CHAPTER 1

Decision Analysis

CONTENTS

PERHAPS THE MOST FUNDAMENTAL AND IMPORTANT TASK THAT A MANAGER FACES is to make decisions in an uncertain environment. For example, a manufacturing manager must decide how much capital to invest in new plant capacity, when future demand for products is uncertain. A marketing manager must decide among a variety of different marketing strategies for a new product, when consumer response to these different marketing strategies is uncertain. An investment manager must decide whether or not to invest in a new venture, or whether or not to merge with another firm in another country, in the face of an uncertain economic and political environment.

In this chapter, we introduce a very important method for structuring and analyzing managerial decision problems in the face of uncertainty, in a systematic and rational manner. The method goes by the name **decision analysis.** The analytical model that is used in decision analysis is called a **decision tree.**

1.1 A DECISION TREE MODEL AND ITS ANALYSIS

Decision analysis is a logical and systematic way to address a wide variety of problems involving decision-making in an uncertain environment. We introduce the method of **decision analysis** and the analytical model of constructing and solving a **decision tree** with the following prototypical decision problem.

BILL SAMPRAS' SUMMER JOB DECISION

Bill Sampras is in the third week of his first semester at the Sloan School of Management at the Massachusetts Institute of Technology (MIT). In addition to spending time preparing for classes, Bill has begun to think seriously about summer employment for the next summer, and in particular about a decision he must make in the next several weeks.

On Bill's flight to Boston at the end of August, he sat next to and struck up an interesting conversation with Vanessa Parker, the Vice President for the Equity Desk of a major investment banking firm. At the end of the flight, Vanessa told Bill directly that she would like to discuss the possibility of hiring Bill for next summer, and that he should contact her directly in mid-November, when her firm starts their planning for summer hiring. Bill felt that she was sufficiently impressed with his experience (he worked in the Finance Department of a Fortune 500 company for four years on short-term investing of excess cash from revenue operations) as well as with his overall demeanor.

When Bill left the company in August to begin studying for his MBA, his boss, John Mason, had taken him aside and also promised him a summer job for the following summer. The summer salary would be $12,000 for twelve weeks back at the company. However, John also told him that the summer job offer would only be good until the end of October. Therefore, Bill must decide whether or not to accept John's summer job offer before he knows any details about Vanessa's potential job offer, as Vanessa had explained that her firm is unwilling to discuss summer job opportunities in detail until mid-November. If Bill were to turn down John's offer, Bill could either accept Vanessa's potential job offer (if it indeed were to materialize), or he could search for a different summer job by participating in the corporate summer recruiting program that the Sloan School of Management offers in January and February.

Bill's Decision Criterion

Let us suppose, for the sake of simplicity, that Bill feels that all summer job opportunities (working for John, working for Vanessa's firm, or obtaining a summer job through corporate recruiting at school) would offer Bill similar learning, networking, and resumé-building experiences. Therefore, we assume that Bill's only criterion on which to differentiate between summer jobs is the summer salary, and that Bill obviously prefers a higher salary to a lower salary.

Constructing a Decision Tree for Bill Sampras' Summer Job Decision Problem

A **decision tree** is a systematic way of organizing and representing the various decisions and uncertainties that a decision-maker faces. Here we construct such a decision tree for Bill Sampras' summer job decision.

Notice that there are, in fact, two decisions that Bill needs to make regarding the summer job problem. First, he must decide whether or not to accept John's summer job offer. Second, if he were to reject John's offer, and Vanessa's firm were to offer him a job in mid-November, he must then decide whether to accept Vanessa's offer or to instead participate in the school's corporate summer recruiting program in January and February.

These decisions are represented chronologically and in a systematic fashion in a drawing called a **decision tree.** Bill's first decision concerns whether to accept or reject John's offer. A decision is represented with a small box that is called a **decision node,** and each possible choice is represented as a line called a **branch** that emanates from the decision node. Therefore, Bill's first decision is represented as shown in Figure 1.1. It is customary to write a brief description of the decision choice on the top of each branch emanating from the decision node. Also, for future reference, we have given the node a label (in this case, the letter "A").

If Bill were to accept John's job offer, then there are no other decisions or uncertainties Bill would need to consider. However, if he were to reject John's job offer, then Bill would face the uncertainty of whether or not Vanessa's firm would subsequently offer Bill a summer job. In a decision tree, an uncertain event is represented with a small circle called an **event node,** and each possible outcome of the event is represented as a line (or branch) that emanates from the event node. Such an event node with its outcome branches is shown in Figure 1.2, and is given the label "B." Again, it is customary to write a brief description of the possible outcomes of the event above each outcome branch.

Unlike a decision node, where the decision-maker gets to select which branch to opt for, at an event node the decision-maker has no such choice. Rather, one can think that at an event node, "nature" or "fate" decides which outcome will take place.

The outcome branches that emanate from an event node must represent a **mutually exclusive** and **collectively exhaustive** set of possible events. By mutually exclusive, we mean that no two outcomes could ever transpire at the same time. By collectively exhaustive, we mean that the set of possible outcomes represents the entire range of possible outcomes. In other words, there is no probability that another non-represented outcome might occur. In our example, at this event node there are two, and only two, distinct outcomes that could occur: one outcome is that Vanessa's firm will offer Bill a summer job, and the other outcome is that Vanessa's firm will not offer Bill a summer job.

If Vanessa's firm were to make Bill a job offer, then Bill would subsequently have to decide to accept or to reject the firm's job offer. In this case, and if Bill were to accept the firm's job offer, then his summer job problem would be resolved. If Bill were to instead reject their offer, then Bill would then have to search for summer employment through the school's corporate summer recruiting program. The decision tree shown in Figure 1.3 represents these further possible eventualities, where the additional decision

FIGURE 1.1
Representation of a decision node.

