APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville
Massachusetts Institute of Technology

McGraw-Hill Publishing Company
New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

15.7. *SIDA Testing*
The incidence of SIDA, a deadly disease, among a certain population is 0.01%. Individuals, randomly selected from this population, are submitted to a SIDA test whose accuracy is 99% both ways. That is to say, the proportion of positive results among people known to be SIDA affected is 99%. Likewise, testing people that are not suffering from the disease yields 99% of negative results. The test gives independent results when repeated.

An individual tests positive.

(*a*) What is the probability that this person is actually affected? (Use both Bayes' Theorem and likelihood ratios.)
(*b*) Discuss the above result as regards the interpretation of the positive result.
(*c*) The test is then repeated twice. What is the probability that the person has SIDA if all three tests are positive? If the two subsequent tests are negative?

15.8. *Weather Expert*
The radio predicts a 60% chance of freezing weather. Your meteorological friend, May Vin, tells you she knows better: it is sure to freeze. From experience you know that she only gets it right 80% of the time.
(*a*) What should your estimate of freezing weather be?
(*b*) What would it be if May had predicted "no freezing weather"?

15.9. *Summer Goods*
Of the summer goods, some are bad. Two percent are defective. Visual inspection is cheap, but only correct half the time. A detailed examination, however, gives a correct diagnosis 90% of the time. Normal procedure is to look the goods over and then to examine in detail the ones that seem defective visually. Goods that fail both tests are rejected.
(*a*) What percent of the goods that pass visual inspection are in fact defective?
(*b*) What percent of defectives are not detected by the total examination process?
(*c*) If the detailed examination were applied to all goods, what percent of defectives would pass? Discuss whether you think this policy would make sense.

15.10. *Championship Playoff*
Before the infinite series, it looks as if either team A or B is equally likely to win the series. Past frequency indicates that "champions" win 70% of their games.
(*a*) If team A wins the first game, what is the probability that it is a "champion"?
(*b*) Use Bayes' Theorem to calculate how many times in a row team A should win so that the probability that it is a "champion" is greater than 90%. Then validate this by likelihood ratios.
(*c*) What is the probability that any team is a "champion" if it wins 2 out of 3 games? 3 out of 5? 4 out of 7?

CHAPTER 16

DECISION ANALYSIS

16.1 OBJECTIVE

This chapter presents the concept and methods of decision analysis, a fundamentally important method of evaluation. This is the approach that should be used whenever the outcomes of potential projects are highly uncertain. Since the planning and design of systems typically must deal with massive uncertainty about the future, as the previous chapter shows, decision analysis is a most valuable tool.

Formally, *decision analysis* is a method of evaluation that leads to three results:

1. It *structures* the problem, which otherwise appears very confusing to most people due to the complexities introduced by uncertainty (Sections 16.3 and 16.4).
2. It *defines optimal choices* for any period, based on a joint consideration of the probabilities and the nature of any outcome of a choice, specifically by calculation of an expected value (Sections 16.5 to 16.7).
3. It *identifies an optimal strategy* over many periods (Sections 16.8 and 16.9).

Decision analysis rests on the simple proposition that a planner or designer should use all the important information available about a problem, specifically the fact that the performance of any system is uncertain. This premise makes decision analysis very different from the traditional economic evaluations, which focus only on the typical or most likely outcome of a situation.

Recognizing the risk of a situation has two important consequences for an evaluation. Mechanically, it makes any evaluation far more complex and requires many more calculations. This fact has inhibited the development of decision analysis, which is only now—with the diffusion of computers—becoming an integral part of systems analysis.

Most importantly, the perception of risk fundamentally changes the nature of the optimal choice over time for any system. Once we acknowledge that we do not know how the future will evolve, we must also accept that it would be presumptuous to define the single best development over time. We must instead determine a *strategy* of choice, a range of possible choices whose final selection will depend on how events occur.

The concept of *insurance* is the equally important corollary to that of a strategy. Since uncertainty prevents us from knowing exactly which series of choices will be best over time, our selection at the first stage of any sequence must provide us with the flexibility to develop as needed later on. It must both position us to take advantage of good opportunities as they arrive, and give us the capability to avoid disaster. This flexibility must be paid for, however. Buying the ability to respond easily to future events is the insurance we must build into an optimal strategy.

16.2 PRIMITIVE MODELS

It is important to understand that decision analysis offers substantial advantages over the other methods for dealing with risk that are now used in practice. Decision analysis is not just a refinement, a more complex version of the primitive approaches that have been used, it represents a real change. The primitive models are not simplistic versions that could be used for first-order approximations; unfortunately, they are wrong.

Perversely, the primitive models of decision analysis are increasingly pervasive. The use of a spreadsheet program on a personal computer makes it easy to lay a problem out for this kind of analysis. As the primitive models seem to be reasonable, if one is not familiar with their fatal flaws, they are also seductive. These two characteristics of ease of use and superficial appeal make the primitive models popular. But they can be completely misleading. That is the reason this section discusses these approaches in detail: it is necessary to warn the reader against error and to motivate the use of the more appropriate decision analysis.

The general problem considered by both decision analysis and the primitive models has three features:

- A set of *possible decisions*, D_i, that the decisionmaker might take
- A set of *uncertain events*, E_j, that may occur
- The *outcomes*, O_{ij}, that result from having chosen D_i and being subjected to E_j

The risky events E_j are often referred to formally as *states of nature*. In practice

decision analysis typically makes an important assumption, which should be carefully noted. The premise is that the probability of any outcome of a choice depends only on passive forces of nature. This excludes the probability of competitors, friends, or enemies who would actively try to influence the outcome of any choices we might make. Game theory deals with this possibility, and is discussed in more advanced texts.

To illustrate the above definitions, consider a typical problem of deciding whether or not to wear a raincoat as you go out. The decisions may be either "wear a raincoat" or "do not wear a raincoat"; the possible events correspond to the weather (fair, misty, rain, downpour, and so on) and the outcomes could be how wet you might become.

The primitive models all differ from decision analysis in that they assume that the probabilities of possible outcomes are not known, indeed cannot be known. They rest on the idea that all that can be assumed is that different outcomes might happen, depending on the situation. Decision analysis assumes that at least some assessment of the probability of events is possible and should be used.

Semantic caution: The primitive models are sometimes said to represent procedures for "decisionmaking under uncertainty," where uncertainty is complete absence of knowledge about the probability outcomes. This is not the meaning used in this text. The view here is that there is never either complete ignorance or knowledge of probabilities. There are only varying degrees of impression about them, that is, uncertainty in the sense of ordinary language.

The primitive models generally describe the problem by a payoff matrix. A *payoff matrix* is a table representing possible outcomes of the available decisions. The decisions are typically arrayed along the vertical axis. The horizontal axis represents the several possible states of nature. The cells in the payoff matrix then each represent the outcome associated with the intersecting decision and state of nature. This is the kind of information an analyst can easily tabulate using a spreadsheet program on a personal computer.

An example situation is depicted by the payoff matrix shown in Table 16.1.

TABLE 16.1
Payoff matrix for example problem

Possible computer system	Annual Profits ($ $\times 10^3$) if jobs require		
	Linear programming	Dynamic programming	Engineering economy
Mainframe	90	40	35
Minis	60	80	10
PCs	0	0	120

This represents the hypothetical outcomes for a consulting firm faced with the choice between three different types of computer systems: a large mainframe, some minicomputers, and individual personal computers (PCs). The states of nature in this case represent the typical kind of analysis needed by their clients over the next several years, a set of requirements presumed to be determined by the vagaries of the general economy. The outcomes are then a measure of net benefits to the company.

This example problem will illustrate the operation and inherent weakness of each of the four traditional primitive models. These are those associated with the Laplace, the Maximin, the Maximax, and the Regret criteria. Recently another primitive approach has become popular, largely due to the use of computers and spreadsheet programs. This is the Weighted Index criterion. It will be illustrated by a different example. Each of the five methods is discussed in turn.

Laplace criterion. The Laplace model proposes that since the probabilities of the states of nature are not known, they should be assumed equal. This simple notion, possibly attributed unjustly to the great French mathematician, has a fair amount of intuitive appeal. Experience shows, for example, that there always are some members of a class of students who spontaneously suggest this model when confronted with an example problem such as that of Table 16.1.

