A PRIMER FOR POLICY ANALYSIS

Edith Stokey and **Richard Zeckhauser**
Kennedy School of Government, Harvard University

W · W · NORTON & COMPANY
New York · London

1 Thinking About Policy Choices

Should the federal government require oil-burning power plants to switch to coal? Given the declining birthrate, should Memorial Hospital convert some of its maternity rooms to a cardiac unit? Should the state university develop a master's program in public policy? Should the Transport Authority extend the subway system to the outer suburbs? Should the United States stockpile grains to diminish extreme price fluctuations? To address these questions one must understand the principles of policy analysis; introducing you to those principles is the goal of this book.

Our Approach to Policy Analysis

The approach to policy analysis throughout this *Primer* is that of the rational decision maker who lays out goals and uses logical processes to explore the best way to reach those goals. He may perform the analysis himself or he may commission others to do parts or all of it for him. The decision maker may be an individual or a group that acts essentially as a unit. We will not consider explicitly the situations in which several decision makers with conflicting objectives participate in a decision. Nonetheless, our approach should prove helpful to an individual who takes part in such a process of shared decision making, whether as a legislator deciding how to vote or as a bureaucrat trying to line up support for a proposal.

In any case, the emphasis in this book is on how decisions ought to be analyzed and made, rather than on the details of the information that should serve as inputs to the decisions. In establishing this framework we rely heavily on the analytic techniques developed in economics, mathematics, operations research, and systems analysis. In actual practice, to be sure, policy analysis is much more broadly eclectic, drawing on a great variety of disciplines, including law, sociology, and political and organizational analysis. We will have little to say about these important complementary disciplines, although you should recognize their relevance for the working analyst. If he is designing a program for welfare reform, he must take account of the capabilities of the state bureaucracies that will implement the program. If he is drawing up safety regulations, he must understand the administrative and judicial processes through which the regulations will ultimately be enforced. Nor will we discuss the natural sciences here, even though understanding how pollutants spewing from tailpipes mix with pollutants escaping up chimneys may be critical for drawing up a set of environmental regulations. In short, understanding and predicting how the world will actually behave is essential for any process of policy formulation. Our concern here, however, is with how the decision maker should structure his thinking about a policy choice and with the analytic models that will aid understanding and prediction, not with all the disciplines that could conceivably provide helpful information.

Most of the materials in this book are equally applicable to a socialist, capitalist, or mixed-enterprise society, to a democracy or a dictatorship, indeed wherever hard policy choices must be made. In deciding whether a vaccine should be used to halt the spread of a threatened epidemic we need not worry about the political or economic ideology of those innoculated. Nor will the optimal scheduling for refuse trucks depend on whether it is capitalist or socialist trash that is being collected.

Questions of values are, nevertheless, a critical and inevitable part of policy analysis. Nothing can be written on the subject without making value judgments, at least implicitly. No specific policies are recommended here; policy issues are used merely for illustration. Still, the very nature of the tools and concepts we expound reflects a philosophical bias and a particular set of ethical concerns. For one thing, the subject itself, policy analysis, is a discipline for working within a political and economic system, not for changing it. For another, we follow in the predominant Western intellectual tradition of recent centuries, which regards the well-being of individuals as the ultimate objective of public policy. We turn in the last few chapters to a further exploration of these points. No doubt those who search for it will find a backward trickle from these later chapters to the more tool-oriented chapters that are the main body of the book.

The Plan of Attack

A Primer for Policy Analysis consists of three major sections. In the first, "Cornerstones," we establish a framework for thinking about policy problems and making choices. The second and much the longest section, "Nuts and Bolts," focuses on the use of models to represent real-world phenomena, and the more general use of analytic methods to assist in the entire process of making decisions. We will work through a toolbox of techniques, starting with fairly simple situations and gradually adding such important complexities as outcomes that are uncertain or that have consequences over future time periods. We deal initially with techniques that help us see clearly what the decision maker's choices are, and then with techniques that assist in identifying and formulating his preferences.

Each technique is to be understood as part of a total structure for thinking about policy choice, as a means of determining some of the pieces of that total structure, and not as an end in itself.

The third section, "Ends and Means," is broader in scope and less technical. To provide a background against which policy analysis can be viewed, it considers critical ethical questions: who should make what policy choices, and on what basis? It lays out the basic criteria for policy choice, identifies the circumstances in which the government should play a role in allocating the resources of society, and reviews briefly the alternative forms that government intervention might take. We might have begun the book with this more philosophical discussion. We did not because this *Primer* is meant to be an essentially practical work emphasizing the structural aspects of policy analysis. Moreover, we want to get you thinking right away in terms of analytic methods, especially if this is a mode of thought you find a bit unfamiliar.

