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

Alternative \ Status of Land	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 ran-

dom 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.

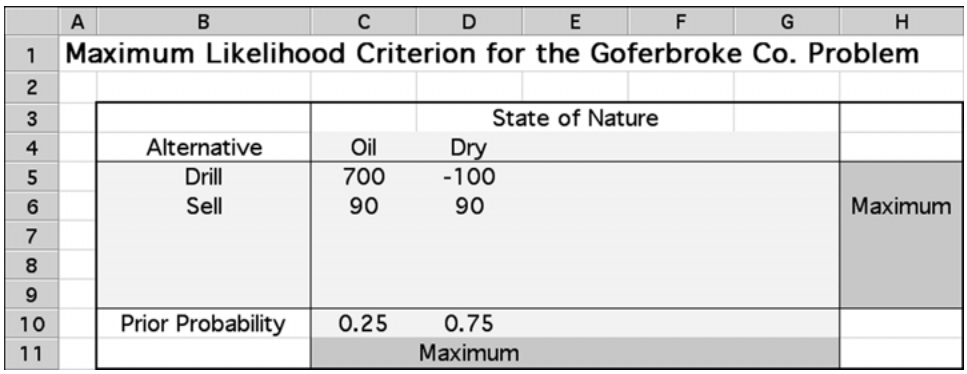

	A	B	C	D	E	F	G	H
1	Maximum Likelihood Criterion for the Goferbroke Co. Problem							
2								
3			State of Nature					
4		Alternative	Oil	Dry				
5		Drill	700	-100				
6		Sell	90	90				Maximum
7								
8								
9								
10		Prior Probability	0.25	0.75				
11			Maximum					

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).

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

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

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

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					

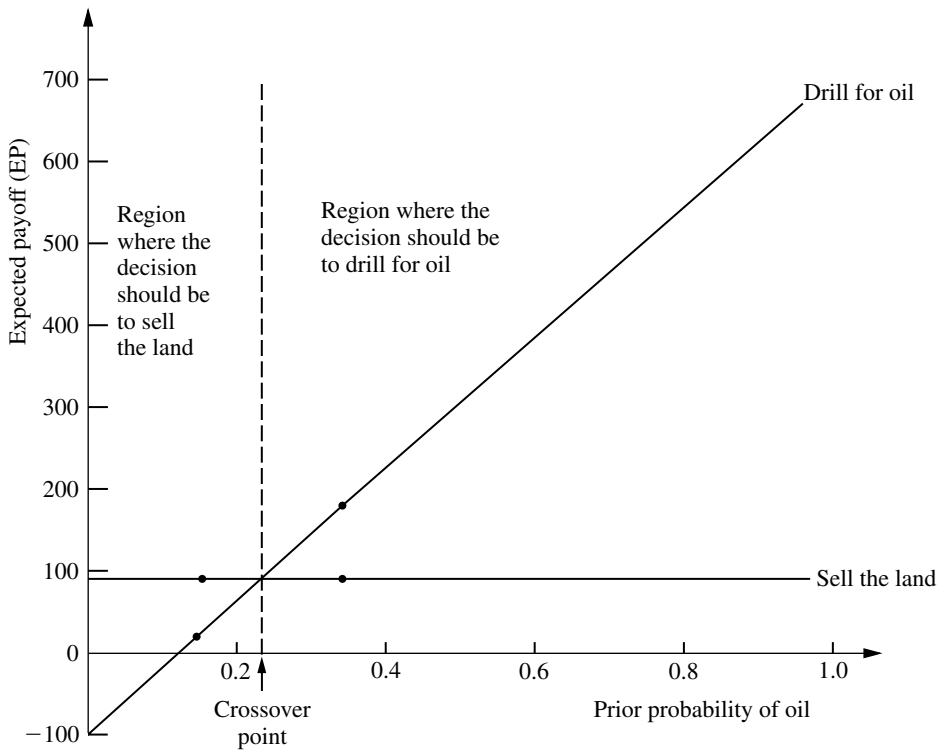

FIGURE 15.5
Graphical display of how the expected payoff for each alternative action changes when the prior probability of oil changes for the first Goferbroke Co. problem.

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,

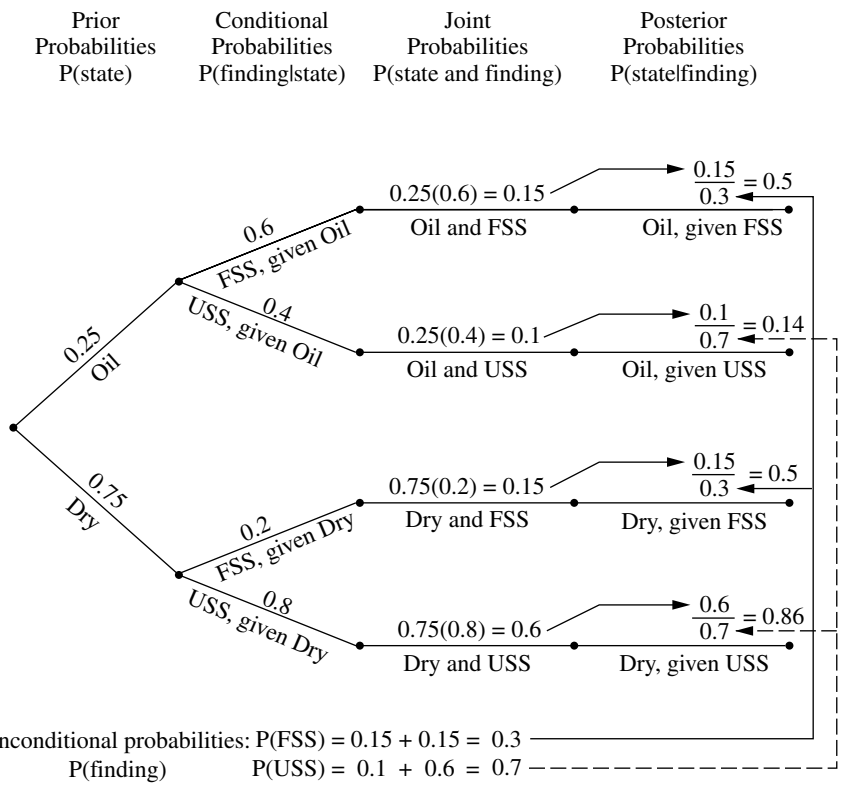

FIGURE 15.6
Probability tree diagram for the full Goferbroke Co. problem showing all the probabilities leading to the calculation of each posterior probability of the state of nature given the finding of the seismic survey.

	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):

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

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

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

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

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

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,

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

	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,

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

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

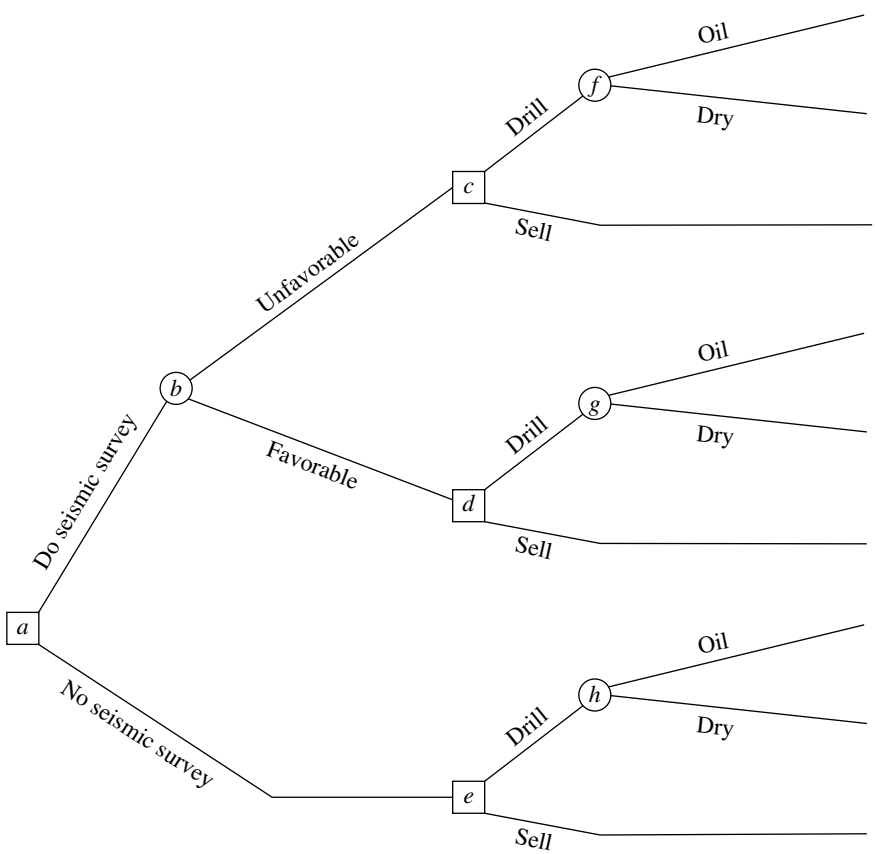

FIGURE 15.9
The decision tree (before including any numbers) for the full Goferbroke Co. problem.

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

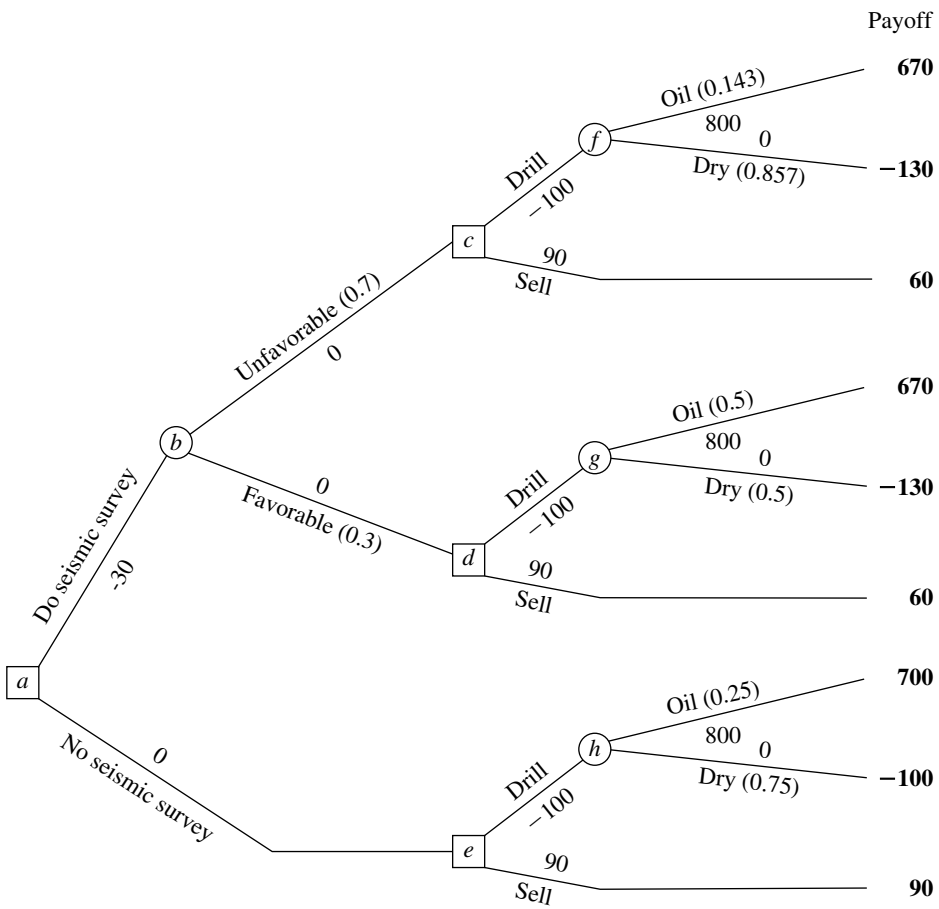

FIGURE 15.10
The decision tree in Fig. 15.9 after adding both the probabilities of random events and the payoffs.

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.

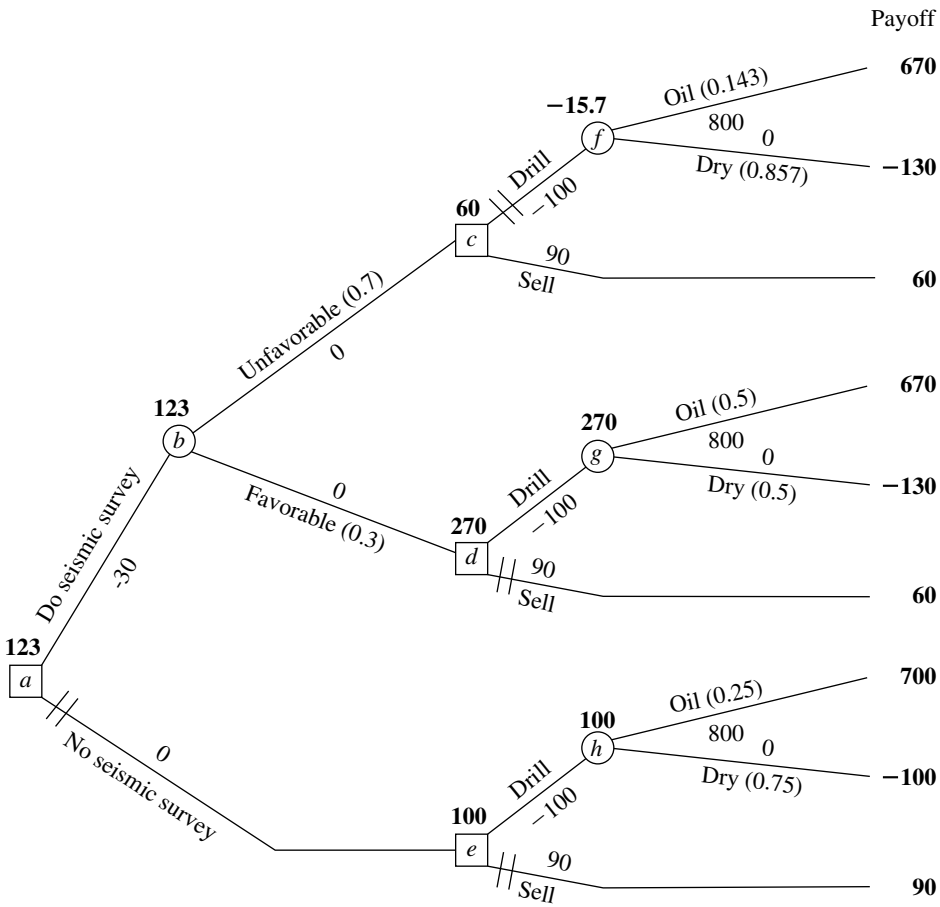

FIGURE 15.11
The final decision tree that records the analysis for the full Goferbroke Co. problem when using monetary payoffs.

