CHAPTER 7

Dealing with Uncertainty: Expected Values, Sensitivity Analysis, and the Value of Information

Cost-benefit analysis often requires us to predict the future. Whether it is desirable to begin a project depends on what we expect will happen after we have begun. But, as mere mortals, we rarely are able to make precise predictions about the future. Indeed, in many situations analysts can be certain that circumstances largely beyond their clients' control, such as epidemics, floods, bumper crops, or fluctuations in international oil prices, will greatly affect the benefits and costs that would be realized from proposed policies. How can analysts reasonably take account of these uncertainties in CBA?

In this chapter, we consider three topics relevant to uncertainty: *expected value* as a measure reflecting risks, *sensitivity analysis* as a way of investigating the robustness of net benefit estimates to different resolutions of uncertainty, and the *value of information* as a benefit category for CBA and as a guide for allocating analytical effort. Expected values take account of the dependence of benefits and costs on the occurrence of specific contingencies, or "states of the world" to which analysts are able to assign probabilities of occurrence. Sensitivity analysis is a way of acknowledging uncertainty about the values of important parameters in our predictions—it should be a component of almost any CBA. When analysts have opportunities for gaining additional information about costs or benefits, they may be able to value the information by explicitly modeling the uncertainty inherent in their decisions. A particular type of information value, called *quasi-option value,* is relevant when assessing currently available alternatives that have different implications for learning about the future.

EXPECTED VALUE ANALYSIS

One can imagine several types of uncertainty about the future. At the most profound level, one might not be able to specify the full range of relevant circumstances that may occur. Indeed, the human and natural worlds are so complex that one cannot hope to anticipate every possible future circumstance. Yet, in many situations of relevance to one's daily life and public policy, it is reasonable to characterize the future in terms of a number of distinct contingencies. For example, in deciding whether to take an umbrella to work, one might reasonably divide the future into two contingencies: Either it will or will not rain sufficiently to make the umbrella useful. Of course, other relevant contingencies can be imagined as well—it will be a dry day, but one may or may not be the victim of an attempted mugging in which the umbrella would prove valuable in self-defense! Yet, if these additional contingencies are very unlikely, it is often reasonable to leave them out of one's model of the future. Modeling the future as a set of relevant contingencies involves yet another narrowing of uncertainty: How likely are each of the contingencies? If one is willing to assign probabilities of occurrence to each of the contingencies, then uncertainty about the future becomes a problem of dealing with risk. In relatively simple situations, risk can be readily incorporated into CBA through expected value analysis.

Contingencies and their Probabilities

Modeling uncertainty as risk begins with the specification of a set of *contingencies* that, within the simplified model of the world being employed, are *exhaustive* and *mutually exclusive*. Contingencies can be thought of as possible events, outcomes, or states of the world such that one and only one of the relevant set of possibilities will actually occur. What makes a set of contingencies the basis of an appropriate model for conducting a CBA of a policy?

One important consideration is that the contingencies capture the full range of likely variation in net benefits of the policy. For example, in evaluating an oil stockpile for use in the event of an oil price shock sometime in the future, one would want to consider at least two contingencies: There never will be an oil price shock (a situation in which the policy is likely to result in net losses), and there will be a major price shock (a situation in which the policy is likely to result in net gains).

Another consideration is how well the contingencies represent the possible outcomes between the extremes. In some circumstances, the possible contingencies can be listed exhaustively so that they are fully representative. More often, however, they sample an infinite number of possibilities. In these circumstances, each contingency can be thought of as a *scenario*, which is just a description of a possible future. Do the specified contingencies provide a sufficient variety of scenarios to convey the possible futures adequately? If so, then the contingencies are representative.

