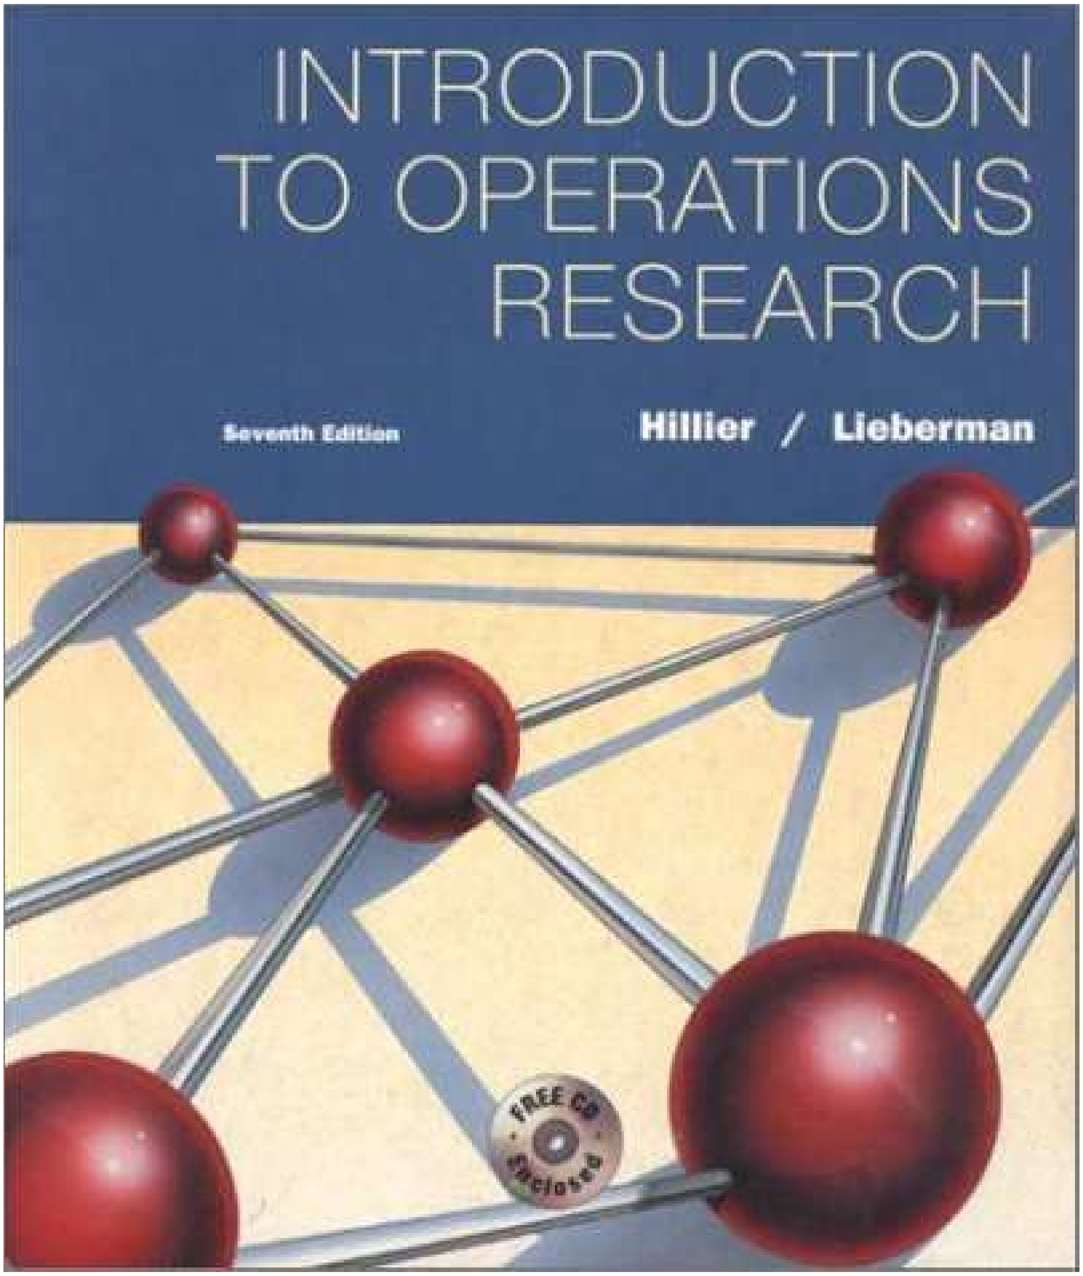

INTRODUCTION TO OPERATIONS RESEARCH

Seventh Edition

Hillier / Lieberman

INTRODUCTION TO OPERATIONS RESEARCH

Seventh Edition

FREDERICK S. HILLIER,
Stanford University

GERALD J. LIEBERMAN,
Late of Stanford University

Cases developed by Karl Schmedders and Molly Stephens

Tutorial software developed by Mark Hillier and Michael O'Sullivan

McGraw Hill

Boston Burr Ridge, IL Dubuque, IA Madison, WI New York
San Francisco St. Louis Bangkok Bogotá Caracas Lisbon London
Madrid Mexico City Milan New Delhi Seoul Singapore Sydney
Taipei Toronto

15

Decision Analysis

The previous chapters have focused mainly on decision making when the consequences of alternative decisions are known with a reasonable degree of certainty. This decision-making environment enabled formulating helpful mathematical models (linear programming, integer programming, nonlinear programming, etc.) with objective functions that specify the estimated consequences of any combination of decisions. Although these consequences usually cannot be predicted with complete certainty, they could at least be estimated with enough accuracy to justify using such models (along with sensitivity analysis, etc.).

However, decisions often must be made in environments that are much more fraught with uncertainty. Here are a few examples.

1. A manufacturer introducing a new product into the marketplace. What will be the reaction of potential customers? How much should be produced? Should the product be test marketed in a small region before deciding upon full distribution? How much advertising is needed to launch the product successfully?
2. A financial firm investing in securities. Which are the market sectors and individual securities with the best prospects? Where is the economy headed? How about interest rates? How should these factors affect the investment decisions?
3. A government contractor bidding on a new contract. What will be the actual costs of the project? Which other companies might be bidding? What are their likely bids?
4. An agricultural firm selecting the mix of crops and livestock for the upcoming season. What will be the weather conditions? Where are prices headed? What will costs be?
5. An oil company deciding whether to drill for oil in a particular location. How likely is oil there? How much? How deep will they need to drill? Should geologists investigate the site further before drilling?

These are the kinds of decision making in the face of great uncertainty that *decision analysis* is designed to address. Decision analysis provides a framework and methodology for rational decision making when the outcomes are uncertain.

The preceding chapter describes how game theory also can be used for certain kinds of decision making in the face of uncertainty. There are some similarities in the approaches used by game theory and decision analysis. However, there also are differences because they are designed for different kinds of applications. We will describe these similarities and differences in Sec. 15.2.

Frequently, one question to be addressed with decision analysis is whether to make the needed decision immediately or to first do some *testing* (at some expense) to reduce the level of uncertainty about the outcome of the decision. For example, the testing might be field testing of a proposed new product to test consumer reaction before making a decision on whether to proceed with full-scale production and marketing of the product. This testing is referred to as performing *experimentation.* Therefore, decision analysis divides decision making between the cases of *without experimentation* and *with experimentation.*

The first section introduces a prototype example that will be carried throughout the chapter for illustrative purposes. Sections 15.2 and 15.3 then present the basic principles of *decision making without experimentation* and *decision making with experimentation.* We next describe *decision trees,* a useful tool for depicting and analyzing the decision process when a series of decisions needs to be made. Section 15.5 introduces *utility theory,* which provides a way of calibrating the possible outcomes of the decision to reflect the true value of these outcomes to the decision maker. We then conclude the chapter by discussing the practical application of decision analysis and summarizing a variety of applications that have been very beneficial to the organizations involved.

15.1 A PROTOTYPE EXAMPLE

The GOFERBROKE COMPANY owns a tract of land that may contain oil. A consulting geologist has reported to management that she believes there is 1 chance in 4 of oil.

Because of this prospect, another oil company has offered to purchase the land for \$90,000. However, Goferbroke is considering holding the land in order to drill for oil itself. The cost of drilling is \$100,000. If oil is found, the resulting expected revenue will be \$800,000, so the company's expected profit (after deducting the cost of drilling) will be \$700,000. A loss of \$100,000 (the drilling cost) will be incurred if the land is dry (no oil).

Table 15.1 summarizes these data. Section 15.2 discusses how to approach the decision of whether to drill or sell based just on these data. (We will refer to this as the *first* Goferbroke Co. problem.)

However, before deciding whether to drill or sell, another option is to conduct a detailed seismic survey of the land to obtain a better estimate of the probability of finding oil. Section 15.3 discusses this case of *decision making with experimentation,* at which point the necessary additional data will be provided.