A Framework for Analysis

What do you do when a complicated policy issue lands on your desk? Suppose it's your first day on the job as a policy analyst in a New York State agency; you are directed to investigate and evaluate alternative pollution control measures for the Hudson River. The problem has so many ramifications you wonder how you will ever sort them out—and even where to begin. You can always muddle along, hoping eventually to develop a feel for the situation, but such a hit-or-miss approach rather goes against the grain. You would prefer to have a standard procedure that will at least help you make a start on digging into a complex policy issue.

Many policy analysts have experimented with a variety of ways to structure complex problems like this one. We suggest the following five-part framework as a starting point. As you gain experience in thinking analytically about policy choices, you will perhaps wish to revise it to suit your own operational style; so much the better.

1. *Establishing the Context.* What is the underlying problem that must be dealt with? What specific objectives are to be pursued in confronting this problem?
2. *Laying Out the Alternatives.* What are the alternative courses of action? What are the possibilities for gathering further information?
3. *Predicting the Consequences.* What are the consequences of each of the alternative actions? What techniques are relevant for predicting these consequences? If outcomes are uncertain, what is the estimated likelihood of each?
4. *Valuing the Outcomes.* By what criteria should we measure success in pursuing each objective? Recognizing that inevitably some alternatives will be superior with respect to certain objectives and inferior with respect to others, how should different combinations of valued objectives be compared with one another?

5. *Making a Choice.* Drawing all aspects of the analysis together, what is the preferred course of action?

We do not mean to imply that an analyst will always proceed in an orderly fashion from one stage of the analysis to the next. Real people—even those who are models of administrative efficiency—can rarely operate so neatly, nor should they try to. But we do insist that each of these five critical areas must be dealt with. The conduct of an analysis will usually turn out in practice to be an iterative process, with the analyst working back and forth among the tasks of identifying problems, defining objectives, enumerating possible alternatives, predicting outcomes, establishing criteria, and valuing tradeoffs, to refine the analysis. This is an entirely sensible approach. We claim only that it is easier to keep track of where you are in this iterative process, and to avoid going around in circles (a disease with which even the best analysts are occasionally afflicted), if you keep in mind a basic framework to which every aspect of the analysis must be related. Furthermore, the consumers of your analysis will thank you, for strict adherence to a clearly visible structure makes for far easier reading and comprehension, and opens up the analysis for evaluation and debate.

We believe you will also find the outline useful as a background for the rest of this book, to help tie together the wide array of methods and concepts that are considered. The techniques described in the following chapters are all aimed at enabling us to provide better answers to one or another of the questions in the outline. At every point as you work your way through the following chapters, ask yourself, "How does this method fit into the overall picture?"

To be sure, not all the questions we bring together here will be addressed in every piece of policy analysis. The analyst will frequently be asked merely to predict outcomes, or will enter the decision process at an intermediate stage, after the range of possible actions has already been delineated. He may be asked to set forth the nature of the tradeoffs that must be made among objectives without making a final choice. This is particularly likely to be the case when a decision revolves around what are sometimes labeled "fragile values," such as risks to health or to the ecosystem. Perhaps the decision maker will be pressed for time, so that waiting for further information (an option that is frequently understressed) is out of the question. And often an analyst will be asked to "suboptimize," to find a best choice for a lower level problem without worrying about the overall problem. Almost all budget decisions are made in this way; the local library trustees are expected to make their expenditure decisions within a given total sum without reference to how the highway department will be spending its funds.

Some Practical Advice

Many of the policy decisions you will encounter will not fit neatly and automatically into the models presented here, for the real world is rich and

complex. Policy analysis is not an assembly line process, where a single-purpose tool can be applied repeatedly to whatever problem comes along. These are a craftsman's tools; you must learn to wield them with skill. Reading about policy analysis is only a beginning. An academic mastery of the tools will hardly prove sufficient; judgment and sophistication in applying them should be your ultimate goal. Therefore our perennial advice to students is "*Practice*!" Practice on all kinds of situations, large and small, public and private. Look regularly at the front page of the newspaper and think hard about one of the policy problems featured. Perhaps a proposed plan for energy conservation is under discussion; see if you can define the immediate objectives of the plan and their relationship to the underlying problem. What procedures would you use to predict the practical outcome of the plan? How would you treat uncertainty? What further information would you want? Should the plan be implemented sequentially? By what criteria would you evaluate the success of a proposed policy? On what basis should the decision be made?

Practice thinking informally in terms of objectives in your day-to-day work. When you are taking part in a budgeting process, say for a committee or a voluntary organization, consider what the organization's objectives might be and what various expenditures would accomplish. For example, suppose you are serving on a committee to allocate limited student aid funds. What are the committee's objectives? How should they be traded off against one another? How do various types and amounts of aid satisfy these objectives?