Figure 7-1 illustrates the representation of a continuous scale with discrete contingencies. The horizontal axis gives the number of inches of summer rainfall in an agricultural region. The vertical axis gives the net benefits of a water storage system, which increase as the amount of rainfall decreases. Imagine that an analyst represents uncertainty about rainfall with only two contingencies: "excessive" and "deficient." The excessive contingency assumes 22 inches of rainfall, which would yield zero net benefits from the storage system. The deficient contingency assumes zero inches of rainfall, which would yield $4.4 million in net benefits. If the relationship between rainfall and net benefits follows the straight line labeled *A*, and all the rainfall amounts between 0 and 22 are equally likely, then the average of net benefits over the full continuous range would be $2.2 million. If the analyst assumed that each of the contingencies were equally likely, then the average over the two contingencies would also be $2.2 million, so that using two scenarios would be adequately representative.[1]

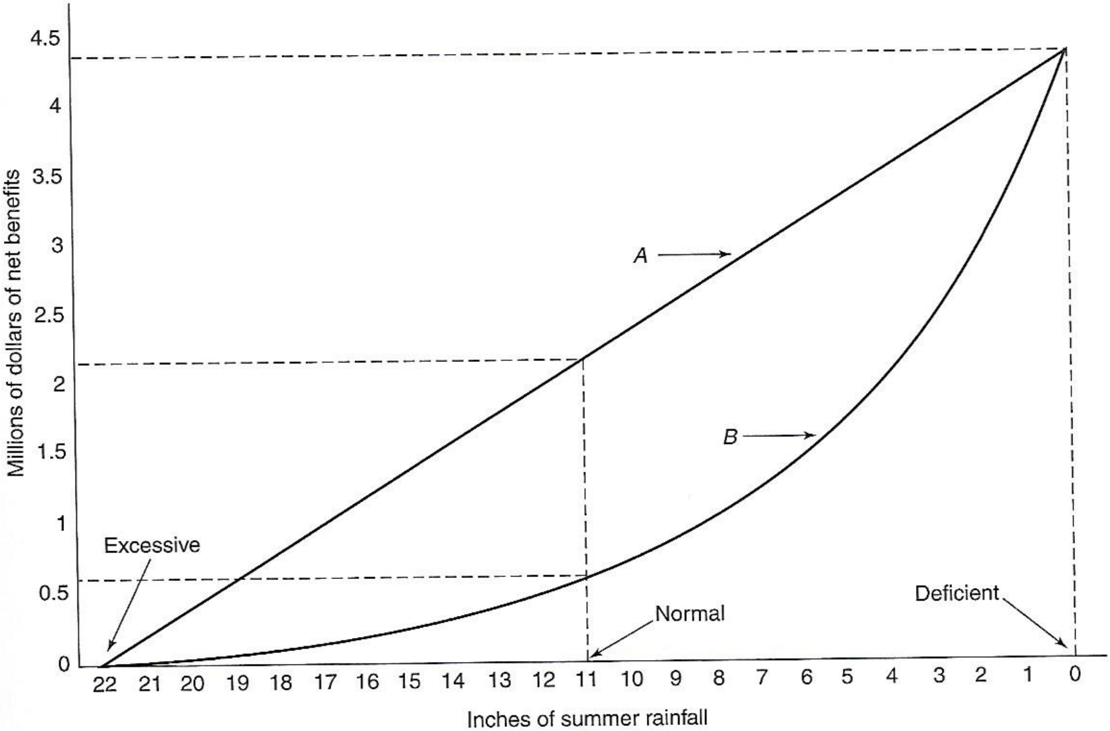

FIGURE 7-1 Representativeness of Contingencies

Now imagine that the net benefits follow the curved line labeled *B*. Again assuming that all rainfall amounts between 0 and 22 inches are equally likely, the average of net benefits over the full continuous range would only be about \$1.1 million, so that using only two contingencies would grossly overestimate the average net benefits from the storage system.[2] Adding "normal" as a contingency that assumes 11 inches of rainfall and averaging net benefits over all three contingencies yields net benefits of \$1.6 million, which is more representative than the average calculated with two contingencies but still considerably larger than the \$1.1 million calculated over the full continuous range. Even more contingencies are desirable. For example, moving to five equally spaced contingencies gives an average of \$1.3 million, which is much closer to the average over the continuous range.[3]

Once we have specified a tractable but representative set of contingencies, the next task is to assign probabilities of occurrence to each of them. To be consistent with the logical requirement that the contingencies taken together are exhaustive and mutually exclusive, the probabilities that we assign must each be nonnegative and sum to exactly 1. Thus, if there are three contingencies, C_1, C_2, and C_3, we must assign corresponding probabilities p_1, p_2, and p_3 such that $p_1 + p_2 + p_3 = 1$.