This company is operating without much capital, so a loss of \$100,000 would be quite serious. In Sec. 15.5, we describe how to refine the evaluation of the consequences of the various possible outcomes.

TABLE 15.1 Prospective profits for the Goferbroke Company

Status of Land / Alternative	Payoff	
	Oil	**Dry**
Drill for oil	\$700,000	−\$100,000
Sell the land	\$ 90,000	\$ 90,000
Chance of status	1 in 4	3 in 4

15.2 DECISION MAKING WITHOUT EXPERIMENTATION

Before seeking a solution to the first Goferbroke Co. problem, we will formulate a general framework for decision making.

In general terms, the decision maker must choose an **action** from a set of possible actions. The set contains all the *feasible alternatives* under consideration for how to proceed with the problem of concern.

This choice of an action must be made in the face of uncertainty, because the outcome will be affected by random factors that are outside the control of the decision maker. These random factors determine what situation will be found at the time that the action is executed. Each of these possible situations is referred to as a possible **state of nature.**

For each combination of an action and a state of nature, the decision maker knows what the resulting payoff would be. The **payoff** is a quantitative measure of the value to the decision maker of the consequences of the outcome. For example, the payoff frequently is represented by the *net monetary gain* (profit), although other measures also can be used (as described in Sec. 15.5). If the consequences of the outcome do not become completely certain even when the state of nature is given, then the payoff becomes an *expected value* (in the statistical sense) of the measure of the consequences. A **payoff table** commonly is used to provide the payoff for each combination of an action and a state of nature.

If you previously studied game theory (Chap. 14), we should point out an interesting analogy between this decision analysis framework and the two-person, zero-sum games described in Chap. 14. The *decision maker* and *nature* can be viewed as the *two players* of such a game. The *alternative actions* and the possible *states of nature* can then be viewed as the available *strategies* for these respective players, where each combination of strategies results in some *payoff* to player 1 (the decision maker). From this viewpoint, the decision analysis framework can be summarized as follows:

1. The *decision maker* needs to choose one of the *alternative actions*.
2. *Nature* then would choose one of the possible *states of nature.*
3. Each combination of an action and state of nature would result in a *payoff,* which is given as one of the entries in a *payoff table.*
4. This payoff table should be used to find an *optimal action* for the decision maker according to an appropriate criterion.

Soon we will present three possibilities for this criterion, where the first one (the maximin payoff criterion) comes from game theory.

However, this analogy to two-person, zero-sum games breaks down in one important respect. In game theory, *both* players are assumed to be *rational* and choosing their strategies to *promote their own welfare.* This description still fits the decision maker, but certainly not nature. By contrast, nature now is a passive player that chooses its strategies (states of nature) in some random fashion. This change means that the game theory criterion for how to choose an optimal strategy (action) will not appeal to many decision makers in the current context.

One additional element needs to be added to the decision analysis framework. The decision maker generally will have some information that should be taken into account about the relative likelihood of the possible states of nature. Such information can usually be translated to a probability distribution, acting as though the state of nature is a random variable, in which case this distribution is referred to as a **prior distribution.** Prior distributions are often subjective in that they may depend upon the experience or intuition of an individual. The probabilities for the respective states of nature provided by the prior distribution are called **prior probabilities.**

Formulation of the Prototype Example in This Framework

As indicated in Table 15.1, the Goferbroke Co. has two possible actions under consideration: drill for oil or sell the land. The possible states of nature are that the land contains oil and that it does not, as designated in the column headings of Table 15.1 by *oil* and *dry.* Since the consulting geologist has estimated that there is 1 chance in 4 of oil (and so 3 chances in 4 of no oil), the prior probabilities of the two states of nature are 0.25 and 0.75, respectively. Therefore, with the payoff in units of thousands of dollars of profit, the payoff table can be obtained directly from Table 15.1, as shown in Table 15.2.

We will use this payoff table next to find the optimal action according to each of the three criteria described below. In each case, we will employ an Excel template provided in this chapter's Excel file for the criterion. These templates expedite entering a payoff table in a spreadsheet format and then applying the criteria.

The Maximin Payoff Criterion

If the decision maker's problem were to be viewed as a *game against nature,* then game theory would say to choose the action according to the *minimax criterion* (as described in Sec. 14.2). From the viewpoint of player 1 (the decision maker), this criterion is more aptly named the *maximin payoff criterion,* as summarized below.

> **Maximin payoff criterion:** For each possible action, find the *minimum payoff* over all possible states of nature. Next, find the *maximum* of these minimum payoffs. Choose the action whose minimum payoff gives this maximum.

The Excel template displayed in Fig. 15.1 shows the application of this criterion to the prototype example. Thus, since the minimum payoff for selling (90) is larger than that for drilling (−100), the former alternative (sell the land) will be chosen as the action to take.

The rationale for this criterion is that it provides the *best guarantee* of the payoff that will be obtained. Regardless of what the true state of nature turns out to be for the example, the payoff from selling the land cannot be less than 90, which provides the best available guarantee. Thus, this criterion takes the pessimistic viewpoint that, regardless of

TABLE 15.2 Payoff table for the decision analysis formulation of the Goferbroke Co. problem

	State of Nature	
Alternative	**Oil**	**Dry**
1. Drill for oil	700	−100
2. Sell the land	90	90
Prior probability	0.25	0.75

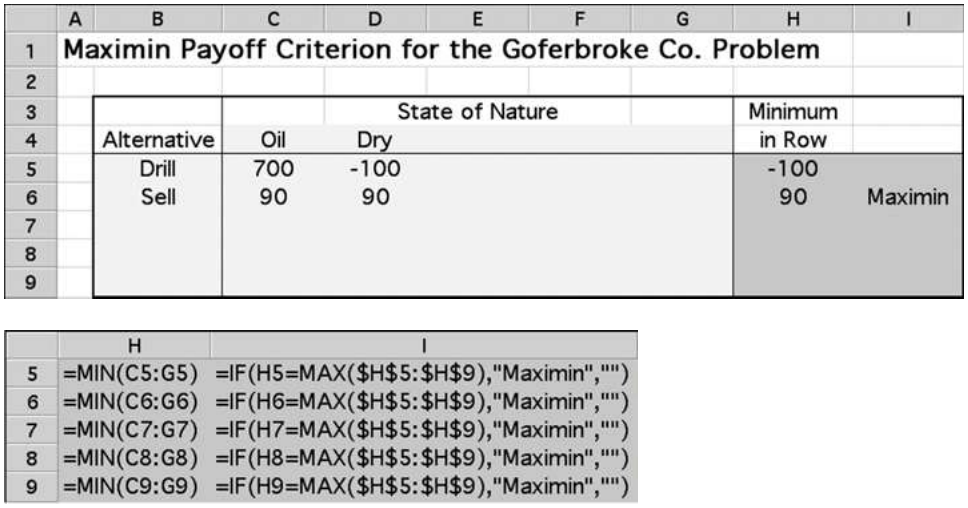

	A	B	C	D	E	F	G	H	I
1	Maximin Payoff Criterion for the Goferbroke Co. Problem								
2									
3				State of Nature				Minimum	
4		Alternative	Oil	Dry				in Row	
5		Drill	700	-100				-100	
6		Sell	90	90				90	Maximin
7									
8									
9									

	H	I
5	=MIN(C5:G5)	=IF(H5=MAX($H$5:$H$9),"Maximin","")
6	=MIN(C6:G6)	=IF(H6=MAX($H$5:$H$9),"Maximin","")
7	=MIN(C7:G7)	=IF(H7=MAX($H$5:$H$9),"Maximin","")
8	=MIN(C8:G8)	=IF(H8=MAX($H$5:$H$9),"Maximin","")
9	=MIN(C9:G9)	=IF(H9=MAX($H$5:$H$9),"Maximin","")

FIGURE 15.1
The application of the Excel template for the *maximin payoff criterion* to the first Goferbroke Co. problem.

which action is selected, the worst state of nature for that action is likely to occur, so we should choose the action which provides the best payoff with its worst state of nature.

This rationale is quite valid when one is competing against a rational and malevolent opponent. However, this criterion is not often used in games against nature because it is an extremely conservative criterion in this context. In effect, it assumes that nature is a conscious opponent that wants to inflict as much damage as possible on the decision maker. Nature is not a malevolent opponent, and the decision maker does not need to focus solely on the worst possible payoff from each action. This is especially true when the worst possible payoff from an action comes from a relatively unlikely state of nature.

Thus, this criterion normally is of interest only to a very cautious decision maker.

The Maximum Likelihood Criterion

The next criterion focuses on the *most likely* state of nature, as summarized below.

> **Maximum likelihood criterion:** Identify the most likely state of nature (the one with the largest prior probability). For this state of nature, find the action with the maximum payoff. Choose this action.

Applying this criterion to the example, Fig. 15.2 indicates that the *Dry* state has the largest prior probability. In the Dry column, the sell alternative has the maximum payoff, so the choice is to sell the land.

