From Edith Stokey & Richard Zeckhauser, A Primer for Policy Analysis

12 Decision Analysis

Many policy problems are hard to sort out. A situation may require a sequence of interrelated decisions that is more complex than the mind can readily encompass. More often it is difficult to decide what is the best choice because the outcomes that result from the required decisions depend in part on chance events; the decision maker is not in full control. Perhaps the uncertainty stems from nature; he does not know, for example, the level to which the spring runoff will carry the river. Perhaps he cannot be sure how many people will take advantage of a new job-training program. Like it or not, he cannot foresee the future; he must take his chances. Since the early 1960s, businessmen have increasingly used a method of analysis called *decision analysis* for tackling problems where decisions must be made sequentially and where uncertainty is a critical element. The application of decision analysis to problems in the public sector lagged by about a decade, mainly because the estimation of probabilities and the valuation of outcomes proved difficult. But decision analysis has increasingly become a valuable tool for analyzing and formulating public policies. It has been employed to address such diverse problems as prescribing appropriate treatment for a sore throat, choosing the site for a new airport for Mexico City, and deciding whether the U.S. should proceed with commercial supersonic flight.

Decision analysis in effect provides us with a road map for picking our way through confusing and uncertain territory. Equally important, it gives us a technique for finding the best route. Without further ado, let's look at the bare bones of the method, so that you will have a general understanding of what decision analysis is all about and how it works. We will then gradually introduce variations on the basic model that will greatly expand the range of situations to which it can be applied.[1]

[1] A lucid introduction to the field is Howard Raiffa's *Decision Analysis* (Reading, Mass.: Addison-Wesley, 1968). We strongly recommend it to those who desire a more comprehensive treatment of decision analysis.

The Decision Tree: A Descriptive Model

Most of the public policy issues for which decision analysis is useful are complex, far too complex to use in introducing the subject. Consequently, we will start with a deliberately trivial problem, culled from the 1974 sample examination for Administrative Officers of the U.S. Foreign Service:

> The officer in charge of a United States Embassy recreation program has decided to replenish the employees club funds by arranging a dinner. It rains nine days out of ten at the post and he must decide whether to hold the dinner indoors or out. An enclosed pavilion is available but uncomfortable, and past experience has shown turnout to be low at indoor functions, resulting in a 60 per cent chance of gaining $100 from a dinner held in the pavilion and a 40 per cent chance of losing $20. On the other hand, an outdoor dinner could be expected to earn $500 unless it rains, in which case the dinner would lose about $10.

(Where this damp and dismal post might be located escapes us.)

The first step in using decision analysis to attack the officer's problem is to diagram the sequence of decisions and chance events that he faces. The particular type of diagram that is used is called a decision tree. In Figure 12–1 we have drawn the decision tree for this situation; you may want to try to figure it out for yourself before reading on.

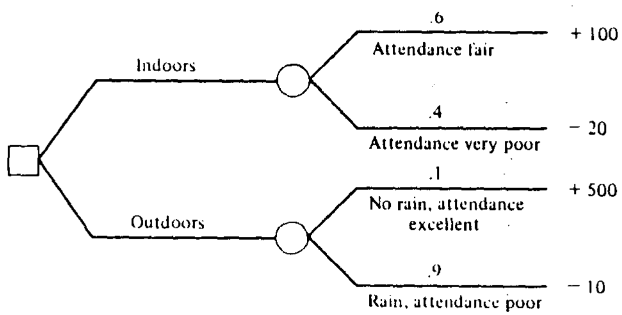

Fig. 12–1

To understand the tree, we begin, quite logically, with the first decision the officer faces: Should he hold the party indoors or outdoors? At the left we draw a square, or *decision node*, to indicate that at this point a decision must be made. We then draw two lines branching out of the decision node to show his two possible choices, and label them *Indoors* and *Outdoors*. Next we ask "What happens if we follow along the upper branch, in other words, if he holds the dinner indoors?" The answer isn't much help: "It all depends." In fact, we are told that it all depends on chance, on how well people turn out for the dinner. Hence at the end of the Indoors branch we draw a circle or *chance node*, to indicate that at this point an uncertainty must be resolved one way or another. At this chance node there are two possibilities: The party may be moderately successful ("attendance fair"), or it may be an utter disaster ("attendance very poor"). Hence we draw two branches for these possible results, and label them accordingly. Moreover, for this decision problem we know the probabilities of these possible outcomes, for we are told that there is a 60 percent chance of the former and a 40 percent chance of the latter. These numbers are recorded along the appropriate branches. Finally, we are also told that the gain from a moderately successful dinner is $100, while very poor attendance will result in the loss of $20. These are the ultimate outcomes or *payoffs* for each possible combination of choice and chance; they are shown at the tips of the branches.

Similarly, the two possibilities for an outdoor affair, their probabilities, and their payoffs are shown emanating from the lower chance node. The tree thus summarizes all the essential information that is available. Note very carefully the order in which events occur: the decision must be made before the decision maker knows what the weather will be. Hence the decision node must precede the chance nodes. The appropriate sequencing of decision and chance nodes always requires close attention when a decision tree is drawn.

The problem as set forth has obviously been drastically simplified; the officer might have additional options, including doing nothing, or other weather possibilities, such as "threatening." He could try to secure a long-range weather forecast, or even pay some extra amount for the option of delaying his decision until a few hours before the affair. It is the formulation and diagramming of the problem that we are concerned with now, not realism.

A decision tree is, then, a flow diagram that shows the logical structure of a decision problem. It contains four elements:

1. *Decision nodes*, which indicate all possible courses of action open to the decision maker;
2. *Chance nodes*, which show the intervening uncertain events and all their possible outcomes;
3. *Probabilities* for each possible outcome of a chance event; and
4. *Payoffs*, which summarize the consequences of each possible combination of choice and chance.

If you fully understand these simple principles for constructing a decision tree, you have already mastered the essence of the method. Yet despite their simplicity, it would be difficult to overemphasize the usefulness and importance of decision trees. As we have found with almost every type of model, the foremost advantage is the discipline imposed by the model. It requires us to structure the problem, break it into manageable pieces, and get all its elements down on paper—tasks that appear deceptively easy. Frequently they turn out to be very difficult. And even when it is hard or perhaps impossible to assign objective probabilities to chance events or to quantify outcomes in physical units, the effort of drawing the tree correctly provides valuable insights into the complexity of a problem. A corollary advantage is that a decision tree helps us communicate assumptions and valuations to others; it is a tool that facilitates sensible policy discussion.

Some Further Clarification

Decision trees must be constructed to show *all* events that can possibly occur at a chance node, and *all* options that might be pursued at a decision node. Moreover, these events and options must be defined in such a way that they don't overlap. (The technical phrase for this state of affairs is "mutually exclusive and collectively exhaustive.") For example, if we wished to consider the weather possibilities for our Foreign Service officer less simplistically, we might add temperature categories "hot" and "cold." But it would be wrong merely to modify the chance nodes by adding hot and cold branches, as in Figure 12–2. Rather we would have to show four possibilities, as in Figure 12–3.

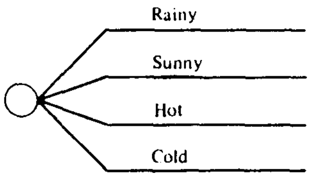

Fig. 12–2

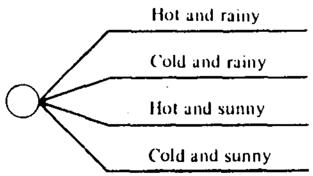

Fig. 12–3

Alternatively, we could diagram this as the two successive chance events of Figure 12–4. Or we could put the "sunny–rainy" chance node on the tree before the "hot–cold" node. In some contexts only one version of the tree will make sense; this is particularly likely to be true when there is a well-defined, chronological sequence. In others there may be little reason to prefer one version to another and the choice will depend on whatever seems most logical in terms of the information available.

Similarly, the possibilities for action must be mutually exclusive and collectively exhaustive. As a homely example, if you are trying to decide whether to wear a raincoat or carry an umbrella, the alternatives presumably are not mutually exclusive. The correct diagram is then Figure 12–5.

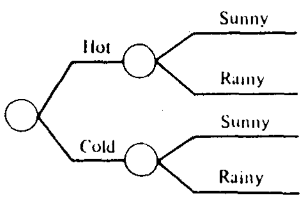

Fig. 12–4

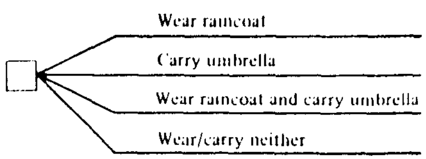

Fig. 12–5

The above examples are deliberately commonplace, yet the point we are making is applicable to all decision problems, however complex. Decision analysis forces you to think carefully about

1. The true nature of the decision problem;
2. The role of chance; and
3. The nature of the sequential interaction of decisions and chance events.