The probabilities may be based solely on historically observed frequencies; on subjective assessments by clients, analysts, or other experts based on a variety of

EXHIBIT 7-1

Being explicit about contingencies, their probabilities, and their consequences can help structure complex decision problems. Consider the following letter that President Abraham Lincoln wrote to Major General George B. McClellan on February 3, 1862:

> My dear Sir:
>
> You and I have distinct, and different plans for a movement of the Army of the Potomac—yours to in be down the Chesapeake, up the Rappahannock to Urbana, and across land to the terminus of the Railroad on the York River—, mine to move directly to a point on the Railroad South West of Manassas.
>
> If you will give me satisfactory answers to the following questions, I shall gladly yield my plan to yours.
>
> First. Does not your plan involve a greatly larger expenditure of time and money than mine?
>
> Second. Wherein is a victory more certain by your plan than mine?
>
> Third. Wherein is a victory more valuable by your plan than mine?
>
> Fourth. In fact, would it not be less valuable, in this, that it would break no great line of the enemy's communications, while mine would?
>
> Fifth. In case of disaster, would not a safe retreat be more difficult by your plan than by mine?
>
> Yours truly, Abraham Lincoln

Source: John G. Nicolay and John Hay, eds., *Abraham Lincoln: Complete Works, Volume Two* (New York: The Century Company, 1894), 120.

information and theory; or on both. For example, return to the contingencies in Figure 7-1: agriculturally "excessive," "normal," and "deficient" precipitation in a river valley for which a water storage system has been proposed. The national weather service may be able to provide data on average annual rainfall over the last century that allows an analyst to estimate the probabilities of the three levels of precipitation from their historical frequencies. If such data were not available, then the analyst would have to base the probabilities on expert opinion, comparison with similar valleys in the region for which data are available, or some other subjective assessment. As such subjective assessments are rarely made with great confidence, it is especially important to investigate the sensitivity of the results to the particular probabilities chosen.

Calculating the Expected Value of Net Benefits

The specification of contingencies and their respective probabilities allows us to calculate the *expected net benefits* of a policy. We do so by first predicting the net benefits of the policy under each contingency and then taking the weighted average of these net benefits over all the contingencies, where the weights are the respective probabilities that the contingencies occur. Specifically, for n contingencies, let B_i be the benefits under contingency i, C_i be the costs under contingency i, and p_i be the probability of contingency i occurring. Then the expected net benefits, $E[NB]$, are given by the formula:

$$E[NB] = p_1(B_1 - C_1) + \cdots + p_n(B_n - C_n) \qquad (7.1)$$

which is just the expected value of net benefits over the n possible outcomes.[4]

EXHIBIT 7-2

In their evaluation of alternative government oil stockpiling programs in the early 1980s, Glen Sweetnam and colleagues at the U.S. Department of Energy modeled the uncertainty surrounding oil market conditions with five contingencies: *slack market*—oil purchases for the U.S. stockpile of up to 1.5 million barrels per day (mmb/d) could be made without affecting the world oil price; *tight market*—oil purchases increase the world price at the rate of $3.60 per mmb/d; *minor disruption*—loss of 1.5 mmb/d to the world market (e.g., caused by a revolution in an oil-exporting country); *moderate disruption*—loss of 6.0 mmb/d to the world market (e.g., caused by a limited war in the Persian Gulf); *major disruption*—loss of 12.0 mmb/d to the world market (e.g., caused by a major war in the Persian Gulf). For each of the 24 years of their planning horizon, they assumed that the probabilities of each of the contingencies occurring depended only on the contingency that occurred in the previous year. For each year, they calculated the social surplus in the U.S. oil market conditional on each of the five market contingencies and change in the size of the stockpile.

