I think that technology has an enormous role to play. It is already playing an enormous role in keeping our books. I don't just mean the family books, I mean our company books. It has changed our organizational lives in all sorts of nonfundamental respects already. But I don't think it has changed them in very many fundamental respects, and likewise our educational practices. Computer-aided instruction just is not here; it is just not important. Does anybody think it is an important part of the scheme? We haven't even incorporated television into instruction, except in a very limited and superficial way. It took only 10 years for the movies to drive vaudeville out of this country (professional performers to

drive out the amateurs). I am at a loss to explain why that has not happened in the universities. That is a whole discussion by itself.

But let me get back to the question about technology. I think technology has an enormous role to play, and it is only going to play it to the extent that the technology is intelligent. We talk about living in an information age, as indeed we do, but the main characteristic of that information age is that information has ceased to be a scarce resource; we are drowning in it. Human attention has become the scarce resource. So any time you bring a machine in, you better ask how much information it is going to take in and how much it is going to give out. And unless it takes in three orders of magnitude more than it gives out, don't buy it, don't let it in the door. To achieve that ratio (that it can take in three times as many orders of magnitude more than it gives out), it has to be intelligent. Because it can't spew anything out at random, it has to have the kind of selectivity that chess players have so they don't look at the 10^{40} possibilities or even 10^8 .

The future of high technology in our educational society—and our productive society generally—is very much tied to the degree to which we understand human intelligence and, hence, are able to build computers that really can collaborate with us and can contribute their intelligence to the total intelligence of the task.

That is happening, and it is not necessarily happening in computer science departments, or even mainly in computer science departments. Most of it is going to be produced by users who don't even think they are doing artificial intelligence, who just have an idea about how to get a computer to do something interesting. You see that happening now in research; you see it happening in chemistry; you see it happening in physics. Increasingly, the computer is becoming the laboratory chemist's or laboratory physicist's right-hand person, and that is going to spread to other areas as well.