The Laplace model defines a simple criterion for choice: select the system with the most desirable expected value. According to this procedure, the preferred choice in the example problem should be the Mainframe computer since the expected value of this option gives the greatest net benefits:

$$\text{E (Mainframe)} = \tfrac{1}{3}(165) = 55 \qquad \text{“Best”}$$

$$\text{E (Minis)} = \tfrac{1}{3}(150) = 50$$

$$\text{E (PCs)} = \tfrac{1}{3}(120) = 40$$

TABLE 16.2
Altered payoff matrix to illustrate the peculiar weakness of the Laplace criterion, its sensitivity to irrelevant states of nature

Possible computer system	Annual profits ($ $\times 10^3$) if jobs require			
	Linear programming	Dynamic programming	Benefit cost	Present value
Mainframe	90	40	35	35
Minis	60	80	10	10
PCs	0	0	120	120

The peculiar flaw of the Laplace model is that it is sensitive to the description of the states of nature, and that its selection can be altered by the introduction of trivial or irrelevant possibilities. To appreciate this phenomenon, consider the payoff matrix in Table 16.2: it is the same as that of Table 16.1 except that the engineering economy jobs have been divided into two components, benefit-cost and present value analysis. This modification to the payoff matrix evidently does not reflect any real change in the actual problem. Yet because of this trivial change the Laplace model identifies a completely different choice as optimal: it now recommends personal computers! Thus:

$$\text{E (Mainframe)} = \tfrac{1}{4}(200) = 50$$

$$\text{E (Minis)} = \tfrac{1}{4}(160) = 40$$

$$\text{E (PCs)} = \tfrac{1}{4}(240) = 60 \qquad \text{“Best”}$$

The Laplace criterion assigns probabilities without regard to the real facts. Thus in the example problem it ascribes a probability to engineering economy problems of $\frac{1}{3}$ in one case and $\frac{1}{2}$ in the other, even though the problem has not changed. This is evidently absurd, and it is surely better to assign probabilities with some reflection, even if these subjective estimates cannot be precise.

Maximin criterion. The maximin criterion indicates that the decisionmaker should choose the option that maximizes the minimum benefit obtainable. This notion also has some intuitive appeal; many people do like the idea of knowing that they have protected themselves against some worst case. For the example problem in Table 16.1, this model thus recommends the Mainframe as the best choice: its minimum profit is 35, greater than 10 or 0 for the other two options.

The difficulty with the maximin criterion is that it totally disregards most of the information we have about the problem. In the example, the procedure ignores five out of the nine entries in the payoff matrix—it only looks at the minima for each possible decision, in this case four cells. Furthermore, it gives absolutely no consideration to whether the different states of nature that lead to the minimum returns are likely or not. This model should therefore not be used; it cannot be wise to disregard relevant information systematically.

Maximax criterion. The maximax criterion is the converse of the maximin: it proposes that the decisionmaker maximize the maximum return available. Logically, it makes as much sense as the maximin: it simply would appeal more to the optimist rather than the pessimist. For the example problem, this model recommends PCs as the best choice since it offers the possibility of the absolute maximum return in the payoff matrix: 120.

This criterion compounds the deficiencies of the maximin criterion. It discards all but one entry in the payoff matrix, focusing only on the highest value. Furthermore, since most individuals do care a great deal about the worst cases that might happen, as Chapters 18 and 19 describe, this model has little appeal.

Regret criterion. The regret criterion is a combination of a criterion similar to the maximin and of the concept of regret. *Regret* is defined, for any option and a specific state of nature, as the difference between the benefit actually received and the maximum benefit that could have been obtained if the appropriate choice had been made. Thus, with reference to Table 16.1, if one had chosen Mainframes and Linear Programming jobs were received, the regret would be zero since one had made the best choice for the state of nature. The regret for the same choice would be 40 if Dynamic Programming jobs were received, since one could have had that much more if one had chosen Minis (i.e., 80 − 40), and the regret would be 85 (i.e., 120 − 35) if one received Engineering Economy jobs. Regret associated with a choice is thus defined with respect to the best outcomes for the states of nature, rather than in comparison with any specific other choice.

The notion of regret has considerable intuitive appeal. Many people define their success relative to others rather than in absolute terms. The concept is reflected by efforts to "keep up with the neighbors" or, conversely, by the "I don't care if I'm mediocre just so long as I'm above class average" attitude. The regret criterion has in fact been used quite considerably in the design of engineering systems, particularly in the development of water resources.

The criterion for selection in dealing with regret is similar to the maximin criterion on benefits. The difference is that one wishes to minimize regret, and specifically to minimize the maximum regret. This is then the *minimax criterion*. It is the inverse of the maximin criterion, and has all its weaknesses and specifically that of ignoring most of the information available about the problem.

In practice, the regret model involves two steps:

1. Transformation of the payoff matrix into a matrix showing "regret"
2. Application of the minimax criterion to this regret matrix

For the example problem, the appropriate regret matrix appears in Table 16.3.

TABLE 16.3
Regret matrix for the example problem of Table 16.1

Possible computer system	Annual "Regret" ($ × 10³) if jobs require: Linear programming	Dynamic programming	Engineering economy
Mainframe	0	40	85
Minis	30	0	110
PCs	90	80	0

The recommended choice is the Mainframe computer since this selection has the smallest maximum regret (85, compared to 110 for Minis and 90 for PCs).

The regret procedure has an additional, serious, peculiar flaw: the ranking of the choices can depend on the type of choices specified. This is because the regret associated with any choice is not an intrinsic property, but is a property defined relative to other possibilities. The consequence is that the ranking of the choices can be biased by the set of choices compared, as well as by irrelevant alternatives, as shown further under the discussion of the weighted index criterion.

The regret procedure can also be intransitive. This phenomenon, a reflection of the sensitivity of the procedure to the set of choices compared, is demonstrated by paired comparisons of the possible decisions in the example problem. The regret matrices associated with each of these three possible combinations appear in Table 16.4. Comparing the Mainframe and Minis choices, the regret criterion indicates that Minis should be preferred:

$$\text{Minis} > \text{Mainframe}$$

Considering Minis and PCs, the latter appear best:

$$\text{PCs} > \text{Minis}$$

When comparing PCs and the Mainframe choices, however, what seemed like the worst choice now appears best:

$$\text{Mainframe} > \text{PCs}$$

This kind of problem should be sufficient reason to reject the regret criterion forever.

TABLE 16.4
Regret matrices for paired comparisons of the options in the example problem

Possible computer system	Annual "Regret" ($ × 10³) if jobs require: Linear programming	Dynamic programming	Engineering economy
Mainframe	0	40	0
Minis	30	0	25
Minis	0	0	110
PCs	60	80	0
Mainframe	0	0	85
PCs	90	40	0

Notation: The symbol ≻ similar to the greater than symbol but with curved legs, should be read as "is preferred to." Facing the other direction, ≺, it means "is preferred by."

Weighted index criterion. The weighted index method has great intuitive appeal and is widely used. In some contexts it has even been the standard approach. This is notably the case in materials engineering, where the issue is to select the material that is best for some application.

The method operates on a table of consequences similar to but different from the payoff matrix previously defined. The difference is that the columns of the table represent known characteristics of each choice rather than possible outcomes associated with different events (see box). In this respect the weighted index method is not strictly comparable to the other criteria. It is discussed here because it is both a primitive method of choice and shares the same kinds of difficulties as the other criteria.

The general concept of the weighted index criterion is that the optimum choice is the one that provides the best weighted average of all the characteristics of the problem. These may be weighted equally or in any ratio that the designer chooses. Thus, for the example shown in the box, it is desirable to keep both the cost and the density of the part down, so that the material with the lowest weighted average of these characteristics will be best.

The first difficulty of the weighted index method lies in the selection of weights. In this regard it has the same drawback of the Laplace criterion in that it is sensitive to the way the columns are defined.

The greater difficulty of this method is associated with the normalization of the characteristics. All the components of the weighted index must indeed be normalized to avoid the dominance of characteristics that happen to be measured in small units. This normalization can be done with respect to a variety of constants such as the extreme values of the characteristics, or average values, or all the values associated with one of the alternatives (see box).

The difficulty is that the ranking of the choices derived from different normalizing constants may easily not be the same. In fact, they can generally be quite different. Table 16.5 shows the dependence of the weighted index on the normalizing constraints for the four materials discussed in the boxes. As Table 16.6 shows, the resulting rankings are dramatically different. In this case, specifically, whether A, B, or C appears best depends on which provides the basis for the normalizing constants. Alternatively, whether C is first, second, or third in the ranking depends on whether one normalizes on the largest, average, or smallest costs.