The model they constructed allowed them to answer the following questions: For any current market condition and stockpile size, what change in stockpile size maximizes the present value of expected net benefits? How much storage capacity should be constructed? How fast should it be added? The model and the answers it provided were influential in policy debates concerning expansion of the U.S. stockpile, the Strategic Petroleum Reserve.

Sources: Adapted from Glen Sweetnam, "Stockpile Policies for Coping with Oil-Supply Disruptions," in George Horwich and Edward J. Mitchell, eds., *Policies for Coping with Oil-Supply Disruptions* (Washington, D.C.: American Enterprise Institute for Public Policy Research, 1982), 82–96. On the role of the model in the policy-making process, see Hank C. Jenkins-Smith and David L. Weimer, "Analysis as Retrograde Action: The Case of Strategic Petroleum Reserves," *Public Administration Review* 45(4) 1985, 485–494.

When facing complicated risk problems, analysts often find it useful to model them as *games against nature*. A game against nature assumes that nature will randomly, and nonstrategically, select a particular state of the world. The random selection of a state of the world is according to assumed probabilities. The selection is nonstrategic in the sense that nature does not alter the probabilities of the states of the world in response to the action selected by the analysts. A game against nature in *normal form* has the following elements: *states of nature* and their *probabilities of occurrence*, *actions* available to the decision maker facing nature, and *payoffs* to the decision maker under each combination of state of nature and action.

Table 7-1 shows the analysis of alternatives for planetary defense against asteroid collisions as a game against nature in normal form. It considers three possible states of nature over the next 100 years: exposure of Earth to collision with an asteroid larger than one kilometer in diameter, which would have enough kinetic energy to impose severe regional or even global effects on society (10 on the Torino Scale); exposure of Earth to collision with an asteroid smaller than one kilometer but larger than 20 meters in diameter, which would have severe local or regional effects on society (8 or 9 on the Torino Scale); and no exposure of Earth to an asteroid larger than 20 meters in diameter. The game shows three actions: build a forward-based asteroid defense, which would station nuclear devices sufficiently deep in space to give a good possibility of

TABLE 7-1 A Game against Nature: Expected Values of Asteroid Defense Alternatives

State of Nature	*Exposure to a Collision with an Asteroid Larger Than One Kilometer in Diameter*	*Exposure to a Collision with an Asteroid between 20 Meters and 1 Kilometer in Diameter*	*No Exposure to Collision with an Asteroid Larger Than 20 Meters in Diameter*	
Probabilities of states of nature (over next century)	.001	.004	.995	
Actions (alternatives)	*Payoffs (net costs in billions of 2000 dollars)*			*Expected Value*
Forward-based asteroid defense	5,060	1,060	60	69
Near-Earth asteroid defense	10,020	2,020	20	**38**
No asteroid defense	30,000	6,000	0	54

Choose near-Earth asteroid defense: Expected net cost = $38 billion.

their timely use in diverting asteroids from collision courses with Earth; build a near-Earth asteroid defense, which would be less expensive but not as effective as the forward-based defense; and build no asteroid defense.

Although actually estimating the payoffs for this game would be a monumental and controversial analytical task, Table 7-1 displays some hypothetical figures. The payoffs, shown as the present value of net costs over the next century, range from $30 trillion (Earth is exposed to a collision with an asteroid larger than one kilometer in diameter in the absence of any asteroid defense) to $0 (Earth is not exposed to collision with an asteroid larger than 20 meters and no defense system is built). Note that estimating the costs of a collision between Earth and an asteroid would itself involve expected value calculations that take account of size, composition, and point of impact of the asteroid. The $30 trillion figure is about half the world's annual sum of gross domestic products.