When the course of action may be carried on at various levels, the use of a decision tree becomes more cumbersome. For example, suppose you are trying to decide what sum of money to carry with you on an extended foreign trip, and how much of it to carry as cash and how much in travellers' checks. A limitless number of combinations are possible. If you are to use decision analysis to tackle this problem, you must, as a practical matter, narrow the problem down to a finite number of discrete choices. You might, for instance, construct the relevant part of the tree as in Figure 12–6. When you are satisfied that you have determined the right branch for

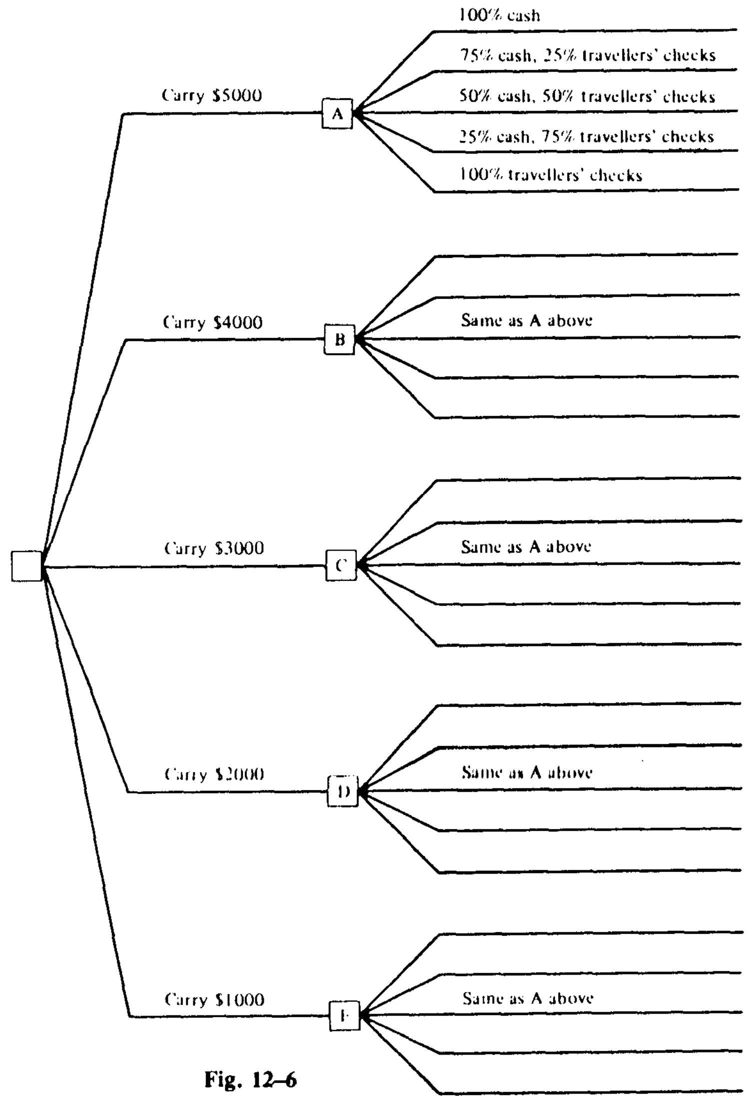

Fig. 12–6

you, you may want to try an amplified version of that branch to refine your choice.

Obviously, an analyst who is determined to keep his blinders on will not find his performance improved by mindlessly sketching a decision tree. But we have discovered that with a little experience, just drawing that first square on a piece of paper encourages most people to think more systematically about the matter at hand. To be sure, the exact shape the tree should take is sometimes far from obvious, as you will discover when you find yourself up against an amorphous problem whose ramifications seem endless. Yet the more difficult it is to develop the model, the more useful it is likely to prove. A decision made without an appropriate model in mind almost certainly will reflect confusion in thinking.

A Representative Decision Problem with Testing: The Choice of Generators

Earlier in this book we considered choices among alternative dam projects, all of which generated electricity and at the same time provided water for irrigation. In the real world the performance characteristics of irrigation systems and power plants are uncertain and variable, and may change from year to year over the extended lifetime of a project. For purposes of illustration we will suppress some of the engineering facts of life and look at a much simplified problem.

Suppose that the choice among dam projects has been narrowed to a particular dam and power plant configuration that will produce a known amount of water for irrigation. The only decision remaining is the choice of generators for the power plant. Assume for convenience that at the end of a year the earth will open and swallow up the whole project.[2] The final stages of development have recently been completed on a new and much more efficient type of generator, but that type has not yet been tested in conjunction with a dam of this kind. The costs of the old type and this new one are identical. A plant equipped with conventional generators will produce electricity worth $5 million a year. Such a plant will supply only a very small proportion of the region's electricity; any excess or deficiency can be sold or bought at a constant price. The output of the new type of generator is uncertain but it is estimated that there is a .3 chance that operating difficulties will develop so that it can be run only at low capacity, in which case it will generate only $3 million of electricity per year. Correspondingly, there is a .7 chance that the new type of generator will work well, in which case it will produce $8 million of electricity. (It might be more realistic to postulate a whole distribution of outcomes, but that would further complicate the problem while offering few compensating gains in insights into the decision analysis technique.) It is possible to build and test a prototype generator at a cost of $.5 million. The tests would

[2] You have already studied benefit–cost analysis and discounting so you know how a stream of benefits over time would be handled in practice. How to aggregate the consequences in future time periods is a separate issue that has no bearing on the point we are now trying to make.

predict the reliability of the new technology with complete accuracy. Before reading on, you should structure this problem for the decision maker in the form of a decision tree. (Don't forget to deduct the cost of the prototype from the payoff wherever appropriate.)

We'll return to this problem shortly; at the moment we want only to give you a little practice in tree drawing. If you wish to check your work, the tree is shown in Figure 12–7. Decision nodes are indicated by numbers, chance nodes by letters. Note carefully the effect of the test on the probabilities at chance nodes C and D, and also that the cost of the test has been included in the payoffs for the appropriate branches.

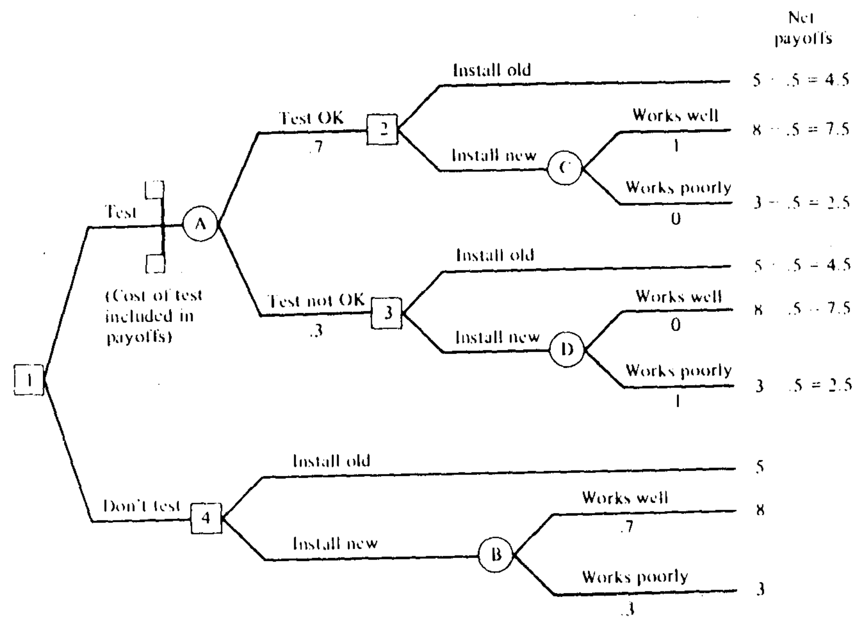

Fig. 12–7 (Payoffs are in millions of dollars.)

You will find that in many discussions of decision analysis costs are entered as "tolls" at the point where they occur. This is one logical way to handle them, and it certainly helps us visualize the problem more accurately. We follow this practice only in part. The symbol is entered on the tree to show the point in the decision process at which an additional cost is incurred; the dollar costs are subtracted to give the net payoffs, which appear on the right at the tips of the branches. It is never wrong to carry the costs out to the tips of the tree in this manner, and in situations where the decision maker is not risk neutral (a matter we'll get to shortly) all costs *must* be carried out to the tips. It therefore seems simpler in this abbreviated treatment of decision analysis to establish uniform rules.

Used as descriptive models, decision trees have great intuitive appeal, for they are easy to understand, readily discussed, and applicable to almost every choice we face in the real world. Trees are also an invaluable normative model, for whenever it is possible to assign values to probabilities and payoffs we can use the tree to determine the preferred course of action. We turn now to consideration of decision analysis as a normative tool that facilitates effective choice.

Decision Analysis: Folding Back and Choosing the Preferred Course of Action

The decision tree lays out the decision process for us so that we may visualize it in its entirety; it does not directly define the preferred course of action. If we try to reach a decision by working along the tree from left to right, following the logical sequence of choice and chance, we can hardly get started. Indeed, beginning that way would undo much of the value of drawing the tree in the first place. The embassy officer can't decide whether to hold his party indoors or out, because he doesn't know how to evaluate the uncertain consequences of either choice. For the same reason, the designers of the power plant can't decide whether or not to test the new generator. To evaluate the choices facing either decision maker, we must start at the right, with the tips of the trees, and work backward along the branches. This process is easier to understand (or at any rate it's easier to describe) if we work through some concrete examples first and talk about general principles afterward.

Working Backward Under Certainty: The Land Bequest

Let's start with a decision problem that has the simplest structure possible. Suppose you are the mayor of a small city. You are trying to decide what to do with a parcel of land recently bequeathed to the city. One group of citizens contends that the land should be used for recreation, another that it should be sold for residential construction. City planners who have studied the situation recommend three alternative recreational uses: (1) as a wildlife refuge; (2) as a municipal pitch-and-putt course; (3) as a public park with tennis courts and a softball diamond. If the parcel is sold for housing, the possibilities are (1) single-family houses; (2) condominiums; (3) mixed-income, federally subsidized apartment buildings. A majority of the City Council is docile and will follow your recommendation, but you find it difficult to make up your mind as to what is the best course for the city. The decision tree is shown in Figure 12–8.

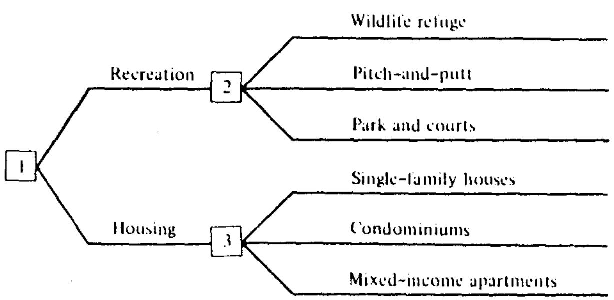

Fig. 12–8

In situations like this, the choice between recreation and housing is clearer if you follow the "work backward" system. Assume for the moment that you will recommend using the parcel for recreation and that you are therefore at decision node 2. You then consider what is the best recreational use for the land. Let's say that you decide that the park and courts are best. Turning next to decision node 3, you decide that if the land is to be used for housing, the best choice is mixed-income apartments. Ordinarily such apartments are not attractive to builders, but several builders have assured you that they are willing to undertake the project, given the sizeable federal subsidy. Your choice is thus narrowed to two contenders, the park and mixed-income apartments. In effect the decision tree has been pruned to that shown in Figure 12–9.

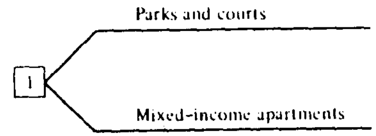

Fig. 12–9

Clearly there is nothing very startling about this process or its end result, and indeed deterministic decision trees such as this one hardly set one's normative blood aboil. Ordinal preferences, which merely provide rankings of alternatives, are always sufficient for finding the best choice.