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

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

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

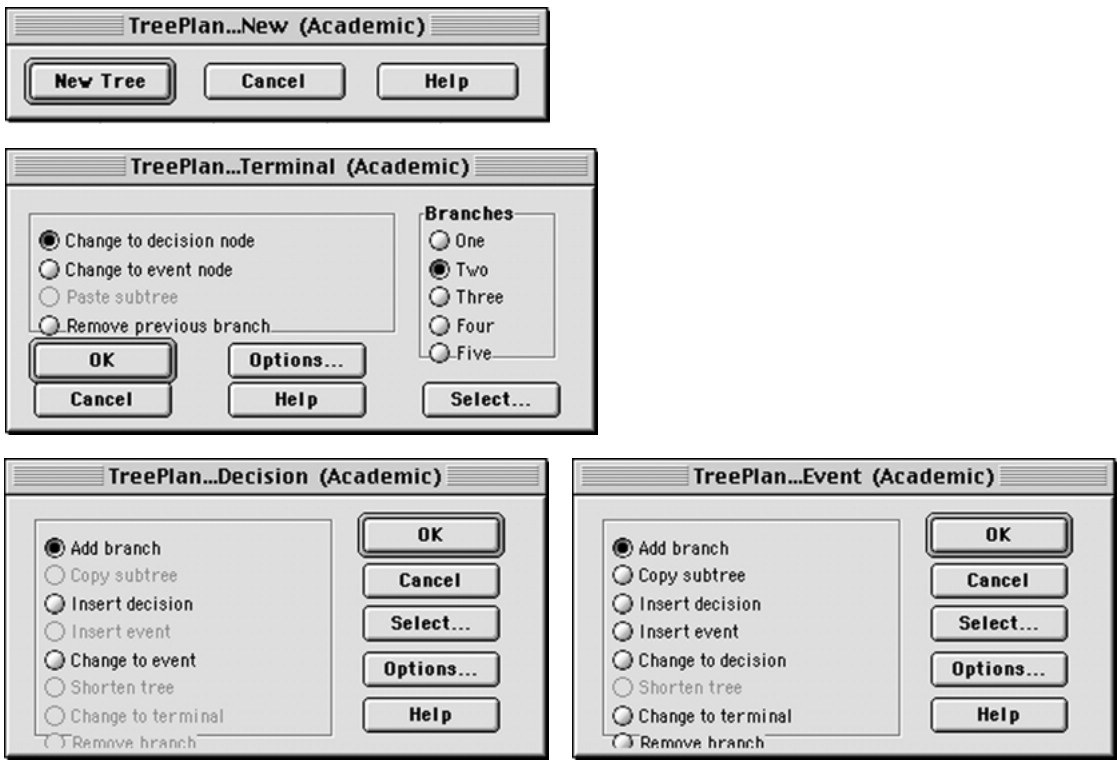

FIGURE 15.12
The dialogue boxes used by TreePlan for constructing a decision tree.

(displayed by a vertical hash mark at the end of a branch) and then choosing Decision Tree from the Tools menu brings up the "TreePlan . . . Terminal" dialogue box (see Fig. 15.12), which enables you to change the terminal fork into either a decision fork or a chance fork with the desired number of branches (between 1 and 5). (TreePlan refers to decision forks as *decision nodes* and to chance forks as *event nodes.*) At any time, you also can click on any existing decision fork (a square) or chance fork (circle) and choose Decision Tree from the Tools menu to bring up the corresponding dialogue box—"TreePlan . . . Decision" or "TreePlan . . . Event"—to make any of the modifications listed in Fig. 15.12 at that fork. To complete the decision tree, the names, cash flows, and probabilities for the various branches are typed directly into the spreadsheet. TreePlan then automatically adds the cash flows to obtain the total cash flows (payoffs) to be shown at the right of each end branch.

15.5 UTILITY THEORY

Thus far, when applying Bayes' decision rule, we have assumed that the expected payoff in *monetary terms* is the appropriate measure of the consequences of taking an action. However, in many situations this assumption is inappropriate.

For example, suppose that an individual is offered the choice of (1) accepting a 50:50 chance of winning \$100,000 or nothing or (2) receiving \$40,000 with certainty. Many people would prefer the \$40,000 even though the expected payoff on the 50:50 chance of winning \$100,000 is \$50,000. A company may be unwilling to invest a large sum of money in a new product even when the expected profit is substantial if there is a risk of losing its investment and thereby becoming bankrupt. People buy insurance even though it is a poor investment from the viewpoint of the expected payoff.

Do these examples invalidate Bayes' decision rule? Fortunately, the answer is no, because there is a way of transforming *monetary values* to an appropriate scale that reflects the decision maker's preferences. This scale is called the *utility function for money.*

Utility Functions for Money

Figure 15.13 shows a typical utility function $u(M)$ for money M. It indicates that an individual having this utility function would value obtaining \$30,000 twice as much as \$10,000 and would value obtaining \$100,000 twice as much as \$30,000. This reflects the fact that

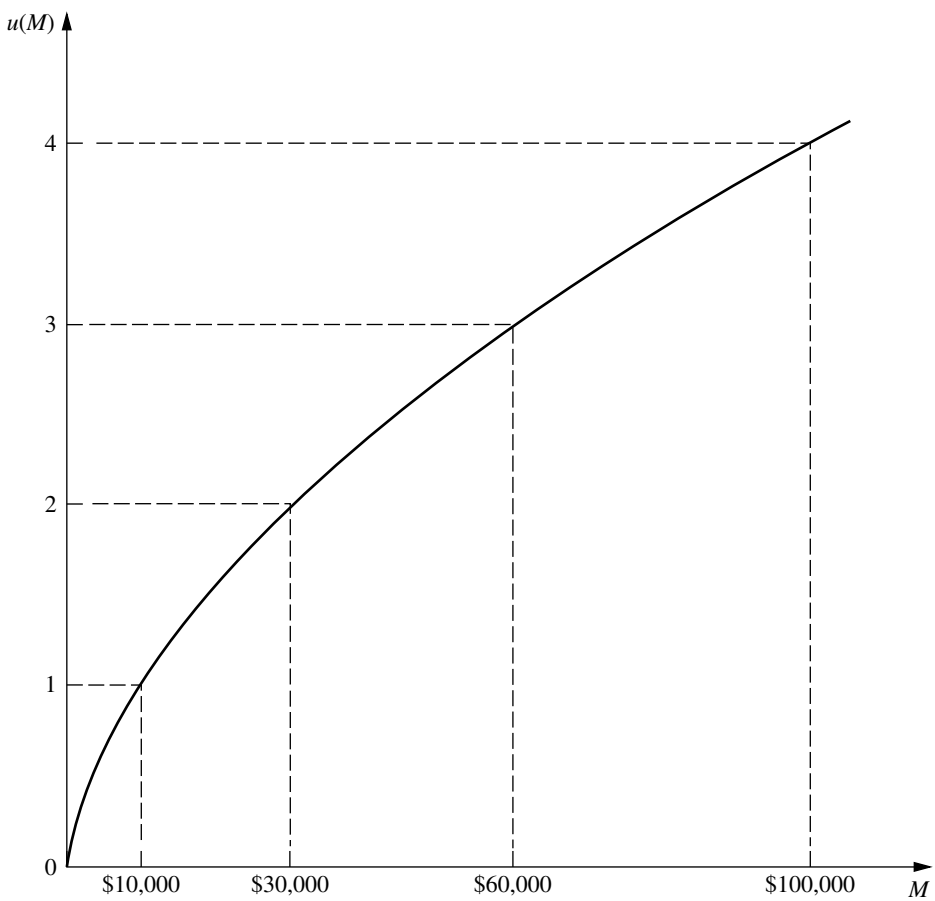

FIGURE 15.13
A typical utility function for money, where $u(M)$ is the utility of obtaining an amount of money M.

the person's highest-priority needs would be met by the first $10,000. Having this decreasing slope of the function as the amount of money increases is referred to as having a **decreasing marginal utility for money.** Such an individual is referred to as being **risk-averse.**

However, not all individuals have a decreasing marginal utility for money. Some people are **risk seekers** instead of *risk-averse,* and they go through life looking for the "big score." The slope of their utility function *increases* as the amount of money increases, so they have an **increasing marginal utility for money.**

The intermediate case is that of a **risk-neutral** individual, who prizes money at its face value. Such an individual's utility for money is simply proportional to the amount of money involved. Although some people appear to be risk-neutral when only small amounts of money are involved, it is unusual to be truly risk-neutral with very large amounts.

It also is possible to exhibit a mixture of these kinds of behavior. For example, an individual might be essentially risk-neutral with small amounts of money, then become a risk seeker with moderate amounts, and then turn risk-averse with large amounts. In addition, one's attitude toward risk can shift over time depending upon circumstances.

An individual's attitude toward risk also may be different when dealing with one's personal finances than when making decisions on behalf of an organization. For example, managers of a business firm need to consider the company's circumstances and the collective philosophy of top management in determining the appropriate attitude toward risk when making managerial decisions.

The fact that different people have different utility functions for money has an important implication for decision making in the face of uncertainty.

> When a *utility function for money* is incorporated into a decision analysis approach to a problem, this utility function must be constructed to fit the preferences and values of the decision maker involved. (The decision maker can be either a single individual or a group of people.)

The key to constructing the utility function for money to fit the decision maker is the following fundamental property of utility functions.

> **Fundamental Property:** Under the assumptions of utility theory, the decision maker's *utility function for money* has the property that the decision maker is *indifferent* between two alternative courses of action if the two alternatives have the *same expected utility.*

To illustrate, suppose that the decision maker has the utility function shown in Fig. 15.13. Further suppose that the decision maker is offered the following opportunity.

> **Offer:** An opportunity to obtain either $100,000 (utility = 4) with probability p or nothing (utility = 0) with probability $(1 - p)$.

Thus,

$E(\text{utility}) = 4p,$ for this offer.

Therefore, for *each* of the following three pairs of alternatives, the decision maker is indifferent between the first and second alternatives:

1. The offer with $p = 0.25$ [$E(\text{utility}) = 1$] or definitely obtaining $10,000 (utility = 1)
2. The offer with $p = 0.5$ [$E(\text{utility}) = 2$] or definitely obtaining $30,000 (utility = 2)
3. The offer with $p = 0.75$ [$E(\text{utility}) = 3$] or definitely obtaining $60,000 (utility = 3)

This example also illustrates one way in which the decision maker's utility function for money can be constructed in the first place. The decision maker would be made the same hypothetical offer to obtain either a large amount of money (for example, \$100,000) with probability p or nothing. Then, for each of a few smaller amounts of money (for example, \$10,000, \$30,000, and \$60,000), the decision maker would be asked to choose a value of p that would make him or her *indifferent* between the offer and definitely obtaining that amount of money. The utility of the smaller amount of money then is p times the utility of the large amount.

The *scale* of the utility function (e.g., utility = 1 for \$10,000) is irrelevant. It is only the *relative values* of the utilities that matter. All the utilities can be multiplied by any positive constant without affecting which alternative course of action will have the largest expected utility.

Now we are ready to summarize the basic role of utility functions in decision analysis.

> When the decision maker's utility function for money is used to measure the relative worth of the various possible monetary outcomes, *Bayes' decision rule* replaces monetary payoffs by the corresponding utilities. Therefore, the optimal action (or series of actions) is the one which *maximizes the expected utility.*

Only utility functions *for money* have been discussed here. However, we should mention that utility functions can sometimes still be constructed when some of or all the important consequences of the alternative courses of action are *not* monetary. (For example, the consequences of a doctor's decision alternatives in treating a patient involve the future health of the patient.) Nevertheless, under these circumstances, it is important to incorporate such value judgments into the decision process. This is not necessarily easy, since it may require making value judgments about the relative desirability of rather intangible consequences. Nevertheless, under these circumstances, it is important to incorporate such value judgments into the decision process.

Applying Utility Theory to the Goferbroke Co. Problem

At the end of Sec. 15.1, we mentioned that the Goferbroke Co. was operating without much capital, so a loss of \$100,000 would be quite serious. The (primary) owner of the company already has gone heavily into debt to keep going. The worst-case scenario would be to come up with \$30,000 for a seismic survey and then still lose \$100,000 by drilling when there is no oil. This scenario would not bankrupt the company at this point, but definitely would leave it in a precarious financial position.

On the other hand, striking oil is an exciting prospect, since earning \$700,000 finally would put the company on a fairly solid financial footing.

To apply the owner's (decision maker's) *utility function for money* to the problem as described in Secs. 15.1 and 15.3, it is necessary to identify the utilities for all the possible monetary payoffs. In units of thousands of dollars, these possible payoffs and the corresponding utilities are given in Table 15.4. We now will discuss how these utilities were obtained.

As a starting point in constructing the utility function, it is natural to let the utility of *zero* money be zero, so $u(0) = 0$. An appropriate next step is to consider the worst scenario and best scenario and then to address the following question.

TABLE 15.4 Utilities for the Goferbroke Co. problem

Monetary Payoff	Utility
−130	−150
−100	−105
60	60
90	90
670	580
700	600

Suppose you have only the following two alternatives. Alternative 1 is to do nothing (payoff and utility = 0). Alternative 2 is to have a probability p of a payoff of 700 and a probability $1 - p$ of a payoff of −130 (loss of 130). What value of p makes you *indifferent* between two alternatives?

The decision maker's choice: $p = \frac{1}{5}$.

If we continue to let $u(M)$ denote the utility of a monetary payoff of M, this choice of p implies that

$$\frac{4}{5}u(-130) + \frac{1}{5}u(700) = 0 \qquad \text{(utility of alternative 1)}.$$

The value of either $u(-130)$ or $u(700)$ can be set arbitrarily (provided only that the first is negative and the second positive) to establish the scale of the utility function. By choosing $u(-130) = -150$ (a convenient choice since it will make $u(M)$ approximately equal to M when M is in the vicinity of 0), this equation then yields $u(700) = 600$.