The last column of Table 7-1 shows expected values for each of the three alternatives. The expected value for each alternative is calculated by summing the products of its payoff conditional on states of nature with the probabilities of those states. For example, the expected value of payoffs (present value of net costs) for no asteroid defense is:

$$(0.001)(\$30{,}000 \text{ billion}) + (0.004)(\$6{,}000 \text{ billion}) + (0.995)(\$0) = \$54 \text{ billion}$$

Similar calculations yield $69 billion for forward-based asteroid defense and $38 billion for near-Earth asteroid defense. As the maximization of expected net benefits is equivalent to minimizing expected net costs, the most efficient alternative is near-Earth asteroid defense. Alternatively, one could think of near-Earth asteroid defense as offering expected net benefits of $16 billion relative to no defense ($54 billion in expected net costs minus $38 billion in expected net costs equals $16 billion in expected net benefits), while forward-based asteroid defense offers negative $15 billion in expected net benefits relative to no defense ($54 billion in expected net costs minus $69 billion in expected net costs equals negative $15 billion in expected net benefits).

In CBA, it is common practice to treat expected values as if they were certain amounts. For example, imagine that a perfect asteroid defense system would have a present value cost of $100 billion under each of the states of nature. In this case, assuming accurate prediction of costs, the $100 billion *would be certain* because it does not depend on which state of nature actually results. CBA generally treats a certain amount such as this as fully commensurate with expected values, even though the latter will not actually result in its expected value. In other words, although the expected net cost of no asteroid defense is $54 billion, assuming an accurate prediction of payoffs, the actually realized net cost will be $30 trillion, $6 trillion, or $0. If the perfect defense system cost $54 billion, then CBA would rank it and no defense as equally efficient.

Treating expected values as if they were certain amounts implies that the person making the comparison has preferences that are *risk neutral*. A person has risk neutral preferences when he or she is indifferent between certain amounts and lotteries with the same expected payoff. A person is *risk averse* if he or she prefers the certain amount and is *risk seeking* if he or she prefers the lottery. Buying insurance, which offers a lower expected payoff than the certain premium charged, indicates risk aversion; buying a lottery ticket, which offers a lower expected value than its price, indicates risk seeking.

Chapter 8 considers the appropriateness of treating expected values and certain equivalents as commensurate (e.g., risk neutrality). It explains that doing so is not conceptually correct in measuring willingness to pay in circumstances in which individuals face uncertainty. Nevertheless, it argues that, in practice, *treating expected values and certain amounts as commensurate is generally reasonable when either the pooling of risk over the collection of policies, or the pooling of risk over the collection of persons affected by a policy, will make the actually realized values of costs and benefits close to their expected values*. For example, a policy that affects the probability of highway accidents involves reasonable pooling of risk across many drivers (some will have accidents, others will not) so that realized values will be close to expected values. In contrast, a policy that affects the risk of asteroid collision does not involve pooling across individuals (either everyone suffers from the global harm if there is a collision or no one does if there is no collision), so that the realized value of costs may be very far from their expected value. As discussed in Chapter 8, such unpooled risk may require an adjustment to expected benefits called *option value*.

Decision Trees and Expected Net Benefits

The basic procedure for expected value analysis, taking weighted averages over contingencies, can be directly extended to situations in which costs and benefits accrue over multiple years, as long as the risks in each year are independent of the realizations of risks in previous years. Consider, for example, a CBA of a dam with a 20-year life. Assume that the costs and benefits of the dam depend only on the contingencies of below-average rainfall and above-average rainfall in the current year. Additionally, if the analyst is willing to make the plausible assumption that the amount of rainfall in any year does not depend on the rainfall in previous years, then the analyst can simply calculate the present value of expected net benefits for each year and calculate the present value of this stream of net benefits in the usual way.

The basic expected value procedure cannot be so directly applied when either the net benefits accruing under contingencies or the probabilities of the contingencies depend on the contingencies that have previously occurred. For example, above-average rainfall in one year may make the irrigation benefits of a dam less in the next year because of accumulated ground water. In the case of a policy to reduce the costs of earthquakes, the probability of a major earthquake may change each year depending on the mix of earthquakes that occurred in the previous year.