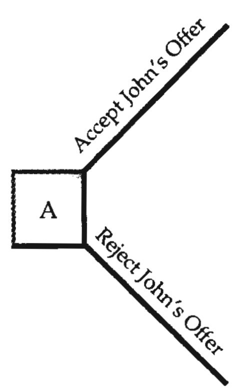

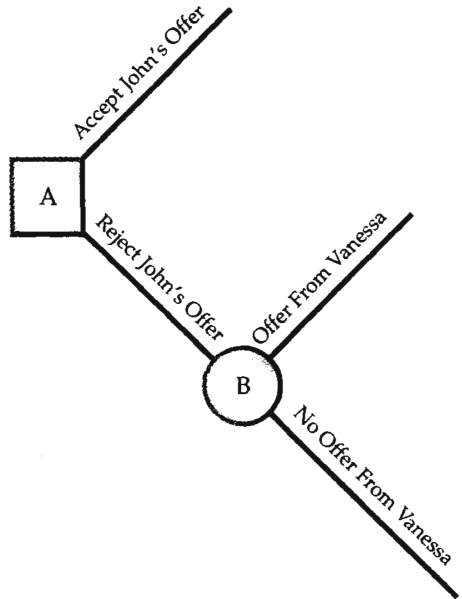

FIGURE 1.2
Representation of an event node.

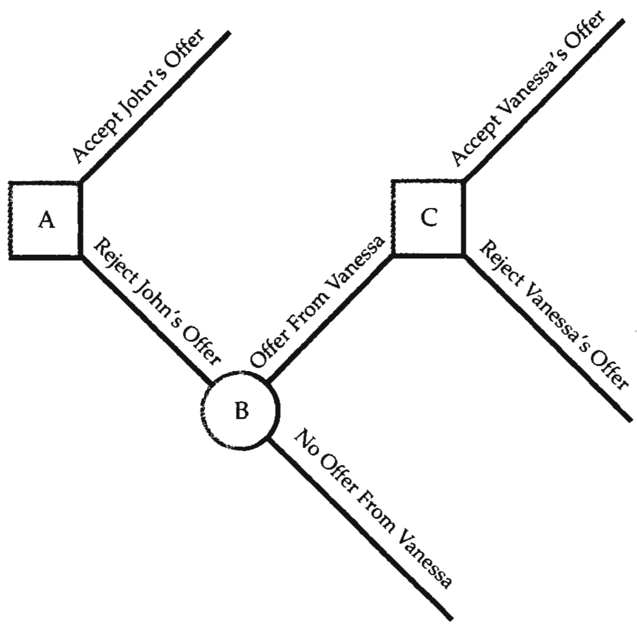

FIGURE 1.3
Further representation of the decision tree.

node C represents the decision that Bill would face if he were to receive a summer job offer from Vanessa's firm.

Assigning Probabilities

Another aspect of constructing a decision tree is the assignment or determination of the probability, i.e., the likelihood, that each of the various uncertain outcomes will transpire.

Let us suppose that Bill has visited the career services center at Sloan and has gathered some summary data on summer salaries received by the previous class of

MBA students. Based on salaries paid to Sloan students who worked in the Sales and Trading Departments at Vanessa's firm the previous summer, Bill has estimated that Vanessa's firm would make offers of $14,000 for twelve weeks' work to summer MBA students this coming summer.

Let us also suppose that we have gathered some data on the salary range for all summer jobs that went to Sloan students last year, and that this data is conveniently summarized in Table 1.1. The table shows five different summer salaries (based on weekly salary) and the associated percentages of students who received this salary. (The school did not have salary information for 5% of the students. In order to be conservative, we assign these students a summer salary of $0.)

Suppose further that our own intuition has suggested that Table 1.1 is a good approximation of the likelihood that Bill would receive the indicated salaries if he were to participate in the school's corporate summer recruiting. That is, we estimate that there is roughly a 5% likelihood that Bill would be able to procure a summer job with a salary of $21,600, and that there is roughly a 25% likelihood that Bill would be able

TABLE 1.1
Distribution of summer salaries.

Weekly Salary	*Total Summer Pay (based on 12 weeks)*	*Percentage of Students Who Received This Salary*
$1,800	$21,600	5%
$1,400	$16,800	25%
$1,000	$12,000	40%
$500	$6,000	25%
$0	$0	5%

FIGURE 1.4
Further representation of the decision tree.