To identify $u(-100)$, a choice of p is made that makes the decision maker indifferent between a payoff of −130 with probability p or definitely incurring a payoff of −100. The choice is $p = 0.7$, so

$$u(-100) = p\,u(-130) = 0.7(-150) = -105.$$

To obtain $u(90)$, a value of p is selected that makes the decision maker indifferent between a payoff of 700 with probability p or definitely obtaining a payoff of 90. The value chosen is $p = 0.15$, so

$$u(90) = p\,u(700) = 0.15(600) = 90.$$

At this point, a smooth curve was drawn through $u(-130)$, $u(-100)$, $u(90)$, and $u(700)$ to obtain the decision maker's *utility function for money* shown in Fig. 15.14. The values on this curve at $M = 60$ and $M = 670$ provide the corresponding utilities, $u(60) = 60$ and $u(670) = 580$, which completes the list of utilities given in the right column of Table 15.4. For contrast, the dashed line drawn at 45° in Fig. 15.14 shows the monetary value M of the amount of money M. This dashed line has provided the values of the payoffs used exclusively in the preceding sections. Note how $u(M)$ essentially equals M for small values

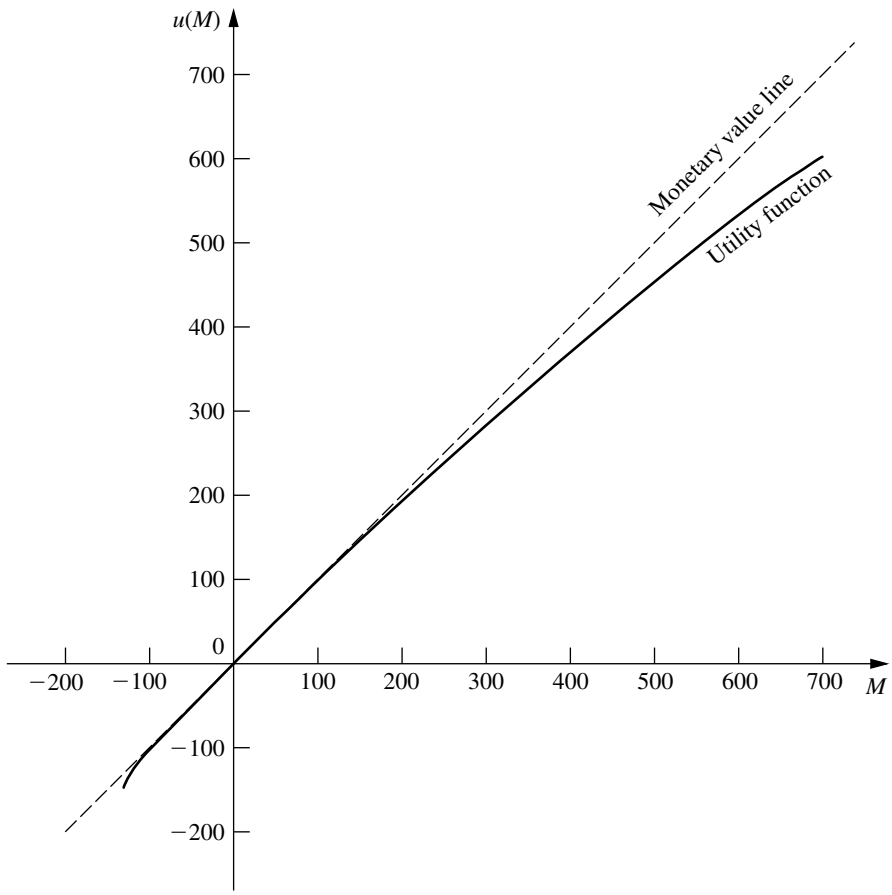

FIGURE 15.14
The utility function for money of the owner of the Goferbroke Co.

(positive or negative) of M, and then how $u(M)$ gradually falls off M for larger values of M. This is typical for a moderately risk-averse individual.

By nature, the owner of the Goferbroke Co. is inclined to be a risk seeker. However, the difficult financial circumstances of his company, which he badly wants to keep solvent, have forced him to adopt a moderately risk-averse stance in addressing his current decisions.

Another Approach for Estimating $u(M)$

The above procedure for constructing $u(M)$ asks the decision maker to repeatedly make a difficult decision about which probability would make him or her indifferent between two alternatives. Many individuals would be uncomfortable with making this kind of decision. Therefore, an alternative approach is sometimes used instead to estimate the utility function for money.

This approach is to assume that the utility function has a certain mathematical form, and then adjust this form to fit the decision maker's attitude toward risk as closely as possible. For example, one particularly popular form to assume (because of its relative simplicity) is the **exponential utility function,**

$$u(M) = R\left(1 - e^{-\frac{M}{R}}\right),$$

where R is the decision maker's *risk tolerance.* This utility function has a decreasing marginal utility for money, so it is designed to fit a *risk-averse* individual. A great aversion to risk corresponds to a small value of R (which would cause the utility function curve to bend sharply), whereas a small aversion to risk corresponds to a large value of R (which gives a much more gradual bend in the curve).

Since the owner of the Goferbroke Co. has a relatively small aversion to risk, the utility function curve in Fig. 15.14 bends quite slowly. The value of R that would give the utilities of $u(670) = 580$ and $u(700) = 600$ is approximately $R = 2{,}250$. On the other hand, the owner becomes much more risk-averse when large losses can occur, since this now would threaten bankruptcy, so the value of R that would give the utility of $u(-130) = -150$ is only about $R = 465$.

Unfortunately, it is not possible to use two different values of R for the same utility function. A drawback of the exponential utility function is that it assumes a constant aversion to risk (a fixed value of R), regardless of how much (or how little) money the decision maker currently has. This doesn't fit the Goferbroke Co. situation, since the current shortage of money makes the owner much more concerned than usual about incurring a large loss.

In other situations where the consequences of the potential losses are not as severe, assuming an exponential utility function may provide a reasonable approximation. In such a case, here is an easy (slightly approximate) way of estimating the appropriate value of R. The decision maker would be asked to choose the number R that would make him (or her) indifferent between the following two alternatives.

A_1: A 50-50 gamble where he would gain R dollars with probability 0.5 and lose $\frac{R}{2}$ dollars with probability 0.5.

A_2: Neither gain nor lose anything.

TreePlan includes the option of using the exponential utility function. All you need to do is click on the Options button in the TreePlan dialogue box and then select "Use Exponential Utility Function." TreePlan uses a different form for the exponential utility function that requires specifying the values of three constants (by choosing Define Name under the Insert menu and entering the values). By choosing the value of R for all three of these constants, this utility function becomes the same as the exponential utility function described above.

Using a Decision Tree to Analyze the Goferbroke Co. Problem with Utilities

Now that the utility function for money of the owner of the Goferbroke Co. has been obtained in Table 15.4 (and Fig. 15.14), this information can be used with a decision tree as summarized next.

> The procedure for using a decision tree to analyze the problem now is *identical* to that described in the preceding section *except* for substituting utilities for monetary payoffs. Therefore, the value obtained to evaluate each fork of the tree now is the *expected utility* there rather than the expected (monetary) payoff. Consequently, the optimal decisions selected by Bayes' decision rule maximize the expected utility for the overall problem.

Thus, our final decision tree shown in Fig. 15.15 closely resembles the one in Fig. 15.11 given in the preceding section. The forks and branches are exactly the same, as are the probabilities for the branches emanating from the chance forks. For informational purposes, the total monetary payoffs still are given to the right of the terminal branches (but we no longer bother to show the individual monetary payoffs next to any of the branches). However, we now have added the utilities on the right side. It is these numbers that have been used to compute the expected utilities given next to all the forks.

These expected utilities lead to the same decisions at forks *a, c,* and *d* as in Fig. 15.11, but the decision at fork *e* now switches to *sell* instead of *drill.* However, the backward induction procedure still leaves fork *e* on a *closed* path. Therefore, the overall optimal policy remains the same as given at the end of Sec. 15.4 (do the seismic survey; sell if the result is unfavorable; drill if the result is favorable).

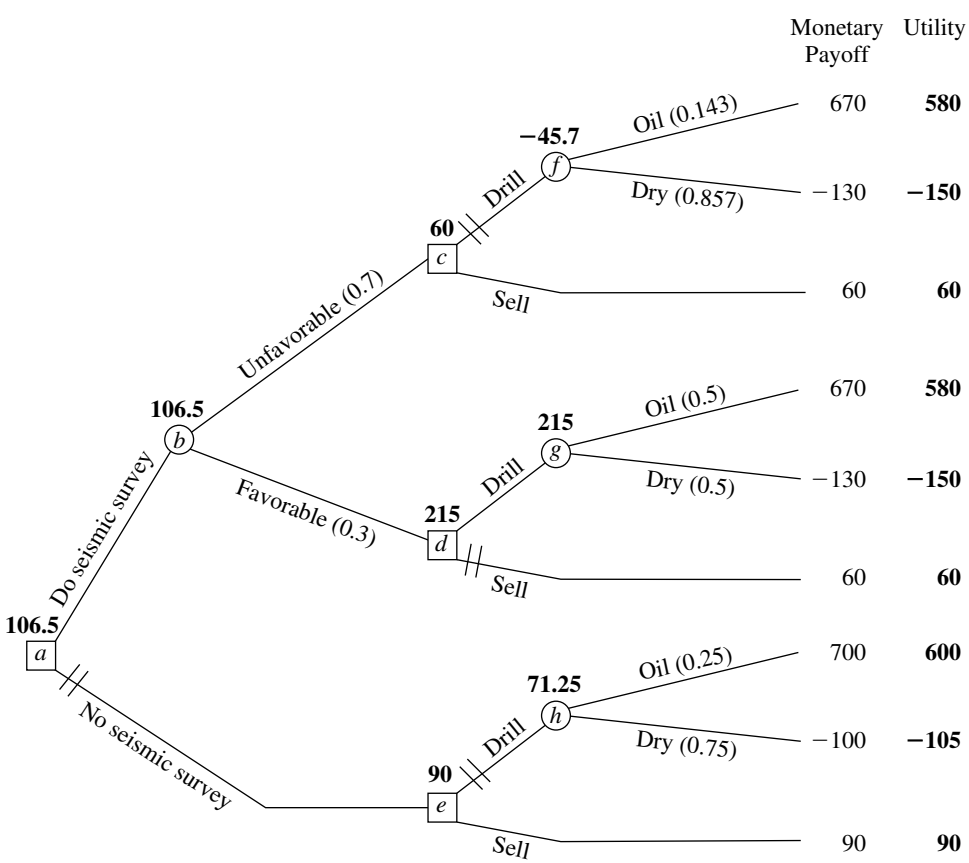

FIGURE 15.15
The final decision tree for the full Goferbroke Co. problem, using the owner's utility function for money to maximize expected utility.

The approach used in the preceding sections of maximizing the expected monetary payoff amounts to assuming that the decision maker is risk-neutral, so that $u(M) = M$. By using utility theory, the optimal solution now reflects the decision maker's attitude about risk. Because the owner of the Goferbroke Co. adopted only a moderately risk-averse stance, the optimal policy did not change from before. For a somewhat more risk-averse owner, the optimal solution would switch to the more conservative approach of immediately selling the land (no seismic survey). (See Prob. 15.5-1.)

The current owner is to be commended for incorporating utility theory into a decision analysis approach to his problem. Utility theory helps to provide a rational approach to decision making in the face of uncertainty. However, many decision makers are not sufficiently comfortable with the relatively abstract notion of utilities, or with working with probabilities to construct a utility function, to be willing to use this approach. Consequently, utility theory is not yet used very widely in practice.

15.6 THE PRACTICAL APPLICATION OF DECISION ANALYSIS

In one sense, this chapter's prototype example (the Goferbroke Co. problem) is a very typical application of decision analysis. Like other applications, management needed to make some decisions (Do a seismic survey? Drill for oil or sell the land?) in the face of great uncertainty. The decisions were difficult because their payoffs were so unpredictable. The outcome depended on factors that were outside management's control (does the land contain oil or is it dry?). Therefore, management needed a framework and methodology for rational decision making in this uncertain environment. These are the usual characteristics of applications of decision analysis.

However, in other ways, the Goferbroke problem is not such a typical application. It was oversimplified to include only two possible states of nature (Oil and Dry), whereas there actually would be a considerable number of distinct possibilities. For example, the actual state might be dry, a small amount of oil, a moderate amount, a large amount, and a huge amount, plus different possibilities concerning the depth of the oil and soil conditions that impact the cost of drilling to reach the oil. Management also was considering only two alternatives for each of two decisions. Real applications commonly involve more decisions, more alternatives to be considered for each one, and many possible states of nature.

When dealing with larger problems, the decision tree can explode in size, with perhaps many thousand terminal branches. In this case, it clearly would not be feasible to construct the tree by hand, including computing posterior probabilities, and calculating the expected payoffs (or utilities) for the various forks, and then identifying the optimal decisions. Fortunately, some excellent software packages (mainly for personal computers) are available specifically for doing this work. Furthermore, special algebraic techniques are being developed and incorporated into the computer solvers for dealing with ever larger problems.[1]

Sensitivity analysis also can become unwieldy on large problems. Although it normally is supported by the computer software, the amount of data generated can easily overwhelm

[1]For example, see C. W. Kirkwood, "An Algebraic Approach to Formulating and Solving Large Models for Sequential Decisions under Uncertainty," *Management Science,* **39:** 900–913, July 1993.

an analyst or decision maker. Therefore, some graphical techniques, such as *tornado diagrams,* have been developed to organize the data in a readily understandable way.[1]

Other kinds of graphical techniques also are available to complement the decision tree in representing and solving decision analysis problems. One that has become quite popular is called the *influence diagram,* and researchers continue to develop others as well.[2]

Many strategic business decisions are made collectively by several members of management. One technique for group decision making is called *decision conferencing*. This is a process where the group comes together for discussions in a decision conference with the help of an analyst and a group facilitator. The facilitator works directly with the group to help it structure and focus discussions, think creatively about the problem, bring assumptions to the surface, and address the full range of issues involved. The analyst uses decision analysis to assist the group in exploring the implications of the various decision alternatives. With the assistance of a computerized group decision support system, the analyst builds and solves models on the spot, and then performs sensitivity analysis to respond to what-if questions from the group.[3]

Applications of decision analysis commonly involve a partnership between the managerial decision maker (whether an individual or a group) and an analyst (whether an individual or a team) with training in OR. Some companies do not have a staff member who is qualified to serve as the analyst. Therefore, a considerable number of management consulting firms specializing in decision analysis have been formed to fill this role. (For example, a few large ones are located in Silicon Valley next to Stanford University, with names such as Applied Decision Analysis and the Strategic Decisions Group.)