Further, the weighted index criterion is sensitive to irrelevant alternatives. If one eliminates material D from the analysis, because it is constantly unattractive and really irrelevant, then the ranking of the remaining materials may be quite turned upside down: C progresses from worst to best! It is not possible to state which (if any!) of these rankings is meaningful or correct. A method so capricious certainly should not form the basis for serious decisions.

Table for Weighted Index Method

Consider the problem of choosing a material for some component of an automobile. In this application, both cost and density (or weight) are important characteristics, the latter because greater weight reduces the fuel economy of the vehicle.

The problem is conveniently described by a table, as generated from a spreadsheet program on a personal computer. For some four materials, it could be as follows:

Material	Cost	Density
A	50	11
B	60	9
C	80	7
D	120	10

Normalization of Characteristics

The normalization defines the values of each characteristic with respect to the reference value. To achieve this, we divide each value of a characteristic for a choice by the reference value for the characteristic.

Refer now to the example choice of materials with the characteristics shown in the previous box. Suppose we wish to normalize all entries in the table of characteristics with respect to material A. The reference value for cost will thus be the cost associated with material A, 50. The normalized cost for material A is thus 1.00 ($=\frac{50}{50}$), for material B it is 1.20 ($=\frac{60}{50}$); and so on. Likewise, the reference value for density is also material A in this case, and the normalized density for A is again 1.00 ($=\frac{11}{11}$); for material B it is 0.8182 ($=\frac{9}{11}$); and so on.

If we normalized by extreme values, the procedure would be the same but the divisor different. For cost it would either be 50 or 120, depending on whether we selected the upper or lower extreme. For density it would be 7 or 11. Note that these reference points need not be connected with the same material.

In fact, the reference point need not be connected with any material. If we normalized on averages, the reference value for cost would be the average cost, 77.5 ($=\frac{310}{4}$), which is not the cost of any material. Indeed, the reference point can be totally arbitrary.

TABLE 16.5
Dependence of weighted index on basis for normalization and alternatives considered, for example materials

	Normalized by A				Normalized by B		
	Normal cost	Normal density	Index		Normal cost	Normal density	Index
A	1.0000	1.0000	1.0000	A	0.8333	1.2222	1.0278
B	1.2000	0.8182	1.0091	B	1.0000	1.0000	1.0000
C	1.6000	0.6364	1.1182	C	1.3333	0.7778	1.0556
D	2.4000	0.9091	1.6545	D	2.0000	1.1111	1.5556

	Normalized by C				Normalized by D		
	Normal cost	Normal density	Index		Normal cost	Normal density	Index
A	0.6250	1.5714	1.0982	A	0.4167	1.1000	0.7583
B	0.7500	1.2857	1.0179	B	0.5000	0.9000	0.7000
C	1.0000	1.0000	1.0000	C	0.6667	0.7000	0.6833
D	1.5000	1.4286	1.4643	D	1.0000	1.0000	1.0000

	Normalized by Maxima				Normalized by Maxima without D		
	Normal cost	Normal density	Index		Normal cost	Normal density	Index
A	0.4167	1.0000	0.7083	A	0.6250	1.0000	0.8125
B	0.5000	0.8182	0.6591	B	0.7500	0.8181	0.7840
C	0.6667	0.6364	0.6515	C	1.0000	0.6363	0.8181
D	1.0000	0.9091	0.9545				

	Normalized by Averages				Normalized by Minima		
	Normal cost	Normal density	Index		Normal cost	Normal density	Index
A	0.6452	1.1892	0.9172	A	1.0000	1.5714	1.2857
B	0.7742	0.9730	0.8736	B	1.2000	1.2857	1.2429
C	1.0323	0.7568	0.8945	C	1.6000	1.0000	1.3000
D	1.5484	1.0811	1.3147	D	2.4000	1.4286	1.9143

TABLE 16.6
Dependence of rankings of materials on the basis of normalization and alternatives considered

Normalizing factors	Order of ranking
Alternative A's cost and weight	A < B < C < D
Alternative B's cost and weight	B < A < C < D
Alternative C's cost and weight	C < B < A < D
Alternative D's cost and weight	C < B < A < D
Smallest cost and weight among all alternatives	B < A < C < D
Average cost and weight of all alternatives	B < C < A < D
Largest cost and weight among all alternatives	C < B < A < D
Largest cost and weight without considering D	B < A < C

Summary. Each of the five primitive models for "decisionmaking under uncertainty" is strongly deficient. Each in its own ways fails to meet the standards we should expect from a sound engineering procedure: intelligent, unambiguous use of all the relevant information about the problem. They should all be rejected.

16.3 COMPLEXITY OF CHOICE

The first contribution of decision analysis is to organize the problem for the analyst or planner. This may seem insignificant but is actually of great practical importance. Without structure to organize one's consideration of the possibilities, people tend to polarize around only a few, and they therefore miss important aspects of the problem. Experience also indicates that the elements of the problem which the structure brings into focus are often the most interesting and desirable solutions.

Correctly stated, any design of a system involves a very great number of alternatives. Even a relatively simple problem involves thousands of different possibilities. This is because each different combination of the design parameters leads to a different configuration, and because the number of combinations of items is an exponential function.

In general, the development of most systems involves three basic parameters:

- *Size*—How big, or what kind of activity should it be?
- *Location*—Where will this system be installed? For whom?
- *Time*—During which period will this project be started?

When each of these factors is part of the design, which is the normal situation for a system that will evolve over time, the number of possible combinations is defined by:

$$\text{Possible Designs} = [(\text{sizes})^{(\text{locations})}]^{(\text{periods})}$$

Mexico City Airport: Possible Designs

This is the first of a series of examples that refer to one of our studies, for the government of Mexico in 1970. The issue was: what is the best evolution of the airport system for Mexico City through the end of the century?

The number of possibilities for such a system is limited, compared to other systems, because airports only make sense in a few different sizes and because there are only a few places where you can place such a facility in a highly populated metropolitan area.

For Mexico City, only four different types or sizes of airports were considered: for domestic, international, general, and military aviation. Likewise, only two locations were available. Finally, we looked at the design for three periods: 1970, 1985, and the year 2000. The number of possible configurations logically possible were thus:

$$\text{Possible Designs} = ((4)^2)^3 = 4096$$

Some of the possible designs are not worth considering in any detail: it would evidently not have been worthwhile to build a complete new airport at a new site in 1985, and then close this down and return to the old site in 2000. Even discarding such combinations leaves us with thousands of possible patterns of development for this fairly simple system.

This is a very large number indeed. Note carefully that it is exponential, not multiplicative as many people would guess it should be (see box and Section 7.1).

Faced with such a large number of choices, most people easily become confused. This difficulty is compounded by uncertainties about what may happen. The mind quickly boggles after a few thoughts along the lines of "If I do this, then that will happen . . . I could then do the other; on the other hand, that might not happen . . . and besides I could always do something else in the first place." Without some structure to help people grasp all these possibilities, they generally focus their attention on a few choices.

This polarization of attention is the normal human way to deal with complexity for people who do not have some structure that will organize the problem. This is unfortunate as it leads, almost inevitably, to inferior choices. The obvious difficulty is that any evaluation that systematically fails to look at a large fraction of the possibilities has a high probability of missing some attractive solutions. This difficulty is increased, moreover, by the way people polarize their attention.

The fact is that people tend to focus on simple, clear-cut solutions. They tend to avoid the ones that are complicated. Specifically, they tend to avoid designs that are combinations of the simpler solutions. This fact is important because, as discussed further in Section 16.5, the best solutions are often precisely those that combine the best features of several simpler designs (see following box).

Mexico City Airport: Polarization of Choices

The question of how Mexico City ought to expand its airport system had preoccupied planners for over a decade before our participation in the process. The issue had actually become quite controversial.

The nature of this debate illustrates well the kind of polarization that often occurs in system planning. When we arrived, just two alternatives were being considered; two out of the over 4000 that could easily be imagined.

Moreover, these two were definitely clear cut. They were based entirely on either of the two possible locations, the existing one at Texcoco or the new site of Zumpango. Specifically, the solutions were:

1. "All Texcoco," all services were to be concentrated at Texcoco for all time.
2. "All Zumpango," all services were to be moved as soon as possible, for all time, to the new site.

No combinations, such as a partial development of a new site or its gradual development, were being considered. After the analysis, these combinations were exactly the ones that turned out to be most desirable.