Such situations require a more flexible framework for handling risk than basic expected value analysis. *Decision analysis* provides the needed framework.[5] Though it takes us too far afield to present decision analysis in any depth here, we sketch its general approach and present simple illustrations that demonstrate its usefulness in CBA. A number of book-length treatments of decision analysis are available for those who wish to pursue this topic in more depth.[6]

Decision analysis can be thought of as a *sequential, or extended form, game* against nature. It proceeds in two basic stages. First, one specifies the logical structure of the decision problem in terms of sequences of decisions and realizations of contingencies using a diagram, called a *decision tree*, that links an initial decision (the trunk) to final outcomes (branches). Second, using *backward induction,* one works from final outcomes back to the initial decision, calculating expected values of net benefits across contingencies and pruning dominated branches (i.e., eliminating branches with lower expected values of net benefits).

Consider a vaccination program against a particular type of influenza that involves various costs.[7] The costs of the program result from immunization expenditures and possible adverse side effects; the benefits consist of the adverse health effects that are avoided if an epidemic occurs. This flu may infect a population over the next two years before sufficient immunity develops worldwide to stop its spread. Figure 7-2 presents a simple decision tree for a CBA of this vaccination program. The tree should be read from left to right to follow the sequence of decisions, denoted by □, and random selections of contingencies, denoted by (○). The tree begins with a decision node, the square labeled 0 at the extreme left. The upper bough represents the decision to implement the vaccination program this year; the lower bough represents the decision not to implement the program this year.

Upper Bough: The Vaccination Program. Follow the upper bough first. If the program is implemented, then it will involve direct administrative costs, C_a, and the costs of adverse side effects, such as contracting the influenza from the vaccine itself, suffered by those who are vaccinated, C_s. Note that C_s, like most of the other costs in this example, is itself an expected cost based on the probability of the side effect, the cost to persons suffering the side effect, and the number of persons vaccinated. The solid vertical line on the bough can be thought of as a toll gate at which point the program costs, $C_a + C_s$, are incurred. A chance node, represented by a circle, appears next. Either the influenza infects the population (the upper branch, which occurs with probability P_1 and results in costs $C_{e|v}$, where the subscript should be read as "epidemic occurs given that the vaccination program has been implemented"), or the influenza does not infect the population (the lower branch, which occurs with probability $1 - P_1$ and results in zero costs at that time). If the influenza does occur, then the population will be immune