Folding Back With Uncertainty: The Hospital Lawn

The more interesting and valuable applications of decision analysis are those in which chance plays an important role in the outcome of a decision. This is the case with the embassy official's decision problem described earlier. Let's look at another deliberately simple problem. A hospital administrator—we'll call her Harriet—must let a contract for reseeding the hospital lawn. She can give the contract to company *A*, which has agreed to do the job for $1500 provided that the weather is good over the next month; the charge will escalate to $2400 if the weather is bad. (Bad weather is defined as less than one or more than four inches of rain.) Or she can give it to company *B*, which has submitted a flat bid of $2000. Meteorological records indicate that there is a 15 percent chance there will be less than one inch of rain and a 25 percent chance of more than four inches. Thus there is a 60 percent chance that the weather will be good. What should Harriet do? The decision tree is shown in Figure 12–10.

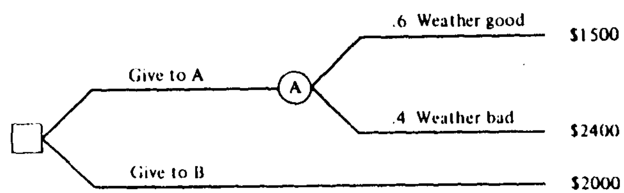

Fig. 12–10

Again we start to work backward from the tips of the tree, or rather Harriet does. Right away she comes to chance node *A*, shown in Figure 12–11. Now if Harriet could pick an outcome, she would choose to have company *A* do the work in good weather for $1500. But she cannot simply choose that outcome since nature is in control. If she gives the contract to company *A* she will have to take whatever comes in the way of weather; in effect she must accept a lottery. This means that she must find a way to assign a value to this node *as a whole*, a measure of what it's worth to her to be in a position where she faces a lottery with a 60 percent chance of spending $1500 on the lawn and a 40 percent chance of spending $2400.

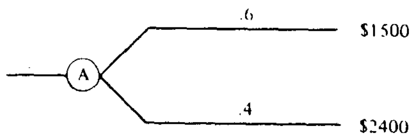

Fig. 12–11

One way she might do this is to determine some sort of an average value for chance node *A*. What do we mean by the average value of a lottery? It is what the average outcome of the lottery would be, in the long run, if the lottery were played again and again. This long-run average is aptly termed the *expected value*. In this case, where dollars are at stake, we usually refer to the *expected monetary value* or *EMV*.

The expected value may be calculated directly. It is found by multiplying the value of each of the possible outcomes at a chance node by its probability and then summing these products. For chance node *A*, we calculate

$$EMV_A = .6(\$1500) + .4(\$2400) = \$1860$$

The expected value if Harriet gives the contract to company *B* is, of course, $2000. The decision tree, thus, in essence reduces to Figure 12–12. If Harriet is willing to play the long-run averages, she should choose the action that offers the best expected value. In this case we're talking about costs; therefore the best EMV is the lowest and Harriet should give the contract to company *A*. A double line (//) is used to indicate that the choice "Give to *B*" has been eliminated.

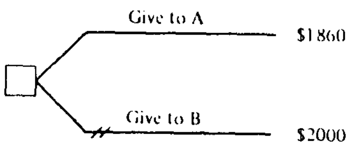

Fig. 12–12

A decision maker who bases his actions on expected values is said to be *risk neutral*. It is reasonable to assume that Harriet is risk neutral, for the amounts at stake in reseeding the hospital lawn are not large, at least not relative to the budget of a hospital in this expensive age. For the time being we will suppose all our decision makers to be risk neutral: this will keep the exposition simple. Later on we'll have more to say about the decision maker who is not risk neutral. In the meantime it may reassure you to know that, even then, the use of expected values usually provides a good first cut at a problem.

Anyway, if Harriet feels the hospital should be risk neutral with respect to the particular lottery it faces, that it should treat a $400 risk the way she herself would treat one for $5, her decision is clear. She gives the contract to company *A* and since hoping is free, if ineffective, she hopes for good weather. Some of you may wonder which path Harriet should follow if the outcomes achieved by following the two paths are equal. This point has been adequately covered in the literature.

> "Would you tell me, please, which way I ought to go from here?"
> "That depends a good deal on where you want to get to," said the Cat.
> "I don't much care where—" said Alice.
> "Then it doesn't matter which way you go," said the Cat.
> —Lewis Carroll, *Alice in Wonderland*

Before returning to the more complex generator problem, let's summarize what we have been doing. Essentially we have assumed that the decision maker is risk neutral, that is, willing to proceed on the basis of expected value. We have then relied on two procedures in working backward from the tips of the tree to the initial decision:

1. At each chance node, we have assigned a value, the expected monetary value or EMV, to the node as a whole.
2. At each decision node, we have chosen the action with the best EMV (the highest or lowest, whichever is appropriate for the case at hand) and have eliminated other possible courses of action.

Howard Raiffa refers to this whole process of working back along the branches of the tree as "averaging out and folding back."

As a check on your understanding of EMV and its use, assume that the embassy officer of Figure 12–1 is risk neutral. Where should he hold the party?[3]

Folding Back with Uncertainty and Learning: The Choice of Generators

Earlier in this chapter we asked you to try to draw the decision tree illustrating the choice between two types of generators for a power plant. We are now ready to go back and determine the preferred choice. The tree was shown in Figure 12–7. We can now average out and fold back from

[3] $EMV_{Indoors} = .6(100) + .4(-20) = 52$. $EMV_{Outdoors} = .1(500) + .9(-10) = 41$. Therefore he should hold the party indoors.

those payoffs to the initial decision. We strongly urge you to try this for yourself before reading on.

If you do indeed stop to work the problem out, you probably will start with the top branch. Working back from its payoff of 4.5, you come to decision node 2—where you can't make a choice until you know the EMV at chance node *C*. Fortunately, that EMV can be determined by inspection to be 7.5, which is entered above the node. The choice at node 2 is then between an outcome of 4.5 with old generators and an expected value of 7.5 with new. Since these are positive payoffs, the preferred choice offers the larger expected value and the old generator branch is crossed out. The same procedure is followed for the other branches; eventually the choice is reduced to Figure 12–13. The decision is clear; all that remains is to cross out the inferior choice, "Don't test." The completed tree is shown in

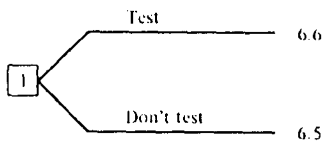

Fig. 12–13

Figure 12–14. The best strategy is to test; if it passes the test, the new type of generator should be installed. If not, the old type should be chosen.

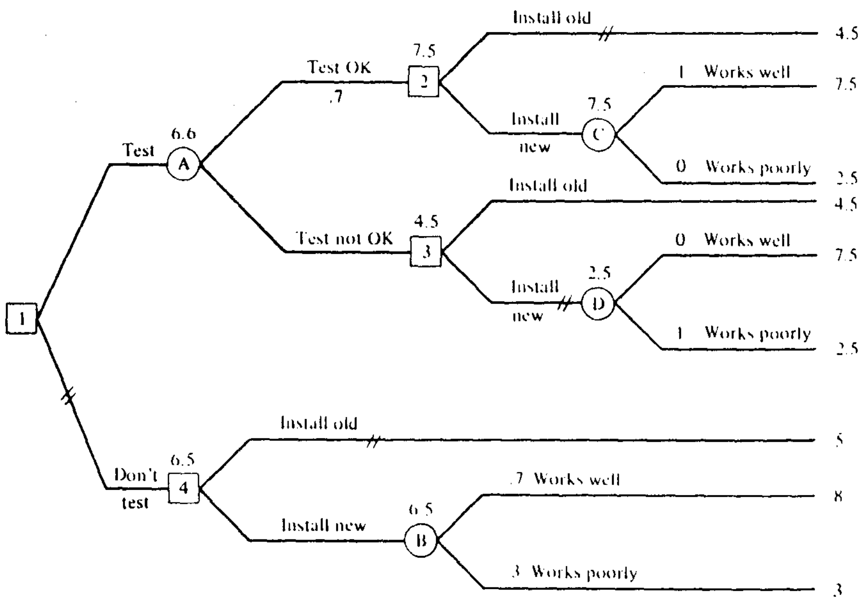

Fig. 12–14

Note that this tree describes a situation where the decision maker must make a first decision, then wait to acquire further information before making a second decision. In other words, he waits until he finds out how the probabilities break, until he knows along which path chance will lead him.

Decision trees thus do more than systematically lay out the opportunities for action and the uncertainties that will affect outcomes; they portray as well the gains from gathering further information. They induce, indeed almost force, the decision maker to consider sequential strategies, where further action will depend on the information observed to date. A major problem with government programs is that once started they are hard to stop, or even to modify. If policy makers used a decision tree at the outset, they might be modest about their ability to predict outcomes and would then be more likely to build information-gathering feedback loops into the decision process.

For example, suppose that you, an educational planner, must determine which of two reading techniques will be the subject of a three-year trial. You discover quickly that you have the option of using one in some classrooms and the second in others. Moreover, you can revamp the trial during its course as information on the two techniques begins to accumulate. An accurate representation of the situation requires drawing the tree so as to reflect *all* the various possibilities for setting up and revising the experiment. In this way decision analysis forces you to structure your choices so that you can react to information as it becomes available. In other words, it emphasizes flexibility in contrast to the construction of an immutable master plan.

Note also, if you haven't already, that although the value at each chance node depends on *all* the possible outcomes, the value at a decision node is the value of the preferred outcome only.[4] This, of course, reflects the fact that the decision maker is in control at a decision node, whereas he must accept whatever fate deals him at a chance node.

Before turning to variations on the basic model, let's look at a more complicated problem of a sort that occasionally appears in the press.

Using a Decision Tree to Structure a Complicated Problem: The Fish Ladder

A hydroelectric power plant on the Connecticut River has been ordered to build a fish ladder so that salmon can swim upstream beyond its dam. Three firms have submitted designs and cost estimates for the ladder.

Design *A* is the most expensive; it will cost $8.4 million and will take three years to build. A ladder of this type is already in successful operation in Oregon.