16.4 STRUCTURE OF DECISIONS

Decision analysis structures the problem to bring out all the relevant choices, as well as all the possible outcomes that can be imagined. The means for doing this is the decision tree.

The *decision tree* is a conceptual device for enumerating each of the possible decisions that can be made, and each of the possible outcomes that may occur according to each of the events or states of nature that may arise. Graphically, it branches out from an initial node, to subsequent nodes where there is further branching, hence the metaphor of a "tree." In practice, decision trees are only drawn for illustrative or textbook purposes; for real problems it is impractical to draw a tree and the concept is embedded in computer files.

The decision tree first of all represents a sequence. This is the pattern of actions that one can take concerning a problem. It consists of the alternation of two elements:

1. *Decision nodes*, which are the moments when possible decisions are considered and a decision made
2. *Chance nodes*, the periods after a decision in which outcomes are determined by prevailing events or states of nature

Graphically, the sequence is:

$$D - C - D - C -$$

where an initial decision is followed by chance events, after which the designers face another set of decisions, and so on.

The decision tree is a "tree" in that each node in the sequence has several branches, as many as needed to show the possible decisions or chance events and resulting outcomes. The decision node thus indicates which specific choices might be made at that time. Figure 16.1 illustrates this for the problem of the Mexico City Airport. For this situation there are 16 different designs that could be implemented during the first period, each of the combinations of four kinds of services at the two separate sites.

The chance nodes differ from the decision nodes. Since we propose to do an analysis, they are in the first place subject to the discipline of the laws of probability. This implies first of all that the several branches define all the events E_j that may occur. Secondly, these events should be distinct; they should not refer to the same thing. If they did, there would be double-counting. Formally, the events referred to by the branches of each chance node should be "mutually exclusive and collectively exhaustive," so that their probabilities, P_j, sum to 1.0 (i.e., something does happen).

The chance nodes differ further from the decision nodes in that they may have several stages of branches from each node. This is a device that allows us to represent joint probabilities of two or more events conveniently (a possibility which is rather awkward in a payoff matrix, for example). Figure 16.2 illustrates this. Section 16.5 shows how to calculate the joint probabilities.

The chance nodes also carry considerable information, in addition to the possible states of nature. A complete branch from a chance node has two tags on it:

1. The probability P_j of the event E_j of that branch occurring (these must sum to 1.0 for any stage).

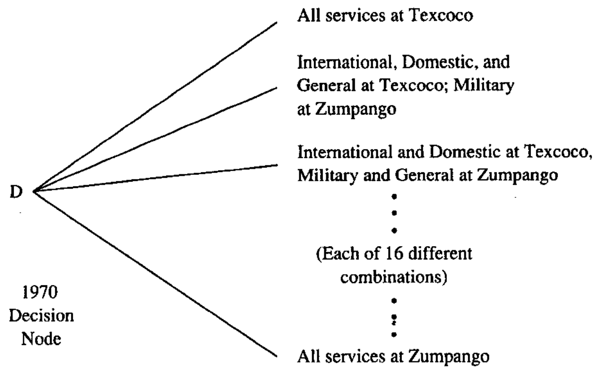

FIGURE 16.1
Initial decision node for the design of the Mexico City airport system: The possible decisions that could be made in the first period of 1970.

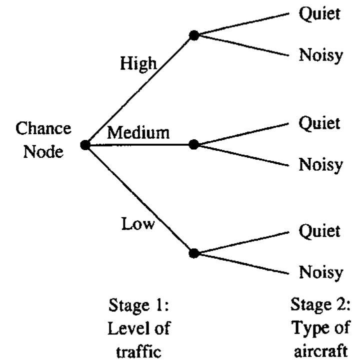

FIGURE 16.2
Structure of a chance node for the design of the Mexico City Airport System: Joint possibilities associated with the two different states of nature—level of traffic and type of aircraft—shown by stages of branches.

2. The outcome O_{ij}, measured in numerical terms, that results from the original decision D_i preceding the chance node and the state of nature or event E_j that prevails afterwards.

Graphically, these items of information are usually placed alongside and at the end of each branch from a chance node, thus

$$C \text{ — } \textit{Probability}, P_j \text{ — Outcome } O_{ij}$$

The complete structure of a decision tree for a real problem is thus quite complex, as might be expected for a complex problem. The beauty of the decision tree is that it can be constructed in a series of simple steps, each of which can be linked in computer files with little difficulty. The analyst can then appreciate the structure without having to draw it completely in detail. Figure 16.3 outlines the complete tree for the design of the airport system for Mexico City.

Semantic caution: The tree metaphor pervades decision analysis. When the tree is complex, as for Mexico City, it may be referred to as a "messy bush." When we discard possible choices, we speak of "pruning" the "branches."

The structure of the decision tree is useful—indeed essential—to the analysis. A first contribution is that it educates the planners as to the real possibilities.

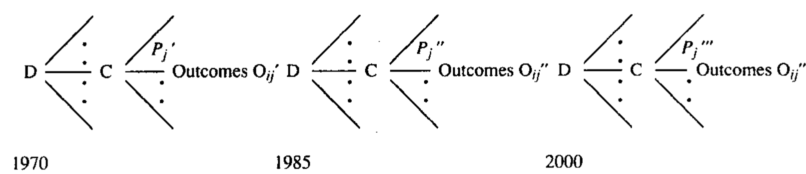

FIGURE 16.3
Sketch of the complete decision tree for the analysis of the Mexico City Airport System.

Thus in the case of the work in Mexico City, simply developing and presenting the decision tree to the planners and the ministers opened up the problem: it helped them realize that their choices were not polar (see Section 16.3), but included a wide range of combinations. The structure thus opened the way to finding the most desirable solution. The decision tree also provides the basis for the actual analysis, as discussed in what follows.

16.5 CALCULATIONS (SINGLE PERIOD)

For ease of presentation, the calculations of decision analysis are first explained and illustrated for a decision covering only a single period. In this case the decisionmaker has only one opportunity to make a choice; there is no follow-on or contingencies that are being taken into account. Such a simple case is, of course, quite unrepresentative of the reality of designing a system. The calculations for the typical case involving many periods are shown in Section 16.7.

Principle for calculations. The essential principle of decision analysis is simple: choose the decision that offers the *best average value.* Each of these terms, "best," "average," and "value" requires careful definition, however.

Value represents the real worth of any outcome to the decisionmaker. This may not correspond in any simple way to size of that outcome. For example, the value of a first plate of food to a starving person might be extremely high because it represents the difference between life and death; the value of the fourth plate on the other hand might be insignificant if the person were now already fed. As Chapters 18 to 20 explain in detail, value can be—and in general is—a complicated, nonlinear function of the outcomes. In other situations, such as those represented by choices whose economic outcomes are relatively small for the company or person concerned, the values can be considered to be equivalent to the outcomes.

To facilitate the discussion of decision analysis itself, this section ignores the possibility that outcomes do not represent their real value to the decisionmaker. This complication is left to Chapters 18 and 19. Once the process of decision analysis and the methods of determining real value are understood, they are easily brought together as needed, as Section 19.8 illustrates.

The "average" value relevant to a risky decision, D_i, is simply its *expected value:* the outcomes weighted by their estimated probability of occurrence. Formally, the average value of a possible decision D_i is then

$$EV(D_i) = \sum_j P_j O_{ij}$$

where $EV(\bullet)$ is a standard notation to indicate expected value.

Semantic caution: The term "expected value" has an established definition that can be confusing in decision analysis. In probability it indicates nothing more than the probabilistic average defined above. It has nothing to do with whether the things being averaged are expressed in outcomes or in their real values, as discussed later in Chapters 18 and 19.

The "best" average value refers to how the decisionmaker sees the outcomes. It is the highest average value of the outcomes that are desirable, and the lowest if they are not.

The calculations necessary to determine the optimal decision are easy. The process requires nothing more than many repetitions of the simple formula for expected value. The calculations may become confusing, however, since risky decisions often involve thousands of possibilities as Section 16.3 indicates. This confusion will be avoided if one organizes the work carefully.