The appeal of this criterion is that the most important state of nature is the most likely one, so the action chosen is the best one for this particularly important state of nature. Basing the decision on the assumption that this state of nature will occur tends to give a better chance of a favorable outcome than assuming any other state of nature. Furthermore, the criterion does not rely on questionable subjective estimates of the probabilities of the respective states of nature other than identifying the most likely state.

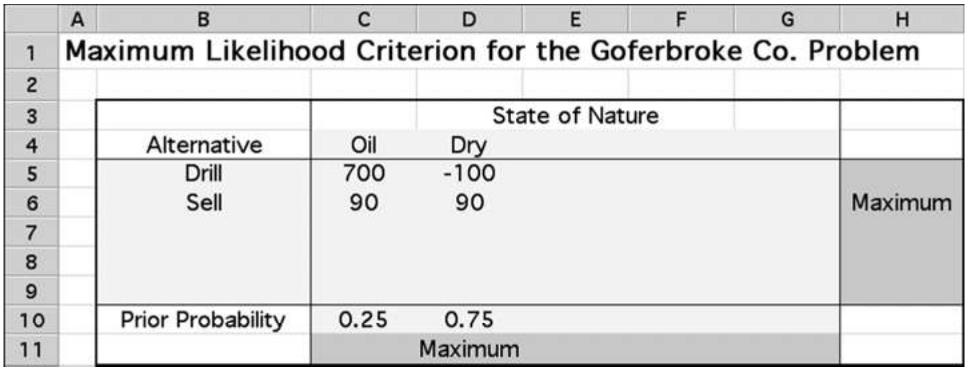

Maximum Likelihood Criterion for the Goferbroke Co. Problem

	State of Nature		
Alternative	Oil	Dry	
Drill	700	-100	Maximum
Sell	90	90	
Prior Probability	0.25	0.75	
	Maximum		

FIGURE 15.2
The application of the Excel template for the *maximum likelihood criterion* to the first Goferbroke Co. problem.

The major drawback of the criterion is that it completely ignores much relevant information. No state of nature is considered other than the most likely one. In a problem with many possible states of nature, the probability of the most likely one may be quite small, so focusing on just this one state of nature is quite unwarranted. Even in the example, where the prior probability of the *Dry* state is 0.75, this criterion ignores the extremely attractive payoff of 700 if the company drills and finds oil. In effect, the criterion does not permit gambling on a low-probability big payoff, no matter how attractive the gamble may be.

Bayes' Decision Rule[1]

Our third criterion, and the one commonly chosen, is *Bayes' decision rule,* described below.

> **Bayes' decision rule:** Using the best available estimates of the probabilities of the respective states of nature (currently the prior probabilities), calculate the expected value of the payoff for each of the possible actions. Choose the action with the maximum expected payoff.

For the prototype example, these expected payoffs are calculated directly from Table 15.2 as follows:

$$\begin{aligned}\text{E[Payoff (drill)]} &= 0.25(700) + 0.75(-100)\\ &= 100.\\ \text{E[Payoff (sell)]} &= 0.25(90) + 0.75(90)\\ &= 90.\end{aligned}$$

Since 100 is larger than 90, the alternative action selected is to drill for oil.

Note that this choice contrasts with the selection of the sell alternative under each of the two preceding criteria.

---

[1] The origin of this name is that this criterion is often credited to the Reverend Thomas Bayes, a nonconforming 18th-century English minister who won renown as a philosopher and mathematician. (The same basic idea has even longer roots in the field of economics.) This decision rule also is sometimes called the *expected monetary value* (*EMF*) criterion, although this is a misnomer for those cases where the measure of the payoff is something other than monetary value (as in Sec. 15.5).

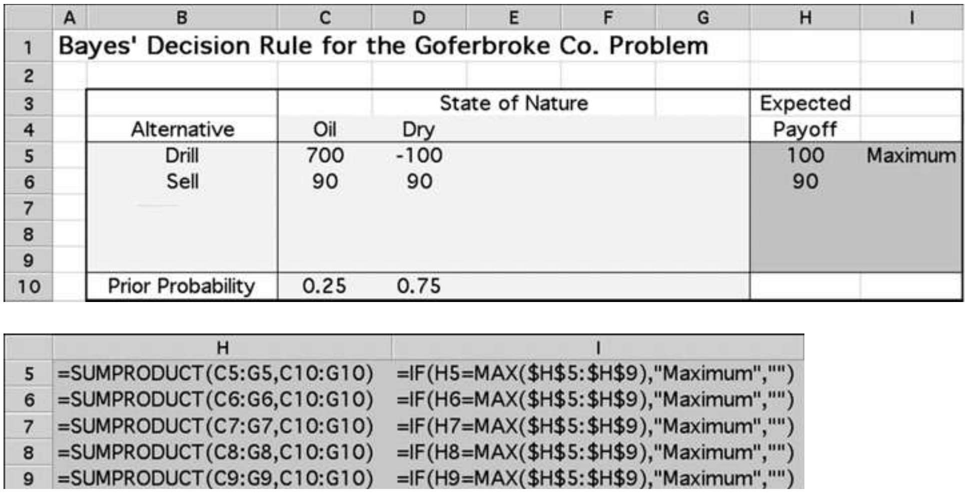

FIGURE 15.3
The application of the Excel template for *Bayes' decision rule* to the first Goferbroke Co. problem.

	A	B	C	D	E	F	G	H	I
1	Bayes' Decision Rule for the Goferbroke Co. Problem								
2									
3				State of Nature				Expected	
4		Alternative	Oil	Dry				Payoff	
5		Drill	700	-100				100	Maximum
6		Sell	90	90				90	
7									
8									
9									
10		Prior Probability	0.25	0.75					

	H	I
5	=SUMPRODUCT(C5:G5,C10:G10)	=IF(H5=MAX($H$5:$H$9),"Maximum","")
6	=SUMPRODUCT(C6:G6,C10:G10)	=IF(H6=MAX($H$5:$H$9),"Maximum","")
7	=SUMPRODUCT(C7:G7,C10:G10)	=IF(H7=MAX($H$5:$H$9),"Maximum","")
8	=SUMPRODUCT(C8:G8,C10:G10)	=IF(H8=MAX($H$5:$H$9),"Maximum","")
9	=SUMPRODUCT(C9:G9,C10:G10)	=IF(H9=MAX($H$5:$H$9),"Maximum","")

Figure 15.3 shows the application of the Excel template for Bayes' decision rule to this problem. The word *Maximum* in cell I5 signifies that the drill alternative in row 5 should be chosen because it has the maximum expected payoff.

The big advantage of Bayes' decision rule is that it incorporates all the available information, including all the payoffs and the best available estimates of the probabilities of the respective states of nature.

It is sometimes argued that these estimates of the probabilities necessarily are largely subjective and so are too shaky to be trusted. There is no accurate way of predicting the future, including a future state of nature, even in probability terms. This argument has some validity. The reasonableness of the estimates of the probabilities should be assessed in each individual situation.

Nevertheless, under many circumstances, past experience and current evidence enable one to develop reasonable estimates of the probabilities. Using this information should provide better grounds for a sound decision than ignoring it. Furthermore, experimentation frequently can be conducted to improve these estimates, as described in the next section. Therefore, we will be using only Bayes' decision rule throughout the remainder of the chapter.

To assess the effect of possible inaccuracies in the prior probabilities, it often is helpful to conduct sensitivity analysis, as described below.

Sensitivity Analysis with Bayes' Decision Rule

Sensitivity analysis commonly is used with various applications of operations research to study the effect if some of the numbers included in the mathematical model are not correct. In this case, the mathematical model is represented by the payoff table shown in Fig. 15.3. The numbers in this table that are most questionable are the prior probabilities in cells C10 and D10. We will focus the sensitivity analysis on these numbers, although a similar approach could be applied to the payoffs given in the table.

The sum of the two prior probabilities must equal 1, so increasing one of these probabilities automatically decreases the other one by the same amount, and vice versa. Goferbroke's management feels that the true chances of having oil on the tract of land are likely to lie somewhere between 15 and 35 percent. In other words, the true prior probability of having oil is likely to be in the range from 0.15 to 0.35, so the corresponding prior probability of the land being dry would range from 0.85 to 0.65.

Sensitivity analysis begins by reapplying Bayes' decision rule twice, once when the prior probability of oil is at the lower end of this range (0.15) and next when it is at the upper end (0.35). Figure 15.4 shows the results from doing this. When the prior probability of oil is only 0.15, the decision swings over to selling the land by a wide margin (an expected payoff of 90 versus only 20 for drilling). However, when this probability is 0.35, the decision is to drill by a wide margin (expected payoff = 180 versus only 90 for selling). Thus, the decision is very *sensitive* to the prior probability of oil. This sensitivity analysis has revealed that it is important to do more, if possible, to pin down just what the true value of the probability of oil is.