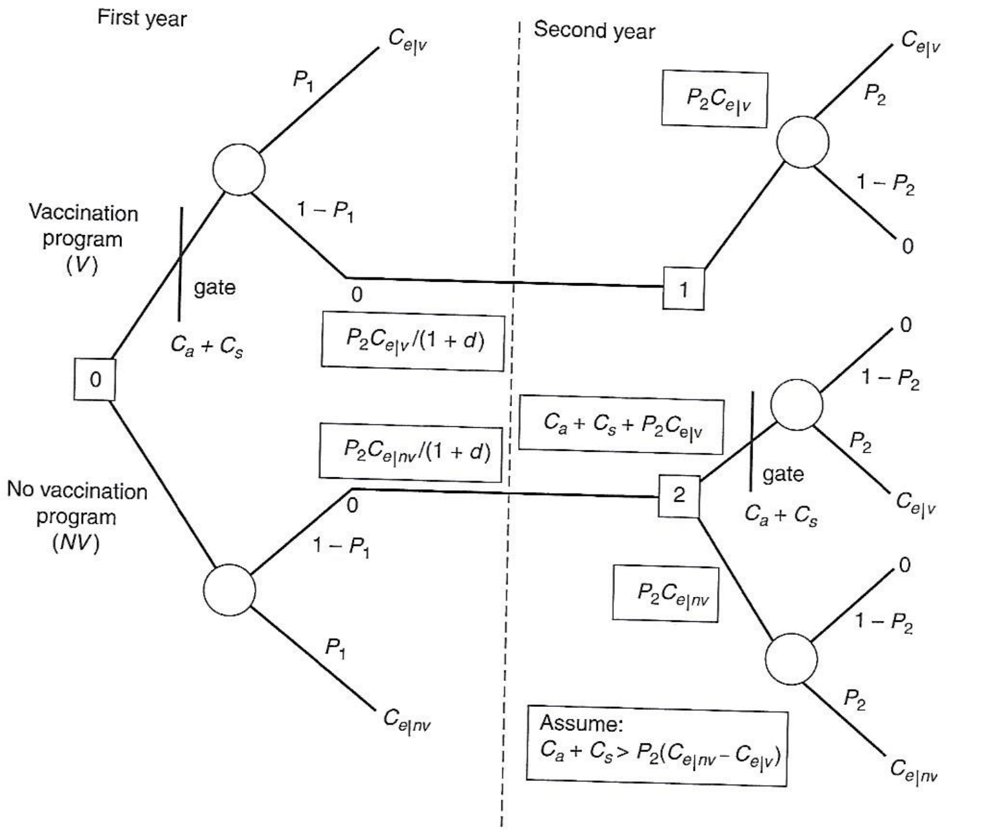

FIGURE 7-2 Decision Tree for Vaccination Program Analysis

in the next year. Thus, the upper branch does not continue. If the influenza does not occur, then there is still a possibility that it might occur in the next year. Therefore, the lower branch continues to the second year, where the square labeled 1 notes the beginning of the second year. It leads directly to another chance node that specifies the two contingencies in the second year: The influenza infects the population (the upper subbranch, which occurs with probability P_2 and results in costs $C_{e|v}$), or the influenza does not infect the population (the lower subbranch, which occurs with probability $1 - P_2$ and results in zero costs).[8] We assume that P_2 is known at the time of the initial decision.[9]

Lower Bough: No Vaccination Program. We now return to the initial decision node and follow the lower bough representing no vaccination program in the first year. Initially there is no cost associated with this decision. A chance node follows with two branches: Either the influenza infects the population (the lower branch, which occurs with probability P_1 and results in costs $C_{e|nv}$), or the influenza does not infect the population (the upper branch, which occurs with probability $1 - P_1$ and results in zero costs).[10] If the influenza does occur, then there is no need to consider the next year. If it does not occur, then the tree continues to decision node 2: Either implement the vaccination program in the second year (the upper subbranch crossing the gate where program costs $C_a + C_s$ are incurred) or do not implement it (the lower subbranch).

If the program is implemented, then a chance node occurs: The influenza infects the population (the lower twig, which occurs with probability P_2 and results in costs $C_{e|v}$), or the influenza does not infect the population (the upper twig, which occurs with probability $1 - P_2$ and results in zero costs). We complete the tree by considering the parallel chance node following the decision not to implement the program in the second year: The influenza infects the population (the lower twig, which occurs with probability P_2 and results in costs $C_{e|nv}$), or the influenza does not infect the population (the upper twig, which occurs with probability $1 - P_2$ and results in zero costs).

Solving the Decision Tree. To solve the decision problem, we work from right to left, replacing chance nodes with their expected costs and pruning off parallel nodes that are dominated. Consider the chance node following decision node 1. Its expected cost, calculated by the expression $P_2C_{e|v} + (1 - P_2)0$, equals $P_2C_{e|v}$.

Now consider the chance nodes following decision node 2. The lower chance node, following a decision not to implement the vaccination program, has an expected cost of $P_2C_{e|nv}$. The upper chance node has an expected cost of $P_2C_{e|v}$, to which must be added the certain payment of program costs so that the full expected cost of implementing the vaccination program in the second year is $C_a + C_s + P_2C_{e|v}$. We can now compare the expected cost of the two possible decisions at node 2: $P_2C_{e|nv}$ versus $C_a + C_s + P_2C_{e|v}$. To illustrate, assume that program costs are greater than the expected cost reduction from the vaccine, that is, $C_a + C_s > P_2(C_{e|nv} - C_{e|v})$, then $P_2C_{e|nv}$ is smaller than $C_a + C_s + P_2C_{e|v}$ so that not implementing the program dominates implementing it. (If this were not the case, then the lower branch would be unequivocally dominated by the upper branch.[11]) We can now prune off the upper subbranch. If we reach decision node 2, then we know that we can obtain expected second-year costs of $P_2C_{e|nv}$.