The calculations will be straightforward if one carefully follows the structure of the decision tree. Specifically, the process should consist of two steps for a single stage:

1. Calculate the probability of the outcomes
2. Calculate the best decision

Calculation of probabilities. When the outcomes are determined by two or more independent events, it is necessary to calculate their probability of occurrence. For example, consider the analysis of a new technology whose success depends both on its cost of production (high, medium, low, etc.) and the overall state of the economy (boom, stagnation, recession, etc.). The estimates of the probabilities of each of these two different kinds of events are typically obtained through a different process and different experts. A production engineer may have provided estimates of the cost of production, for instance, and an economist the economic forecasts. The actual outcomes will depend on the events that occur as combinations of these two events such as "high cost and boom economy," "high cost and recession," and so on. The probability of these outcomes is the *joint probability* of the relevant events.

The joint probabilities can be easily calculated only if the events are *independent*, that is, if their probabilities of occuring do not depend on each other. (In the previous example the probabilities might not be independent: the probability of getting high costs could conceivably depend on whether the economy were in a boom, or a recession with plentiful labor.) The joint probabilities of independent events are simply the factor of their individual probabilities. Thus the probability of "high cost and boom" is simply the probability of high cost of production times the probability of a boom economy.

In practice, the calculation of joint probabilities may be a really important step in the process. This is because the outcomes may depend on many factors leading to a great many possible combinations. The calculations then usually indicate that the distribution of the probability of the outcomes is much broader than decisionmakers otherwise imagine. The detailed calculations, naturally done by computer, thus educate the decisionmakers as to the real risks involved in any decision (see box).

Calculation of Probabilities

This example comes from the files of an electric power utility that was trying to identify the best way to conserve energy. Their analysis indicated that the outcomes, in terms of net benefits, were determined by five independent factors: the lifetime of the device; the company's future discount rate; and the escalation rates or rates of increase in required capacity, production costs, and revenues.

Schematically, their chance node after any possible decision was as in Figure 16.4(*a*). The total number of branches from the chance node were $(3\times3\times3\times3\times2) = 162$. Each branch, being a joint probability of all five different independent events, had a specific, distinct outcome. These constituted a probability distribution that could be summarized graphically for easy display. Figure 16.4(*b*) gives an example. This kind of result helped the directors of the utility appreciate that any decision could have a really wide range of outcomes.

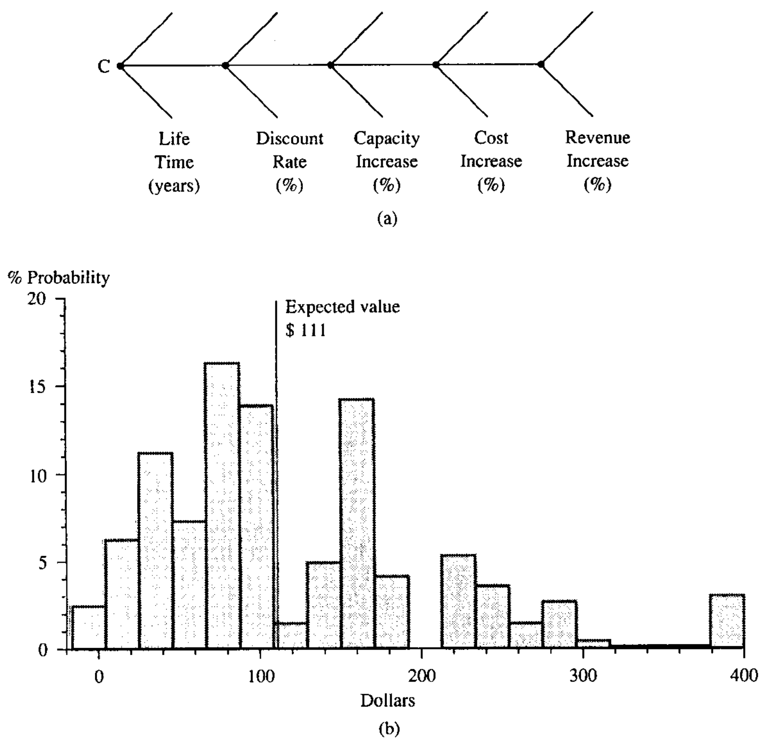

FIGURE 16.4
(*a*) Sequence of difference chance events. (*b*) Joint distribution of outcomes resulting from sequence of chance events.

Optimal decision for single period. The best decision is the one with the best average value. Insofar as the outcomes correspond to their real worth, the best decision is simply the one with the best expected value for the outcomes.

The expected value for each possible decision is easily calculated if there is only one period of possible decisions and chance nodes in the decision tree. It is simply the probabilistic average as defined previously:

$$EV(D_i) = \sum_j P_j O_{ij}$$

The optimal decision D* is thus the decision with the best expected value.

To illustrate the calculations, consider this simple decision: should you take your raincoat or not, the next day you go to work or school? Suppose for simplicity that the risky events E_j are that it may either rain or be clear, and further that

- the probability of rain, $P(\text{Rain}) = 0.4$, so that the probability of no rain is $P(\text{Clear}) = 0.6$.
- the values of the possible outcomes, O_{ij} (such as being caught in the rain without a raincoat) are the numbers indicated on the right in the following decision tree.

Your decision tree for this single period is thus:

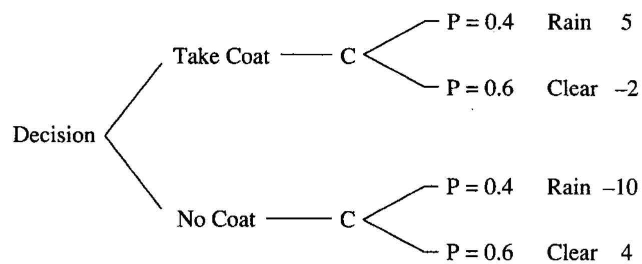

To calculate the optimal decision, we calculate the expected value of each choice:

$$\begin{aligned} EV(\text{Take Coat}) &= P(\text{Rain})(5) + P(\text{Clear})(-2) \\ &= (0.4)(5) + (0.6)(-2) = 0.8 \end{aligned}$$

$$EV(\text{No Coat}) = (0.4)(-10) + (0.6)(4) = -1.6$$

The best choice would be to take the raincoat:

$$EV(D^*) = EV(\text{Raincoat}) = 0.8 > EV(\text{No Coat})$$

16.6 CONTINUOUS PROBABILITY DISTRIBUTIONS

For ease of presentation, the decision trees have so far been shown with a finite number of branches emerging from each chance node. This should not be taken to imply that decision analysis can only handle a finite number of outcomes. On the contrary, the method can easily deal with the infinite number of possibilities associated with a continuous probability distribution.

The application of decision analysis to situations involving continuous probability distributions is essentially identical to the procedure just described. The only difference is that we must use integrals to evaluate the expected values. Thus, if the probability distribution is some function of the chance event, E_j:

$$\text{probability distribution} = \text{pdf}\,(E_j)$$

then the expected value of any decision is simply the integral of this distribution times the possible outcomes:

$$EV(D_i) = \int O_{ij} \bullet \text{pdf}\,(E_j)\, dE_j$$

See box for an example application.

Many kinds of risky situations are best represented by continuous probability distributions, as indicated in Section 15.4. These include all kinds of physical and social phenomena that are to occur repeatedly and for which a continuous distribution can be estimated. They include, for example,

- rainfall, needed for the planning of water resources systems.
- earthquakes, whose intensity must be estimated for the design of structures.
- traffic intensity, by time of day and season, useful for the management of all kinds of networks, from communications to highways.
- economic conditions, which play an important part in the success of business investments.

16.7 CALCULATIONS (MULTIPLE PERIODS)

The realistic situation for any system is that there are many stages of possible decision and chance nodes, as for the problem of the Mexico City Airport illustrated in Figure 16.3. To find the optimal choice, one must calculate the best decisions stage by stage, starting with the last period. The best decision for this last period is calculated as for a single period (Section 16.5).

The process of finding the optimal decision for multiple periods is similar to that employed in dynamic programming (see Chapter 7). The result of this procedure is a great economy in the number of calculations necessary, just as there is through the use of dynamic programming.

The problem of finding the best decision in a sequence of possible decisions is equivalent to finding the best paths through the decision tree. It thus has the same form as the sequential route problems discussed in Section 7.2, and can be solved similarly. In practice, a formal dynamic program is essentially never used, as it is easier to proceed through the decision tree than to set up the formal structure of a dynamic program. But the overall effect on the number of calculations is the same.