Letting

p = prior probability of oil,

the expected payoff from drilling for any p is

$$\begin{aligned} E[\text{Payoff (drill)}] &= 700p - 100(1 - p) \\ &= 800p - 100. \end{aligned}$$

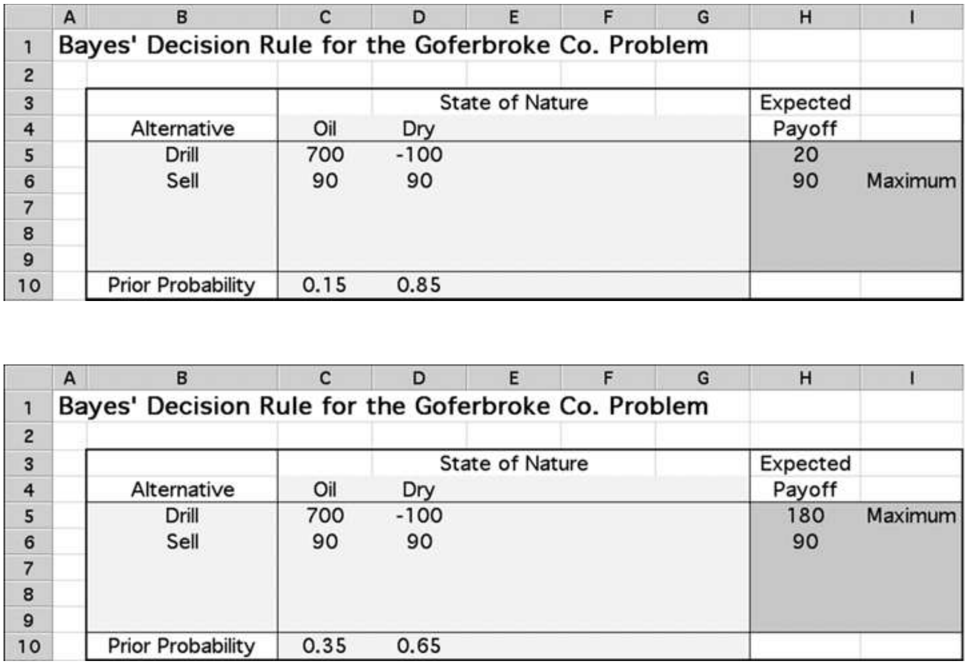

FIGURE 15.4
Performing sensitivity analysis by trying alternative values of the prior probability of oil.

	A	B	C	D	E	F	G	H	I
1	Bayes' Decision Rule for the Goferbroke Co. Problem								
2									
3					State of Nature			Expected	
4		Alternative	Oil	Dry				Payoff	
5		Drill	700	-100				20	
6		Sell	90	90				90	Maximum
7									
8									
9									
10		Prior Probability	0.15	0.85					

	A	B	C	D	E	F	G	H	I
1	Bayes' Decision Rule for the Goferbroke Co. Problem								
2									
3					State of Nature			Expected	
4		Alternative	Oil	Dry				Payoff	
5		Drill	700	-100				180	Maximum
6		Sell	90	90				90	
7									
8									
9									
10		Prior Probability	0.35	0.65					

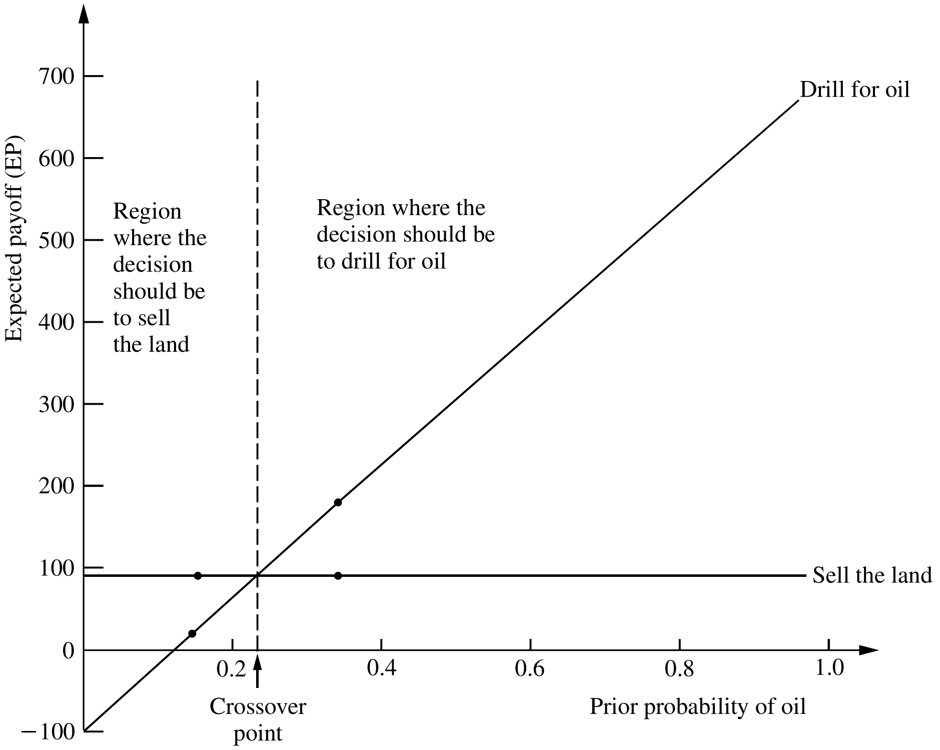

FIGURE 15.5
Graphical display of how the expected payoff for each alternative action changes when the prior probability of oil changes for the first Goferbroke Co. problem.

The slanting line in Fig. 15.5 shows the plot of this expected payoff versus *p,* which is just the line passing through the two points given by cells C10 and H5 in the two spreadsheets in Fig. 15.4. Since the payoff from selling the land would be 90 for any *p,* the flat line in Fig. 15.5 gives *E*[Payoff (sell)] versus *p*.

The point in Fig. 15.5 where the two lines intersect is the **crossover point** where the decision shifts from one alternative (sell the land) to the other (drill for oil) as the prior probability increases. To find this point, we set

$$E[\text{Payoff (drill)}] = E[\text{Payoff (sell)}]$$
$$800p - 100 = 90$$
$$p = \frac{190}{800} = 0.2375$$

Conclusion: Should sell the land if $p < 0.2375$.
Should drill for oil if $p > 0.2375$.

For other problems that have more than two alternative actions, the same kind of analysis can be applied. The main difference is that there now would be more than two lines (one per alternative) in the graphical display corresponding to Fig. 15.5. However, the top line for any particular value of the prior probability still indicates which alternative should be chosen. With more than two lines, there might be more than one crossover point where the decision shifts from one alternative to another.

For a problem with more than two possible states of nature, the most straightforward approach is to focus the sensitivity analysis on only two states at a time as described above. This again would involve investigating what happens when the prior probability of one state increases as the prior probability of the other state decreases by the same amount, holding fixed the prior probabilities of the remaining states. This procedure then can be repeated for as many other pairs of states as desired.

Practitioners sometimes use software to assist them in performing this kind of sensitivity analysis, including generating the graphs. For example, an Excel add-in in your OR Courseware called SensIt is designed specifically for conducting sensitivity analysis with probabilistic models such as when applying Bayes' decision rule. Complete documentation for SensIt is included on your CD-ROM.

Because the decision the Goferbroke Co. should make depends so critically on the true probability of oil, serious consideration should be given to conducting a seismic survey to estimate this probability more closely. We will explore this option in the next two sections.

15.3 DECISION MAKING WITH EXPERIMENTATION

Frequently, additional testing (experimentation) can be done to improve the preliminary estimates of the probabilities of the respective states of nature provided by the prior probabilities. These improved estimates are called **posterior probabilities.**

We first update the Goferbroke Co. example to incorporate experimentation, then describe how to derive the posterior probabilities, and finally discuss how to decide whether it is worthwhile to conduct experimentation.

Continuing the Prototype Example

As mentioned at the end of Sec. 15.1, an available option before making a decision is to conduct a detailed seismic survey of the land to obtain a better estimate of the probability of oil. The cost is $30,000.

A seismic survey obtains seismic soundings that indicate whether the geological structure is favorable to the presence of oil. We will divide the possible findings of the survey into the following two categories:

USS: Unfavorable seismic soundings; oil is fairly unlikely.
FSS: Favorable seismic soundings; oil is fairly likely.

Based on past experience, if there is oil, then the probability of unfavorable seismic soundings is

$$P(\text{USS} \mid \text{State} = \text{Oil}) = 0.4, \qquad \text{so} \qquad P(\text{FSS} \mid \text{State} = \text{Oil}) = 1 - 0.4 = 0.6.$$

Similarly, if there is no oil (i.e., the true state of nature is *Dry*), then the probability of unfavorable seismic soundings is estimated to be

$$P(\text{USS} \mid \text{State} = \text{Dry}) = 0.8, \qquad \text{so} \qquad P(\text{FSS} \mid \text{State} = \text{Dry}) = 1 - 0.8 = 0.2.$$

We soon will use these data to find the posterior probabilities of the respective states of nature *given* the seismic soundings.