Decision analysis is widely used around the world. For proprietary reasons (among others), companies usually do not publish articles in professional journals to describe their applications of OR techniques, including decision analysis. Fortunately, such articles do filter out once in awhile, with some of them appearing in the journal called *Interfaces.* The articles about decision analysis provide valuable insights about the practical application of this technique in practice.

Table 15.5 briefly summarizes the nature of some of the applications of decision analysis that have appeared in *Interfaces.* The rightmost column identifies the specific issue of the journal for each application. Note in the other columns the wide diversity of organizations and applications (with public utilities as the heaviest users). For each specific application, think about how uncertainties in the situation make decision analysis a natural technique to use.

[1]For further information, see T. G. Eschenbach, "Spiderplots versus Tornado Diagrams for Sensitivity Analysis," *Interfaces,* **22:** 40–46, Nov.–Dec. 1992.

[2]For example, see P. P. Schnoy, "A Comparison of Graphical Techniques for Decision Analysis," *European Journal of Operational Research,* **78:** 1–21, Oct. 13, 1994. Also see Z. Covaliu and R. M. Oliver, "Representation and Solution of Decision Problems Using Sequential Decision Diagrams," *Management Science,* **41:** 1860–1881, Dec., 1995, as well as Chaps. 4 and 9 in K. T. Marshall and R. M. Oliver, *Decision Making and Forecasting,* McGraw-Hill, New York, 1995.

[3]For further information, see the two articles on decision conferencing in the November–December 1992 issue of *Interfaces,* where one describes an application in Australia and the other summarizes the experience of 26 decision conferences in Hungary.

TABLE 15.5 Some applications of decision analysis

Organization	Nature of Application	Issue of *Interfaces*
Amoco Oil Co.	Used utilities to evaluate strategies for marketing its products through full-facility service stations.	Dec., 1982
Ohio Edison Co.	Evaluated and selected particulate emission control equipment for a coal-fired power plant.	Feb., 1983
New England Electric System	Determined an appropriate bid for the salvage rights to a grounded ship.	March–April, 1984
National Weather Service	Developed a plan for responding to flood forecasts and warnings.	May–June, 1984
National Forest Administrations	Planned prescribed fires to improve forest and rangeland ecosystems.	Sept.–Oct., 1984
Tomco Oil Corp.	Chose between two site locations for drilling an oil well, with 74 states of nature.	March–April, 1986
Personal decision	Used decision criteria without probabilities to choose between adjustable-rate and fixed-rate mortgages.	May–June, 1986
U.S. Postal Service	Chose between six alternatives for a postal automation program, saving $200 million.	March–April, 1987; Jan.–Feb., 1988
Santa Clara University	Evaluated whether to implement a drug-testing program for their intercollegiate athletes.	May–June, 1990
Independent Living Center (Australia)	A decision conference developed a strategic plan for reorganizing the center.	Nov.–Dec., 1992
DuPont Corp.	Many applications to strategic planning; one added $175 million in value.	Nov.–Dec., 1992
British Columbia Hydro and Power Authority	Elicited a utility function for clarifying value trade-offs for many strategic issues.	Nov.–Dec., 1992
U.S. Department of Defense	Improved the decision process for the acquisition of weapon systems.	Nov.–Dec., 1992
Electric utility industry	Considered health and environmental risks in dealing with utility-generated solid wastes and air emissions.	Nov.–Dec., 1992
An anonymous international bank	Developed a contingency-planning program against fire and power failure for all services.	Nov.–Dec., 1992
General Motors	More than 40 major decision analysis projects over 5 years.	Nov.–Dec., 1992
Southern Company (electric utility)	Evaluated alternative preventive maintenance programs for motor vehicle and construction equipment fleets.	May–June, 1993

Organization	Nature of Application	Issue of *Interfaces*
ICI Americas	Selected research and development projects with little data available for assessing them.	Nov.–Dec., 1993
Federal National Mortgage Association	Used utilities to select the composition of a portfolio of home mortgage assets.	May–June, 1994
Oglethorpe Power Corp.	Evaluated whether to invest in a major transmission system and how to finance it.	March–April, 1995
Phillips Petroleum Co.	Evaluated oil exploration opportunities with a consistent risk-taking policy.	Nov.–Dec., 1995
Energy Electric System	Evaluated schedules for preventive maintenance for electrical generator units.	July–Aug., 1996

If you would like to do more reading about the practical application of decision analysis, a good place to begin would be the November–December 1992 issue of *Interfaces.* This is a special issue devoted entirely to decision analysis and the related area of risk analysis. It includes many interesting articles, including descriptions of basic methods, sensitivity analysis, and decision conferencing. Also included are several of the articles on applications that are listed in Table 15.5.

15.7 CONCLUSIONS

Decision analysis has become an important technique for decision making in the face of uncertainty. It is characterized by enumerating all the available courses of action, identifying the payoffs for all possible outcomes, and quantifying the subjective probabilities for all the possible random events. When these data are available, decision analysis becomes a powerful tool for determining an optimal course of action.

One option that can be readily incorporated into the analysis is to perform experimentation to obtain better estimates of the probabilities of the possible states of nature. Decision trees are a useful visual tool for analyzing this option or any series of decisions.

Utility theory provides a way of incorporating the decision maker's attitude toward risk into the analysis.

Good software (including TreePlan and SensIt in your OR Courseware) is becoming widely available for performing decision analysis.

SELECTED REFERENCES

1. Clemen, R. T.: *Making Hard Decisions: An Introduction to Decision Analysis,* 2d ed., Duxbury Press, Belmont, CA, 1996.

2. Edwards, W. (ed.): *Utility Theories: Measurements and Applications,* Kluwer Academic Publishers, Boston, 1992.

3. Fishburn, P. C.: "Foundations of Decision Analysis: Along the Way," *Management Science,* **35:** 387–405, 1989.

4. Fishburn, P. C.: *Nonlinear Preference and Utility Theory,* The Johns Hopkins Press, Baltimore, MD, 1988.
5. Hillier, F. S., M. S. Hillier, and G. J. Lieberman: *Introduction to Management Science: A Modeling and Case Studies Approach with Spreadsheets,* Irwin/McGraw-Hill, Burr Ridge, IL, 2000, chap. 10.
6. Marshall, K. T., and R. M. Oliver: *Decision Making and Forecasting,* McGraw-Hill, New York, 1995.

LEARNING AIDS FOR THIS CHAPTER IN YOUR OR COURSEWARE



"Ch. 15—Decision Analysis" Excel File:

TreePlan Decision Trees for Goferbroke Problems (3)
Template for Maximin Payoff Criterion
Template for Maximum Likelihood Criterion
Template for Bayes' Decision Rule
Decision Analysis Spreadsheets for Goferbroke Problems (2)
Template for Expected Payoff with Perfect Information
Template for Posterior Probabilities

Excel Add-Ins:

TreePlan (academic version)
SensIt (academic version)

PROBLEMS

The symbols to the left of some of the problems (or their parts) have the following meaning:

T: The corresponding Excel template listed above can be helpful.
A: The corresponding Excel add-in listed above can be used.

An asterisk on the problem number indicates that at least a partial answer is given in the back of the book.

15.2-1.* Silicon Dynamics has developed a new computer chip that will enable it to begin producing and marketing a personal computer if it so desires. Alternatively, it can sell the rights to the computer chip for $15 million. If the company chooses to build computers, the profitability of the venture depends upon the company's ability to market the computer during the first year. It has sufficient access to retail outlets that it can guarantee sales of 10,000 computers. On the other hand, if this computer catches on, the company can sell 100,000 machines. For analysis purposes, these two levels of sales are taken to be the two possible outcomes of marketing the computer, but it is unclear what their prior probabilities are. The cost of setting up the assembly line is $6 million. The difference between the selling price and the variable cost of each computer is $600.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table.
(b) Develop a graph that plots the expected payoff for each of the alternative actions versus the prior probability of selling 10,000 computers.
(c) Referring to the graph developed in part (*b*), use algebra to solve for the *crossover point.* Explain the significance of this point.
A **(d)** Develop a graph that plots the expected payoff (when using Bayes' decision rule) versus the prior probability of selling 10,000 computers.
T **(e)** Assuming the prior probabilities of the two levels of sales are both 0.5, which alternative action should be chosen?

15.2-2. Jean Clark is the manager of the Midtown Saveway Grocery Store. She now needs to replenish her supply of strawberries. Her regular supplier can provide as many cases as she wants. However, because these strawberries already are very ripe, she will need to sell them tomorrow and then discard any that remain unsold. Jean estimates that she will be able to sell 10, 11, 12, or 13 cases tomorrow. She can purchase the strawberries for $3 per case and

sell them for $8 per case. Jean now needs to decide how many cases to purchase.

Jean has checked the store's records on daily sales of strawberries. On this basis, she estimates that the prior probabilities are 0.2, 0.4, 0.3, and 0.1 for being able to sell 10, 11, 12, and 13 cases of strawberries tomorrow.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table.

T **(b)** How many cases of strawberries should Jean purchase if she uses the maximin payoff criterion?

T **(c)** How many cases should be purchased according to the maximum likelihood criterion?

T **(d)** How many cases should be purchased according to Bayes' decision rule?

T **(e)** Jean thinks she has the prior probabilities just about right for selling 10 cases and selling 13 cases, but is uncertain about how to split the prior probabilities for 11 cases and 12 cases. Reapply Bayes' decision rule when the prior probabilities of 11 and 12 cases are (i) 0.2 and 0.5, (ii) 0.3 and 0.4, and (iii) 0.5 and 0.2.

T **15.2-3.*** Warren Buffy is an enormously wealthy investor who has built his fortune through his legendary investing acumen. He currently has been offered three major investments and he would like to choose one. The first one is a *conservative investment* that would perform very well in an improving economy and only suffer a small loss in a worsening economy. The second is a *speculative investment* that would perform extremely well in an improving economy but would do very badly in a worsening economy. The third is a *countercyclical investment* that would lose some money in an improving economy but would perform well in a worsening economy.

Warren believes that there are three possible scenarios over the lives of these potential investments: (1) an improving economy, (2) a stable economy, and (3) a worsening economy. He is pessimistic about where the economy is headed, and so has assigned prior probabilities of 0.1, 0.5, and 0.4, respectively, to these three scenarios. He also estimates that his profits under these respective scenarios are those given by the following table:

	Improving Economy	Stable Economy	Worsening Economy
Conservative investment	$30 million	$ 5 million	−$10 million
Speculative investment	$40 million	$10 million	−$30 million
Countercyclical investment	−$10 million	0	$15 million
Prior probability	0.1	0.5	0.4

Which investment should Warren make under each of the following criteria?

(a) Maximin payoff criterion.

(b) Maximum likelihood criterion.

(c) Bayes' decision rule.

15.2-4. Reconsider Prob. 15.2-3. Warren Buffy decides that Bayes' decision rule is his most reliable decision criterion. He believes that 0.1 is just about right as the prior probability of an improving economy, but is quite uncertain about how to split the remaining probabilities between a stable economy and a worsening economy. Therefore, he now wishes to do sensitivity analysis with respect to these latter two prior probabilities.

T **(a)** Reapply Bayes' decision rule when the prior probability of a stable economy is 0.3 and the prior probability of a worsening economy is 0.6.

T **(b)** Reapply Bayes' decision rule when the prior probability of a stable economy is 0.7 and the prior probability of a worsening economy is 0.2.

(c) Graph the expected profit for each of the three investment alternatives versus the prior probability of a stable economy (with the prior probability of an improving economy fixed at 0.1). Use this graph to identify the crossover points where the decision shifts from one investment to another.

(d) Use algebra to solve for the crossover points identified in part (*c*).

A **(e)** Develop a graph that plots the expected profit (when using Bayes' decision rule) versus the prior probability of a stable economy.

15.2-5.* Consider a decision analysis problem whose payoffs (in units of thousands of dollars) are given by the following payoff table:

	State of Nature	
Alternative	S_1	S_2
A_1	80	25
A_2	30	50
A_3	60	40
Prior probability	0.4	0.6

T **(a)** Which alternative should be chosen under the maximin payoff criterion?

T **(b)** Which alternative should be chosen under the maximum likelihood criterion?

T **(c)** Which alternative should be chosen under Bayes' decision rule?

A **(d)** Using Bayes' decision rule, do sensitivity analysis graphically with respect to the prior probabilities to determine the crossover points where the decision shifts from one alternative to another.

(e) Use algebra to solve for the crossover points identified in part (*d*).

15.2-6. You are given the following payoff table (in units of thousands of dollars) for a decision analysis problem:

	State of Nature		
Alternative	S_1	S_2	S_3
A_1	220	170	110
A_2	200	180	150
Prior probability	0.6	0.3	0.1

T **(a)** Which alternative should be chosen under the maximin payoff criterion?

T **(b)** Which alternative should be chosen under the maximum likelihood criterion?

T **(c)** Which alternative should be chosen under Bayes' decision rule?

(d) Using Bayes' decision rule, do sensitivity analysis graphically with respect to the prior probabilities of states S_1 and S_2 (without changing the prior probability of state S_3) to determine the crossover point where the decision shifts from one alternative to the other. Then use algebra to calculate this crossover point.

(e) Repeat part (*d*) for the prior probabilities of states S_1 and S_3.

(f) Repeat part (*d*) for the prior probabilities of states S_2 and S_3.

(g) If you feel that the true probabilities of the states of nature are within 10 percent of the given prior probabilities, which alternative would you choose?

15.2-7. Dwight Moody is the manager of a large farm with 1,000 acres of arable land. For greater efficiency, Dwight always devotes the farm to growing one crop at a time. He now needs to make a decision on which one of four crops to grow during the upcoming growing season. For each of these crops, Dwight has obtained the following estimates of crop yields and net incomes per bushel under various weather conditions.