Practice on your own problems and decisions, using models to get your thinking straight or to illuminate commonplace events. For example, when you find yourself waiting in line, ask yourself what could be accomplished with additional service capacity, and what the benefits of such a move would be. When the local school committee advocates an inexpensive building with high maintenance costs, think about the tradeoff between present and future spending that is implied.

Above all, practice presenting your conclusions systematically; you don't need to become a gifted and sophisticated analyst before you can upgrade your output. Make up your mind that at least once every day you will deliberately apply the outline set forth above to a problem you face. You'll be amazed at what it will do for your reputation for perceptiveness and good judgment.

2 Models: A General Discussion

Of all the terms that policy analysts like to toss about, *model* is perhaps the most confusing to the layman, for models are many things to many people. A model is a simplified representation of some aspect of the real world, sometimes of an object, sometimes of a situation or a process. It may be an actual physical representation—a globe, for instance—or a diagram, a concept, or even a set of equations. It is a purposeful reduction of a mass of information to a manageable size and shape, and hence is a principal tool in the analyst's workbox. Indeed, we will be employing models throughout this book.

Models are of particular importance for public policy analysts, who are frequently forced to make policy recommendations in the face of a bewildering conglomeration of facts and estimates. The analyst must strip away the nonessentials that cloud a problem to expose the structural relationships among the important variables, so that the consequences of a particular policy choice may be predicted. The statement "No one can predict what will happen," while perhaps correct in a very narrow sense, is extraordinarily dangerous. It may allow the status quo to persist unquestioned, or it may lead more or less inevitably to the selection of the alternative that is most popular at the moment. We prefer the risks of decisions based on predictions, including predictions that admit to uncertainties, to decisions made by default.

We all rely on models in the conduct of our daily lives. A parent marks the heights of his growing children on the wall; you draw a map so that your friends may find their way to your house; we frequently use an average value for some sequence of events in the past (the weather, for example) in planning for the future. The models of policy analysts are merely extensions—albeit sometimes highly sophisticated—of such simple models as these. In justifying the use of any model, whether we use it to describe Mrs. Smith's behavior in the fruit market or to inform the federal government's deliberations about the appropriate policies for energy research and development, we must always come back to the same fundamental question: Do the gains in insight and manageability outweigh the sacrifice in realism that we incur by stripping away descriptive details?

Types of Models

It is customary to start off a discussion of models with a taxonomic listing. We will not attempt to draw up such a formal tabulation, for there is an infinite number of dimensions in which models may differ. Rather, we will list some of the kinds of models in common use, more or less in order of increasing difficulty or abstractness.

A model may be a model in the sense that we have used the term since childhood: a physical model that is a more or less accurate physical representation or image of some real-life phenomenon. Sometimes the scale is very much reduced, as with a model airplane. On the other hand, a model of the human eye might be many times enlarged. Urban planners frequently use models of proposed projects to show how things will look and facilitate public discussion. A still more abstract model of classroom space is used by the registrars of many schools in assigning rooms. Typically, a large wall chart will have slots for each classroom and each hour of the week; availability of space is seen by a glance at the open slots.

A step removed from these models are the diagrammatic models; a road map is a familiar example. The traditional map tries to capture the essential features of a road network to assist us in getting from one place to another. It regards as nonessential such details as which route has the best restaurants, although that might be given high priority in a guide for the traveling gourmet. It enables us only to guess at which route is the least congested, although it distinguishes superhighways from country roads. Yet it is usually adequate for the purpose at hand. Blueprints of a house, another example of this type of model, tell the builder what he needs to know. The painter would find them less helpful.

Flow charts are a particularly valuable sort of diagrammatic model, especially in situations where some commodity or some portion of the population passes at a regular rate of flow from one condition to another. In recent years we have often seen flow charts indicating the processes by which crude oil becomes gasoline or some other petroleum product. A flow diagram may be useful in describing how an individual apprehended for a crime passes through the stages of a criminal justice system. Figure 2–1 illustrates a flow chart designed to model the dynamics of heroin use. The population is divided into three major categories: nonusers of heroin, unsupervised users, and those in supervised programs. Each of these categories is in turn subdivided. In addition, the model includes the possibility of death. The arrows indicate the movement of people from one category to another. The government's objective may then be thought of as influencing favorably the rates at which these movements occur.