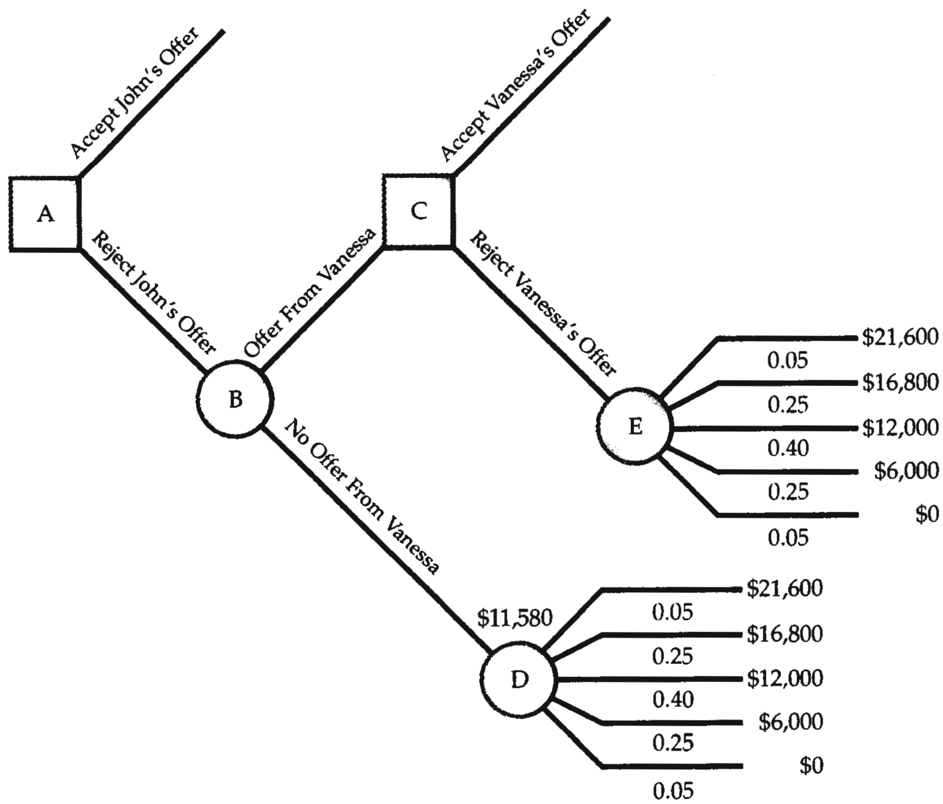

to procure a summer job with a salary of $16,800, etc. The now-expanded decision tree for the problem is shown in Figure 1.4, which includes event nodes D and E for the eventuality that Bill would participate in corporate summer recruiting if he were not to receive a job offer from Vanessa's firm, or if he were to reject an offer from Vanessa's firm. It is customary to write the probabilities of the various outcomes underneath their respective outcome branches, as is done in the figure.

Finally, let us estimate the likelihood that Vanessa's firm will offer Bill a job. Without much thought, we might assign this outcome a probability of 0.50, that is, there is a 50% likelihood that Vanessa's firm would offer Bill a summer job. On further reflection, we know that Vanessa was very impressed with Bill, and she sounded certain that she wanted to hire him. However, very many of Bill's classmates are also very talented (like him), and Bill has heard that competition for investment banking jobs is in fact very intense. Based on these musings, let us assign the probability that Bill would receive a summer job offer from Vanessa's firm to be 0.60. Therefore, the likelihood that Bill would not receive a job offer from Vanessa's firm would then be 0.40. These two numbers are shown in the decision tree in Figure 1.5.

Valuing the Final Branches

The next step in the decision analysis modeling methodology is to assign numerical values to the outcomes associated with the "final" branches of the decision tree, based on the decision criterion that has been adopted. As discussed earlier, Bill's decision criterion is his salary. Therefore, we assign the salary implication of each final branch and write this down to the right of the final branch, as shown in Figure 1.6.

FIGURE 1.5
Further representation of the decision tree.

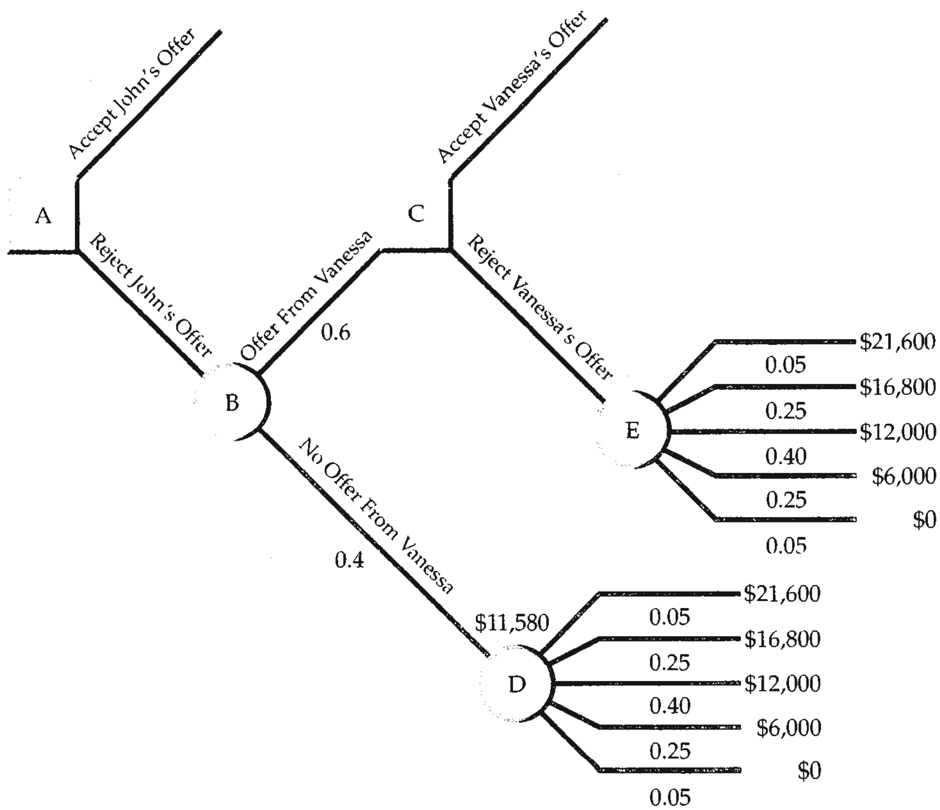

FIGURE 1.6
The completed decision tree.