	Expected Yield, Bushels/Acre			
Weather	**Crop 1**	**Crop 2**	**Crop 3**	**Crop 4**
Dry	20	15	30	40
Moderate	35	20	25	40
Damp	40	30	25	40
Net income per bushel	\$1.00	\$1.50	\$1.00	\$0.50

After referring to historical meteorological records, Dwight also estimated the following prior probabilities for the weather during the growing season:

Dry	0.3
Moderate	0.5
Damp	0.2

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table.

T **(b)** Use Bayes' decision rule to determine which crop to grow.

T **(c)** Using Bayes' decision rule, do sensitivity analysis with respect to the prior probabilities of moderate weather and damp weather (without changing the prior probability of dry weather) by re-solving when the prior probability of moderate weather is 0.2, 0.3, 0.4, and 0.6.

T **15.2-8.*** A new type of airplane is to be purchased by the Air Force, and the number of spare engines to be ordered must be determined. The Air Force must order these spare engines in batches of five, and it can choose among only 15, 20, or 25 spares. The supplier of these engines has two plants, and the Air Force must make its decision prior to knowing which plant will be used. However, the Air Force knows from past experience that two-thirds of all types of airplane engines are produced in Plant *A,* and only one-third are produced in Plant *B.* The Air Force also knows that the number of spare engines required when production takes place at Plant *A* is approximated by a Poisson distribution with mean $\theta = 21$, whereas the number of spare engines required when production takes place at Plant *B* is approximated by a Poisson distribution with mean $\theta = 24$. The cost of a spare engine purchased now is \$400,000, whereas the cost of a spare engine purchased at a later date is \$900,000. Spares must always be supplied if they are demanded, and unused engines will be scrapped when the airplanes become obsolete. Holding costs and interest are to be neglected. From these data, the total costs (negative payoffs) have been computed as follows:

	State of Nature	
Alternative	$\theta = 21$	$\theta = 24$
Order 15	1.155×10^7	1.414×10^7
Order 20	1.012×10^7	1.207×10^7
Order 25	1.047×10^7	1.135×10^7

Determine the optimal action under Bayes' decision rule.

15.2-9. An individual makes decisions according to Bayes' decision rule. For her current problem, she has constructed the fol-

lowing payoff table, and she now wishes to maximize the expected payoff.

Alternative	State of Nature		
	θ_1	θ_2	θ_3
a_1	$2x$	60	10
a_2	25	40	90
a_3	35	$3x$	30
Prior probability	0.4	0.2	0.4

The value of x currently is 50, but there is an opportunity to increase x by spending some money now.

What is the maximum amount that should be spent to increase x to 75?

15.3-1.* Reconsider Prob. 15.2-1. Management of Silicon Dynamics now is considering doing full-fledged market research at a cost of $1 million to predict which of the two levels of demand is likely to occur. Previous experience indicates that such market research is correct two-thirds of the time.

T **(a)** Find EVPI for this problem.

(b) Does the answer in part (*a*) indicate that it might be worthwhile to perform this market research?

(c) Develop a probability tree diagram to obtain the posterior probabilities of the two levels of demand for each of the two possible outcomes of the market research.

T **(d)** Use the corresponding Excel template to check your answers in part (*c*).

(d) Find EVE. Is it worthwhile to perform the market research?

15.3-2. You are given the following payoff table (in units of thousands of dollars) for a decision analysis problem:

Alternative	State of Nature		
	S_1	S_2	S_3
A_1	4	0	0
A_2	0	2	0
A_3	3	0	1
Prior probability	0.2	0.5	0.3

T **(a)** According to Bayes' decision rule, which alternative should be chosen?

T **(b)** Find EVPI.

(c) You are given the opportunity to spend $1,000 to obtain more information about which state of nature is likely to occur. Given your answer to part (*b*), might it be worthwhile to spend this money?

15.3-3.* Betsy Pitzer makes decisions according to Bayes' decision rule. For her current problem, Betsy has constructed the following payoff table (in units of dollars):

Alternative	State of Nature		
	S_1	S_2	S_3
A_1	50	100	−100
A_2	0	10	−10
A_3	20	40	−40
Prior probability	0.5	0.3	0.2

T **(a)** Which alternative should Betsy choose?

T **(b)** Find EVPI.

(c) What is the most that Betsy should consider paying to obtain more information about which state of nature will occur?

15.3-4. Using Bayes' decision rule, consider the decision analysis problem having the following payoff table (in units of thousands of dollars):

Alternative	State of Nature		
	S_1	S_2	S_3
A_1	−100	10	100
A_2	−10	20	50
A_3	10	10	60
Prior probability	0.2	0.3	0.5

T **(a)** Which alternative should be chosen? What is the resulting expected payoff?

(b) You are offered the opportunity to obtain information which will tell you with certainty whether the first state of nature S_1 will occur. What is the maximum amount you should pay for the information? Assuming you will obtain the information, how should this information be used to choose an alternative? What is the resulting expected payoff (excluding the payment)?

(c) Now repeat part (*b*) if the information offered concerns S_2 instead of S_1.

(d) Now repeat part (*b*) if the information offered concerns S_3 instead of S_1.

T **(e)** Now suppose that the opportunity is offered to provide information which will tell you with certainty which state of nature will occur (perfect information). What is the maximum amount you should pay for the information? Assuming you will obtain the information, how should this information be used to choose an alternative? What is the resulting expected payoff (excluding the payment)?

(f) If you have the opportunity to do some testing that will give you partial additional information (not perfect information) about the state of nature, what is the maximum amount you should consider paying for this information?

15.3-5. Reconsider the Goferbroke Co. prototype example, including its analysis in Sec. 15.3. With the help of a consulting geologist, some historical data have been obtained that provide more precise information on the likelihood of obtaining favorable seismic soundings on similar tracts of land. Specifically, when the land contains oil, favorable seismic soundings are obtained 80 percent of the time. This percentage changes to 40 percent when the land is dry.

(a) Revise Fig. 15.6 to find the new posterior probabilities.

T **(b)** Use the corresponding Excel template to check your answers in part (*a*).

(c) What is the resulting optimal policy?

15.3-6. You are given the following payoff table (in units of dollars):

	State of Nature	
Alternative	S_1	S_2
A_1	400	−100
A_2	0	100
Prior probability	0.4	0.6

You have the option of paying \$100 to have research done to better predict which state of nature will occur. When the true state of nature is S_1, the research will accurately predict S_1 60 percent of the time (but will inaccurately predict S_2 40 percent of the time). When the true state of nature is S_2, the research will accurately predict S_2 80 percent of the time (but will inaccurately predict S_1 20 percent of the time).

T **(a)** Given that the research is not done, use Bayes' decision rule to determine which decision alternative should be chosen.

T **(b)** Find EVPI. Does this answer indicate that it might be worthwhile to do the research?

(c) Given that the research is done, find the joint probability of each of the following pairs of outcomes: (i) the state of nature is S_1 and the research predicts S_1, (ii) the state of nature is S_1 and the research predicts S_2, (iii) the state of nature is S_2 and the research predicts S_1, and (iv) the state of nature is S_2 and the research predicts S_2.

(d) Find the unconditional probability that the research predicts S_1. Also find the unconditional probability that the research predicts S_2

(e) Given that the research is done, use your answers in parts (*c*) and (*d*) to determine the posterior probabilities of the states of nature for each of the two possible predictions of the research.

T **(f)** Use the corresponding Excel template to obtain the answers for part (*e*).

T **(g)** Given that the research predicts S_1, use Bayes' decision rule to determine which decision alternative should be chosen and the resulting expected payoff.

T **(h)** Repeat part (*g*) when the research predicts S_2.

T **(i)** Given that research is done, what is the expected payoff when using Bayes' decision rule?

(j) Use the preceding results to determine the optimal policy regarding whether to do the research and the choice of the decision alternative.

15.3-7. You are given the opportunity to guess whether a coin is fair or two-headed, where the prior probabilities are 0.5 for each of these possibilities. If you are correct, you win \$5; otherwise, you lose \$5. You are also given the option of seeing a demonstration flip of the coin before making your guess. You wish to use Bayes' decision rule to maximize expected profit.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, states of nature, and payoff table.

T **(b)** What is the optimal action, given that you decline the option of seeing a demonstration flip?

T **(c)** Find EVPI.

(d) Use the procedure presented in Sec. 15.3 to calculate the posterior distribution if the demonstration flip is a tail. Do the same if the flip is a head.

T **(e)** Use the corresponding Excel template to confirm your results in part (*d*).

(f) Determine your optimal policy.

(g) Now suppose that you must pay to see the demonstration flip. What is the most that you should be willing to pay?

T **15.3-8.*** Reconsider Prob. 15.2-8. Suppose now that the Air Force knows that a similar type of engine was produced for an earlier version of the type of airplane currently under consideration. The order size for this earlier version was the same as for the current type. Furthermore, the probability distribution of the number of spare engines required, given the plant where production takes place, is believed to be the same for this earlier airplane model and the current one. The engine for the current order will be produced in the same plant as the previous model, although the Air Force does not know which of the two plants this is. The Air Force does

have access to the data on the number of spares actually required for the older version, but the supplier has not revealed the production location.

(a) How much money is it worthwhile to pay for perfect information on which plant will produce these engines?

(b) Assume that the cost of the data on the old airplane model is free and that 30 spares were required. You are given that the probability of 30 spares, given a Poisson distribution with mean θ, is 0.013 for $\theta = 21$ and 0.036 for $\theta = 24$. Find the optimal action under Bayes' decision rule.

15.3-9.* Vincent Cuomo is the credit manager for the Fine Fabrics Mill. He is currently faced with the question of whether to extend $100,000 credit to a potential new customer, a dress manufacturer. Vincent has three categories for the credit-worthiness of a company: poor risk, average risk, and good risk, but he does not know which category fits this potential customer. Experience indicates that 20 percent of companies similar to this dress manufacturer are poor risks, 50 percent are average risks, and 30 percent are good risks. If credit is extended, the expected profit for poor risks is −$15,000, for average risks $10,000, and for good risks $20,000. If credit is not extended, the dress manufacturer will turn to another mill. Vincent is able to consult a credit-rating organization for a fee of $5,000 per company evaluated. For companies whose actual credit record with the mill turns out to fall into each of the three categories, the following table shows the percentages that were given each of the three possible credit evaluations by the credit-rating organization.

	Actual Credit Record		
Credit Evaluation	**Poor**	**Average**	**Good**
Poor	50%	40%	20%
Average	40%	50%	40%
Good	10%	10%	40%

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table when the credit-rating organization is not used.

T **(b)** Assuming the credit-rating organization is not used, use Bayes' decision rule to determine which decision alternative should be chosen.

T **(c)** Find EVPI. Does this answer indicate that consideration should be given to using the credit-rating organization?

(d) Assume now that the credit-rating organization is used. Develop a probability tree diagram to find the posterior probabilities of the respective states of nature for each of the three possible credit evaluations of this potential customer.

T **(e)** Use the corresponding Excel template to obtain the answers for part (*d*).

(f) Determine Vincent's optimal policy.

15.3-10. An athletic league does drug testing of its athletes, 10 percent of whom use drugs. This test, however, is only 95 percent reliable. That is, a drug user will test positive with probability 0.95 and negative with probability 0.05, and a nonuser will test negative with probability 0.95 and positive with probability 0.05.

Develop a probability tree diagram to determine the posterior probability of each of the following outcomes of testing an athlete.

(a) The athlete is a drug user, given that the test is positive.

(b) The athlete is not a drug user, given that the test is positive.

(c) The athlete is a drug user, given that the test is negative.

(d) The athlete is not a drug user, given that the test is negative.

T **(e)** Use the corresponding Excel template to check your answers in the preceding parts.

15.3-11. Management of the Telemore Company is considering developing and marketing a new product. It is estimated to be twice as likely that the product would prove to be successful as unsuccessful. It it were successful, the expected profit would be $1,500,000. If unsuccessful, the expected loss would be $1,800,000. A marketing survey can be conducted at a cost of $300,000 to predict whether the product would be successful. Past experience with such surveys indicates that successful products have been predicted to be successful 80 percent of the time, whereas unsuccessful products have been predicted to be unsuccessful 70 percent of the time.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table when the market survey is not conducted.

T **(b)** Assuming the market survey is not conducted, use Bayes' decision rule to determine which decision alternative should be chosen.

T **(c)** Find EVPI. Does this answer indicate that consideration should be given to conducting the market survey?

T **(d)** Assume now that the market survey is conducted. Find the posterior probabilities of the respective states of nature for each of the two possible predictions from the market survey.

(e) Find the optimal policy regarding whether to conduct the market survey and whether to develop and market the new product.

15.3-12. The Hit-and-Miss Manufacturing Company produces items that have a probability p of being defective. These items are produced in lots of 150. Past experience indicates that p for an entire lot is either 0.05 or 0.25. Furthermore, in 80 percent of the lots produced, p equals 0.05 (so p equals 0.25 in 20 percent of the lots). These items are then used in an assembly, and ultimately their quality is determined before the final assembly leaves the plant. Initially the company can *either* screen each item in a lot at a cost of

$10 per item and replace defective items *or* use the items directly without screening. If the latter action is chosen, the cost of rework is ultimately $100 per defective item. Because screening requires scheduling of inspectors and equipment, the decision to screen or not screen must be made 2 days before the screening is to take place. However, one item can be taken from the lot and sent to a laboratory for inspection, and its quality (defective or nondefective) can be reported before the screen/no screen decision must be made. The cost of this initial inspection is $125.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table if the single item is not inspected in advance.

T **(b)** Assuming the single item is not inspected in advance, use Bayes' decision rule to determine which decision alternative should be chosen.

T **(c)** Find EVPI. Does this answer indicate that consideration should be given to inspecting the single item in advance?

T **(d)** Assume now that the single item is inspected in advance. Find the posterior probabilities of the respective states of nature for each of the two possible outcomes of this inspection.

(e) Find EVE. Is inspecting the single item worthwhile?

(f) Determine the optimal policy.

T **15.3-13.*** Consider two weighted coins. Coin 1 has a probability of 0.3 of turning up heads, and coin 2 has a probability of 0.6 of turning up heads. A coin is tossed once; the probability that coin 1 is tossed is 0.6, and the probability that coin 2 is tossed is 0.4. The decision maker uses Bayes' decision rule to decide which coin is tossed. The payoff table is as follows:

	State of Nature	
Alternative	**Coin 1 Tossed**	**Coin 2 Tossed**
Say coin 1 tossed	0	−1
Say coin 2 tossed	−1	0
Prior probability	0.6	0.4

(a) What is the optimal action before the coin is tossed?