Design *B* is apparently similar, although it cuts a few corners to save construction time and money. It will cost $7.4 million and will take two years to build. We say "apparently" because although to the human eye it appears to capture the essential design features, the ichthyologists are reluctant to guarantee absolutely that the salmon will agree; they (the

[4] This is another of those points that are so obvious we're embarrassed to mention them. Yet experience has shown that students unfamiliar with decision analysis are likely to trip themselves up on just this point.

ichthyologists, that is) estimate the probability of success at .9. If for some reason the fish refuse to climb the ladder, the problem will become obvious by the end of the first year of operation. Modifications that will unquestionably satisfy the salmon can then be carried out at an additional cost of $2.8 million and a further delay of a year.

Design *C* is an altogether different type of ladder; the ichthyologists believe that it has only a .7 chance of success. It is far less expensive—$5.6 million—and will take only a year to build. It will take an additional year to determine whether or not it works. If it does not work, it will have to be abandoned and a ladder of type *A* or *B* will then have to be built. Although no fish ladders of type *C* are now in operation, one is currently under construction in New Brunswick, and by a year from now it will be known whether or not it is successful. If it works there, we can confidently expect it to work here.

Environmental and recreational benefits from a successful ladder are estimated at $1 million a year, whichever type of ladder it may be.

The design of the ladder must be approved by a state agency. The agency quite reasonably decides that its goal should be to minimize total cost, on the theory that costs to the utility will eventually wind up on the consumers' bills by way of rate increases, and the losses in environmental and recreational benefits due to construction delays will be borne by substantially the same group. The agency discounts both costs and benefits at 8 percent per year. It is anticipated that construction costs will rise at the rate of 10 percent per year.

Assuming that the three ladders will last equally long, and that the agency is risk neutral, what is the agency's best course of action? (For convenience, assume that construction costs are incurred when a ladder comes on line.)

This problem is typical of a situation where the decision maker desperately needs to keep track of what's going on and where he is in the process. It is, of course, an armchair case, with all the loose ends tidied up. But it hints at several problems that would be encountered in the real world, in particular the evaluation of intangibles (environment and recreation), allowance for the passage of time, and the possibility of acquiring further information. The probabilities are not more than "best guesses"—but if this is the best information you have, you should use it.

The decision tree for the fish ladder is shown in Figure 12–15, but before you turn to it you should try to formulate the tree and work out the

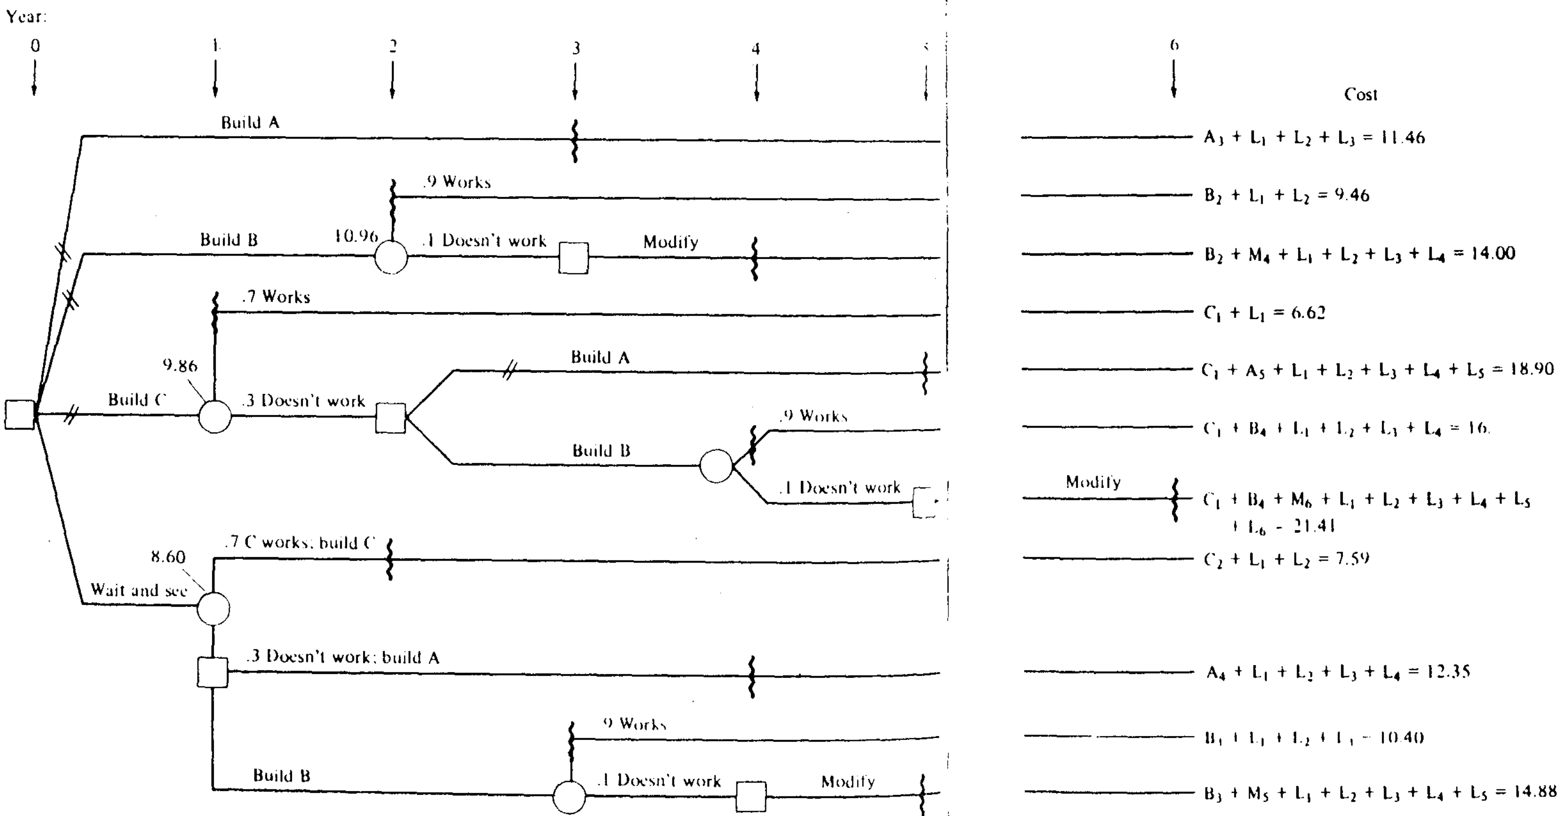

Fig. 12–15 *A*, *B*, *C* are costs ($ million). Subscripts indicate time at which incurred. *L* is loss of environmental benefits for subscripted year.

preferred choice for yourself. A new feature has been added, and we suggest that you incorporate it in your tree: the entire tree is put in a time frame, partly for clarity of exposition and partly because it makes it easier to figure the time lost along each route. A zigzag line (}) is used to indicate the point at which a fish ladder comes on line.

Allowing for Risk Aversion

If this were as far as decision analysis could take us, it would still be valuable as a conceptual framework for thinking about decision problems, as well as a direct guide to choice whenever the decision maker operates on an expected value basis. Fortunately, decision analysis also points us toward a systematic approach to situations in which the decision maker will not be risk neutral. In this section we take an intuitive look at this approach.

We saw above that it was plausible to assume that Harriet, the hospital administrator, was risk neutral, for the amounts involved were not large. But let's look at a more dramatic example. Suppose an individual—let's call him Henry—is compelled to make a decision that in effect amounts to choosing between the two lotteries shown in Figure 12–16. If he is risk neutral, he will choose lottery *A* because it has a larger EMV, despite the fact that he could lose a bundle—$10,000, to be exact. At this point you

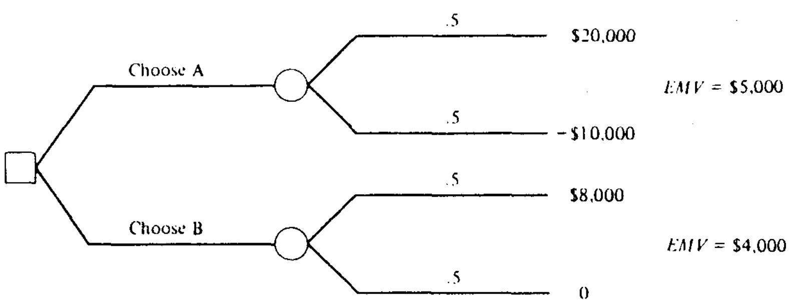

Fig. 12–16

may feel that risk neutrality is fine for Henry if that's the way he wants it, but it's not your cup of tea. In fact, the structure of your preferences is quite different. This is indeed a sensible reaction. Most people who are risk neutral when relatively small amounts are at stake are *risk averse* if the sums involved are large enough. If you prefer lottery *B* to *A*, you too are risk averse; most of your friends would probably react the same way. Yet the same people may on occasion exhibit *risk seeking* behavior. Anyone who has spent $1 for a raffle ticket that gives him a one-in-a-thousand shot at a $500 television set has been a risk seeker, although one might argue that the thrill of the game is an additional positive consideration.

A Problem That Introduces Risk Aversion: The Desalinization Plant

In our examples thus far, we have relied on expected value, in other words on the long-run averages, to see which of the available alternatives we should undertake. Let's take a closer look at what this implies. Consider the problem faced by the government of a small island that has insufficient fresh water. A desalinization plant is to be built, with a choice of technologies available. Technology *E* is well established; it will produce 6 million gallons of fresh water a day. Technology *A* is more advanced but less certain. There is a 50–50 chance that it will work well, in which case it will have an output of 10 million gallons a day. If it works poorly, the output will be 3 million gallons. The decision tree is shown in Figure 12–17.

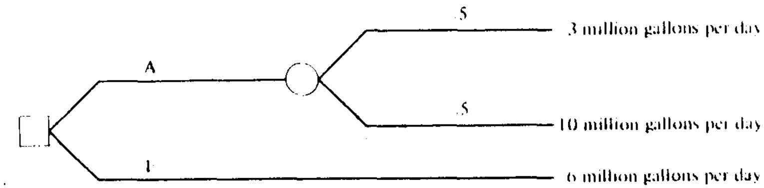

Fig. 12–17

If the government is interested in the greatest expected number of gallons, then it should choose Technology *A*. Indeed this would be the sensible action if, whatever the outcome, the island will be regularly purchasing additional supplies from some convenient source such as a pipeline from the mainland. (Thus far, this decision problem sounds like the electricity-generating problem discussed earlier.) The number of gallons at risk with the advanced technology may then be translated directly into dollars at risk should an additional purchase be necessary. That dollar amount will presumably be small relative to the total budget for the island. An expected value or averaging process would therefore be reasonable.