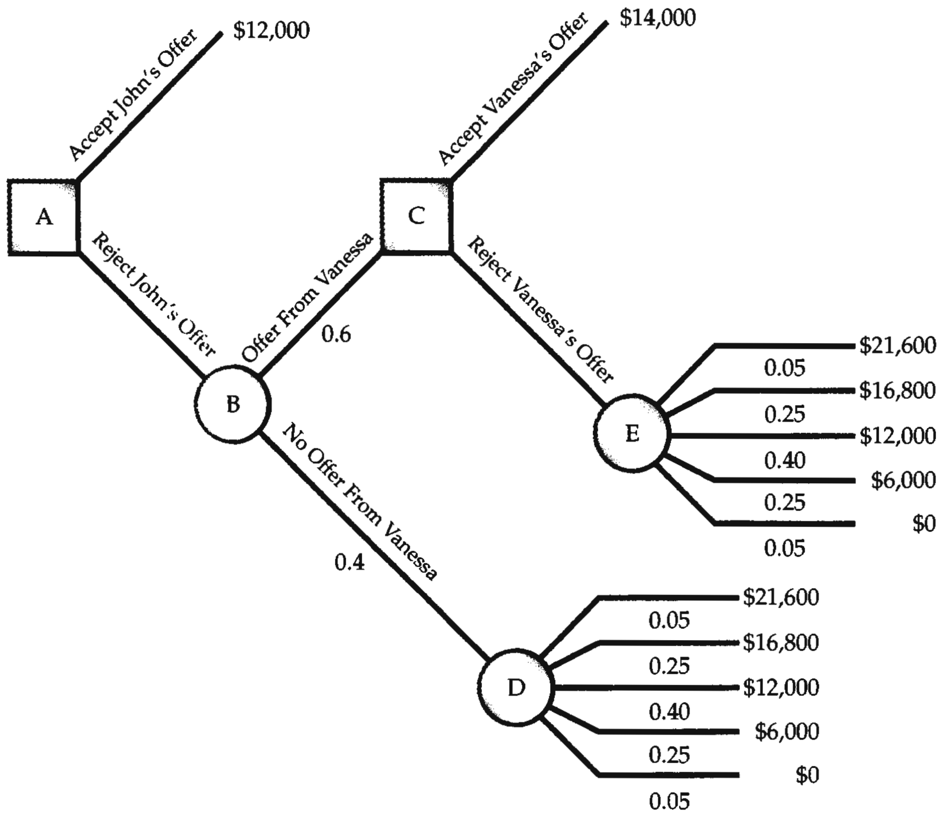

Fundamental Aspects of Decision Trees

Let us pause and look again at the decision tree as shown in Figure 1.6. Notice that time in the decision tree flows from left to right, and the placement of the decision nodes and the event nodes is logically consistent with the way events will play out in reality. Any event or decision that must logically precede certain other events and decisions is appropriately placed in the tree to reflect this logical dependence.

The tree has two decision nodes, namely node A and node C. Node A represents the decision Bill must make soon: whether to accept or reject John's offer. Node C represents the decision Bill might have to make in late November: whether to accept or reject Vanessa's offer. The branches emanating from each decision node represent all of the possible decisions under consideration at that point in time under the appropriate circumstances.

There are three event nodes in the tree, namely nodes B, D, and E. Node B represents the uncertain event of whether or not Bill will receive a job offer from Vanessa's firm. Node D (and also Node E) represents the uncertain events governing the school's corporate summer recruiting salaries. The branches emanating from each event node represent a set of mutually exclusive and collectively exhaustive outcomes from the event node. Furthermore, the sum of the probabilities of each outcome branch emanating from a given event node must sum to one. (This is because the set of possible outcomes is collectively exhaustive.)

These important characteristics of a decision tree are summarized as follows:

Key Characteristics of a Decision Tree

1. Time in a decision tree flows from left to right, and the placement of the decision nodes and the event nodes is logically consistent with the way events will play out in reality. Any event or decision that must logically precede certain other events and decisions is appropriately placed in the tree to reflect this logical dependence.
2. The branches emanating from each decision node represent all of the possible decisions under consideration at that point in time under the appropriate circumstances.
3. The branches emanating from each event node represent a set of mutually exclusive and collectively exhaustive outcomes of the event node.
4. The sum of the probabilities of each outcome branch emanating from a given event node must sum to one.
5. Each and every "final" branch of the decision tree has a numerical value associated with it. This numerical value usually represents some measure of monetary value, such as salary, revenue, cost, etc.

Notice that in the case of Bill's summer job decision, all of the numerical values associated with the final branches in the decision tree are dollar figures of salaries, which are inherently objective measures to work with. However, Bill might also wish to consider subjective measures in making his decision. We have conveniently assumed for simplicity that the intangible benefits of his summer job options, such as opportunities to learn, networking, resumé-building, etc., would be the same at either his former employer, Vanessa's firm, or in any job offer he might receive through the school's corporate summer recruiting. In reality, these subjective measures would not be the same for all of Bill's possible options. Of course, another important subjective factor, which Bill might also consider, is the value of the time he would have to spend in corporate summer recruiting. Although we will analyze the decision tree ignoring all of these subjective measures, the value of Bill's time should at least be considered when reviewing the conclusions afterward.

Solution of Bill's Problem by Folding Back the Decision Tree

If Bill's choice were simply between accepting a job offer of $12,000 or accepting a different job offer of $14,000, then his decision would be easy: he would take the higher salary offer. However, in the presence of uncertainty, it is not necessarily obvious how Bill might proceed.

Suppose, for example, that Bill were to reject John's offer, and that in mid-November he were to receive an offer of $14,000 from Vanessa's firm. He would then be at node C of the decision tree. How would he go about deciding between obtaining a summer salary of $14,000 with certainty, and the distribution of possible salaries he might obtain (with varying degrees of uncertainty) from participating in the school's corporate summer recruiting? The criterion that most decision-makers feel is most appropriate to use in this setting is to convert the distribution of possible salaries to a single numerical value using the expected monetary value (EMV) of the possible outcomes:

> The **expected monetary value** or EMV of an uncertain event is the weighted average of all possible numerical outcomes, with the probabilities of each of the possible outcomes used as the weights.

Therefore, for example, the EMV of participating in corporate summer recruiting is computed as follows:

$$\text{EMV} = 0.05 \times \$21{,}600 + 0.25 \times \$16{,}800 + 0.40 \times \$12{,}000 + 0.25 \times \$6{,}000 + 0.05 \times \$0 = \$11{,}580.$$

The EMV of a certain event is defined to be the monetary value of the event. For example, suppose that Bill were to receive a job offer from Vanessa's firm, and that he were to accept the job offer. Then the EMV of this choice would simply be $14,000.