(b) What is the optimal action after the coin is tossed if the outcome is heads? If it is tails?

15.3-14. A new type of photographic film has been developed. It is packaged in sets of five sheets, where each sheet provides an instantaneous snapshot. Because this process is new, the manufacturer has attached an additional sheet to the package, so that the store may test one sheet before it sells the package of five. In promoting the film, the manufacturer offers to refund the entire purchase price of the film if one of the five is defective. This refund must be paid by the camera store, and the selling price has been fixed at $2 if this guarantee is to be valid. The camera store may sell the film for $1 if the preceding guarantee is replaced by one that pays $0.20 for each defective sheet. The cost of the film to the camera store is $0.40, and the film is not returnable. The store may choose any one of three actions:

1. Scrap the film.

2. Sell the film for $2.

3. Sell the film for $1.

(a) If the six states of nature correspond to 0, 1, 2, 3, 4, and 5 defective sheets in the package, complete the following payoff table:

	State of Nature					
Alternative	**0**	**1**	**2**	**3**	**4**	**5**
1	−0.40					
2	1.60		−0.40			
3	0.60	0.40		0.00		

T **(b)** The store has accumulated the following information on sales of 60 such packages:

	Defectives in Package					
Quality of Attached Sheet	**0**	**1**	**2**	**3**	**4**	**5**
Good	10	8	6	4	2	0
Bad	0	2	4	6	8	10
Total	10	10	10	10	10	10

These data indicate that each state of nature is equally likely, so that this prior distribution can be assumed. What is the optimal action under Bayes' decision rule (before the attached sheet is tested) for a package of film?

(c) Now assume that the attached sheet is tested. Use a probability tree diagram to find the posterior probabilities of the state of nature for each of the two possible outcomes of this testing.

T **(d)** What is the optimal expected payoff for a package of film if the attached sheet is tested? What is the optimal action if the sheet is good? If it is bad?

15.3-15. There are two biased coins with probabilities of landing heads of 0.8 and 0.4, respectively. One coin is chosen at random (each with probability $\frac{1}{2}$) to be tossed twice. You are to receive $100 if you correctly predict how many heads will occur in two tosses.

T **(a)** Using Bayes' decision rule, what is the optimal prediction, and what is the corresponding expected payoff?

T **(b)** Suppose now that you may observe a practice toss of the

chosen coin before predicting. Use the corresponding Excel template to find the posterior probabilities for which coin is being tossed.

T **(c)** Determine your optimal prediction after observing the practice toss. What is the resulting expected payoff?

(d) Find EVE for observing the practice toss. If you must pay \$30 to observe the practice toss, what is your optimal policy?

15.4-1.* Reconsider Prob. 15.3-1. The management of Silicon Dynamics now wants to see a decision tree displaying the entire problem.

(a) Construct and solve this decision tree by hand.

A **(b)** Use TreePlan to construct and solve this decision tree.

15.4-2. You are given the decision tree to the right, where the numbers in parentheses are probabilities and the numbers on the far right are payoffs at these terminal points.

(a) Analyze this decision tree to obtain the optimal policy.

A **(b)** Use TreePlan to construct and solve the same decision tree.

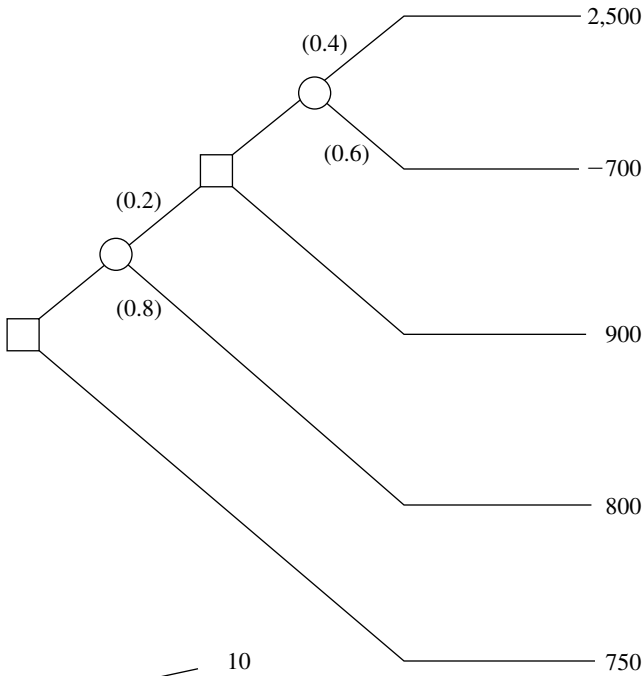

15.4-3. You are given the decision tree below, with the probabilities at chance forks shown in parentheses and with the payoffs at terminal points shown on the right. Analyze this decision tree to obtain the optimal policy.

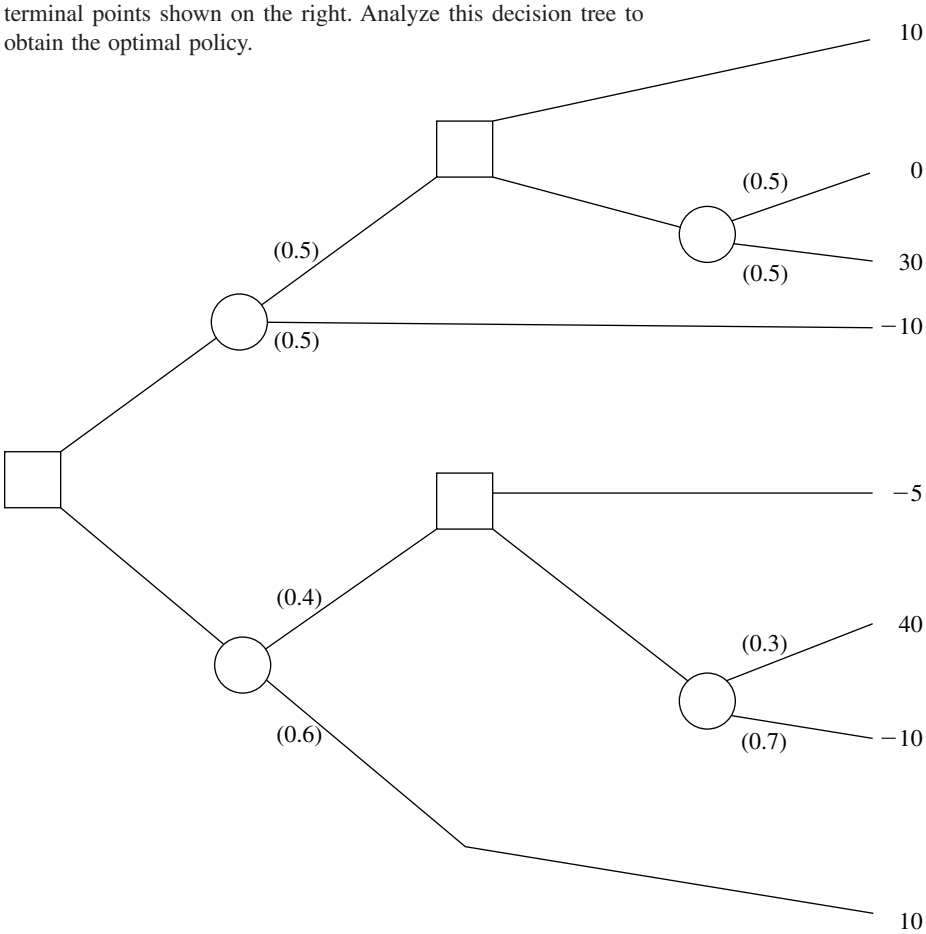

15.4-4.* The Athletic Department of Leland University is considering whether to hold an extensive campaign next year to raise funds for a new athletic field. The response to the campaign depends heavily upon the success of the football team this fall. In the past, the football team has had winning seasons 60 percent of the time. If the football team has a winning season (W) this fall, then many of the alumnae and alumni will contribute and the campaign will raise $3 million. If the team has a losing season (L), few will contribute and the campaign will lose $2 million. If no campaign is undertaken, no costs are incurred. On September 1, just before the football season begins, the Athletic Department needs to make its decision about whether to hold the campaign next year.

(a) Develop a decision analysis formulation of this problem by identifying the alternative actions, the states of nature, and the payoff table.

T **(b)** According to Bayes' decision rule, should the campaign be undertaken?

T **(c)** What is EVPI?

(d) A famous football guru, William Walsh, has offered his services to help evaluate whether the team will have a winning season. For $100,000, he will carefully evaluate the team throughout spring practice and then throughout preseason workouts. William then will provide his prediction on September 1 regarding what kind of season, W or L, the team will have. In similar situations in the past when evaluating teams that have winning seasons 50 percent of the time, his predictions have been correct 75 percent of the time. Considering that this team has more of a winning tradition, if William predicts a winning season, what is the posterior probability that the team actually will have a winning season? What is the posterior probability of a losing season? If Williams predicts a losing season instead, what is the posterior probability of a winning season? Of a losing season? Show how these answers are obtained from a probability tree diagram.

T **(e)** Use the corresponding Excel template to obtain the answers requested in part (*d*).

(f) Draw the decision tree for this entire problem by hand. Analyze this decision tree to determine the optimal policy regarding whether to hire William and whether to undertake the campaign.

A **(g)** Use TreePlan to construct and solve this decision tree.

15.4-5. The comptroller of the Macrosoft Corporation has $100 million of excess funds to invest. She has been instructed to invest the entire amount for 1 year in either stocks or bonds (but not both) and then to reinvest the entire fund in either stocks or bonds (but not both) for 1 year more. The objective is to maximize the expected monetary value of the fund at the end of the second year.

The annual rates of return on these investments depend on the economic environment, as shown in the following table:

	Rate of Return	
Economic Environment	**Stocks**	**Bonds**
Growth	20%	5%
Recession	−10%	10%
Depression	−50%	20%

The probabilities of growth, recession, and depression for the first year are 0.7, 0.3, and 0, respectively. If growth occurs in the first year, these probabilities remain the same for the second year. However, if a recession occurs in the first year, these probabilities change to 0.2, 0.7, and 0.1, respectively, for the second year.

(a) Construct the decision tree for this problem by hand.

(b) Analyze the decision tree to identify the optimal policy.

A **(c)** Use TreePlan to construct and solve the decision tree.

15.4-6. On Monday, a certain stock closed at $10 per share. On Tuesday, you expect the stock to close at $9, $10, or $11 per share, with respective probabilities 0.3, 0.3, and 0.4. On Wednesday, you expect the stock to close 10 percent lower, unchanged, or 10 percent higher than Tuesday's close, with the following probabilities:

Today's Close	10% Lower	Unchanged	10% Higher
$ 9	0.4	0.3	0.3
$10	0.2	0.2	0.6
$11	0.1	0.2	0.7

On Tuesday, you are directed to buy 100 shares of the stock before Thursday. All purchases are made at the end of the day, at the known closing price for that day, so your only options are to buy at the end of Tuesday or at the end of Wednesday. You wish to determine the optimal strategy for whether to buy on Tuesday or defer the purchase until Wednesday, given the Tuesday closing price, to minimize the expected purchase price.

(a) Develop and evaluate a decision tree by hand for determining the optimal strategy.

A **(b)** Use TreePlan to construct and solve the decision tree.

15.4-7. Use the scenario given in Prob. 15.3-7.

(a) Draw and properly label the decision tree. Include all the payoffs but not the probabilities.

T **(b)** Find the probabilities for the branches emanating from the chance forks.

(c) Apply the backward induction procedure, and identify the resulting optimal policy.

A **15.4-8.** Jose Morales manages a large outdoor fruit stand in one of the less affluent neighborhoods of San Jose, California. To replenish his supply, Jose buys boxes of fruit early each morning from a grower south of San Jose. About 90 percent of the boxes of fruit turn out to be of satisfactory quality, but the other 10 percent are unsatisfactory. A satisfactory box contains 80 percent excellent fruit and will earn $200 profit for Jose. An unsatisfactory box contains 30 percent excellent fruit and will produce a loss of $1,000. Before Jose decides to accept a box, he is given the opportunity to sample one piece of fruit to test whether it is excellent. Based on that sample, he then has the option of rejecting the box without paying for it. Jose wonders (1) whether he should continue buying from this grower, (2) if so, whether it is worthwhile sampling just one piece of fruit from a box, and (3) if so, whether he should be accepting or rejecting the box based on the outcome of this sampling.

Use TreePlan (and the Excel template for posterior probabilities) to construct and solve the decision tree for this problem.

15.4-9. Use the scenario given in Prob. 15.3-9.

(a) Draw and properly label the decision tree. Include all the payoffs but not the probabilities.

T **(b)** Find the probabilities for the branches emanating from the chance forks.

(c) Apply the backward induction procedure, and identify the resulting optimal policy.

15.4-10.* The Morton Ward Company is considering the introduction of a new product that is believed to have a 50-50 chance of being successful. One option is to try out the product in a test market, at a cost of $5 million, before making the introduction decision. Past experience shows that ultimately successful products are approved in the test market 80 percent of the time, whereas ultimately unsuccessful products are approved in the test market only 25 percent of the time. If the product is successful, the net profit to the company will be $40 million; if unsuccessful, the net loss will be $15 million.

T **(a)** Discarding the option of trying out the product in a test market, develop a decision analysis formulation of the problem by identifying the alternative actions, states of nature, and payoff table. Then apply Bayes' decision rule to determine the optimal decision alternative.

T **(b)** Find EVPI.

A **(c)** Now including the option of trying out the product in a test market, use TreePlan (and the Excel template for posterior probabilities) to construct and solve the decision tree for this problem.

A **(d)** There is some uncertainty in the stated profit and loss figures ($40 million and $15 million). Either could vary from its base by as much as 25 percent in either direction. Use SensIt to generate a graph for each that plots the expected payoff over this range of variability.

15.4-11. Use the scenario given in Prob. 15.3.-11.

(a) Draw and properly label the decision tree. Include all the payoffs but not the probabilities.

T **(b)** Find the probabilities for the branches emanating from the chance forks.

(c) Apply the backward induction procedure, and identify the resulting optimal policy.

15.4-12. Use the scenario given in Prob. 15.3-12.

(a) Draw and properly label the decision tree. Include all the payoffs but not the probabilities.

T **(b)** Find the probabilities for the branches emanating from the chance forks.