But what if this desalinization plant is to be a primary source of water for the island (and this is where this example departs from the generating plant example), and if additional water can be brought in only by an expensive process, say, by tanker? Then we will need to adopt a different approach. Although there will certainly be gains in producing 10 million gallons a day rather than 6, the uses to which the last gallons are put will be less valued than earlier uses. Perhaps the last million gallons a day will be used on lawns and fairways. In contrast, a drop to 3 million gallons a day would mean curtailing much more valuable uses of the water, such as irrigation of vegetable crops.

Suppose, for example, that experience with droughts in the past indicates that the first 3 million gallons to be produced by the plant will be worth $10,000 to the islanders, a supply of 6 million gallons will be worth $18,000, and one of 10 million gallons is worth a total of $25,000. The decision tree now becomes Figure 12–18. Looked at in this light, the traditional technology is superior. In other words, the loss in value if output is reduced from 6 million to 3 million gallons a day is $8000, which is greater than the gain of $7000 that would result from an increase in output from 6 million to 10 million gallons a day. Therefore the average dollar value achieved by pursuing Alternative A is less than the average dollar value achieved by pursuing Alternative E. With average dollars as our criterion of choice, Alternative E should be chosen. Moreover, if the island government is in addition risk averse with respect to money lotteries, the strength of this preference for Alternative E would be increased, for Alternative E is a sure thing.

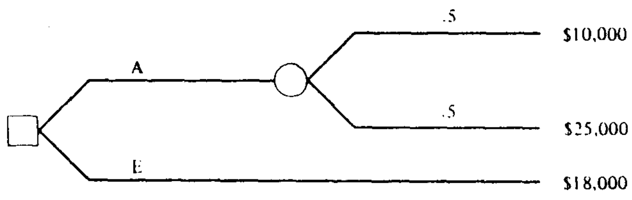

Fig. 12–18

This example shows that relying on the average output of some quantity such as gallons per day may be inappropriate. The gains from the same gallon increment may not be the same across different ranges of values. In this case it was legitimate to convert gallons to dollar equivalents and to use expected dollars as our criterion of choice. The islanders were willing to be risk neutral with respect to dollars, at least for the magnitude of dollars involved, but they were not willing to be risk neutral with respect to quantities of water.

But just as the value of gallons may diminish as we get more of them, so too the real worth to us of a dollar may not be the same at all dollar levels. Suppose you are offered a choice between $10,000 for sure, and the flip of a coin to determine whether you get $0 or $25,000. You might well prefer to receive $10,000, despite the fact that the lottery offers an expected dollar value of $12,500. The simple lesson is that when a gamble may substantially alter one's endowment of a vital commodity, whether dollars or gallons of water or health or environmental amenities, it may be inappropriate to use expected values in making a choice.

Economists and decision theorists have developed methods for handling situations where expected dollars are not a suitable criterion for choice because there is a substantial variation in the value of dollars achieved under alternative outcomes. The basic approach involves converting dollar outcomes to an artificially constructed scale of values called utilities. The methodology is called *utility theory*. It can appear deceptively simple, but in fact the scaling process is rigorously defined and ensures that the individual will be risk neutral with respect to lotteries that are expressed in terms of these utilities. As we converted gallons to dollars in the desalinization plant example, so utility theory converts dollars, or any other measurable unit, to utility units. We shall not deal with utility theory here; it is left to an appendix at the end of this chapter. The lazy reader will take comfort from our confession that we have encountered few actual public policy analyses where utility theory was explicitly employed, though some might claim that its use is implicit in structuring a graduated income tax. The budding decision analyst will master that appendix, hoping perhaps to be among the first to apply this important technique to the assessment of critical policy issues. One major application of the utility theory approach, was addressed to the appropriate location of the new Mexico City airport.[5] It had a significant impact on the Mexican government's ultimate policy choice.

The Value of Information

We return again to the problem in which a decision maker was required to choose between two types of generators for a power plant. The output characteristics for the old type of generator were known with certainty; for the new they were not. The decision maker also had the option of acquiring accurate information about the new technology by having a prototype built and tested. We found that the latter was the best choice, with the final decision relying on the results of the test. (The complete decision tree is shown in Figure 12–7.)

Is it always advisable to gather additional information before making a decision? If you think for a minute, the answer clearly must be no. First and most simply, it may be that the information, whatever it is, cannot possibly change the decision. Second, even when a decision might be altered, the cost of acquiring the information in terms of resources and delays may be more than the information is worth. For example, the director of the dam project would not have found it worthwhile to spend $2 million to test the prototype generator, because the expected value of the "Test" branch would then be only $5.1 million, significantly less than the $6.5 million EMV of the "Don't test" branch. Just how much should the decision maker be willing to pay?

We can work this out using our decision tree. Let T be the cost of the test in millions (we will get back to specific values of T in a moment). Then the payoffs (reading from the top of the tree down) are $5 - T$, $8 - T$, and so on. If we average out and fold back, the tree reduces to Figure 12–19. Note

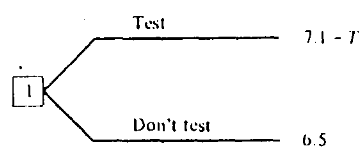

Fig. 12–19
(All payoffs are in millions of dollars.)

[5] See Ralph L. Keeney, "A Decision Analysis with Multiple Objectives: The Mexico City Airport," *Bell Journal of Economics and Management Science* 4 (1973): 101–17.

that we don't need to know the value of T in order to make a choice at decision node 2. The choice is between "Install old" for a payoff of $5 - T$ and "Install new" for a payoff of $8 - T$. Whatever the value of T, "Install new" is better. But at decision node 1, the choice depends on what T is. Clearly the decision maker should choose "Test" as long as the cost of the test is less than \$.6 million; if it is greater than \$.6 million, he should install the new type of generator without testing. Thus for any amount up to \$.6 million, the test is a good buy; anything more would be a mistake. Hence the test is worth \$.6 million to the decision maker. We call this amount the Expected Value of Perfect Information, or *EVPI*. A decision maker who can undertake a costly perfect test should be willing to pay up to his EVPI for it. In general, the EVPI is the difference between the EMV of the "Test" branch for a perfectly accurate test and the EMV of the next best branch.

Another Illustration of the Value of Information: The Metropole Subway System

Metropole, a large subway system, is considering the installation of automatic train-speed controls on its new branch. The savings in personnel costs will be significant. Two types of controls are available. The type sold by Venerable Engineering has been thoroughly tested in use. It will cost \$14 million. A much cheaper alternative is produced by Innovative Technology. Unfortunately, it has not been proven in practice. Metropole believes that it is 60 percent likely that Innovative's controls will work. If that system is purchased and it then fails, all payments made to Innovative will be refunded. But Metropole will have to fall back on a nonautomated system for six months while the Innovative controls are removed and the Venerable system installed. Extra personnel costs during these six months will amount to \$3 million. Innovative Technology is eager to pull off a sale to a big subway system and offers its speed controls to Metropole for \$10 million. The decision tree (for pedagogic simplicity, we ignore discounting considerations) is shown in Figure 12–20.

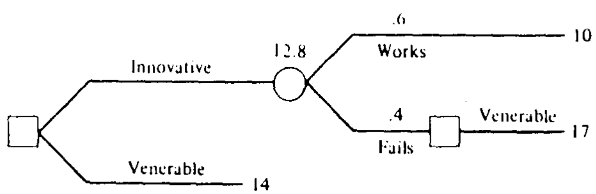

Fig. 12–20
(All payoffs are in millions of dollars.)

Note that in this decision problem we are seeking to minimize expected costs, and hence should choose the lowest expected value at each decision node. If Metropole is risk neutral, and we shall assume that it is on gambles of this size, then it should adopt the Innovative system.

Another possibility is available. A simulated system could be developed to create the equivalent of a field test for the Innovative controls. Metropole has not yet received firm bids for this field test. (Before it does, you might want to work out how much it should pay at the most.) Assume that the test could be run in the next few months, that it is completely reliable, that if it succeeds the Innovative system will be installed on schedule, and that if it fails the Venerable system will be available at the time it is needed. The decision tree has now acquired another branch. Before looking at Figure 12-21, try to work out for yourself how much Metropole should be willing to pay for the test.

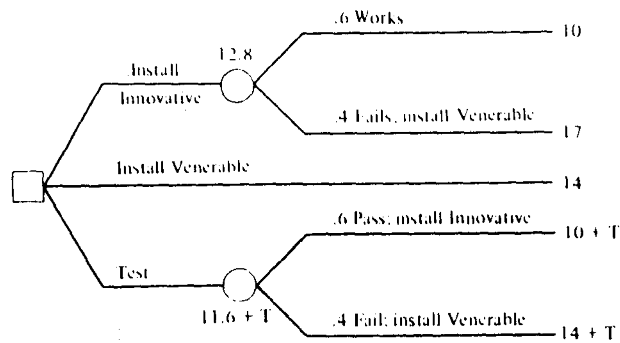

Fig. 12–21

The EMV of the "Test" branch is \$11.6 million plus the cost of the test, T. The EMV of the "Don't test" branch is \$12.8 million. Therefore the EVPI is $(12.8) - (11.6) = 1.2$ million dollars. Metropole should be willing to pay up to that amount for the test.

Perfect information is rarely to be had through a straightforward expenditure of dollars; we can acquire it, if at all, only by waiting, and waiting carries its own costs, costs that must be considered as part of the test costs. For example, in our subway example, if the conditions were different, the test might have delayed the installation of automated speed controls by one month. This then would have added \$.5 million in personnel costs to the system, whatever the outcome of the test. Even when they occur in intangible form, it is usually helpful to think about information costs and how much the information is worth.