Notice that the EMV of the choice to participate in corporate recruiting is $11,580, which is less than $14,000 (the EMV of accepting the offer from Vanessa's firm), and so under the EMV criterion, Bill would prefer the job offer from Vanessa's firm to the option of participating in corporate summer recruiting.

The EMV is one way to convert a group of possible outcomes with monetary values and probabilities to a single number that weighs each possible outcome by its probability. The EMV represents an "averaging" approach to uncertainty. It is quite intuitive, and is quite appropriate for a wide variety of decision problems under uncertainty. (However, there are cases where it is not necessarily the best method for converting a group of possible outcomes to a single number. In Section 1.5, we discuss several aspects of the EMV criterion further.)

Using the EMV criterion, we can now "solve" the decision tree. We do so by evaluating every event node using the EMV of the event node, and evaluating every decision node by choosing that decision which has the best EMV. This is accomplished by starting at the final branches of the tree, and then working "backwards" to the starting node of the decision tree. For this reason, the process of solving the decision tree is called **folding back the decision tree.** It is also occasionally referred to as **backwards induction.** This process is illustrated in the following discussion.

Starting from any one of the "final" nodes of the decision tree, we proceed backwards. As we have already seen, the EMV of node E is $11,580. It is customary to write the EMV of an event node above the node, as is shown in Figure 1.7. Similarly, the EMV of node D is also $11,580, which we write above node D. This is also displayed in Figure 1.7.

We next examine decision node C, which corresponds to the event that Bill receives a job offer from Vanessa's firm. At this decision node, there are two choices. The first choice is for Bill to accept the offer from Vanessa's firm, which has an EMV of $14,000. The second choice is to reject the offer, and instead to participate in corporate summer recruiting, which has an EMV of $11,580. As the EMV of $11,580 is less than the EMV of $14,000, it is better to choose the branch corresponding to accepting Vanessa's offer. Pictorially, we show this by crossing off the inferior choice by drawing two lines through the branch, and by writing the monetary value of the best choice above the decision node. This is shown in Figure 1.7 as well.

We continue by evaluating event node B, which is the event node corresponding to the event where Vanessa's firm either will or will not offer Bill a summer job. The methodology we use is the same as evaluating the salary distributions from participating in corporate summer recruiting. We compute the EMV of the node by computing the

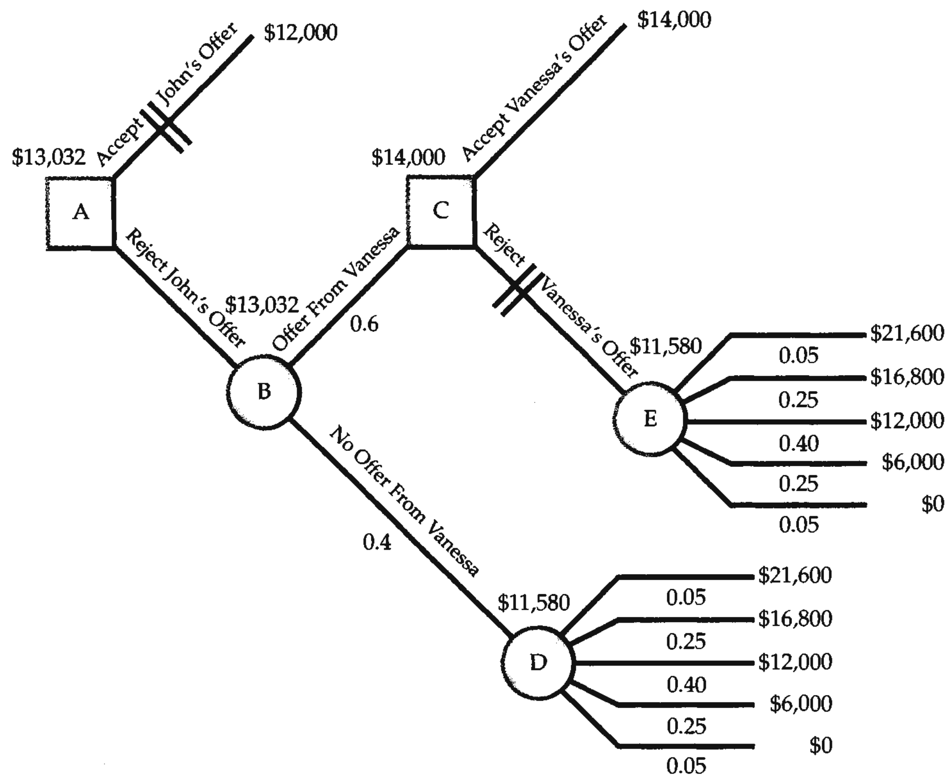

FIGURE 1.7 Solution of the decision tree.

weighted average of the EMVs of each of the outcomes, weighted by the probabilities corresponding to each of the outcomes. In this case, this means multiplying the probability of an offer (0.60) by the $14,000 value of decision node C, then multiplying the probability of not receiving an offer from Vanessa's firm (0.40) times the EMV of node D, which is $11,580, and then adding the two quantities. The calculations are:

$$\text{EMV} = 0.60 \times \$14{,}000 + 0.40 \times \$11{,}580 = \$13{,}032.$$

This number is then placed above the node, as shown in Figure 1.7.

The last step in solving the decision tree is to evaluate the remaining node, which is the first node of the tree. This is a decision node, and its evaluation is accomplished by comparing the better of the two EMV values of the branches that emanate from it. The upper branch, which corresponds to accepting John's offer, has an EMV of $12,000. The lower branch, which corresponds to rejecting John's offer and proceeding onward, has an EMV of $13,032. As this latter value is the highest, we cross off the branch corresponding to accepting John's offer, and place the EMV value of $13,032 above the initial node. The completed solution of the decision tree is shown in Figure 1.7.