Posterior Probabilities

Proceeding now in general terms, we let

n = number of possible states of nature;
$P(\text{State} = \text{state } i)$ = prior probability that true state of nature is state i, for $i = 1, 2, \ldots, n$;
Finding = finding from experimentation (a random variable);
Finding j = one possible value of finding;
$P(\text{State} = \text{state } i \mid \text{Finding} = \text{finding } j)$ = posterior probability that true state of nature is state i, given that Finding = finding j, for $i = 1, 2, \ldots, n$.

The question currently being addressed is the following:

> Given $P(\text{State} = \text{state } i)$ and $P(\text{Finding} = \text{finding } j \mid \text{State} = \text{state } i)$, for $i = 1, 2, \ldots, n$, what is $P(\text{State} = \text{state } i \mid \text{Finding} = \text{finding } j)$?

This question is answered by combining the following standard formulas of probability theory:

$$P(\text{State} = \text{state } i \mid \text{Finding} = \text{finding } j) = \frac{P(\text{State} = \text{state } i, \text{Finding} = \text{finding } j)}{P(\text{Finding} = \text{finding } j)}$$

$$P(\text{Finding} = \text{finding } j) = \sum_{k=1}^{n} P(\text{State} = \text{state } k, \text{Finding} = \text{finding } j)$$

$$P(\text{State} = \text{state } i, \text{Finding} = \text{finding } j) = P(\text{Finding} = \text{finding } j \mid \text{State} = \text{state } i)\, P(\text{State} = \text{state } i).$$

Therefore, for each $i = 1, 2, \ldots, n$, the desired formula for the corresponding posterior probability is

$$P(\text{State} = \text{state } i \mid \text{Finding} = \text{finding } j) = \frac{P(\text{Finding} = \text{finding } j \mid \text{State} = \text{state } i)P(\text{State} = \text{state } i)}{\sum_{k=1}^{n} P(\text{Finding} = \text{finding } j \mid \text{State} = \text{state } k)P(\text{State} = \text{state } k)}$$

(This formula often is referred to as **Bayes' theorem** because it was developed by Thomas Bayes, the same 18th-century mathematician who is credited with developing Bayes' decision rule.)

Now let us return to the prototype example and apply this formula. If the finding of the seismic survey is unfavorable seismic soundings (USS), then the posterior probabilities are

$$P(\text{State} = \text{Oil} \mid \text{Finding} = \text{USS}) = \frac{0.4(0.25)}{0.4(0.25) + 0.8(0.75)} = \frac{1}{7},$$

$$P(\text{State} = \text{Dry} \mid \text{Finding} = \text{USS}) = 1 - \frac{1}{7} = \frac{6}{7}.$$

Similarly, if the seismic survey gives favorable seismic soundings (FSS), then

$$P(\text{State} = \text{Oil} \mid \text{Finding} = \text{FSS}) = \frac{0.6(0.25)}{0.6(0.25) + 0.2(0.75)} = \frac{1}{2},$$

$$P(\text{State} = \text{Dry} \mid \text{Finding} = \text{FSS}) = 1 - \frac{1}{2} = \frac{1}{2}.$$

The **probability tree diagram** in Fig. 15.6 shows a nice way of organizing these calculations in an intuitive manner. The prior probabilities in the first column and the conditional probabilities in the second column are part of the input data for the problem. Multiplying each probability in the first column by a probability in the second column gives the corresponding joint probability in the third column. Each joint probability then becomes the numerator in the calculation of the corresponding posterior probability in the fourth column. Cumulating the joint probabilities with the same finding (as shown at the bottom of the figure) provides the denominator for each posterior probability with this finding.

Your OR Courseware also includes an Excel template for computing these posterior probabilities, as shown in Fig. 15.7.

After these computations have been completed, Bayes' decision rule can be applied just as before, with the posterior probabilities now replacing the prior probabilities. Again,

FIGURE 15.6
Probability tree diagram for the full Goferbroke Co. problem showing all the probabilities leading to the calculation of each posterior probability of the state of nature given the finding of the seismic survey.

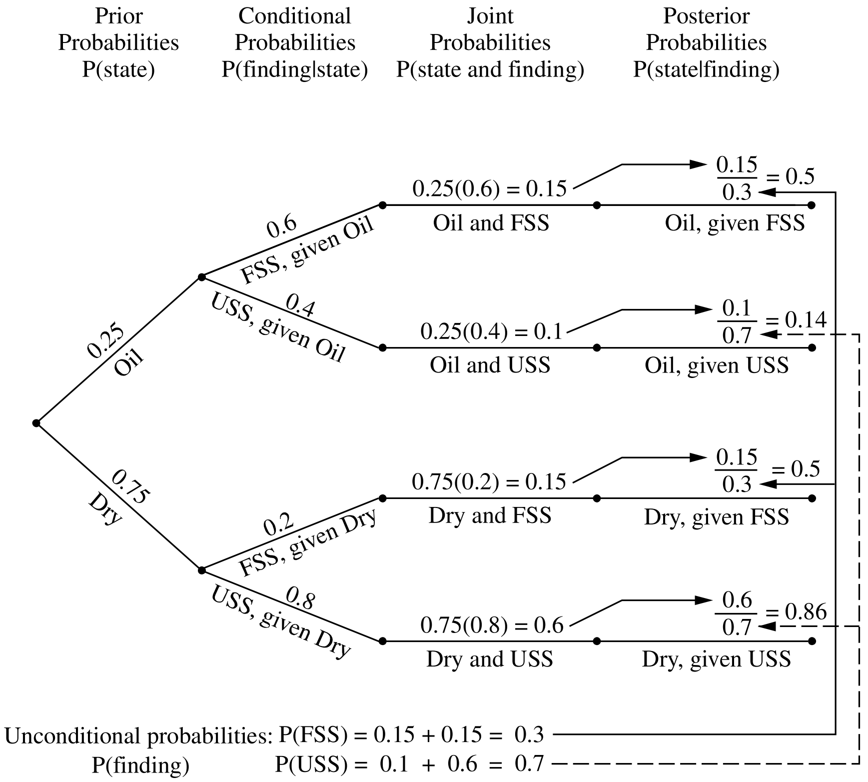

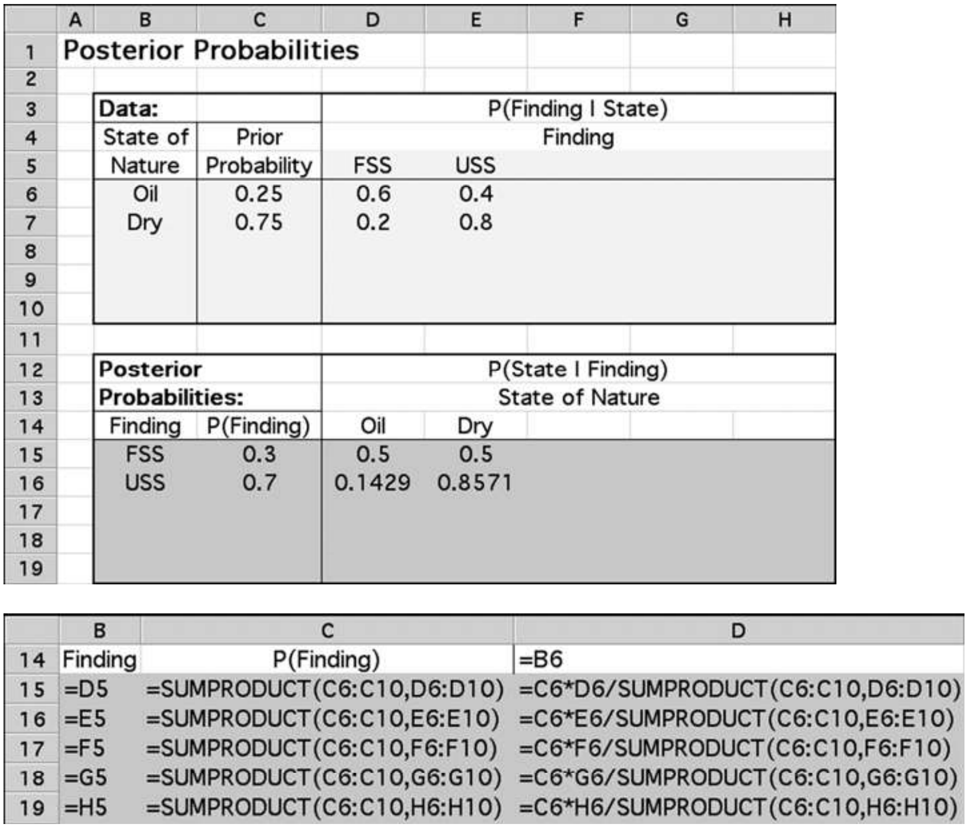

	A	B	C	D	E	F	G	H
1	Posterior Probabilities							
2								
3		Data:		P(Finding I State)				
4		State of	Prior	Finding				
5		Nature	Probability	FSS	USS			
6		Oil	0.25	0.6	0.4			
7		Dry	0.75	0.2	0.8			
8								
9								
10								
11								
12		Posterior		P(State I Finding)				
13		Probabilities:		State of Nature				
14		Finding	P(Finding)	Oil	Dry			
15		FSS	0.3	0.5	0.5			
16		USS	0.7	0.1429	0.8571			
17								
18								
19								