Drawing Inferences from Imperfect Tests

It was easy to trace the implications of the tests we considered for the prototype generator and the speed control system, for we assumed that the tests were perfectly accurate. The real world is usually less cooperative, and is less precise about revealing the true situation. Our devices for gathering information turn out to be tests that are imperfect in a variety of ways. Must a test be wholly reliable in order to be of value? Any test whose results may change a decision has some value. (Poker players make mighty efforts to catch the slightest clue or hint of information from their fellow players.) Naturally, the more imperfect the test, the less we should be willing to pay for it. The EVPI serves as an upper bound for testing costs; if information is not completely reliable, we should pay at a maximum somewhat less than that amount.

Let's stick with our Metropole example a bit longer; we now assume that the simulated field test is imperfect. If the test result is "Fail," the controls will surely fail in operation. But "Pass" is less reliable. The controls are expected to pass the test 80 percent of the time, but the pass result is accurate only 75 percent of the time; one time in four when the controls pass the test they will fail in practice. This implies that the controls will work in practice 80 percent × 75 percent = 60 percent of the time. How much should Metropole be willing to pay for this imperfect test?

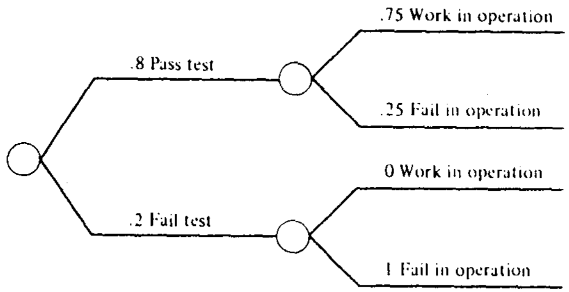

Fig. 12–22
(All payoffs in millions of dollars; IT is the cost of the imperfect test.)

We can summarize the test information in an *event tree*, as in Figure 12–22. (An event tree has only chance nodes; such trees have been used notably in studies of nuclear reactor safety.) This minitree is then incorporated in a full decision tree, as is shown in Figure 12–23. Working backward from the tips of the tree, we immediately see that "Install Innovative without testing" has an EMV of 12.8, and "Install Venerable " an EMV of 14. These figures are the same as in the earlier version of the problem. As before, Venerable controls are eliminated from further consideration because the expected cost is higher than for Innovative. The "Test Innovative" branch is more complicated, but if we average out and fold back we find that its EMV is $12.2 million plus the cost of the imperfect test, *IT*. Metropole should be willing to pay up to $.6 million for the imperfect test.

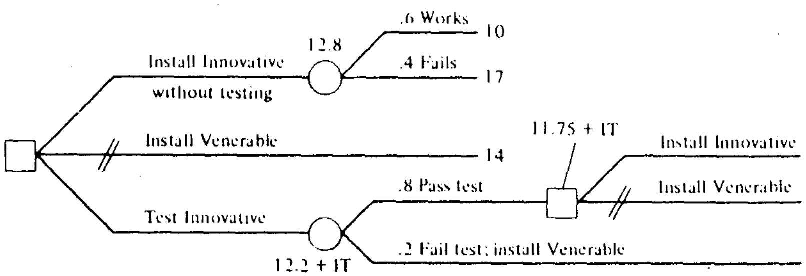

Fig. 12–23

Putting Imperfect Test Information into Usable Form: Abused Children

In the example just considered, we were supplied with all the test information we needed. But frequently the information isn't as neatly arranged as this. Consider the following situation, based on an investigation of the problems in identifying abused children.[6]

School officials believe that 3 percent of a city's 10,000 school children are physically abused. Measures can be taken to help these children, but first they must be located. It is proposed to carry out a preliminary screening of all children; when evidence of abuse (such as bruises of a certain type) is found, interviews with parents will follow. Unfortunately the screening process is not entirely reliable. If a child is actually abused, the chance is 95 percent that the test used for the screening will be positive, i.e., will indicate that the child is abused. On the other hand, if the child is not abused, the chance is only 10 percent that the test will be positive.[7] School officials are most anxious to identify cases of abuse, yet must proceed cautiously given the enormous stigma that attaches to parents who are falsely accused. The tradeoffs are painful, and it is imperative that school officials understand the implications of a positive or negative test.

The difficulty is that the information is not in a shape that's readily usable. However, we have all the numbers we need; we just have to perform a few calculations and rearrange them. Figure 12–24 sets forth what we now know.

But it doesn't serve any purpose to screen a child *after* we know he is

[6] Richard Light, "Abused and Neglected Children in America: A Study of Alternative Policies," *Harvard Educational Review* 43, no. 4 (November 1973).
[7] For purposes of illustration, abuse and the test for it are assumed to be yes–no variables (the technical term is "binary"). We all know that this is an oversimplification; there are many degrees of abuse, and evidence of abuse varies in its degree of ambiguity. Introducing this additional complexity would cause no conceptual difficulty, though more arithmetic would be required to get the information into usable form. Note, moreover, that when it comes to taking action on behalf of abused children, both the screening test and the actual existence of abuse may well be viewed as binary. Authorities may determine that even though the results of the screening may be reported along a continuous scale, scores above a certain cutoff point will be regarded as indicating abuse and will result in direct intervention with the family.

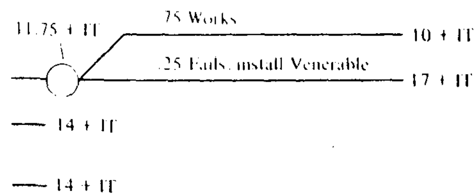

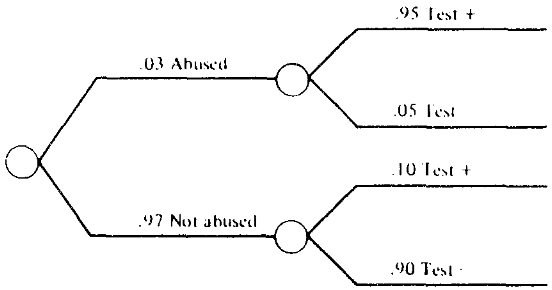

Fig. 12–24

abused. We need to know what the likelihood is that a positive test means that a child is indeed abused, or that a negative test means that he is not. In other words, we want the information in the form of Figure 12–25.

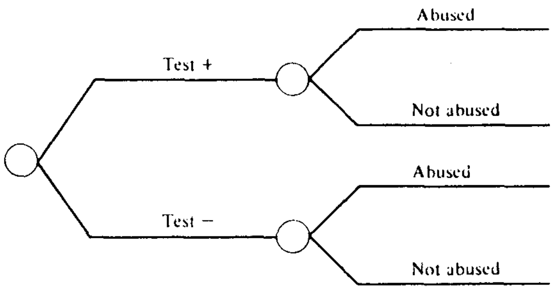

Fig. 12–25

Look again at the tree shown in Figure 12–24, and think about what it means to follow along the topmost branch: it means that the child is abused and at the same time the test is positive. The probability of this combined event—the joint probability—is .03 × .95 = .0285. It is called the *path probability* because it's the probability of following the whole course of one particular path on the decision tree. (Granted, in the interest of consistency it might be called the branch probability, but it isn't.) Look next at the topmost path in Figure 12–25; if we follow along this path it means that the child's test is positive and he also is abused. But this combined event is just the same as the combined event "is abused and has positive test"; the order is immaterial. Hence the joint probability must be the same, .0285. Similarly, each of the other paths on the second tree is the counterpart of a path on the first tree, although the order from top to bottom is not the same on both trees. The probabilities for each path are shown at the tips of the branches in Figures 12–26*a* and 12–26*b*.

Now comes the crucial step. Notice that *if* a child's test is positive it must be that either he is abused (with probability .0285 for the joint event "test positive, is abused") or he isn't (with probability .097 for the joint event "test positive, is not abused"); the total probability of a positive test is the sum of these, .1255. Similarly, we find that the probability of a negative test is the sum of the probability that the child is abused and has a

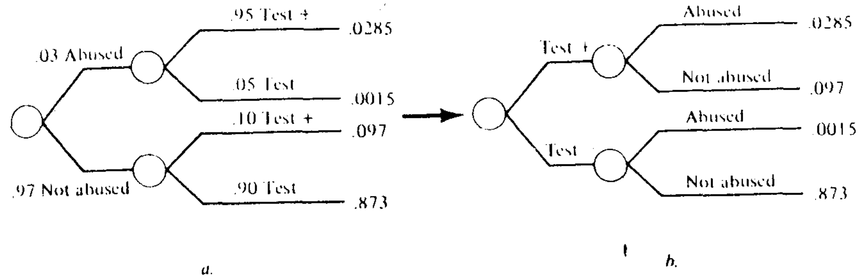

Fig. 12–26

negative test (.0015) and the probability that he is not abused and has a negative test (.873), or .8745. Entering these results on Figure 12–25's target tree, we have Figure 12–27.

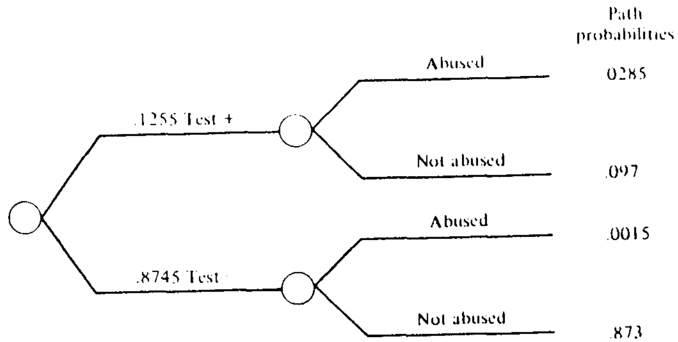

Fig. 12–27

It remains only to fill in the probabilities at the final two chance nodes, and this is a matter of simple arithmetic. If the probability of a positive test is .1255, and the probability of the path "Test +; abused" is .0285, then the probability that the child is abused, *once his test is positive*, is .0285/.1255 = .2271. Following the same procedure, we fill in the rest of the tree, as in Figure 12–28.

This event tree is in a form that is of use to us; it is what we have been aiming for. The whole process is called "tree flipping"; it is the intuitive version of Bayes' formula.[8] The tree flipping technique offers many advantages. It's simple; there's no formula to forget or foul up; and most

[8] For those who are more comfortable with formal mathematical notation we print the usual Bayes' formula:

$$p(A \mid +) = \frac{p(A)p(+ \mid A)}{[p(A)p(+ \mid A)] + [p(\text{Not } A)p(+ \mid \text{Not } A)]}$$

In this notation $p(A)$ is the overall probability (the "prior" probability) of abuse. The vertical line | indicates a conditional probability; $p(A \mid +)$, for example, is to be read "the probability of abuse, given that the test is positive." The formula was originally set forth by the Reverend Thomas Bayes, an 18th-century English cleric. We strongly urge you to rely on tree flipping for updating probabilities.