Let us now look again at the solved decision tree and examine the "optimal decision strategy" under uncertainty. According to the solved tree, Bill should not accept John's job offer, i.e., he should reject John's job offer. This is shown at the first decision node. Then, if Bill receives a job offer from Vanessa's firm, he should accept this offer. This is shown at the second decision node. Of course, if he does not receive a job offer from Vanessa's firm, he would then participate in the school's corporate summer recruiting program. The EMV of John's optimal decision strategy is $13,032.

Summarizing, Bill's optimal decision strategy can be stated as follows:

Bill's Optimal Decision Strategy:

- Bill should reject John's offer in October.
- If Vanessa's firm offers him a job, he should accept it. If Vanessa's firm does not offer him a summer job, he should participate in the school's corporate summer recruiting.
- The EMV of this strategy is $13,032.

Note that the output from constructing and solving the decision tree is a very concrete plan of action, which states what decisions should be made under each possible uncertain outcome that might prevail.

The procedure for solving a decision tree can be formally stated as follows:

Procedure for Solving a Decision Tree

1. Start with the final branches of the decision tree, and evaluate each event node and each decision node, as follows:
   - For an event node, compute the EMV of the node by computing the weighted average of the EMV of each branch weighted by its probability. Write this EMV number above the event node.
   - For a decision node, compute the EMV of the node by choosing that branch emanating from the node with the best EMV value. Write this EMV number above the decision node, and cross off those branches emanating from the node with inferior EMV values by drawing a double line through them.
2. The decision tree is solved when all nodes have been evaluated.
3. The EMV of the optimal decision strategy is the EMV computed for the starting branch of the tree.

As we mentioned already, the process of solving the decision tree in this manner is called **folding back the decision tree.** It is also sometimes referred to as **backwards induction.**

Sensitivity Analysis of the Optimal Decision

If this were an actual business decision, it would be naive to adopt the optimal decision strategy derived above, without a critical evaluation of the impact of the key data assumptions that were made in the development of the decision tree model. For example, consider the following data-related issues that we might want to address:

- **Issue 1: The probability that Vanessa's firm would offer Bill a summer job.** We have subjectively assumed that the probability that Vanessa's firm would offer Bill a summer job to be 0.60. It would be wise to test how changes in this probability might affect the optimal decision strategy.
- **Issue 2: The cost of Bill's time and effort in participating in the school's corporate summer recruiting.** We have implicitly assumed that the cost of Bill's time and effort in participating in the school's corporate summer recruiting would be zero. It would be wise to test how high the implicit cost of participating

in corporate summer recruiting would have to be before the optimal decision strategy would change.

- **Issue 3: The distribution of summer salaries that Bill could expect to receive.** We have assumed that the distribution of summer salaries that Bill could expect to receive is given by the numbers in Table 1.1. It would be wise to test how changes in this distribution of salaries might affect the optimal decision strategy.

The process of testing and evaluating how the solution to a decision tree behaves in the presence of changes in the data is referred to as **sensitivity analysis.** The process of performing sensitivity analysis is as much an art as it is a science. It usually involves choosing several key data values and then testing how the solution of the decision tree model changes as each of these data values are modified, one at a time. Such a process is very important for understanding what data are driving the optimal decision strategy and how the decision tree model behaves under changes in key data values. The exercise of performing sensitivity analysis is important in order to gain confidence in the validity of the model and is necessary before one bases one's decisions on the output from a decision tree model. We illustrate next the art of sensitivity analysis by performing the three data changes suggested previously.

Note that in order to evaluate how the optimal decision strategy behaves as a function of changes in the data assumptions, we will have to solve and re-solve the decision tree model many times, each time with slightly different values of certain data. Obviously, one way to do this would be to re-draw the tree each time and perform all of the necessary arithmetic computations by hand each time. This approach is obviously very tedious and repetitive, and in fact we can do this much more conveniently with the help of a computer spreadsheet. We can represent the decision tree problem and its solution very conveniently on a spreadsheet, illustrated in Figure 1.8 and explained in the following discussion.

FIGURE 1.8 Spreadsheet representation of Bill Sampras' summer job problem.

	Spreadsheet Representation of Bill Sampras' Decision Problem		
Data			
Value of John's offer	$12,000		
Value of Vanessa's offer	$14,000		
Probability of offer from Vanessa's firm	0.60		
Cost of participating in Recruiting	$0		
	Distribution of Salaries from Recruiting		
	Weekly Salary	**Total Summer Pay**	**Percentage of Students**
		(based on 12 weeks)	**who Received this Salary**
	$1,800	$21,600	5%
	$1,400	$16,800	25%
	$1,000	$12,000	40%
	$500	$6,000	25%
	$0	$0	5%
	EMV of Nodes		
	Nodes	**EMV**	
	A	$13,032	
	B	$13,032	
	C	$14,000	
	D	$11,580	
	E	$11,580	

In the spreadsheet representation of Figure 1.8, the data for the decision tree is given in the upper part of the spreadsheet, and the "solution" of the spreadsheet is computed in the lower part in the "EMV of Nodes" table. The computation of the EMV of each node is performed automatically as a function of the data. For example, we know that node E of the spreadsheet has its EMV computed as follows:

$$\text{EMV of node E} = 0.05 \times \$21{,}600 + 0.25 \times \$16{,}800 + 0.40 \times \$12{,}000 + 0.25 \times \$6{,}000 + 0.05 \times \$0 = \$11{,}580.$$

The EMV of node D is computed in an identical manner. As presented earlier, the EMV of node C is the maximum of the EMV of node E and the value of an offer from Vanessa's firm, and is computed as

$$\text{EMV of node C} = \text{MAX}\{\text{EMV of node E}, \$14{,}000\}.$$

Similarly, the EMV of nodes B and A are given by

$$\text{EMV of node B} = (0.60) \times (\text{EMV of node C}) + (1 - 0.60) \times (\text{EMV of node D})$$

and

$$\text{EMV of node A} = \text{MAX}\{\text{EMV of node B}, \$12{,}000\}.$$

All of these formulas can be conveniently represented in a spreadsheet, and such a spreadsheet is shown in Figure 1.8. Note that the EMV numbers for all of the nodes in the spreadsheet correspond exactly to those computed "by hand" in the solution of the decision tree shown in Figure 1.7.