	B	C	D
14	Finding	P(Finding)	=B6
15	=D5	=SUMPRODUCT(C6:C10,D6:D10)	=C6*D6/SUMPRODUCT(C6:C10,D6:D10)
16	=E5	=SUMPRODUCT(C6:C10,E6:E10)	=C6*E6/SUMPRODUCT(C6:C10,E6:E10)
17	=F5	=SUMPRODUCT(C6:C10,F6:F10)	=C6*F6/SUMPRODUCT(C6:C10,F6:F10)
18	=G5	=SUMPRODUCT(C6:C10,G6:G10)	=C6*G6/SUMPRODUCT(C6:C10,G6:G10)
19	=H5	=SUMPRODUCT(C6:C10,H6:H10)	=C6*H6/SUMPRODUCT(C6:C10,H6:H10)

FIGURE 15.7
This *posterior probabilities* template in your OR Courseware enables efficient calculation of posterior probabilities, as illustrated here for the full Goferbroke Co. problem.

by using the payoffs (in units of thousands of dollars) from Table 15.2 and subtracting the cost of the experimentation, we obtain the results shown below.

Expected payoffs if finding is unfavorable seismic soundings (USS):

$$\begin{aligned} E[\text{Payoff (drill} \mid \text{Finding = USS)}] &= \frac{1}{7}(700) + \frac{6}{7}(-100) - 30 \\ &= -15.7. \\ E[\text{Payoff (sell} \mid \text{Finding = USS)}] &= \frac{1}{7}(90) + \frac{6}{7}(90) - 30 \\ &= 60. \end{aligned}$$

Expected payoffs if finding is favorable seismic soundings (FSS):

$$\begin{aligned} E[\text{Payoff (drill} \mid \text{Finding = FSS)}] &= \frac{1}{2}(700) + \frac{1}{2}(-100) - 30 \\ &= 270. \end{aligned}$$

TABLE 15.3 The optimal policy with experimentation, under Bayes' decision rule, for the Goferbroke Co. problem

Finding from Seismic Survey	Optimal Action	Expected Payoff Excluding Cost of Survey	Expected Payoff Including Cost of Survey
USS	Sell the land	90	60
FSS	Drill for oil	300	270

$$\begin{aligned} E[\text{Payoff (sell} \mid \text{Finding = FSS)}] &= \frac{1}{2}(90) + \frac{1}{2}(90) - 30 \\ &= 60. \end{aligned}$$

Since the objective is to maximize the expected payoff, these results yield the optimal policy shown in Table 15.3.

However, what this analysis does not answer is whether it is worth spending $30,000 to conduct the experimentation (the seismic survey). Perhaps it would be better to forgo this major expense and just use the optimal solution without experimentation (drill for oil, with an expected payoff of $100,000). We address this issue next.

The Value of Experimentation

Before performing any experiment, we should determine its potential value. We present two complementary methods of evaluating its potential value.

The first method assumes (unrealistically) that the experiment will remove *all* uncertainty about what the true state of nature is, and then this method makes a very quick calculation of what the resulting *improvement in the expected payoff* would be (ignoring the cost of the experiment). This quantity, called the *expected value of perfect information,* provides an *upper bound* on the potential value of the experiment. Therefore, if this upper bound is less than the cost of the experiment, the experiment definitely should be forgone.

However, if this upper bound exceeds the cost of the experiment, then the second (slower) method should be used next. This method calculates the *actual* improvement in the expected payoff (ignoring the cost of the experiment) that would result from performing the experiment. Comparing this improvement with the cost indicates whether the experiment should be performed.

Expected Value of Perfect Information. Suppose now that the experiment could definitely identify what the true state of nature is, thereby providing "perfect" information. Whichever state of nature is identified, you naturally choose the action with the maximum payoff for that state. We do not know in advance which state of nature will be identified, so a calculation of the expected payoff with perfect information (ignoring the cost of the experiment) requires weighting the maximum payoff for each state of nature by the prior probability of that state of nature.

Figure 15.8 shows the Excel template in your OR Courseware that can be used to organize and perform this calculation. Using the equation given for cell F13,

$$\begin{aligned} \text{Expected payoff with perfect information} &= 0.25(700) + 0.75(90) \\ &= 242.5. \end{aligned}$$

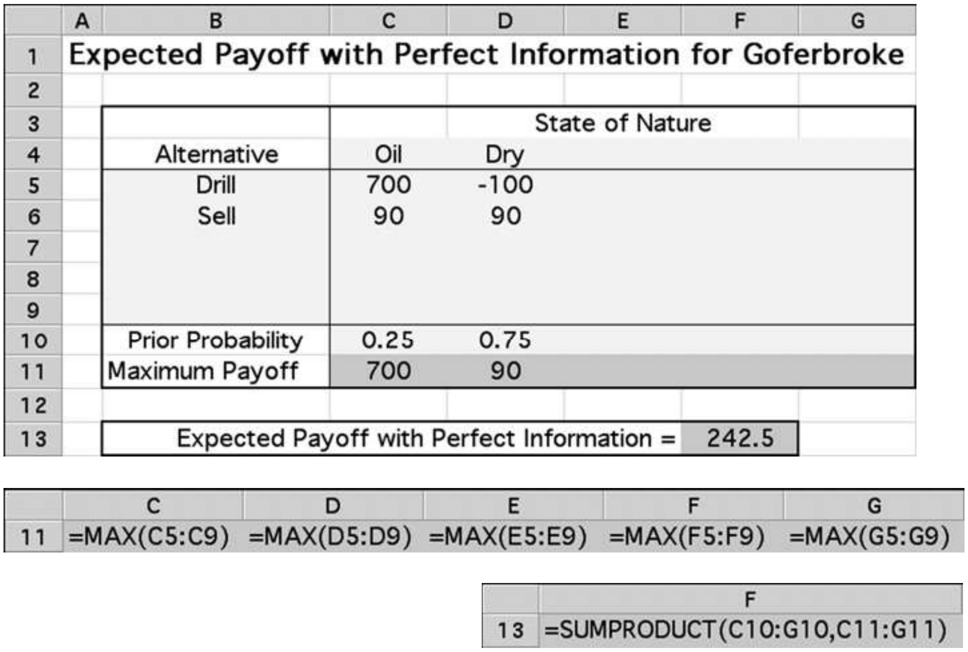

	A	B	C	D	E	F	G
1	Expected Payoff with Perfect Information for Goferbroke						
2							
3					State of Nature		
4		Alternative	Oil	Dry			
5		Drill	700	-100			
6		Sell	90	90			
7							
8							
9							
10		Prior Probability	0.25	0.75			
11		Maximum Payoff	700	90			
12							
13		Expected Payoff with Perfect Information =				242.5	

	C	D	E	F	G
11	=MAX(C5:C9)	=MAX(D5:D9)	=MAX(E5:E9)	=MAX(F5:F9)	=MAX(G5:G9)

	F
13	=SUMPRODUCT(C10:G10,C11:G11)

FIGURE 15.8
This Excel template for obtaining the expected payoff with perfect information is applied here to the first Goferbroke Co. problem.

Thus, if the Goferbroke Co. could learn before choosing its action whether the land contains oil, the expected payoff as of now (before acquiring this information) would be \$242,500 (excluding the cost of the experiment generating the information.)

To evaluate whether the experiment should be conducted, we now use this quantity to calculate the expected value of perfect information.

The **expected value of perfect information,** abbreviated **EVPI,** is calculated as

EVPI = expected payoff with perfect information − expected payoff without experimentation.[1]

Thus, since experimentation usually cannot provide perfect information, EVPI provides an upper bound on the expected value of experimentation.

For the prototype example, we found in Sec. 15.2 that the expected payoff without experimentation (under Bayes' decision rule) is 100. Therefore,

$$\text{EVPI} = 242.5 - 100 = 142.5.$$

Since 142.5 far exceeds 30, the cost of experimentation (a seismic survey), it may be worthwhile to proceed with the seismic survey. To find out for sure, we now go to the second method of evaluating the potential benefit of experimentation.

---

[1] The *value of perfect information* is a random variable equal to the payoff with perfect information *minus* the payoff without experimentation. EVPI is the expected value of this random variable.