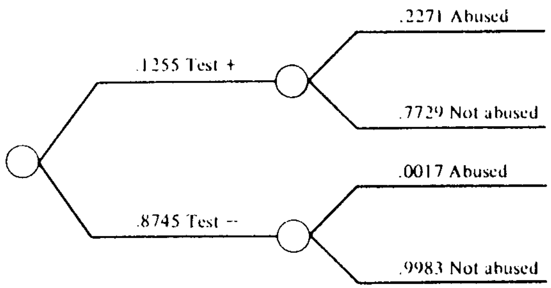

Fig. 12–28

important, it works well for more complex situations, such as those in which there are more than two underlying conditions. Suppose, for example, that the underlying conditions are no evidence of abuse, strong evidence, and ambiguous evidence, while the test reports are +, −, and ?. The formula is a mess in such cases, but the tree is crystal clear.

Practice in Tree Flipping: A Medical Test

A doctor must treat a patient who has a tumor. He knows that 70 percent of similar tumors are benign. He can perform a test, but the test is not perfectly accurate. If the tumor is malignant, long experience with the test indicates that the probability is 80 percent that the test will be positive, and 10 percent that it will be negative; 10 percent of the tests are inconclusive. If the tumor is benign, the probability is 70 percent that the test will be negative, 20 percent that it will be positive; again, 10 percent of the tests are inconclusive. What is the significance of a positive or negative test?

The information immediately available to the doctor is described succinctly in the event tree in Figure 12–29. But it doesn't serve any

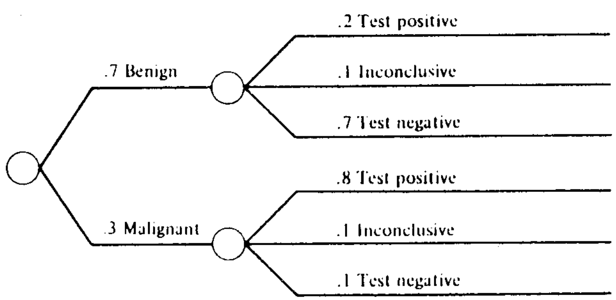

Fig. 12–29

purpose to perform the test after we know whether the tumor is benign or malignant. The doctor needs to know the probabilities for the event tree in Figure 12–30.

Try to work out the answer for yourself before looking at the complete conversion in Figure 12–31.

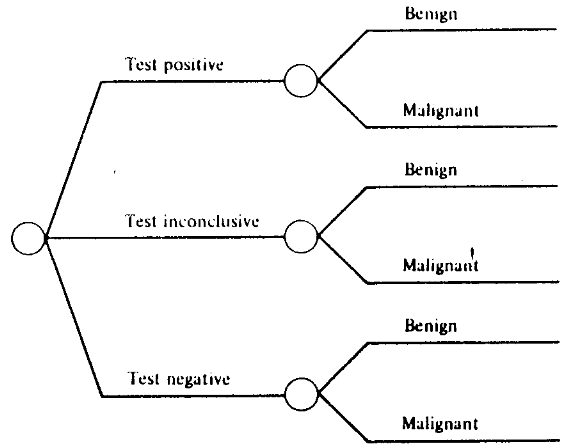

Fig. 12–30

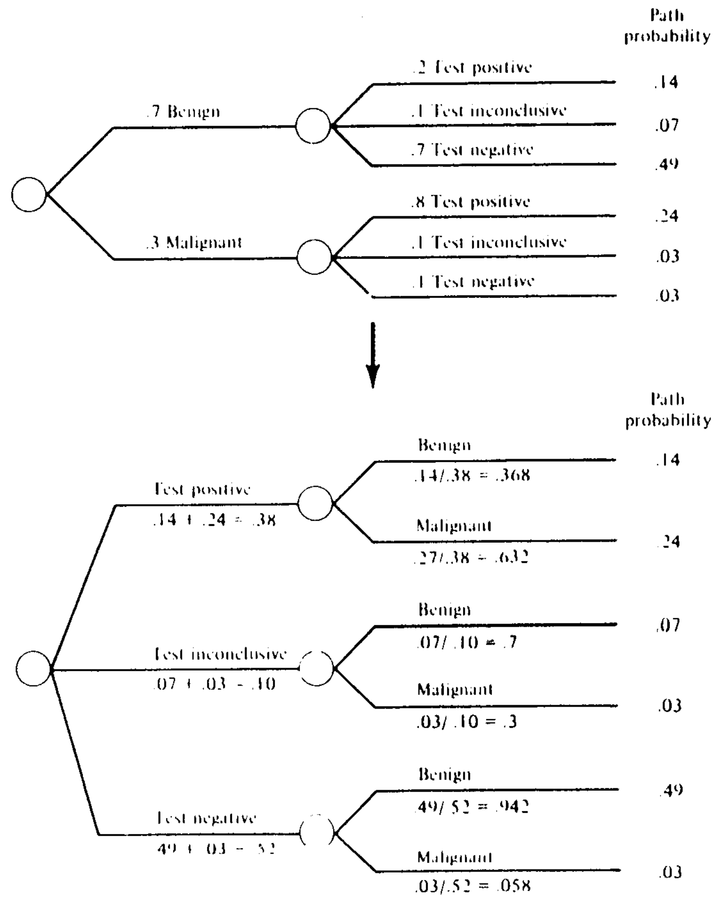

Fig. 12–31

Imperfect Tests and Intangible Payoffs

Both the child abuse and the tumor tests ultimately lead to decisions about subsequent courses of action that cannot be quantified in any clearcut fashion. Nevertheless decision analysis poses the nature of the decision maker's dilemma very clearly. For example, suppose for simplicity that good medical practice requires that a malignant tumor be treated surgically and a benign tumor left alone. Then four general types of intermediate outcome are possible: (1) the patient has a necessary operation; (2) the patient undergoes surgery unnecessarily; (3) the patient receives no treatment and none is needed; (4) the patient fails to receive necessary treatment. For each of these intermediate outcomes, the patient will have certain probabilities of recovering and dying. (In the long run, of course, the probability of dying is 1.0. We are talking about the short run.)

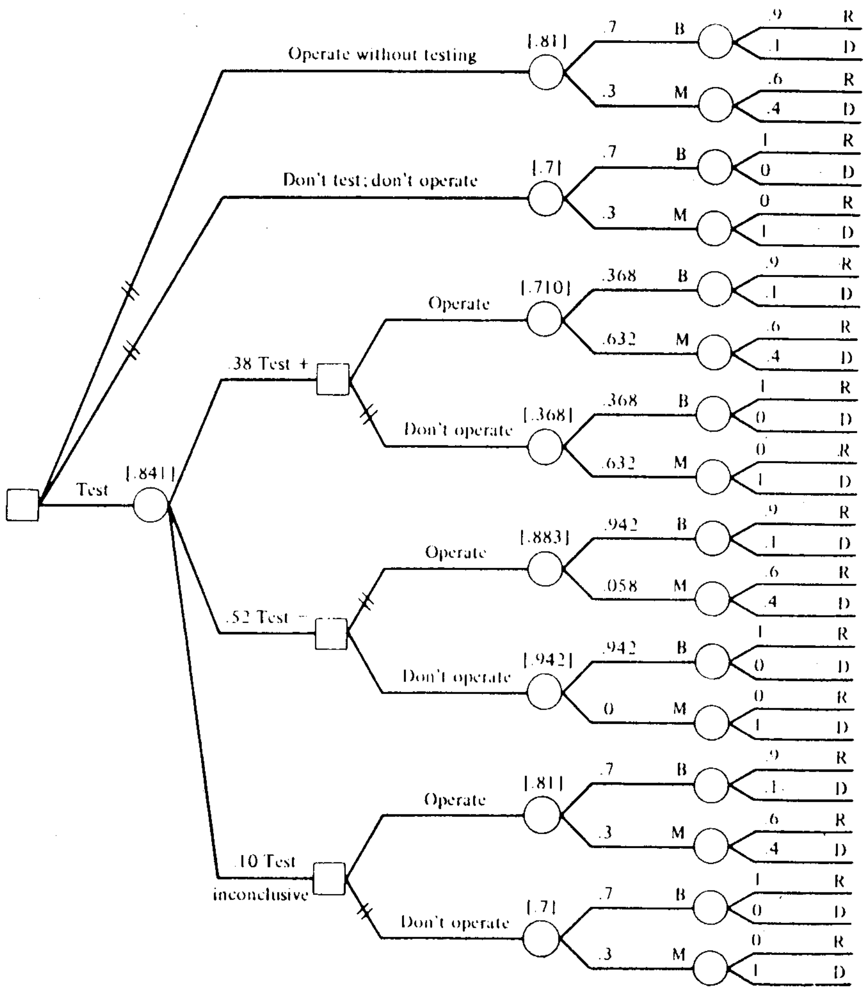

Fig. 12–32

Decision analysis cannot help the doctor determine the value of each of these outcomes. But it does remind him starkly that whatever decision he makes, given what he knows about the probability that a tumor is benign or malignant and the likelihood of recovery or death in each situation, he will implicitly set boundaries on the values he assigns to these outcomes. The decision tree is shown in Figure 12–32. The benign/malignant probabilities are transposed from Figure 12–31. If the recover/die probabilities can be estimated, the course that gives the patient the best chance of recovery can be directly determined. We have used conjectural probabilities for recover (*R*) and die (*D*) in Figure 12–32; the numbers in brackets [ ] are the resulting chances of recovery. Death of the patient due to other causes is excluded from consideration, in order to simplify the exposition; additional chance nodes could easily be added to allow for such contingencies. Similarly, "recover" is of course much too broad a category; in practice we might well insist that the prospective quality of the patient's life be taken into consideration.