Consider the analysis of the development of the Mexico City Airport discussed in Section 16.3. At each period there were 16 possible decisions as Figure 16.1 shows. For the three periods there were thus $(16)^3 = 4096$ different combinations, and thus paths through the decision tree. This is the number of expected value calculations that would be necessary if we examined each possible sequence explicitly. By analyzing the problem period by period, however, we only considered 16 possible decisions at each of the three periods, and thus $16 \times 3 = 48$ in all. The size of the problem was thus multiplicative rather than exponential in the number of periods—(16×3) versus $(16)^3$.

The reason the analysis is so much easier when done stage by stage is that we determine a single best decision after the chance outcome of a previous period, and thus drop from consideration all the other possible decisions at that period—

Optimal Design Against Earthquakes

The definition of the optimal design against earthquakes was a classic application of decision analysis to a situation involving continuous probability distributions. Some years ago, we did this for the United States.

There were—and are in general—a discrete number of possible choices. These are represented by building codes that define consistent standards of strength for a structure. Naturally, designing according to a more stringent building code both provides more strength and costs more money. The issue is: which is the right code for any given locality?

The optimal choice depends on the kinds of earthquakes that occur. Where there are only insignificant earthquakes, it is not worthwhile to spend much money on protection against them, and vice-versa.

The possible intensity of earthquakes in any region is conveniently summarized by a probability distribution giving the frequency of occurrence of earthquakes of any given intensity. This information is routinely collected in earthquake zones, and reported either on the Richter scale or the Modified Mercalli Index.

The damage that may occur as a result of an earthquake, O_{ij}, depends both on the choice of the building code, D_i, and the intensity of the earthquake, I_j. The expected cost associated with any code is thus the cost of its implementation plus the expected damage due to earthquakes. This is:

$$\text{Expected Cost of Building Code} = \text{Construction Cost} + \int (\text{Damage due to } I_j)(\text{Distribution of } I_j)\,dI_j$$

as well as the combinations they would represent. This is the step known as "pruning" the decision tree.

The procedure is best illustrated by a simple example with only a very few branches. This kind of problem does not demonstrate the way the analysis reduces the number of calculations; it has few branches to begin with and thus few to prune. But one can fully graph a simple problem as one cannot a realistic problem such as that of the Mexico City Airport!

The example concerns the problem of deciding how best to plan for development. The idea is, as often occurs in practice, that one must first decide whether to reserve a site for expansion and then later decide whether to build. In between these decisions one gets a chance to see how fast ones' activities grow. The final outcome depends on this growth and some other chance events occurring after the decision as to whether to build. This situation is, in fact, a simplified version of the Mexico City case, as Figure 16.5 shows.

The analysis of the problem begins with the last period and consists of calculating the expected value associated with each of the possible decisions at the available decision nodes. From these analyses we can conclude which decision is best at each decision node in the last period. Knowing what is best at that time and what we could do if we were ever in the position defined by that node, we can drop from further analysis all the possible decisions we could avoid. Figure 16.6 shows the result so far. Note that since we now know what we would do as a result of each chance outcome in the preceding stage (for example, if we had reserved a site and there were growth, we would build), and knowing its

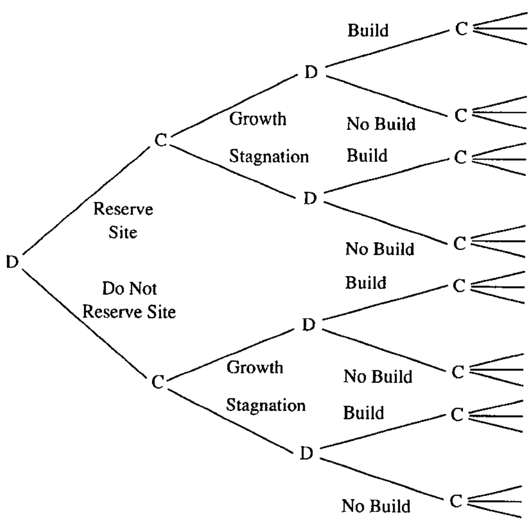

FIGURE 16.5 A decision tree for the example problem of trying to determine the best pattern of development for some group.

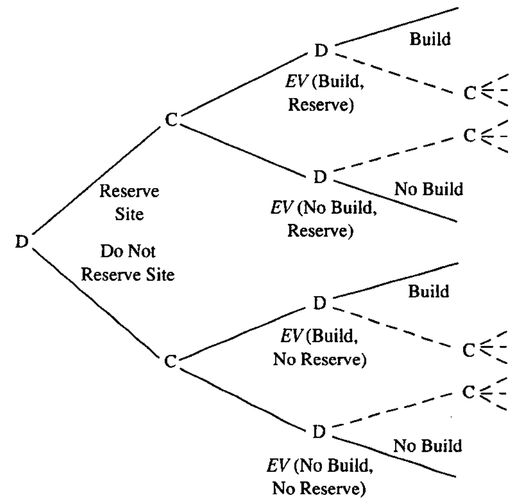

FIGURE 16.6 Example decision tree after analysis of last stage in which nonoptimal branches have been eliminated from further consideration.

expected value, we simply define the value of that outcome as the expected value of the best decision at that point.

The process continues with the next, that is, the preceding, period. At this point the problem is as shown in Figure 16.7. Expected values of each possible decision are again calculated and the decision with the best expected value is selected. Note how this analysis looks at simple outcomes of the chance nodes, rather than the really complex ones represented by all the possible decisions and branches that could be taken later.

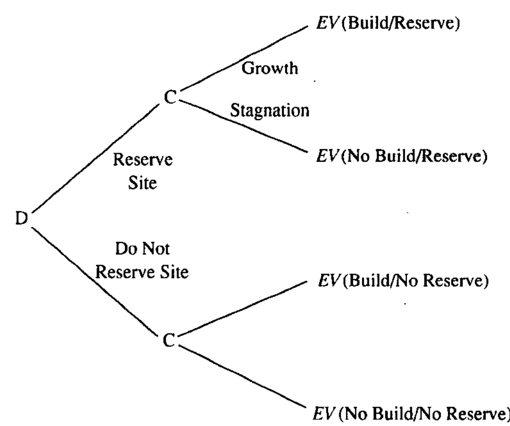

FIGURE 16.7 The example problem viewed from the next, preceding stage.

16.8 DECISION ANALYSIS AS STRATEGY

Results. The final result of a multistage decision tree is a strategy. Formally, a *strategy* can be considered a sequence of decisions contingent on all the possible outcomes. It is not a single sequence of decisions best for all occasions.

For example, let us say that the optimal decision at the first stage of a multistage site selection problem is to reserve a site. The decision at the next stage is not specified by this first decision; it depends on the intervening chance outcome. For our example, the final recommendation would be: First reserve a site; then build if there is growth or do not build if there is stagnation. This was, in fact, our final recommendation to the President of Mexico and his decision in 1971 concerning the capital's airport. Zumpango was reserved, but ultimately not built since traffic stagnated as a joint result of economic recession and development of direct flights to coastal resorts. We considered this a most successful decision: it preserved the options when needed, but avoided massive costs that ultimately were not necessary.

This kind of result is completely different than the "master plans" frequently encountered in practice. A *master plan* specifies a single sequence of development in each phase (such as to build a new airport at Zumpango, close the old airport, transfer all traffic, etc.). Master plans have been the typical result of most efforts to define the development of large-scale engineering products.

The definition of an optimal strategy of development is one of the real contributions of decision analysis. It permits and encourages the designer to deal effectively with the inevitable uncertainties. Master plans by contrast promote inflexibility and a resistance to doing the right thing for prevailing circumstances. Master plans are also inherently far less than optimal since there is always a reasonable chance that the proposed master plan (e.g., reserve site and build) is disadvantageous compared to some other pattern of development (e.g., reserve site but do not build—if there is no growth).

16.9 APPLICATION

The use of decision analysis was recently very effective in facilitating the design of a major national system. We used it to help the Australian government decide how to develop airports for their major hub, Sydney.

The advantage of using decision analysis in this case is particularly evident when contrasted to the standard alternative of benefit-cost analysis, which had been repeatedly used quite unsuccessfully on the same problem. This case also nicely illustrates how decision analysis

- educates the client, by structuring the problem to bring into focus the real range of risk and options.
- defines an optimal choice, particularly for the immediate first stage.
- identifies a preferred strategy for dealing with the inevitable eventual changes in the situation.

Background. Australian governments have long recognized the need to provide additional airport capacity for Sydney. The existing commercial airport, Kingsford-Smith, had—when we were called in to help in 1984—only two runways to serve a city of three million. Comparable cities, such as Boston, will have three or four major runways.