We now show how the spreadsheet representation of the decision tree can be used to study how the optimal decision strategy changes relative to the three key data issues discussed above at the start of this subsection. To begin, consider the first issue, which concerns the sensitivity of the optimal decision strategy to the value of the probability that Vanessa's firm will offer Bill a summer job. Denote this probability by p, i.e.,

$$p = \text{probability that Vanessa's firm will offer Bill a summer job.}$$

If we test a variety of values of p in the spreadsheet representation of the decision tree, we will find that the optimal decision strategy (which is to reject John's job offer, and to accept a job offer from Vanessa's firm if it is offered) remains the same for all values of p greater than or equal to $p = 0.174$. Figure 1.9 shows the output of the spreadsheet when $p = 0.18$, for example, and notice that the EMV of node B is \$12,016, which is just barely above the threshold value of \$12,000. For values of p at or below $p = 0.17$, the EMV of node B becomes less than \$12,000, which results in a new optimal decision strategy of accepting John's job offer. We can conclude the following:

- As long as the probability of Vanessa's firm offering Bill a job is 0.18 or larger, then the optimal decision strategy will still be to reject John's offer and to accept a summer job with Vanessa's firm if they offer it to him.

This is reassuring, as it is reasonable for Bill to be very confident that the probability of Vanessa's firm offering him a summer job is surely greater than 0.18.

We next use the spreadsheet representation of the decision tree to study the second data assumption issue, which concerns the sensitivity of the optimal decision strategy to the implicit cost to Bill (in terms of his time) of participating in the school's corporate summer recruiting program. Denote this cost by c, i.e.,

$$c = \text{implicit cost to Bill of participating in the school's corporate summer recruiting program.}$$

FIGURE 1.9 Output of the spreadsheet of Bill Sampras' summer job problem when the probability that Vanessa's firm will make Bill an offer is 0.18.

	Spreadsheet Representation of Bill Sampras' Decision Problem		
Data			
Value of John's offer	$12,000		
Value of Vanessa's offer	$14,000		
Probability of offer from Vanessa's firm	0.18		
Cost of participating in Recruiting	$0		
	Distribution of Salaries from Recruiting		
	Weekly Salary	**Total Summer Pay**	**Percentage of Students**
		(based on 12 weeks)	**who Received this Salary**
	$1,800	$21,600	5%
	$1,400	$16,800	25%
	$1,000	$12,000	40%
	$500	$6,000	25%
	$0	$0	5%
	EMV of Nodes		
	Nodes	**EMV**	
	A	$12,016	
	B	$12,016	
	C	$14,000	
	D	$11,580	
	E	$11,580	

If we test a variety of values of c in the spreadsheet representation of the decision tree, we will notice that the current optimal decision strategy (which is to reject John's job offer, and to accept a job offer from Vanessa's firm if it is offered) remains the same for all values of c less than $c = \$2,578$. Figure 1.10 shows the output of the spreadsheet when $c = \$2,578$. For values of c above $c = \$2,578$, the EMV of node B becomes less than $12,000, which results in a new optimal decision strategy of accepting John's job offer. We can conclude the following:

- As long as the implicit cost to Bill of participating in summer recruiting is less than $2,578, then the optimal decision strategy will still be to reject John's offer and to accept a summer job with Vanessa's firm if they offer it to him.

This is also reassuring, as it is reasonable to estimate that the implicit cost to Bill of participating in the school's corporate summer recruiting program is much less than $2,578.

We next use the spreadsheet representation of the decision tree to study the third data issue, which concerns the sensitivity of the optimal decision strategy to the distribution of possible summer job salaries from participating in corporate recruiting. Recall that Table 1.1 contains the data for the salaries Bill might possibly realize by participating in corporate summer recruiting. Let us explore the consequences of changing all of the possible salary offers of Table 1.1 by an amount S. That is, we will explore modifying Bill's possible summer salaries by an amount S. If we test a variety of values of S in the spreadsheet representation of the model, we will notice that the current optimal decision strategy remains optimal for all values of S less than $S = \$2,419$. Figure 1.11 shows the output of the spreadsheet when $S = \$2,419$. For values of S above $S = \$2,420$, the EMV of node E will become greater than or equal to $14,000, and consequently Bill's optimal decision strategy will change: he would reject an offer from Vanessa's firm if it materialized, and instead would participate in the school's corporate summer recruiting program. We can conclude:

FIGURE 1.10 Output of the spreadsheet of Bill Sampras' summer job problem if the cost of Bill's time spent participating in corporate summer recruiting is $2,578.

	Spreadsheet Representation of Bill Sampras' Decision Problem		
Data			
Value of John's offer	$12,000		
Value of Vanessa's offer	$14,000		
Probability of offer from Vanessa's firm	0.60		
Cost of participating in Recruiting	$2,578		
	Distribution of Salaries from Recruiting		
	Weekly Salary	**Total Summer Pay**	**Percentage of Students**
		(based on 12 weeks)	**who Received this Salary**
	$1,800	$21,600	5%
	$1,400	$16,800	25%
	$1,000	$12,000	40%
	$500	$6,000	25%
	$0	$0	5%
	EMV of Nodes		
	Nodes	**EMV**	
	A	$12,001	
	B	$12,001	
	C	$14,000	
	D	$9,002	
	E	$9,002	

FIGURE 1.11 Output of the spreadsheet of Bill Sampras' summer job problem if summer salaries from recruiting were $2,419 higher.