Drawing boxes and showing the connections between them is a good

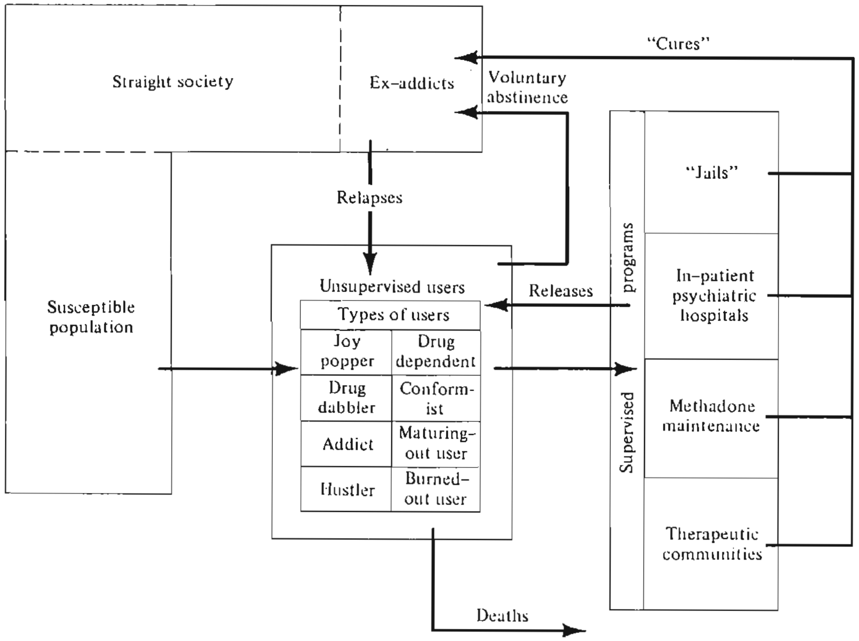

Fig. 2–1. *A Dynamic Model of the Heroin-Using Population*
Source: Mark H. Moore, "Anatomy of the Heroin Problem: An Exercise in Problem Definition," *Policy Analysis* 2, no. 4 (Fall 1976): 656.

way to begin to attack numerous problems. Many skeptics, once into the process, have discovered how tough and challenging it is and how much is learned in trying to do it. If you find it difficult to draw up such a model, you probably do not understand all aspects of the system you are trying to model.

Decision trees are first cousins to the flow chart; both identify distinct stages in a complex process. Figure 2–2 shows the simple decision tree (which is read from left to right) for a power authority that must decide whether to install conventional generators or a new, untested type in its projected power plant. If it takes action *B* and installs the new type, it must then await the outcome of a chance event (designated by the circle): either the new type of generator works or it doesn't. Decision trees move beyond flow charts, however, in pointing us toward the best choice. They have been developed primarily to aid decisions that must allow for chance

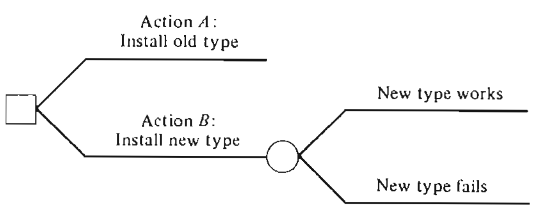

Fig. 2–2

events in the outside world. Chapter 12 is devoted to models of this type, which fall under the general heading of decision analysis.

Graphs and charts are another type of diagrammatic model. We will see in the next chapter that the economist uses an indifference map as a graphical model of preferences, and a possibility frontier to trace out the combinations of available outputs.

The term *model* is also used for what are essentially conceptual models. In a simple way, the person slicing a pie has the model of a circle in the back of his mind. He knows that the easiest way to divide a circle evenly is to cut it into wedge-shaped pieces. Rarely do we see anyone try to cut a round pie into concentric circles.

Long division sometimes serves as a conceptual model. It is easily applied and almost always helpful. Most of us instinctively employ it when we want to find out how much we are getting for a given level and type of expenditure. Large numbers are sometimes confusing; we can't think as intelligently as we would like when we are confronted by all those zeros in front of the decimal point. But we can try to convert them to terms we can come to grips with. The original estimate of the cost to the federal government of the swine flu immunization program was $130 million. Most of us have had little acquaintance with that kind of money or that type of program. But most of us also instinctively performed a little long division. If half the people in the country are immunized, $130 million comes to a bit more than $1 per person. A $2 million housing project that provides homes for 40 families comes out to $50,000 per family dwelling unit, a number we can understand and readily compare with alternatives.

On a more complex level, such terms as *feedback* or *contagion* are often used as conceptual models of processes that are widely understood. A thermostat is a familiar example of a feedback mechanism. Many suburban fire departments have an analogous mutual aid mechanism to reassign engines routinely from one town to another when one department is overburdened. The heroin study from which Figure 2–1 was reprinted includes a contagion model to show how heroin use spreads among a susceptible population. Such concepts can serve as appropriate metaphors in a variety of situations that differ drastically on the surface.