(c) Apply the backward induction procedure, and identify the resulting optimal policy.

15.4-13. Use the scenario given in Prob. 15.3-13.

(a) Draw and properly label the decision tree. Include all the payoffs but not the probabilities.

T **(b)** Find the probabilities for the branches emanating from the chance forks.

(c) Apply the backward induction procedure, and identify the resulting optimal policy.

15.4-14. Chelsea Bush is an emerging candidate for her party's nomination for President of the United States. She now is considering whether to run in the high-stakes Super Tuesday primaries. If she enters the Super Tuesday (S.T.) primaries, she and her advisers believe that she will either do well (finish first or second) or do poorly (finish third or worse) with probabilities 0.4 and 0.6, respectively. Doing well on Super Tuesday will net the candidate's campaign approximately $16 million in new contributions, whereas a poor showing will mean a loss of $10 million after numerous TV ads are paid for. Alternatively, she may choose not to run at all on Super Tuesday and incur no costs.

Chelsea's advisers realize that her chances of success on Super Tuesday may be affected by the outcome of the smaller New Hampshire (N.H.) primary occurring 3 weeks before Super Tuesday. Political analysts feel that the results of New Hampshire's primary are correct two-thirds of the time in predicting the results of the Super Tuesday primaries. Among Chelsea's advisers is a decision analysis expert who uses this information to calculate the following probabilities:

$$P\{\text{Chelsea does well in S.T. primaries, given she does well in N.H.}\} = \frac{4}{7}$$

$$P\{\text{Chelsea does well in S.T. primaries, given she does poorly in N.H.}\} = \frac{1}{4}$$

$$P\{\text{Chelsea does well in N.H. primary}\} = \frac{7}{15}$$

The cost of entering and campaigning in the New Hampshire primary is estimated to be $1.6 million.

Chelsea feels that her chance of winning the nomination depends largely on having substantial funds available after the Super Tuesday primaries to carry on a vigorous campaign the rest of the way. Therefore, she wants to choose the strategy (whether to run in the New Hampshire primary and then whether to run in the Super Tuesday primaries) that will maximize her expected funds after these primaries.

(a) Construct and solve the decision tree for this problem.

A **(b)** There is some uncertainty in the estimates of a gain of $16 million or a loss of $10 million depending on the showing on Super Tuesday. Either amount could differ from this estimate by as much as 25 percent in either direction. Develop a graph for each amount that plots the expected payoff over this range of variability.

A **15.4-15.** The executive search being conducted for Western Bank by Headhunters Inc. may finally be bearing fruit. The position to be filled is a key one—Vice President for Information Processing—because this person will have responsibility for developing a state-of-the-art management information system that will link together Western's many branch banks. However, Headhunters feels they have found just the right person, Matthew Fenton, who has an excellent record in a similar position for a midsized bank in New York.

After a round of interviews, Western's president believes that Matthew has a probability of 0.7 of designing the management information system successfully. If Matthew is successful, the company will realize a profit of $2 million (net of Matthew's salary, training, recruiting costs, and expenses). If he is not successful, the company will realize a net loss of $400,000.

For an additional fee of $20,000, Headhunters will provide a detailed investigative process (including an extensive background check, a battery of academic and psychological tests, etc.) that will further pinpoint Matthew's potential for success. This process has been found to be 90 percent reliable; i.e., a candidate who would successfully design the management information system will pass the test with probability 0.9, and a candidate who would not successfully design the system will fail the test with probability 0.9.

Western's top management needs to decide whether to hire Matthew and whether to have Headhunters conduct the detailed investigative process before making this decision.

(a) Construct the decision tree for this problem.

T **(b)** Find the probabilities for the branches emanating from the chance nodes.

(c) Analyze the decision tree to identify the optimal policy.

(d) Now suppose that the Headhunters' fee for administering its detailed investigative process is negotiable. What is the maximum amount that Western Bank should pay?

15.5-1. Reconsider the Goferbroke Co. prototype example, including the application of utilities in Sec. 15.5. The owner now has decided that, given the company's precarious financial situation, he needs to take a much more risk-averse approach to the problem. Therefore, he has revised the utilities given in Table 15.4 as follows: $u(-130) = -200$, $u(-100) = -130$, $u(60) = 60$, $u(90) = 90$, $u(670) = 440$, and $u(700) = 450$.

(a) Analyze the revised decision tree corresponding to Fig. 15.15 by hand to obtain the new optimal policy.

A **(b)** Use TreePlan to construct and solve this revised decision tree.

15.5-2.* You live in an area that has a possibility of incurring a massive earthquake, so you are considering buying earthquake insurance on your home at an annual cost of $180. The probability of an earthquake damaging your home during 1 year is 0.001. If this happens, you estimate that the cost of the damage (fully covered by earthquake insurance) will be $160,000. Your total assets (including your home) are worth $250,000.

T **(a)** Apply Bayes' decision rule to determine which alternative (take the insurance or not) maximizes your expected assets after 1 year.

(b) You now have constructed a utility function that measures how much you value having total assets worth x dollars ($x \geq 0$). This utility function is $u(x) = \sqrt{x}$. Compare the utility of reducing your total assets next year by the cost of the earthquake insurance with the expected utility next year of not taking the earthquake insurance. Should you take the insurance?

15.5-3. For your graduation present from college, your parents are offering you your choice of two alternatives The first alternative is to give you a money gift of $19,000. The second alternative is to make an investment in your name. This investment will quickly have the following two possible outcomes:

Outcome	Probability
Receive $10,000	0.3
Receive $30,000	0.7

Your utility for receiving M thousand dollars is given by the utility function $u(M) = \sqrt{M + 6}$. Which choice should you make to maximize expected utility?

15.5-4.* Reconsider Prob. 15.5-3. You now are uncertain about what your true utility function for receiving money is, so you are in the process of constructing this utility function. So far, you have found that $u(19) = 16.7$ and $u(30) = 20$ are the utility of receiving $19,000 and $30,000, respectively. You also have concluded that you are indifferent between the two alternatives offered to you by your parents. Use this information to find $u(10)$.

15.5-5. You wish to construct your personal utility function $u(M)$ for receiving M thousand dollars. After setting $u(0) = 0$, you next set $u(1) = 1$ as your utility for receiving \$1,000. You next want to find $u(10)$ and then $u(5)$.

(a) You offer yourself the following two hypothetical alternatives:

A_1: Obtain \$10,000 with probability p.
Obtain 0 with probability $(1 - p)$.
A_2: Definitely obtain \$1,000.

You then ask yourself the question: What value of p makes you indifferent between these two alternatives? Your answer is $p = 0.125$. Find $u(10)$.

(b) You next repeat part (*a*) except for changing the second alternative to definitely receiving \$5,000. The value of p that makes you indifferent between these two alternatives now is $p = 0.5625$. Find $u(5)$.

(c) Repeat parts (*a*) and (*b*), but now use *your* personal choices for p.

15.5-6. You are given the following payoff table:

	State of Nature	
Alternative	S_1	S_2
A_1	25	36
A_2	100	0
A_3	0	49
Prior probability	p	$1 - p$

(a) Assume that your utility function for the payoffs is $u(x) = \sqrt{x}$. Plot the expected utility of each alternative action versus the value of p on the same graph. For each alternative action, find the range of values of p over which this alternative maximizes the expected utility.

A **(b)** Now assume that your utility function is the exponential utility function with a risk tolerance of $R = 5$. Use TreePlan to construct and solve the resulting decision tree in turn for $p = 0.25$, $p = 0.5$, and $p = 0.75$.

15.5-7. Dr. Switzer has a seriously ill patient but has had trouble diagnosing the specific cause of the illness. The doctor now has narrowed the cause down to two alternatives: disease *A* or disease *B*. Based on the evidence so far, she feels that the two alternatives are equally likely.

Beyond the testing already done, there is no test available to determine if the cause is disease *B*. One test is available for disease *A*, but it has two major problems. First, it is very expensive. Second, it is somewhat unreliable, giving an accurate result only 80 percent of the time. Thus, it will give a positive result (indicating disease *A*) for only 80 percent of patients who have disease *A*, whereas it will give a positive result for 20 percent of patients who actually have disease *B* instead.

Disease *B* is a very serious disease with no known treatment. It is sometimes fatal, and those who survive remain in poor health with a poor quality of life thereafter. The prognosis is similar for victims of disease *A* if it is left untreated. However, there is a fairly expensive treatment available that eliminates the danger for those with disease *A*, and it may return them to good health. Unfortunately, it is a relatively radical treatment that always leads to death if the patient actually has disease *B* instead.

The probability distribution for the prognosis for this patient is given for each case in the following table, where the column headings (after the first one) indicate the disease for the patient.

	Outcome Probabilities			
	No Treatment		Receive Treatment for Disease *A*	
Outcome	***A***	***B***	***A***	***B***
Die	0.2	0.5	0	1.0
Survive with poor health	0.8	0.5	0.5	0
Return to good health	0	0	0.5	0

The patient has assigned the following utilities to the possible outcomes:

Outcome	**Utility**
Die	0
Survive with poor health	10
Return to good health	30

In addition, these utilities should be incremented by -2 if the patient incurs the cost of the test for disease *A* and by -1 if the patient (or the patient's estate) incurs the cost of the treatment for disease *A*.

Use decision analysis with a complete decision tree to determine if the patient should undergo the test for disease *A* and then how to proceed (receive the treatment for disease *A*?) to maximize the patient's expected utility.

15.5-8. Consider the following decision tree, where the probabilities for each chance fork are shown in parentheses.

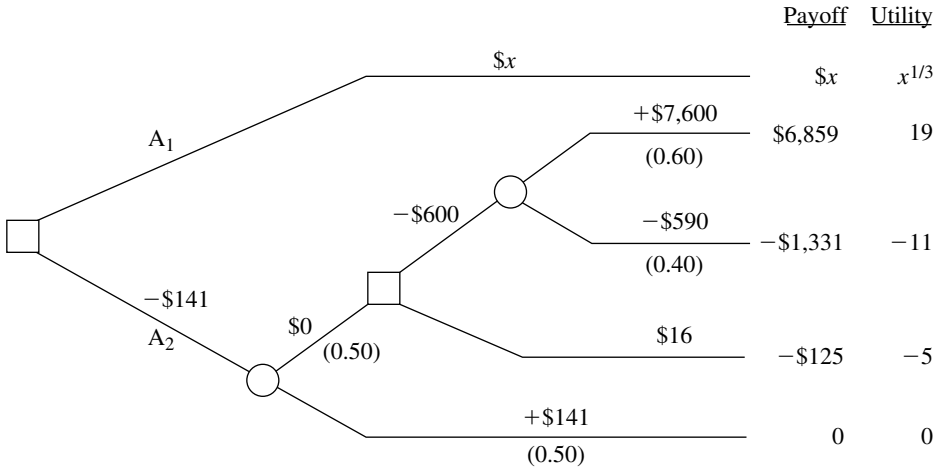

The dollar amount given next to each branch is the cash flow generated along that branch, where these intermediate cash flows add up to the total net cash flow shown to the right of each terminal branch. (The unknown amount for the top branch is represented by the variable x.) The decision maker has a utility function $u(y) = y^{\frac{1}{3}}$ where y is the total net cash flow after a terminal branch. The resulting utilities for the various terminal branches are shown to the right of the decision tree.

Use these utilities to analyze the decision tree. Then determine the value of x for which the decision maker is indifferent between alternative actions A_1 and A_2.

15.5-9. You want to choose between actions A_1 and A_2 in the following decision tree, but you are uncertain about the value of the probability p, so you need to perform sensitivity analysis of p as well.

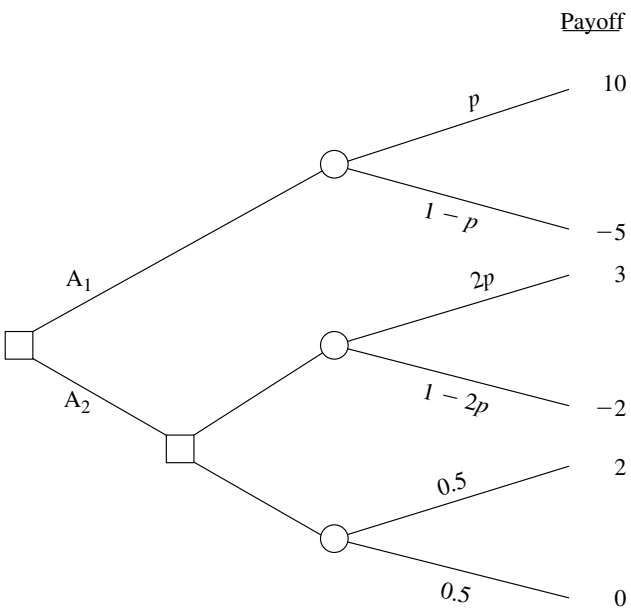

Your utility function for money (the payoff received) is

$$u(M) = \begin{cases} M^2 & \text{if } M \geq 0 \\ M & \text{if } M < 0. \end{cases}$$

(a) For $p = 0.25$, determine which action is optimal in the sense that it maximizes the expected utility of the payoff.

(b) Determine the range of values of the probability p ($0 \leq p \leq 0.5$) for which this same action remains optimal.

15.6-1. Select one of the applications of decision analysis listed in Table 15.5. Read the article describing the application in the indicated issue of *Interfaces*. Write a two-page summary of the application and the benefits it provided.

15.6-2. Select three of the applications of decision analysis listed in Table 15.5. Read the articles describing the applications in the indicated issues of *Interfaces*. For each one, write a one-page summary of the application and the benefits it provided.

CASE 15.1 BRAINY BUSINESS

While El Niño is pouring its rain on northern California, Charlotte Rothstein, CEO, major shareholder and founder of Cerebrosoft, sits in her office, contemplating the decision she faces regarding her company's newest proposed product, Brainet. This has been a particularly difficult decision. Brainet might catch on and sell very well. However, Charlotte is concerned about the risk involved. In this competitive market, marketing Brainet also could lead to substantial losses. Should she go ahead anyway and start the marketing campaign? Or just abandon the product? Or perhaps buy additional marketing research information from a local market research company before deciding whether to launch the product? She has to make a decision very soon and so, as she slowly drinks from her glass of high protein-power multivitamin juice, she reflects on the events of the past few years.