	Spreadsheet Representation of Bill Sampras' Decision Problem		
Data			
Value of John's offer	$12,000		
Value of Vanessa's offer	$14,000		
Probability of offer from Vanessa's firm	0.60		
Cost of participating in Recruiting	$0		
	Distribution of Salaries from Recruiting		
	Weekly Salary	**Total Summer Pay**	**Percentage of Students**
		(based on 12 weeks)	**who Received this Salary**
	$1,800	$24,019	5%
	$1,400	$19,219	25%
	$1,000	$14,419	40%
	$500	$8,419	25%
	$0	$2,419	5%
	EMV of Nodes		
	Nodes	**EMV**	
	A	$14,000	
	B	$14,000	
	C	$14,000	
	D	$13,999	
	E	$13,999	

- In order for Bill's optimal decision strategy to change, all of the possible summer corporate recruiting salaries of Table 1.1 would have to increase by more than $2,419.

This is also reassuring, as it is reasonable to anticipate that summer salaries from corporate summer recruiting in general would not be $2,419 higher this coming summer than they were last summer.

We can summarize our findings as follows:

- For all three of the data issues that we have explored (the probability p of Vanessa's firm offering Bill a summer job, the implicit cost c of participating in corporate summer recruiting, and an increase S in all possible salary values from corporate summer recruiting), we have found that the optimal decision strategy does not change unless these quantities take on unreasonable values. Therefore, it is safe to proceed with confidence in recommending to Bill Sampras that he adopt the optimal decision strategy found in the solution to the decision tree model. Namely, he should reject John's job offer, and he should accept a job offer from Vanessa's firm if such an offer is made.

In some applications of decision analysis, the decision-maker might discover that the optimal decision strategy is very sensitive to a key data value. If this happens, it is then obviously important to spend some effort to determine the most reasonable value of that data. For instance, in the decision tree we have constructed, suppose that in fact the optimal decision was very sensitive to the probability p that Vanessa's firm would offer Bill a summer job. We might then want to gather data on how many offers Vanessa's firm made to Sloan students in previous years, and in particular we might want to look at how students with Bill's general profile fared when they applied for jobs with Vanessa's firm. This information could then be used to develop a more exact estimate of the probability p that Bill would receive a job offer from Vanessa's firm.

Note that in this sensitivity analysis exercise, we have only changed one data value at a time. In some problem instances, the decision-maker might want to test how the model behaves under simultaneous changes in more than one data value. This is a bit more difficult to analyze, of course.

1.2 SUMMARY OF THE GENERAL METHOD OF DECISION ANALYSIS

The example of Bill Sampras' summer job decision problem illustrates the format of the general method of **decision analysis** to systematically analyze a decision problem. The format of this general method is as follows:

Principal Steps of Decision Analysis

1. Structure the decision problem. List all of the decisions that have to be made. List all of the uncertain events in the problem and all of their possible outcomes.
2. Construct the basic decision tree by placing the decision nodes and the event nodes in their chronological and logically consistent order.
3. Determine the probability of each of the possible outcomes of each of the uncertain events. Write these probabilities on the decision tree.
4. Determine the numerical values of each of the final branches of the decision tree. Write these numerical values on the decision tree.
5. Solve the decision tree using the folding-back procedure:

(a) Start with the final branches of the decision tree, and evaluate each event node and each decision node, as follows:

- For an event node, compute the EMV of the node by computing the weighted average of the EMV of each branch weighted by its probability. Write this EMV number above the event node.
- For a decision node, compute the EMV of the node by choosing that branch emanating from the node with the best EMV value. Write this EMV number above the decision node and cross off those branches emanating from the node with inferior EMV values by drawing a double line through them.

(b) The decision tree is solved when all nodes have been evaluated.

(c) The EMV of the optimal decision strategy is the EMV computed for the starting branch of the tree.

6. Perform sensitivity analysis on all key data values. For each data value for which the decision-maker lacks confidence, test how the optimal decision strategy will change relative to a change in the data value, one data value at a time.

As mentioned earlier, the solution of the decision tree and the sensitivity analysis procedure typically involve a number of mechanical arithmetic calculations. Unless the decision tree is small, it might be wise to construct a spreadsheet version of the decision tree in order to perform these calculations automatically and quickly. (And of course, a spreadsheet version of the model will also eliminate the likelihood of making arithmetical errors!)

1.3 ANOTHER DECISION TREE MODEL AND ITS ANALYSIS

In this section, we continue to illustrate the methodology of decision analysis by considering a strategic development decision problem encountered by a new company called Bio-Imaging, Incorporated.

BIO-IMAGING DEVELOPMENT STRATEGIES

In 2004, the company Bio-Imaging, Incorporated was formed by James Bates, Scott Tillman, and Michael Ford, in order to develop, produce, and market a new and potentially extremely beneficial tool in medical diagnosis. Scott Tillman and James Bates were each recent graduates from Massachusetts Institute of Technology (MIT), and Michael Ford was a professor of neurology at Massachusetts General Hospital (MGH). As part of his graduate studies at MIT, Scott had developed a new technique and a software package to process MRI (magnetic resonance imaging) scans of brains of patients using a personal computer. The software, using state of the art computer graphics, would construct a three-dimensional picture of a patient's brain and could be used to find the exact location of a brain lesion or a brain tumor, estimate its volume and shape, and even locate the centers in the brain that would be affected by the tumor. Scott's work was an extension of earlier two-dimensional image processing