One of the most famous conceptual models of recent years is Garrett Hardin's model of the commons, the common grazing ground of the medieval English village.[1] Cattle owners holding the rights to pasture their animals on the commons, like car owners who have the right to use a common expressway, ignore the cost that their own use imposes on others. The inevitable result is overgrazing or overcongestion that is costly to all. And as Hardin observed, the problem of the commons arises in countless contexts that have nothing to do with cows or cars.

Conceptual models are often used to make nonquantified predictions about the behavior of individuals or institutions. Many of the theoretical

---

[1] "The Tragedy of the Commons," *Science* 162, no. 3059 (December 13, 1968): 1243–48.

constructs of sociology and political science are designed to help us understand such behavior.

A Simple Model: Compound Interest

Our definition of model seems to be getting broader and broader, so perhaps it is time to turn to a more restricted type that is in fact what many analysts mean when they speak of models. These are the formal mathematical models that describe explicitly the quantitative changes in a particular variable or system in response to various stimuli. For example, suppose we wish to investigate what happens to a sum of money that is left in a savings bank for a number of years. (We deliberately choose this simple situation because it is one with which all our readers will be thoroughly familiar.) What facts do we need? We need to know the initial sum of money (call it S_0), the rate of interest paid on the sum on deposit (call it r), and how long the money is to be left in the bank (let n be the number of years). These are the parameters of the problem. The term *parameter* is used to suggest that the particular variable in question is fixed, at least for the duration of the present exercise. The decision maker is to take it as given; he cannot change it. It implies, however, that the model is formulated in such a way that it can handle a parameter change should one occur.

Following this notation, S_n will be the sum at the end of n years. We can illustrate the progression from the original deposit to the sum after n years with the following diagram:

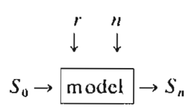

What does this model actually look like? Suppose $r = .05$. At the end of the first year we have a sum S_1, which is related to the original sum by the formula:

$$S_1 = (1.05)S_0$$

That is, $1 will yield just $1.05. (The rather backward bank of this example compounds interest annually, though its real-world competitors have been driven to daily compounding.) At the end of the second year we have

$$S_2 = (1.05)S_1 = (1.05)(1.05)S_0 = (1.05)^2S_0$$

And at the end of n years we have an amount

$$S_n = (1.05)^nS_0$$

This, then, is the analytical model for compound interest at 5 percent. The general model for any rate r is

$$S_n = (1 + r)^nS_0$$

The mathematical models in the chapters that follow are more complex, but we will continue to make use of the simple compound interest model for purposes of definition and exposition. One class of mathematical

models, those involving statistical inference, are not discussed in this book, since they are well covered in many elementary texts.[2]

Frequently the study of a complex problem will make use of several subsidiary mathematical models to predict the changes in important variables, or the consequences of different policy choices or of different assumptions about the true state of the world. Many people find that the use of physical or conceptual models comes naturally to them, although they may have to remind themselves, say, to try a flow chart in a particular situation. It is the mathematical models and their quantitative implications that seem more difficult, and for this reason disproportionate attention is devoted to them in the subsequent discussion.

Before moving on, we should mention one further use of the term *model* that has proven confusing to some students of policy analysis. Some people in the field have come to use the term broadly to refer to several separate models addressed to a common issue. As an example, suppose Alice Smith has analyzed the problem of public transportation in the city of Brownbury, using a number of subsidiary models to predict population growth, changing land-use patterns, or what have you. Even though she may not provide a comprehensive model of the complete transportation system, her entire product will often be referred to as "the Smith model." This is not a strictly accurate use of the term, but most of us have learned to live with it.

Formal models are used increasingly in the analysis of public systems. Hence, merely for purposes of speaking the language, a familiarity with them is of growing importance for people in associated fields. In acquiring this familiarity, the student must necessarily start with narrow and specific applications of various models. He thus runs the risk that he will overemphasize the solution aspect of modeling, which is really its least significant feature. Much more important is learning how to use models to formulate problems better, especially those broad, fuzzy problems where it is hard to know how to begin, and to get others to discuss what they feel to be their essential elements. As a further payoff, you will find that as you use models more and more, you begin to develop analytic insights in other situations where you do not consciously engage in a modeling process.

What we have presented here is by no means an exhaustive list of all the kinds of models you may encounter or use. To reiterate, however, all models have one basic feature in common—they aim at reducing the complexity of the problem at hand by eliminating nonessential features so that we may concentrate on the features that describe the primary behavior of the significant variables.