The situation was becoming increasingly urgent because the few plausible sites were being engulfed by metropolitan development. For Sydney, this situation was particularly acute. The Kingsford-Smith Airport sits on restricted landfill in Botany Bay, surrounded by both hills and densely inhabited communities. Furthermore, the Sydney region is ringed by substantial mountains and cut up by rivers. The options were quite limited as the need pressed.

Previous governments had made several major attempts to decide how to expand the airport system for Sydney. Each had tried to define the best master plan for the region, either to extend the existing site by filling in a major portion of the Bay or to acquire another site and build a new airport. None of the previous efforts had been successful in defining an acceptable solution, despite millions of dollars worth of first-class work in engineering and planning.

The preceding studies had each evaluated the alternatives using some form of benefit-cost analysis. Basically they would attempt to estimate future costs and benefits of any candidate design over the next 20 to 30 years; based on some forecasts of growth. Inevitably, by the time the final recommendations came to be debated, some two years or so after the forecasts were factored into the analysis, these predictions were already wrong. These analyses were thus simply not credible.

The proponents of each master plan were always in the unfortunate position of "defending the indefensible." To demonstrate that their chosen plan was the best alternative they had to assert that their forecasts were correct. This is quite impossible; no amount of technical knowledge can predict popular trends over 20 or 30 years (see Section 15.1). Air travel itself changes so rapidly that the fallacy of their logic was evident almost immediately.

Education of client. The first step in any evaluation is to pick the right form of analysis for the situation at hand, as Chapter 10 stresses. Specifically, it is necessary to determine whether risk is an important factor for the problem.

Previous analysts had in effect assumed that risk was insignificant. They had based their evaluations on specific forecasts of what would happen, and had gone to great pains to defend these predictions when it came to justifying their recommended master plan. At the time our team arrived on the scene, the Australian government had an official forecasting bureau. This agency was confidently predicting an annual rate of growth of air traffic of 2%, maximum, and the government wished us to base our evaluation on this forecast.

The fact is that 20 to 30 year forecasts must be regarded as highly inaccurate. Section 15.1 makes the point. Long-term predictions concerning developing industries, such as air transport, are especially uncertain since they both grow rapidly and fluctuate substantially. This reality had to be brought home to the client.

Demonstrating that the official forecasts were unreliable was easy, as it generally should be. All we had to do was to contrast the current predictions with the previous ones made 5 and 10 years earlier, with the historical record and, for extra reinforcement, with experience in comparable situations elsewhere. These comparisons showed that growth rates could be as high as 10% for many years, were about 7% over the previous 20 years, and that previous forecasts had recently asserted that long-term rates would be about 5%! Naturally, the professional forecasters were upset that this evidence was brought forward: it effectively dismissed their work. Yet the client had to recognize, as anyone should, that the future outcomes are highly uncertain.

Secondly, we had to help the client appreciate the complexity of the decisions. Previously the problem had been presented as a stark, polar alternative: build a major facility either at the current or at a new site. The reality of course is that any large system requires many stages, each involving many possible decisions, leading to a very great number of alternatives. Drawing out the decision tree graphically illustrates the issue. We stressed the point verbally by emphasizing that the only immediate decision the government had to make was to select a site; decisions about whether any airport should be built, or at what size, only had to be made many years later, after the actual growth in traffic had been observed.

Evaluation. Decision analysis was the appropriate form of evaluation because of the many possible outcomes of any decision. It is the only method that can deal adequately with risky alternatives. It is unquestionably better for this problem than the benefit-cost analyses, which have no reasonable way of dealing with contingencies.

The analysis itself proceeded normally. It clearly showed that the best immediate step was to acquire a site for a new airport—however the traffic might develop and whatever decisions the government might eventually want to take. This is because securing the site provides insurance against the chance of growth, and good decisions almost inevitably provide means to deal with risks.

We did not use a decision tree when it came time to present the results to the Minister for Aviation, the government generally, and ultimately the public-at-large through the environmental impact statement. One cannot expect that this concept, which requires considerable study, will mean anything to most people. We therefore used a simple payoff matrix to summarize the analysis. Table 16.7 reproduces our presentation.

This display makes it obvious that the optimal choice, for the first stage of the process, is to secure a site for a new airport. Whereas a new site is not the better choice if we absolutely knew that air traffic would stagnate (at 2% a year), the reality is that we cannot be so sure. By acquiring a new site the government has at least an adequate solution whatever happens.

Strategy. A strategy of development is the only reasonable solution when there is risk, as decision analysis indicates. Once one recognizes that long-term predictions are unreliable, master plans for what should be done are indefensible—as

TABLE 16.7
Consequences of alternative first stage decisions for Sydney

Action	High forecast	Median forecast	Low forecast
Select site	Good decision: site available to meet demand	Good decision: site available to meet demand	Neutral decision: temporary sterilization of site; land put to alternative use until required
No action	Poor decision: congestion at Kingsford-Smith Airport; site at higher cost or greater distance subsequently selected	Poor decision: congestion at Kingsford-Smith Airport; site at higher cost or greater distance subsequently selected	Neutral decision: proposed sites put to higher priority use: site at higher cost or greater distance subsequently selected

previous planners for Sydney's airport had discovered. Our recommended strategy was simple:

- First, acquire a new site.
- Then, depending on how fast traffic grew, either build to the appropriate size or not.

Results. The most important result is that decision analysis led to an acceptable decision. Within a year after we submitted our recommendation and environmental impact statement, uneventful public hearings had been concluded and the government began to acquire the site. Three years later the job was more than half done.

The reason decision analysis succeeded where benefit-cost had failed so miserably is that the recommendation made simple common sense. Our key argument was that we could not predict the future and *therefore* had to take adequate precautions. Rather than trying to "defend the undefensible" idea of omniscience, we made a strength of the reality of risk.

REFERENCES

Behn, R. D., and Vaupel, J. W., (1982). *Quick Analysis for Busy Decision-Makers*, Basic Books, New York.

de Neufville, R., and Keeney, R. (1972, 1974). "Use of Decision Analysis in Airport Development for Mexico City," Chapter 23 in *Analysis of Public Systems*, A. W. Drake, R. Keeney, and P. M. Morse, eds., MIT Press, Cambridge, MA; and Chapter 24 in *Systems Planning and Design*, R. de Neufville, and D. H. Marks, eds., Prentice-Hall, Englewood Cliffs, NJ.

de Neufville, R., (1990). "Successful Siting of Airports: the Sydney Example," *Journal of Transportation Engineering*, American Society of Civil Engineers, Vol. 116.
Raiffa, H., (1970). *Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley, Reading, MA.

PROBLEMS

16.1 *Election Campaign*
As a politician you can end your campaign for election either in City A or B. You estimate that a speech in A will gain you either 2000 votes or, if local students protest, 4000 votes more than your opponent. However, if your opponent also comes to A you will have to cancel the speech and debate, in which case you will get 1000 votes less than your opponent. Your manager estimates that the probability of your opponent showing up is 0.5. The probability that the students will protest is 0.3.

Alternatively, you can go to B, where there are no students and your opponent has already been, where you may either gain 2000 votes (probability 0.3) or 600 votes (probability 0.7).

(*a*) Which city would you go to if you are trying to maximize your votes?
(*b*) Which city should you go to if current polls show you are losing by 800 votes?

16.2 *Money Bags, Take 2*
Welcome back (see Problem 15.1) to the "Money Bags" TV show. You have won $600. Then Monty, the MC, offers you, for your $600, a wallet in which there are either ten $100 bills (making $1000) or one $20 bill and three $100 bills (making $320). He also gives you the option, for a cost of $100, to pull one of the bills out of the wallet before choosing whether or not to take the contents (which will be either $1000 or $320) at the additional cost of $600.

(*a*) Structure the decision tree for determining whether you should keep your $600, take the wallet without pulling a bill (at a cost of $600), or look at a bill before deciding whether or not you should buy the wallet. Calculate the appropriate consequences and probabilities for the tree.
(*b*) What is the strategy that maximizes the expected monetary value?

16.3 *Assembly Robot*
An industrial assembly robot is being used to join two machine parts. A peg from Part 127 must be inserted into a hole in Part 34A. The robot may be programmed either to pound Part 127 into Part 34A, or to nudge it in gently. The pounding process is faster if Part 34A was presented to the robot in the proper alignment, but slower otherwise.