Expected Value of Experimentation. Rather than just obtain an upper bound on the *expected increase in payoff* (excluding the cost of the experiment) due to performing experimentation, we now will do somewhat more work to calculate this expected increase directly. This quantity is called the *expected value of experimentation.*

Calculating this quantity requires first computing the expected payoff with experimentation (excluding the cost of the experiment). Obtaining this latter quantity requires doing all the work described earlier to find all the posterior probabilities, the resulting optimal policy with experimentation, and the corresponding expected payoff (excluding the cost of the experiment) for each possible finding from the experiment. Then each of these expected payoffs needs to be weighted by the probability of the corresponding finding, that is,

$$\text{Expected payoff with experimentation} = \sum_j P(\text{Finding} = \text{finding } j) \\ E[\text{payoff} \mid \text{Finding} = \text{finding } j],$$

where the summation is taken over all possible values of j.

For the prototype example, we have already done all the work to obtain the terms on the right side of this equation. The values of $P(\text{Finding} = \text{finding } j)$ for the two possible findings from the seismic survey—unfavorable (USS) and favorable (FSS)—were calculated at the bottom of the probability tree diagram in Fig. 15.6 as

$$P(\text{USS}) = 0.7, \qquad P(\text{FSS}) = 0.3.$$

For the optimal policy with experimentation, the corresponding expected payoff (excluding the cost of the seismic survey) for each finding was obtained in the third column of Table 15.3 as

$$E(\text{Payoff} \mid \text{Finding} = \text{USS}) = 90,$$
$$E(\text{Payoff} \mid \text{Finding} = \text{FSS}) = 270.$$

With these numbers,

$$\begin{aligned}\text{Expected payoff with experimentation} &= 0.7(90) + 0.3(300)\\ &= 153.\end{aligned}$$

Now we are ready to calculate the expected value of experimentation.

The **expected value of experimentation,** abbreviated **EVE,** is calculated as

EVE = expected payoff with experimentation − expected payoff without experimentation.

Thus, EVE identifies the potential value of experimentation.

For the Goferbroke Co.,

$$\text{EVE} = 153 - 100 = 53.$$

Since this value exceeds 30, the cost of conducting a detailed seismic survey (in units of thousands of dollars), this experimentation should be done.

15.4 DECISION TREES

Decision trees provide a useful way of *visually displaying* the problem and then *organizing the computational work* already described in the preceding two sections. These trees are especially helpful when a *sequence of decisions* must be made.

Constructing the Decision Tree

The prototype example involves a sequence of two decisions:

1. Should a seismic survey be conducted before an action is chosen?
2. Which action (drill for oil or sell the land) should be chosen?

The corresponding decision tree (before adding numbers and performing computations) is displayed in Fig. 15.9.

The nodes of the decision tree are referred to as **forks,** and the arcs are called **branches.**

> A **decision fork,** represented by a square, indicates that a decision needs to be made at that point in the process. A **chance fork,** represented by a circle, indicates that a random event occurs at that point.

Thus, in Fig. 15.9, the first decision is represented by decision fork *a*. Fork *b* is a chance fork representing the random event of the outcome of the seismic survey. The two branches emanating from fork *b* represent the two possible outcomes of the survey. Next

FIGURE 15.9
The decision tree (before including any numbers) for the full Goferbroke Co. problem.

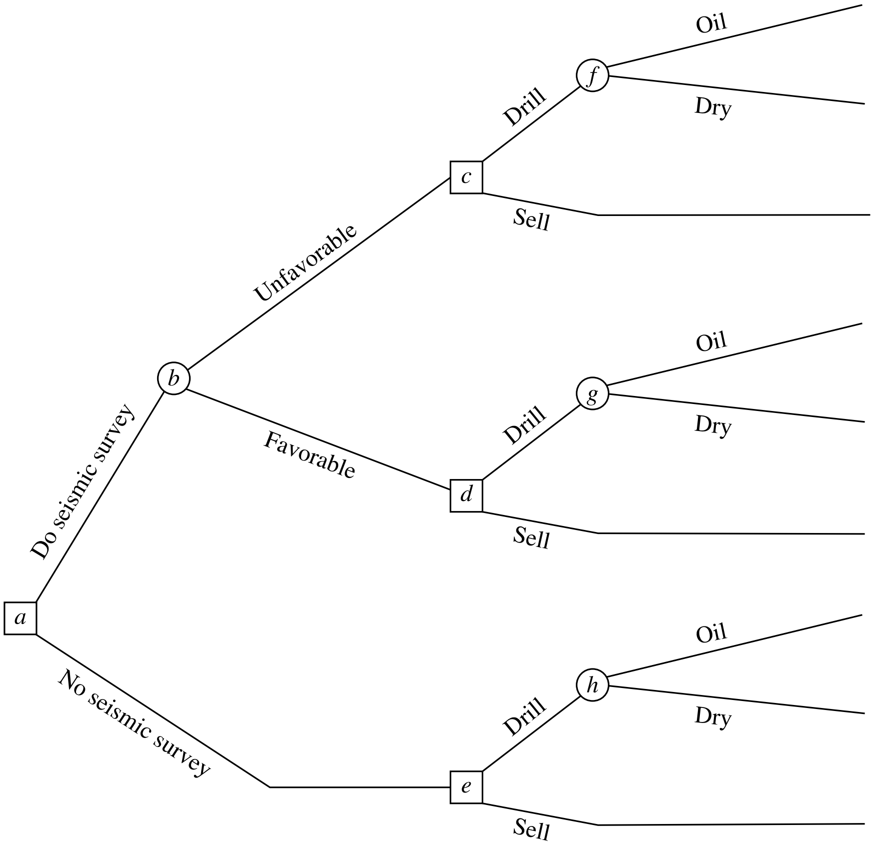

comes the second decision (forks *c, d,* and *e*) with its two possible choices. If the decision is to drill for oil, then we come to another chance fork (forks *f, g,* and *h*), where its two branches correspond to the two possible states of nature.

Note that the path followed from fork *a* to reach any terminal branch (except the bottom one) is determined both by the decisions made and by random events that are outside the control of the decision maker. This is characteristic of problems addressed by decision analysis.

The next step in constructing the decision tree is to insert numbers into the tree as shown in Fig. 15.10. The numbers under or over the branches that are *not* in parentheses are the cash flows (in thousands of dollars) that occur at those branches. For each path through the tree from node *a* to a terminal branch, these same numbers then are added to obtain the resulting total payoff shown in boldface to the right of that branch. The last set of numbers is the probabilities of random events. In particular, since each branch emanating from a chance fork represents a possible random event, the probability of this event occurring from this fork has been inserted in parentheses along this branch. From chance

FIGURE 15.10
The decision tree in Fig. 15.9 after adding both the probabilities of random events and the payoffs.

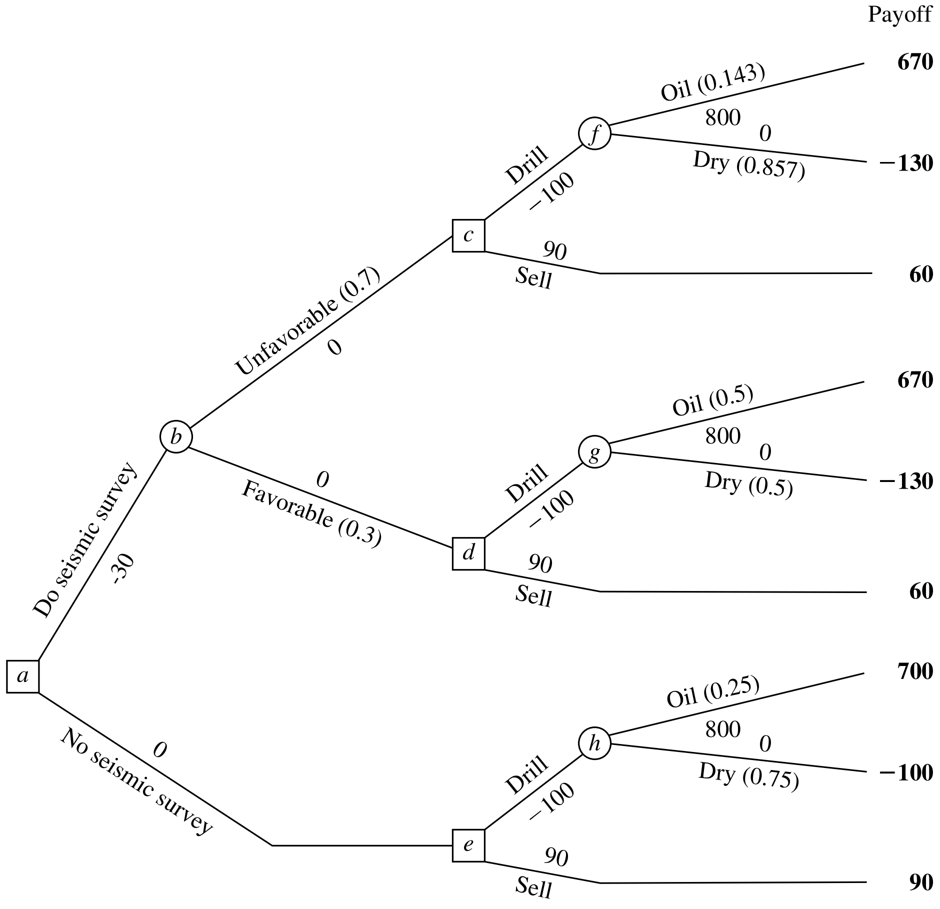

fork *h,* the probabilities are the *prior probabilities* of these states of nature, since no seismic survey has been conducted to obtain more information in this case. However, chance forks *f* and *g* lead out of a decision to do the seismic survey (and then to drill). Therefore, the probabilities from these chance forks are the *posterior probabilities* of the states of nature, given the finding from the seismic survey, where these numbers are given in Figs. 15.6 and 15.7. Finally, we have the two branches emanating from chance fork *b.* The numbers here are the probabilities of these findings from the seismic survey, Favorable (FSS) or Unfavorable (USS), as given underneath the probability tree diagram in Fig. 15.6 or in cells C15:C16 of Fig. 15.7.