Descriptive vs. Prescriptive Models

Improved decision making is the goal of model building. The ultimate justification for models must thus rest on their usefulness in aiding decisions. Some models illuminate choices by showing us more clearly

[2] See, for example, Ralph E. Beals, *Statistics for Economists* (Chicago: Rand McNally, 1970).

what those choices entail, what outcomes will result from what actions. These models are called *descriptive*; they describe the way the world operates. Others go further and provide rules for making the optimal choice. They are categorized as *prescriptive*; they help prescribe courses of action. The distinction is worth pinning down, so let's try to be more precise.

A descriptive model is just what you would expect: it attempts to describe or explain something, or to predict how some variables will respond to changes in other parts of a system. The "something" may be an entire system or only a piece of it. We mentioned earlier that the relation between action and consequence may be complex indeed. It is this relationship that the descriptive model focuses on, perhaps on its entirety, perhaps on only a part of it. For example, an environmentalist may model the flow of pollutants through a particular body of water; an econometrician may construct a model of the entire economy in an attempt to predict the effects of alternative tax proposals.

A prescriptive or normative or optimizing model—take your pick—consists of two parts. The first part is a descriptive model that encompasses all the choices open to the decision maker and predicts the outcome of each. (If he has no choice as to the actions he may take, or if all actions lead to the same outcome, then he has no problem and there is no need for a prescriptive model.) The second part of a prescriptive model is a set of procedures for choosing among alternative actions, given the decision maker's preferences among the outcomes. Prescriptive modeling includes procedures to help the decision maker sort out those preferences.

We may illustrate these definitions with the descriptive model of accumulating compound interest that we discussed earlier. Generalizing the model so that it holds for any interest rate r, we find that

$$S_n = (1 + r)^n S_0$$

where S_0 is the initial sum of money and S_n the amount to which S_0 accumulates after n years. As it stands, this is a simple descriptive model; it predicts the amount to which a given sum of money will accumulate if it is left in the bank for a given number of years at a given annual interest rate. It is not a prescriptive model because (1) as it is constructed, no indication is given that the decision maker has any choice as to the values of r or n or S_0. In other words, the model is not formulated in terms of alternatives. The choices may be there, but if they are we have not been told what they are. And (2) even if we were to redraft the model so that it relates the outcome S_n to some sort of a choice of r or n or S_0, we have no criteria for evaluating different S_ns.

Let's construct an artificial situation in which we can convert this into a prescriptive model. Suppose, first, that our decision maker has a fixed sum, \$10,000, which he has already decided will be left in a savings bank for 2 years. He can deposit the money in any of five banks, which are essentially identical in every respect that is important to the decision maker except one—they do not all pay the same rate of interest. The first bank pays r_1 percent, the second r_2, and so on. Now our compound interest model becomes

$$S_{2i} = (1 + r_i)^2(10{,}000), \qquad i = 1, \ldots, 5$$

where S_{2i} is the accumulation at the end of two years if \$10,000 is deposited in the ith bank, and r_i is the rate paid by the ith bank. Now the decision maker has a choice among actions (i.e., among banks) and these actions lead to different outcomes. This takes care of the first part of the prescriptive model discussed above.

To satisfy the second part of a prescriptive model, we need a rule for choosing among actions that reflects the decision maker's preferences as to the outcomes or S_{2i}s. Let us assume that our decision maker is a straightforward person who prefers more to less. The decision rule in this case is obvious: to maximize S_{2i}, choose the bank with the largest r. The model is now prescriptive, albeit trivial.

How might this model become more complicated? Well, for one thing, different banks may have different kinds of service charges associated with their accounts, or may offer different borrowing privileges to depositors. Some may reward a new depositor with an electric blanket; others offer him tickets to the World Series or a collapsible dinghy. The interest rates may be guaranteed for different periods of time, or the notice of withdrawal requirements may differ. Most important, different types and degrees of risk may be involved, especially if one is choosing not just among savings banks but among different types of investment instruments. Such complications would require a more elaborate descriptive model, as well as a far more complex decision rule to deal with the decision maker's preferences for income, convenience, risk, and the like.

If all this sounds easy, don't be misled; it isn't. Constructing a descriptive model may be extremely difficult, especially when the relation between action and outcome is complex or remote. Determining objectives and finding ways to evaluate outcomes in terms of those objectives is likely to be even harder. If complications can conveniently be introduced one by one, however, we may learn to deal with increasingly complex situations.

Deterministic vs. Probabilistic Models

We have defined models as simplified representations of reality, and it is natural to think of reality as consisting of hard facts and sure things. Many models deal with situations where each action has a certain outcome. You flick the switch and the light goes on; you take an automobile plant and add prescribed quantities of materials and labor, and a given number of cars come off the assembly line. This does not mean that nothing ever goes wrong. Now and then something obviously does. What it does mean is that the outcome of an action is so close to being a sure thing that the model may take it as certain.