Cerebrosoft was founded by Charlotte and two friends after they had graduated from business school. The company is located in the heart of Silicon Valley. Charlotte and her friends managed to make money in their second year in business and continued to do so every year since. Cerebrosoft was one of the first companies to sell software over the World Wide Web and to develop PC-based software tools for the multimedia sector. Two of the products generate 80 percent of the company's revenues: Audiatur and Videatur. Each product has sold more than 100,000 units during the past year. Business is done over the Web: customers can download a trial version of the software, test it, and if they are satisfied with what they see, they can purchase the product (by using a password that enables them to disable the time counter in the trial version). Both products are priced at $75.95 and are exclusively sold over the Web.

Although the World Wide Web is a network of computers of different types, running different kinds of software, a standardized protocol between the computers enables them to communicate. Users can "surf" the Web and visit computers many thousand miles away, accessing information available at the site. Users can also make files available on the Web, and this is how Cerebrosoft generates its sales. Selling software over the Web eliminates many of the traditional cost factors of consumer products: packaging, storage, distribution, sales force, etc. Instead, potential customers can download a trial version, take a look at it (that is, use the product) before its trial period expires, and then decide whether to buy it. Furthermore, Cerebrosoft can always make the most recent files available to the customer, avoiding the problem of having outdated software in the distribution pipeline.

Charlotte is interrupted in her thoughts by the arrival of Jeannie Korn. Jeannie is in charge of marketing for on-line products and Brainet has had her particular attention from the beginning. She is more than ready to provide the advice that Charlotte has requested. "Charlotte, I think we should really go ahead with Brainet. The software engineers have convinced me that the current version is robust and we want to be on the market with this as soon as possible! From the data for our product launches during the past two years we can get a rather reliable estimate of how the market will respond to the new product, don't you think? And look!" She pulls out some presentation slides. "During that time period we launched 12 new products altogether and 4 of them sold more than 30,000 units during the first 6 months alone! Even better: the

last two we launched even sold more than 40,000 copies during the first two quarters!" Charlotte knows these numbers as well as Jeannie does. After all, two of these launches have been products she herself helped to develop. But she feels uneasy about this particular product launch. The company has grown rapidly during the past three years and its financial capabilities are already rather stretched. A poor product launch for Brainet would cost the company a lot of money, something that isn't available right now due to the investments Cerebrosoft has recently made.

Later in the afternoon, Charlotte meets with Reggie Ruffin, a jack-of-all-trades and the production manager. Reggie has a solid track record in his field and Charlotte wants his opinion on the Brainet project.

"Well, Charlotte, quite frankly I think that there are three main factors that are relevant to the success of this project: competition, units sold, and cost—ah, and of course our pricing. Have you decided on the price yet?"

"I am still considering which of the three strategies would be most beneficial to us. Selling for $50.00 and trying to maximize revenues—or selling for $30.00 and trying to maximize market share. Of course, there is still your third alternative; we could sell for $40.00 and try to do both."

At this point Reggie focuses on the sheet of paper in front of him. "And I still believe that the $40.00 alternative is the best one. Concerning the costs, I checked the records; basically we have to amortize the development costs we incurred for Brainet. So far we have spent $800,000 and we expect to spend another $50,000 per year for support and shipping the CDs to those who want a hardcopy on top of their downloaded software." Reggie next hands a report to Charlotte. "Here we have some data on the industry. I just received that yesterday, hot off the press. Let's see what we can learn about the industry here." He shows Charlotte some of the highlights. Reggie then agrees to compile the most relevant information contained in the report and have it ready for Charlotte the following morning. It takes him long into the night to gather the data from the pages of the report, but in the end he produces three tables, one for each of the three alternative pricing strategies. Each table shows the corresponding probability of various amounts of sales given the level of competition (high, medium, or low) that develops from other companies.

The next morning Charlotte is sipping from another power drink. Jeannie and Reggie will be in her office any moment now and, with their help, she will have to decide what to do with Brainet. Should they launch the product? If so, at what price?

TABLE 1 Probability distribution of unit sales, given a high price ($50)

	Level of Competition		
Sales	**High**	**Medium**	**Low**
50,000 units	0.2	0.25	0.3
30,000 units	0.25	0.3	0.35
20,000 units	0.55	0.45	0.35

TABLE 2 Probability distribution of unit sales, given a medium price ($40)

	Level of Competition		
Sales	**High**	**Medium**	**Low**
50,000 units	0.25	0.30	0.40
30,000 units	0.35	0.40	0.50
20,000 units	0.40	0.30	0.10

TABLE 3 Probability distribution of unit sales, given a low price ($30)

	Level of Competition		
Sales	**High**	**Medium**	**Low**
50,000 units	0.35	0.40	0.50
30,000 units	0.40	0.50	0.45
20,000 units	0.25	0.10	0.05

When Jeannie and Reggie enter the office, Jeannie immediately bursts out: "Guys, I just spoke to our marketing research company. They say that they could do a study for us about the competitive situation for the introduction of Brainet and deliver the results within a week."

"How much do they want for the study?"

"I knew you'd ask that, Reggie. They want $10,000 and I think it's a fair deal."

At this point Charlotte steps into the conversation. "Do we have any data on the quality of the work of this marketing research company?"

"Yes, I do have some reports here. After analyzing them, I have come to the conclusion that the marketing research company is not very good in predicting the competitive environment for medium or low pricing. Therefore, we should not ask them to do the study for us if we decide on one of these two pricing strategies. However, in the case of high pricing, they do quite well: given that the competition turned out to be high, they predicted it correctly 80 percent of the time, while 15 percent of the time they predicted medium competition in that setting. Given that the competition turned out to be medium, they predicted high competition 15 percent of the time and medium competition 80 percent of the time. Finally, for the case of low competition, the numbers were 90 percent of the time a correct prediction, 7 percent of the time a 'medium' prediction and 3 percent of the time a 'high' prediction."

Charlotte feels that all these numbers are too much for her. "Don't we have a simple estimate of how the market will react?"

"Some prior probabilities, you mean? Sure, from our past experience, the likelihood of facing high competition is 20 percent, whereas it is 70 percent for medium competition and 10 percent for low competition," Jeannie has her numbers always ready when needed.

All that is left to do now is to sit down and make sense of all this. . . .

(a) For the initial analysis, ignore the opportunity of obtaining more information by hiring the marketing research company. Identify the alternative actions and the states of nature. Construct the payoff table. Then formulate the decision problem in a decision tree. Clearly distinguish between decision and chance forks and include all the relevant data.
(b) What is Charlotte's decision if she uses the maximum likelihood criterion? The maximin payoff criterion?
(c) What is Charlotte's decision if she uses Bayes' decision rule?
(d) Now consider the possibility of doing the market research. Develop the corresponding decision tree. Calculate the relevant probabilities and analyze the decision tree. Should Cerebrosoft pay the $10,000 for the marketing research? What is the overall optimal policy?

CASE 15.2 SMART STEERING SUPPORT

On a sunny May morning, Marc Binton, CEO of Bay Area Automobile Gadgets (BAAG), enters the conference room on the 40th floor of the Gates building in San Francisco, where BAAG's offices are located. The other executive officers of the company have already gathered. The meeting has only one item on its agenda: planning a research and development project to develop a new driver support system (DSS). Brian Huang, Manager of Research and Development, is walking around nervously. He has to inform the group about the R&D strategy he has developed for the DSS. Marc has identified DSS as the strategic new product for the company. Julie Aker, Vice President of Marketing, will speak after Brian. She will give detailed information about the target segment, expected sales, and marketing costs associated with the introduction of the DSS.

BAAG builds electronic nonaudio equipment for luxury cars. Founded by a group of Stanford graduates, the company sold its first product—a car routing system relying on a technology called global positioning satellites (GPS)—a few years ago. Such routing systems help drivers to find directions to their desired destinations using satellites to determine the exact position of the car. To keep up with technology and to meet the wishes of their customers, the company has added a number of new features to its router during the last few years. The DSS will be a completely new product, incorporating recent developments in GPS as well as voice recognition and display technologies. Marc strongly supports this product, as it will give BAAG a competitive advantage over its Asian and European competitors.

Driver support systems have been a field of intense research for more than a decade. These systems provide the driver with a wide range of information, such as directions, road conditions, traffic updates, etc. The information exchange can take place verbally or via projection of text onto the windscreen. Other features help the driver avoid obstacles that have been identified by cars ahead on the road (these cars transmit the information to the following vehicles). Marc wants to incorporate all these features and other technologies into one support system that would then be sold to BAAG's customers in the automobile industry.

After all the attendees have taken their seats, Brian starts his presentation: "Marc asked me to inform you about our efforts with the driver support system, particularly the road scanning device. We have reached a stage where we basically have to make a go or no-go decision concerning the research for this device, which, as you all know by now, is a key feature in the DSS. We have already integrated the other devices, such as the PGS-based positioning and direction system. The question we have to deal with is whether to fund basic research into the road scanning device. If this research were successful, we then would have to decide if we want to develop a product based on these results—or if we just want to sell the technology without developing a product. If we do decide to develop the product ourselves, there is a chance that the product development process might not be successful. In that case, we could still sell the technology. In the case of successful product development, we would have to decide whether to market the product. If we decide not to market the developed product, we could at least sell the product concept that was the result of our successful research and development efforts. Doing so would earn more than just selling the technology prematurely. If, on the other hand, we decide to market the driver support system, then we are faced with the uncertainty of how the product will be received by our customers."

"You completely lost me." snipes Marc.

Max, Julie's assistant, just shakes his head and murmurs, "those techno-nerds. . . ."

Brian starts to explain: "Sorry for the confusion. Let's just go through it again, step by step."

"Good idea—and perhaps make smaller steps!" Julie obviously dislikes Brian's style of presentation.

"OK, the first decision we are facing is whether to invest in research for the road scanning device."

"How much would that cost us?" asks Marc.

"Our estimated budget for this is $300,000. Once we invest that money, the outcome of the research effort is somewhat uncertain. Our engineers assess the probability of successful research at 80 percent."

"That's a pretty optimistic success rate, don't you think?" Julie remarks sarcastically. She still remembers the disaster with Brian's last project, the fingerprint-based car security system. After spending half a million dollars, the development engineers concluded that it would be impossible to produce the security system at an attractive price.

Brian senses Julie's hostility and shoots back: "In engineering we are quite accustomed to these success rates—something we can't say about marketing. . . ."

"What would be the next step?" intervenes Marc.

"Hm, sorry. If the research is not successful, then we can only sell the DSS in its current form."

"The profit estimate for that scenario is $2 million," Julie throws in.

"If, however, the research effort is successful, then we will have to make another decision, namely, whether to go on to the development stage."

"If we wouldn't want to develop a product at that point, would that mean that we would have to sell the DSS as it is now?" asks Max.

"Yes, Max. Except that additionally we would earn some $200,000 from selling our research results to GM. Their research division is very interested in our work and they have offered me that money for our findings."

"Ah, now that's good news," remarks Julie.

Brian continues, "If, however, after successfully completing the research stage, we decide to develop a new product then we'll have to spend another $800,000 for that task, at a chance of 35 percent of not being successful."

"So you are telling us we'll have to spend $800,000 for a ticket in a lottery where we have a 35 percent chance of not winning anything?" asks Julie.

"Julie, don't focus on the losses, but on the potential gains! The chance of winning in this lottery, as you call it, is 65 percent. I believe that that's much more than with a normal lottery ticket," says Marc.

"Thanks, Marc," says Brian. "Once we invest that money in development, we have two possible outcomes: either we will be successful in developing the road scanning device or we won't. If we fail, then once again we'll sell the DSS in its current form and cash in the $200,000 from GM for the research results. If the development process is successful, then we have to decide whether to market the new product."

"Why wouldn't we want to market it after successfully developing it?" asks Marc.

"That's a good question. Basically what I mean is that we could decide not to sell the product ourselves but instead give the right to sell it to somebody else, to GM, for example. They would pay us $1 million for it."

"I like those numbers!" remarks Julie.

"Once we decide to build the product and market it, we will face the market uncertainties and I'm sure that Julie has those numbers ready for us. Thanks."

At this point, Brian sits down and Julie comes forward to give her presentation. Immediately some colorful slides are projected on the wall behind her as Max operates the computer.

"Thanks, Brian. Well, here's the data we have been able to gather from some marketing research. The acceptance of our new product in the market can be high, medium, or low," Julie is pointing to some figures projected on the wall behind her. "Our estimates indicate that high acceptance would result in profits of $8.0 million, and that medium acceptance would give us $4.0 million. In the unfortunate case of a poor reception by our customers, we still expect $2.2 million in profit. I should mention that these profits do not include the additional costs of marketing or R&D expenses."

"So, you are saying that in the worst case we'll make barely more money than with the current product?" asks Brian.

"Yes, that's what I am saying."

"What budget would you need for the marketing of our DSS with the road scanner?" asks Marc.

"For that we would need an additional $200,000 on top of what has already been included in the profit estimates," Julie replies.

"What are the chances of ending up with a high, medium, or low acceptance of the new DSS?" asks Brian.

"We can see those numbers at the bottom of the slide," says Julie, while she is turning toward the projection behind her. There is a 30 percent chance of high market acceptance and a 20 percent chance of low market acceptance.

At this point, Marc moves in his seat and asks: "Given all these numbers and bits of information, what are you suggesting that we do?"

(a) Organize the available data on cost and profit estimates in a table.
(b) Formulate the problem in a decision tree. Clearly distinguish between decision and chance forks.
(c) Calculate the expected payoffs for each fork in the decision tree.
(d) What is BAAG's optimal policy according to Bayes' decision rule?
(e) What would be the expected value of perfect information on the outcome of the research effort?
(f) What would be the expected value of perfect information on the outcome of the development effort?
(g) Marc is a risk-averse decision maker. In a number of interviews, his utility function for money was assessed to be

$$u(M) = \frac{1 - e^{-\frac{M}{12}}}{1 - e^{-\frac{1}{12}}},$$

where M is the company's net profit in units of hundreds of thousands of dollars (e.g., $M = 8$ would imply a net profit of \$800,000). Using Marc's utility function, calculate the utility for each terminal branch of the decision tree.
(h) Determine the expected utilities for all forks in the decision tree.
(i) Based on Marc's utility function, what is BAAG's optimal policy?
(j) Based on Marc's utility function, what would be the expected value of perfect information on the outcome of the research effort?
(k) Based on Marc's utility function, what would be the expected value of perfect information on the outcome of the development effort?