Performing the Analysis

Having constructed the decision tree, including its numbers, we now are ready to analyze the problem by using the following procedure.

1. Start at the right side of the decision tree and move left one column at a time. For each column, perform either step 2 or step 3 depending upon whether the forks in that column are chance forks or decision forks.
2. For each chance fork, calculate its *expected payoff* by multiplying the expected payoff of each branch (shown in boldface to the right of the branch) by the probability of that branch and then summing these products. Record this expected payoff for each decision fork in boldface next to the fork, and designate this quantity as also being the expected payoff for the branch leading to this fork.
3. For each decision fork, compare the expected payoffs of its branches and choose the alternative whose branch has the largest expected payoff. In each case, record the choice on the decision tree by inserting a double dash as a barrier through each rejected branch.

To begin the procedure, consider the rightmost column of forks, namely, chance forks *f, g,* and *h.* Applying step 2, their expected payoffs (EP) are calculated as

$$\text{EP} = \frac{1}{7}(670) + \frac{6}{7}(-130) = -15.7, \qquad \text{for fork } f,$$

$$\text{EP} = \frac{1}{2}(670) + \frac{1}{2}(-130) = 270, \qquad \text{for fork } g,$$

$$\text{EP} = \frac{1}{4}(700) + \frac{3}{4}(-100) = 100, \qquad \text{for fork } h.$$

These expected payoffs then are placed above these forks, as shown in Fig. 15.11.

Next, we move one column to the left, which consists of decision forks *c, d,* and *e.* The expected payoff for a branch that leads to a chance fork now is recorded in boldface over that chance fork. Therefore, step 3 can be applied as follows.

Fork *c:* Drill alternative has EP $= -15.7$.
Sell alternative has EP $= 60$.
$60 > -15.7$, so choose the Sell alternative.

Fork *d:* Drill alternative has EP $= 270$.
Sell alternative has EP $= 60$.
$270 > 60$, so choose the Drill alternative.

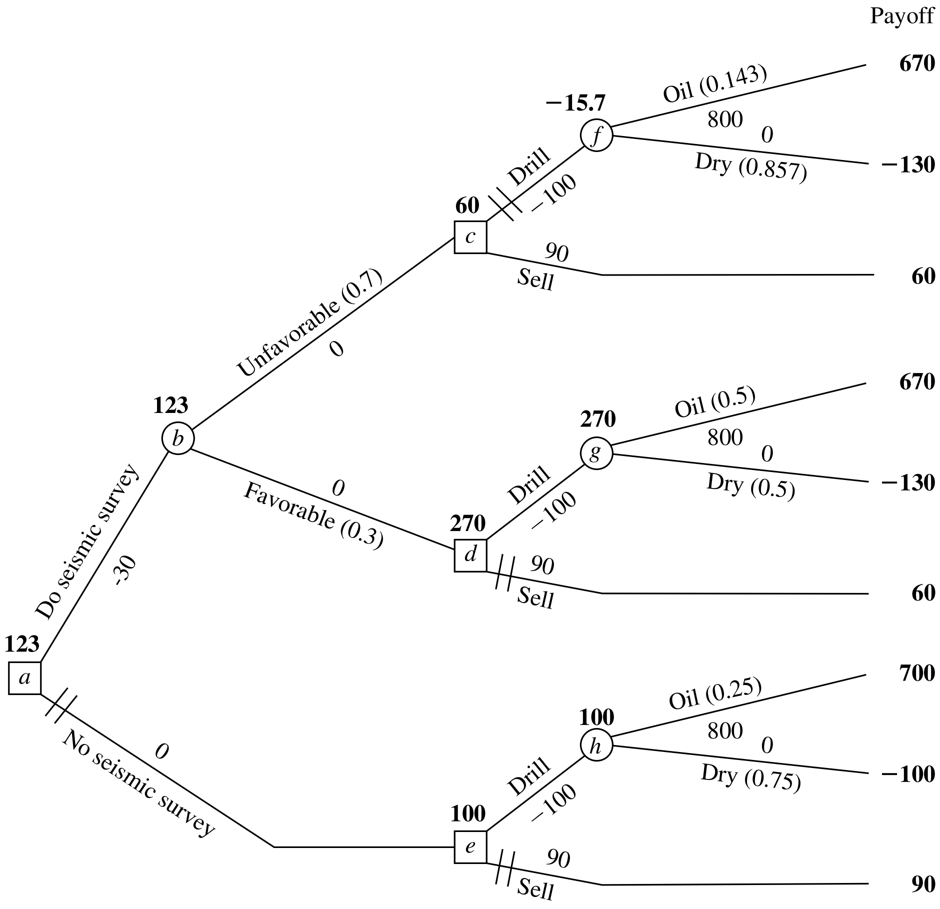

FIGURE 15.11
The final decision tree that records the analysis for the full Goferbroke Co. problem when using monetary payoffs.

Fork *e:* Drill alternative has EP = 100.
Sell alternative has EP = 90.
100 > 90, so choose the Drill alternative.

The expected payoff for each chosen alternative now would be recorded in boldface over its decision node, as already shown in Fig. 15.11. The chosen alternative also is indicated by inserting a double dash as a barrier through each rejected branch.

Next, moving one more column to the left brings us to fork *b*. Since this is a chance fork, step 2 of the procedure needs to be applied. The expected payoff for each of its branches is recorded over the following decision fork. Therefore, the expected payoff is

$$EP = 0.7(60) + 0.3(270) = 123, \qquad \text{for fork } b,$$

as recorded over this fork in Fig. 15.11.

Finally, we move left to fork *a,* a decision fork. Applying step 3 yields

Fork *a:* Do seismic survey has EP = 123.
No seismic survey has EP = 100.
123 > 100, so choose Do seismic survey.

This expected payoff of 123 now would be recorded over the fork, and a double dash inserted to indicate the rejected branch, as already shown in Fig. 15.11.

This procedure has moved from right to left for analysis purposes. However, having completed the decision tree in this way, the decision maker now can read the tree from left to right to see the actual progression of events. The double dashes have closed off the undesirable paths. Therefore, given the payoffs for the final outcomes shown on the right side, *Bayes' decision rule* says to follow only the open paths from left to right to achieve the largest possible expected payoff.

Following the open paths from left to right in Fig. 15.11 yields the following optimal policy, according to Bayes' decision rule.

Optimal policy:
Do the seismic survey.
If the result is unfavorable, sell the land.
If the result is favorable, drill for oil.
The expected payoff (including the cost of the seismic survey) is 123 ($123,000).

This (unique) optimal solution naturally is the same as that obtained in the preceding section without the benefit of a decision tree. (See the optimal policy with experimentation given in Table 15.3 and the conclusion at the end of Sec. 15.3 that experimentation is worthwhile.)

For any decision tree, this **backward induction procedure** always will lead to the *optimal policy* (or policies) after the probabilities are computed for the branches emanating from a chance fork.

Helpful Software

Practitioners sometimes use special software to help construct and analyze decision trees. This software often is in the form of an Excel add-in. One popular add-in of this type is *TreePlan,* which is shareware developed by Professor Michael Middleton. The academic version of TreePlan is included in your OR Courseware, along with Professor Middleton's companion shareware SensIt mentioned at the end of Sec. 15.2.

It is straightforward to use TreePlan to quickly construct a decision tree equivalent to the one in Fig. 15.11, as well as much larger ones. In the process, TreePlan also will automatically solve the decision tree. The Excel file for this chapter includes the TreePlan decision trees for three versions of the Goferbroke Co. problem. Complete documentation for TreePlan also is included on the CD-ROM.

To construct a decision tree with TreePlan, go to its Tools menu and choose *Decision Tree,* which brings up the "TreePlan . . . New" dialogue box shown in Fig. 15.12. Clicking on New Tree then adds a tree to the spreadsheet that initially consists of a single (square) decision fork with two branches. Clicking just to the right of a terminal fork