In other situations the true state of affairs is not at all certain; nevertheless, you are satisfied to treat it as certain. Perhaps some element in your model behaves randomly, but you know that using an average value will provide a sufficiently good approximation. Per pupil expenditures, for example, are frequently used in models of school systems. Or perhaps you would like to test the implications of various alternative assumptions. In population forecasts, for example, a model of how population changes over time may make use of several alternative sets of assumptions about birth and death rates, without any commitment as to which set is the most likely. If you assess the likelihood that each of these sets of assumptions is correct, then you may be able to predict the likelihood that each corresponding outcome will be observed. Note, however, that for each set of assumptions the outcome is taken as certain.

Models such as these, in which the outcome is assumed to be certain, are called *deterministic*. Given the relationships, the initial conditions, and the actions (which may be simply the passage of time, as in a population model), the outcome is uniquely determined.

You might think that this is as far as we can go with models, but it isn't. In some situations the outcome of a particular action is not unique. Instead there is a range or a number of possible outcomes, for which the probabilities may be estimated. Many models relating to health policies are of this sort. We might wish to model, for example, the progress of an epidemic. Suppose that, on average, 30 percent of the people coming in contact with the victim of a certain disease will themselves catch it. If 10 people are exposed, on average 3 will get sick. But sometimes 2 or 4 will catch the disease, and there is a small chance that all 10 will. A model that considers the various possible outcomes can help health planners estimate the probabilities of experiences of various sizes.

In such cases it is frequently possible to construct a model that will help us to trace the consequences of various actions that have probabilistic outcomes. Such a probabilistic model will illuminate the kinds of choices we are facing far better than a model that relies simply on an average value.[3] Since some of the inputs or some of the processes are probabilistic in nature, we will necessarily end up with a probability distribution of possible outcomes.

Let us milk one more illustration out of the compound interest model. Suppose you have decided to deposit $10,000 in the local bank and let it sit there for 2 years. At present the bank pays 5 percent per year, compounded annually, but realistically you know that next year's interest rate is subject to change. You estimate that the probability that it will remain at 5 percent is .6; that it will increase to 5¼ percent, .3; and that it will drop to 4¾ percent, .1.

[3] The distinction between a deterministic and a probabilistic model isn't very important in itself. But you should know that there are more ways to cope with variables that behave randomly than just taking an average and letting it go at that. Looked at from the proper angle, a deterministic model is simply a probabilistic model in which the probabilities of various outcomes happen to be either 0 or 100 percent. For example, a thrown ball has a 100 percent chance of coming down, a 0 percent chance of continuing upward indefinitely.

In tabular form,

$$p(.05) = .6$$
$$p(.0525) = .3$$
$$p(.0475) = .1$$

where $p(.05)$ means "the probability that the interest rate in the second year is .05."

The model that predicts how we'll make out with our money is

$$S_2 = (\$10{,}000)(1.05)(1 + r_2)$$

where S_2 is the sum on deposit at the end of 2 years and r_2 is the interest rate in the second year.

We might simply find the average of the possible values that S_2 could achieve. But in fact we can do better than that. Given our probability estimates for each interest rate, we can predict that

$$p(S_2 = \$11{,}025) = .6$$
$$p(S_2 = \$11{,}051.25) = .3$$
$$p(S_2 = \$10{,}998.75) = .1$$

Thus in this primitive probabilistic model, the probability distribution for S_2 coming out of the model happens to be the same as that for r_2 going in.

As an example of a more complex probabilistic model, consider the IRS office where people queue up to wait for assistance in preparing their income tax returns. They arrive irregularly, according to some probability distribution; the time it takes to serve each taxpayer is also irregular, and follows another probability distribution. If we model the queuing process, we will be able to predict the waiting times, or the length of the queues, in the form of still more probability distributions that will hardly be identical to or as simple as those with which we started. This distribution will be valuable in helping the director of the office to balance personnel costs against waiting times for the customers.

These probabilistic models are descriptive; without additional inputs they tell us nothing about how to make choices among chancy outcomes. Later in this book we will consider an important prescriptive model for decision making under uncertainty; it is generally known as decision analysis.

Choosing the "Right" Model

Our choice among models depends on the type of situation we confront, what we want to know, the level of detail we need, and the variables that we can control. The model will be judged by how well it works, or by how accurately it predicts, whether it is a complex set of equations fed into a computer or merely a few pencil scratches on the back of an envelope. This means that the assumptions that drive a good model must be accurate. We must take great care with assumptions that play an important role in value of different types of information. It may focus our attention on how little we know about a situation. If more information may be available in the near future, the desirability of postponing all or part of a decision in order to maintain flexibility becomes obvious. It may even be advisable to spend time and money to get more information before making a decision.