At decision node 0 the expected costs of implementing the vaccination program (i.e., following the upper bough) consist of direct costs plus the expected costs of the following chance node, which now has the payoffs $C_{e|v}$ if there is an epidemic and the discounted expected value of node 1, $P_2C_{e|v}/(1 + d)$ if there is not an epidemic. Note that because this latter cost occurs in the second year, it is discounted using rate d. Thus, the present value of expected costs from implementing the vaccination program is given by:

$$E[C_v] = C_a + C_s + P_1C_{e|v} + (1 - P_1)P_2C_{e|v}/(1 + d) \qquad (7.2)$$

where the last term incorporates the expected costs from the second year.

The expected costs of not implementing the vaccination program are calculated in the same way: The payoff if there is not an epidemic becomes the discounted expected costs from decision node 2, $P_2C_{e|nv}/(1 + d)$; the payoff if there is an epidemic is still $C_{e|nv}$. Therefore, the expression:

$$E[C_{nv}] = P_1C_{e|nv} + (1 - P_1)P_2C_{e|nv}/(1 + d) \qquad (7.3)$$

gives the present value of expected costs of not implementing the program.

The final step is to compare the present values of expected costs for the two possible decisions at node 0. We prune the bough with the larger present value of expected costs. The remaining bough is the optimal decision.

As an illustration, suppose that we have gathered data suggesting the following values for parameters in the decision tree: $P_1 = .4$, $P_2 = .2$, $d = .05$, $C_{e|v} = .5C_{e|nv}$ (the vaccination program cuts the costs of influenza by half), $C_a = .1C_{e|nv}$ (the vaccination costs 10 percent of the costs of the influenza), and $C_s = .01C_{e|nv}$ (the side-effect costs are 1 percent of the costs of the influenza). For these values, $E[C_v] = .367C_{e|nv}$ and $E[C_{nv}] = .514C_{e|nv}$. Therefore, the vaccination program should be implemented in the first year because $E[C_v] < E[C_{nv}]$.

Calculating Expected Net Benefits of the Vaccination Program. Returning explicitly to CBA, we can recognize the benefits of the vaccination program as the costs it avoids. Thus, the present value of expected net benefits of the vaccination program is simply $E[C_{nv}] - E[C_v]$, which in the numerical illustration presented in the preceding paragraph equals $0.147C_{e|nv}$. In Chapter 8, we return to the question of the appropriateness of expected net benefits as a generalization of net benefits in CBA.

Extending Decision Analysis. Decision analysis can be applied to both public- and private-sector issues, and it can be used to structure much more complicated analyses than the CBA of the vaccination program. Straightforward extensions include more than two alternatives at decision nodes, more than two contingencies at chance nodes, more than two periods of time, and different probabilities of events in different periods. For example, analyses of the U.S. oil stockpiling program typically involve trees so large that they can only be fully represented and solved by computers.[12] Even in less complex situations, however, decision analysis can be very helpful in showing how risk should be incorporated into the calculation of expected net benefits.

SENSITIVITY ANALYSIS

Whether or not we structure a CBA explicitly in terms of contingencies and their probabilities, we always face some uncertainty about the magnitude of the impacts we predict and the values we assign to them. Our basic analysis usually submerges this uncertainty by using our most plausible estimates of these unknown quantities. These estimates comprise what is called the *base case*. The purpose of sensitivity analysis is to acknowledge the underlying uncertainty. In particular, it should convey how sensitive predicted net benefits are to changes in assumptions. If the sign of net benefits does not change when we consider the range of reasonable assumptions, then our results are robust, and we can have greater confidence in them.

Large numbers of unknown quantities, the usual situation in CBA, make the brute force approach of looking at all combinations of assumptions unfeasible. For example, the vaccination program analysis, which we further develop in the next section, involves 17 different uncertain numerical assumptions. If we looked at just three different values for each assumption, there would still be over 129 million different combinations of assumptions to consider.[13] Even if we could compute net benefits for all