	Part 34A presentation	
Robot action	**Aligned**	**Not aligned**
Pound	0.2 s	1 s
Nudge	0.4 s	0.6 s

The conveyor belt seems to align correctly about 70% of Part 34As that are fed to this robot.

The robot is also equipped with a laser and a set of photocells that can be used to test that alignment "visually." Correct alignments will give "okay" test results 80% of the time, while misalignments yield "not okay" readings 90% of the time. The laser-bearing head takes 0.001 s to make the test itself, but 0.05 s are required for the head to swing into test position and then back out of the way, so that Part 127 may approach without hitting the laser.

(*a*) Structure the decision tree for the choice of pound, nudge, or test.
(*b*) Find *P*(aligned/"okay"), *P*(aligned/"not okay") using Bayes' Theorem.
(*c*) Assign appropriate probabilities and outcomes to the branches of the tree. Find the expected time of each alternative strategy by folding back the tree. Which strategy will be the fastest?

16.4 *Mars Probe*
A group of researchers has leased space on a Mars probe for an experiment to be run once the probe lands. They have limited grant money available, but if their experiment succeeds, they will receive a grant for $10 million to continue their research. If their experiment fails they will get nothing. The experiment to be conducted can be run in one of two ways, but can only be run once. Scientists in the group have two theories about the Martian surface which they deem equally likely. The validity of the theories affects the probabilities of success of the two methods as indicated in the following table.

	Probability of success	
True theory	**Method X**	**Method Y**
A	0.8	0.4
B	0.4	0.7

(*a*) Calculate the probability of success if method X is used; if method Y is used.
(*b*) Draw the decision tree for deciding how to conduct the experiment.
(*c*) Which method should be used?

16.5 *Software Development*
You are president of Celestial Software company, and have enough money to buy an exclusive license to produce software for one of the two personal computers due to be released next summer. You are convinced that only one of these will be a hit, but you do not know which. You must act soon or another company will purchase the license before you do. If you own the license for the PC that becomes a hit, you will make a profit, but if your PC flops, then you will suffer a loss:

	Financial outcome ($1M)	
License purchased	**Hit**	**Flop**
A	10	−1
B	8	−0.5

(*a*) Assuming that both PCs are equally likely to become hits, draw the decision tree and label its branches with the appropriate probabilities and outcomes.
(*b*) Which license should you buy?

16.6 *Traffic Department*
As head of the Traffic Department you are planning a new system of traffic lights. Your experts are divided into two groups. One believes in theory A and one in theory B (A or B is true, not both). At present you consider the two theories equally plausible. You have two options for your light system—systems X and Y.

If theory A is true and you adopt system X, then your payoff will be 1 with probability 0.8, and 0 with probability 0.2. If theory B is true and you adopt system X, then your payoff will be 1 with probability 0.1, and 0 with probability 0.9. Your payoff to system Y is 0.5.

You have a two-period time horizon. You can choose system Y now for the two periods or you can experiment for the first period (i.e., use system X) and then choose a system for the second period depending on your payoff in period 1.
(*a*) Draw the decision tree. Label carefully.
(*b*) Put the payoffs at the ends of the branches.
(*c*) Put the probabilities you can write down immediately on the appropriate branches.
(*d*) Explain how to compute the other probabilities.
(*e*) Assuming that you wish to maximize the expected value of the payoff, what strategy should you follow?

16.7 *Oil Exploration*
An oil company has the option of exploring for oil at only one of three sites. The company estimates that the three sites have approximately the same productive lifetimes and costs and that their probable production is as in the table:

	P(output), barrels per day			
Site	**0**	**20K**	**100K**	**1000K**
A	0.05	0.4	0.5	0.05
B	0.1	0.75	0.1	0.05
C	0.7	0.05	0.05	0.2

At which site should the company explore?

16.8 *Earthquake Protection*
Your friend, Phil Vallee, is going to buy a $300,000 house on the West Coast. However, the site is on a 45° slope and is susceptible to landslides caused by earthquakes. During the 10 years Phil plans to live in the house, there is a 10% chance that a damaging earthquake (which would cause a slide and destroy the house) will occur. The probability of having more than one earthquake during the next 10 years can be neglected. If such an earthquake occurs, it will be "moderate" with 90% probability and "severe" with 10% probability. Your friend can take one of three actions: (1) do not provide earthquake protection (and spend exactly $300,000 initially), (2) design the foundations of the house, at an additional cost of $15,000, so that it can take a moderate earthquake but not a severe one, or (3) spend an additional $120,000 and design the house so that it will survive even a severe earthquake.
(*a*) Draw a decision tree showing all possible actions, outcomes, and consequences (in dollar terms).
(*b*) If money were the only consideration, which alternative should you recommend?

16.9 *Sonny's PV's*
Sonny Reyes (see Problem 15.6) must decide how to manufacture the PV panels. He has three choices:

- Develop a new method
- Alter existing methods
- Get an outside firm to produce them

Developing the new method would yield a profit of $11 million if successful, and if it cannot be developed, the outside firm must be used for a profit of $2M. There is a 70% chance that the new method will be successfully developed.

Altering the existing method successfully will yield a profit of $7M, and there is a 90% chance of the alterations being successful. If not successful, the outside firm must be used for a profit of $3M. If used immediately, the outside firm will definitely be able to produce the panels and this would lead to a profit of $5M.

Neglecting any time considerations,
(*a*) Structure the decision tree.
(*b*) Solve and select the best strategy.

16.10 *PY-RIC Enterprises*
M.T. "Vic" Torrey heads the division of PY-RIC that launches communications satellites. Vic has to decide whether to go ahead with today's launch or delay it for more favorable conditions. At present the probability of a successful launch is estimated at 0.90. The resulting profits would be $20M. If the launch fails, the loss of the payload costs $120M. A month's delay would increase the probability of success to 0.95, but the profit would drop to $10 M due to penalty fees paid to clients. If the launch is totally aborted, PY-RIC must refund clients $10M.
(*a*) Draw the decision tree.
(*b*) What should Vic do if he decides only according to expected values of profit?

16.11 *Marian Haste*
Marian Haste is a painter under contract to paint the exterior of a building for $225,000. Unfortunately for her, the outside temperature may drop below freezing. If it does, the paint will not stick well and may peel off. The paint may also peel off even it does not freeze. Either way, Marian would have to repaint the building, at a cost of $150,000.

The radio forecasts a 60% chance of freezing. Marian also believes there is a $\frac{1}{3}$ chance the paint may peel if it freezes, but a 10% chance if it does not. Her choices now are to go ahead and paint, or to defer until there is no chance of frost. Deferring the job would require her to pay $30,000 in overtime and penalties. What would you advise Marian Haste to do? (For your own amusement, the question is should Marian Haste go ahead, and repaint at leisure?)

16.12 *P. O'Toole*
Paddy O'Toole makes outdoor furniture that sells in the summer (the Paddy O. Furniture line). Production starts in January, however, right after the holidays. That is when Paddy must decide on whether to hire 1, 2, or 3 assistants at $14,000 each for the season. He could then produce Low (1200 units), Medium (1800 units), or High (2400 units). His product sells for $30/unit on average in June. Units not sold in June must be sold at half price. Based on experience, he estimates P(N), the probability that the demand for his furniture next June will be N, as:

N	:	1000	1500	2000	2500
P(N)	:	0.2	0.3	0.4	0.1

How many assistants should Paddy hire to maximize expected profit?

CHAPTER 17

INFORMATION GATHERING

17.1 THE ISSUE

New information about a situation provides us with the opportunity to base our decisions on more current, better estimates of the consequences of any choice. We may expect that new information thus leads to better decisions. This logic suggests that any decision analysis should consider spending some effort gathering information.

Any decision problem has opportunities for gathering new information. It is collected through deliberate efforts to test the situation. For example

- An oil company may run detailed geological surveys before deciding whether to drill a test well.
- A manufacturer can build a prototype machine or even a factory before committing to full-scale production.
- A distributor may test market a new product before deciding if and how to sell it.
- A planner may use forecasting models to project future demand for a service.

In general, professionals are always engaged in some explicit forms of information gathering, so that they can improve their decisions.

Information can also be obtained simply by waiting. Indeed, the best way to determine what will happen N years hence is to wait until then and find out. If a decision must be taken earlier, one can at least wait awhile and see how some of the uncertainties about the future are resolved. There is thus a good reason to defer decisions.

In practical terms, the decisionmaker must decide how much information to obtain, if any. We must expect that additional information will be less and