A further advantage is the possibility of experimenting with the model rather than with the system itself. This is likely to be particularly valuable in planning a new public facility, such as a hospital. Planners need to know the implications for design of assumptions about a host of variables such as intensive care load, emergency arrival rate, and the like. In such a case it may be impossible to experiment with the hospital itself, and it certainly would be prohibitively expensive in terms of time, money, morale, and health. The model highlights the critical tradeoffs, such as longer stays for chronic-care patients versus an extra day for maternity cases.

In other situations, experiments with a model may suggest potential beneficial changes in a system that warrant cautious moves in a particular direction. In 1965 a famous case of this sort resulted in a dramatic improvement in the service provided by the Northeast Frontier Railway in India. On this single-line railway, trains moving in opposite directions could cross only at crossing stations, which necessarily involved halting at least one of the trains. Delays were monumental. Analysts experimented with a simplified mathematical model and found that the number of such crossing delays rises sharply with the number of trains using the line. In fact, the number of delays increases more rapidly than the square of the number of trains. Thus an increase from 10 to 11 trains per day would bring about not just a 10 percent increase in the number of crossing delays, but rather an increase of $(11^2 - 10^2)/10^2$, or 21 percent. This finding suggested that the railroad should try running 2 of its 15 trains coupled. When the experiment resulted in a better than 20 percent improvement in the average number of kilometers each diesel engine covered in a day, the scheme was extended.[4]

Constructing a model frequently facilitates communication among those concerned with a policy issue. It makes specific the definition of the problem (or at least of a piece of it); it sets out on paper what might otherwise remain buried in the analyst's mind. (This seems to be the one point on which harmony prevails among admirers and critics of that famous volume modeling world economic behavior, *The Limits to Growth*.)[5] Moreover, when we refer to frequently employed classical models—linear programming, for instance—we employ a common language that permits shortcuts in communication. To be sure, this degree of precision makes it harder to smudge over some factors, and in some situations (especially those political in nature) this may be a disadvantage as well as an advantage.

Perhaps most important of all, experience with modeling helps us develop general insights that can be applied even to unfamiliar situations. For example, we soon learn that systems involving queues often behave in an almost counterintuitive fashion; even a very small increase in service capacity offered may radically reduce waiting time at a facility.

[4] For a description of the model, see Jagjit Singh, *Great Ideas of Operations Research* (New York: Dover Publications, Inc., 1968), p. 165.

[5] Donella H. Meadows, Dennis L. Meadows, Jørgen Randers, and William W. Behrens III (New York: Universe Books, 1972).

How much should the individual who desires a broad, rather than deep, comprehension of policy analysis know about actual models? We suggest the following five goals, in no particular order of priority:

1. You should become aware of a few conceptual models that are widely useful because they provide a shorthand description of situations or processes that pervade our society. Earlier we cited feedback and contagion as widely used conceptual models, and mentioned Hardin's famous model of the commons. You should become knowledgeable about the existence of the formal models most commonly used in statistics, in project evaluation, and in operations research. Some of these, such as benefit–cost analysis, linear programming, and decision analysis, are discussed in some detail later in this volume.

2. You should acquire the habit of asking yourself, when faced with a complex problem, if constructing a model would shed any light on the problem or subject you to the discipline of testing your own understanding. You should develop both an ability to construct simple models of your own and the habit of relying on an expert to refine the model and develop empirical data as a problem becomes more technical.

3. You should learn to use some of the more basic formal models in thinking about the structure of a problem even when you do not carry your thinking to the point of quantification. The insights gained through this nonrigorous use of modeling concepts may well be the most valuable outcome of a study of modeling.

4. To accompany your enhanced technical proficiency, you should develop a healthy and informed skepticism about models, and become aware of their limitations. A good way to start is by making sure that you understand all the assumptions about the relationships that the model implies and the data that it uses. Test the model's plausibility by seeing how it performs under simple conditions. Ask some probing questions: What are the critical features of the model that is being presented to you? Where may it diverge from reality in a manner that might lead our reasoning astray? What about causality—are you dealing with a situation where it is essential to understand why the variables in a system behave the way they do, or is simply modeling the observed relationship sufficient for your purposes? What parameter values are significant, in the sense that changing them would drastically change the predictions made or choices suggested by the model? Indeed, you should be careful to develop your own models in such a way that you can answer these questions.

5. You should keep the real world of policy making in mind as you build or use a model. A model so vast and so complex that no one can understand how it works will benefit neither the modeler nor the client of the analysis.

In the next chapter we will begin to discuss types of mathematical models in greater detail. We will then introduce a variety of widely used modeling techniques that have proven particularly useful in practice.