The child abuse problem is similar. What should school officials do once the preliminary screening indicates that a child is abused? Perhaps their first instinct is to pursue vigorously each suspected case of abuse. The relevant part of this event tree is shown in Figure 12–33. In other words, the chances are greater than three in four that, if a suspected case is followed up, the parents will turn out to be innocent. The school department may well feel that this is hardly the way to win friends for its programs, or do good for its community. It must weigh the severe harm of permitting an occasional case of child abuse to continue against the serious costs of falsely implicating parents. In practice this analysis would probably lead the department to schedule careful follow-ups after symptoms of abuse are detected. Or if test outcomes were in the form of several gradations from negative to positive, as might well be the case in the real world, perhaps only the strongest positives would be pursued.

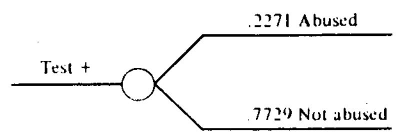

Fig. 12–33

These treat–don't treat dilemmas are encountered in many guises. Even when unnecessary treatment is not positively harmful, it is likely to be expensive and time-consuming. Utility theory offers each decision maker a way to quantify the value of these risky and intangible outcomes in terms of his own personal preferences; it is discussed in the appendix to this chapter.

Decision Analysis and a Contemporary Policy Issue

One of the most complex issues currently facing the United States is the question of how electricity is to be generated for the next half century. The choice between continued reliance on the light water nuclear reactor and development of the breeder (more accurately, the liquid metal fast breeder reactor) is one aspect of that controversy. A study of the optimal timing of research and development (R&D) for the U.S. breeder reactor program illustrates the tremendous capabilities that decision analysis offers for clarifying such a problem.[9]

The Clinch River breeder reactor, a demonstration plant located near Oak Ridge, Tennessee, is scheduled for completion in 1983, provided Congress does not cut off its funding. The next step planned for the breeder program is to build a large breeder reactor that would serve as a commercial prototype, to be followed in turn by the first commercial breeder reactor. The study examines four options: (1) concurrent development of Clinch River and the commercial prototype, with a decision to be made in 1986 as to whether we should then proceed with a full-scale commercial breeder; (2) sequential development, with Clinch River to be developed now, a decision on the prototype to be made in 1986, and a decision on the commercial breeder in 2005; (3) waiting until 1986 before making a decision on both Clinch River and the prototype, with the concurrent development route then pursued if information is favorable; and (4) stopping development now and ceasing R&D. The first stage of the decision tree is shown in Figure 12–34. The dates at the tips of the branches indicate when the first commercial prototype would come into use.

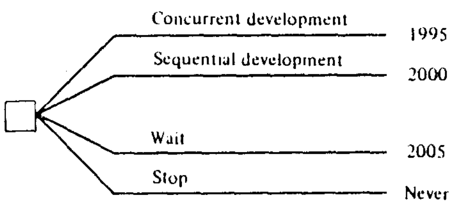

Fig. 12–34

In this analysis, several key uncertainties affect the outcome of each decision and hence are included in the complete decision tree. These uncertainties relate to the success of the initial breeders, the supply of uranium, future energy demands, capital cost differentials between the light water reactor and the breeder, the future availability of other advanced technologies, breeder R&D costs, and public reaction to the issues of nuclear safety and potential environmental damage. Benefits take the form of lower energy costs in the future.

The Atomic Energy Commission (AEC, later ERDA, the Energy Research and Development Administration), on the other hand, favors a deterministic model to investigate the ramifications of the R&D decision, using a number of possible scenarios for the future. Net benefits are then calculated for each scenario. The AEC thus relies on a once-and-for-all masterplan, with a commitment to an entire timetable. In contrast, the decision tree approach would permit the AEC to set up a decision process that enables it to take advantage of new information as it becomes available. As a result of the scenario approach, those opposed to the further development of the breeder have focused on the scenarios that make it look bad, while the proponents have stressed those that lead to a recommendation of concurrent development. No sense of the likelihood of each scenario emerges from the analysis, and hence no estimation of expected net benefits is possible. Both types of scenario are plausible; the debate continues.

One of the interesting features of the decision analysis study of this issue is its treatment of environmental and safety issues. Rather than addressing them directly, all such issues are subsumed under a single uncertainty, the likelihood of a moratorium on further use of nuclear power, assuming implicitly that a moratorium will be invoked when environmental and safety costs are high. This approach has several virtues. It limits the scope of the analysis; in effect it says only, "*If* we decide to go ahead with the breeder, this is the economically effective way to go." In thus separating dollar outlay and benefit questions from the crucially important intangible considerations, the tradeoffs between dollars and potential threats to safety and the environment are made explicit. Focusing on the likelihood of a nuclear moratorium also recognizes that, since the parties to the nuclear debate will never reach agreement on the safety of the breeder, the critical issue is what people believe about its safety.

Figure 12–35 shows the essential features of the complete decision tree. A programming model is used to derive estimates of the relevant costs and benefits, and probabilities are carefully assessed. The tree is then folded back and we find that if the decision is to be made on economic grounds alone, the concurrent development strategy is preferable, though not by a wide margin. Not to leave you in suspense, the expected values (discounted at a 10 percent annual rate) of the four alternatives are shown in Table 12–1.

You may argue that the analysis is all well and good, but given the political mood in this country and the widely expressed concerns about reactor safety and spent fuel disposal, the large outlay implied by concurrent or even sequential development is no longer feasible. Decision analysis makes it relatively easy to evaluate altered situations; of course, it may be necessary to fold back the tree afresh. Sometimes it's simply a matter of introducing revised probabilities or refining the estimates of the

Table 12–1

Concurrent development	$13.8 billion
Sequential development	$12.5 billion
Wait	$11.5 billion
Stop	$ 0

[9] Richard G. Richels, *R & D under Uncertainty: A Study of the U.S. Breeder Reactor Program* (Cambridge, Mass.: Energy and Environmental Policy Center, Harvard University, 1976).

payoffs. At other times whole branches must be struck off the tree. Since the nuclear R&D analysis was completed, for example, the breeder has been put on hold. Sometimes the exploration of new options will require us to add a bit of vegetation. Here again, the required revisions in the analysis

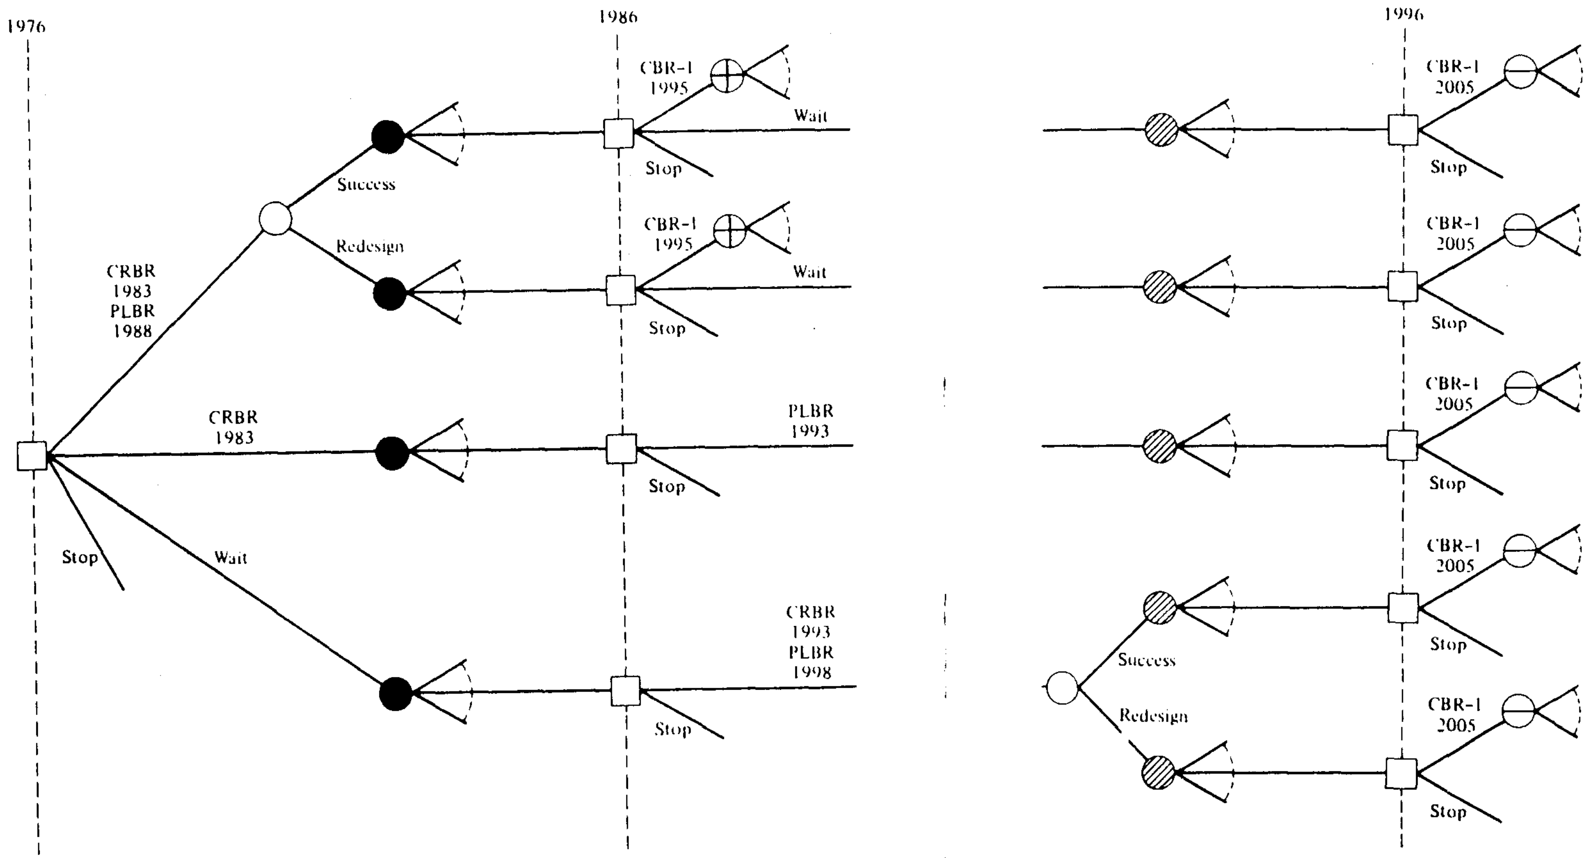

Fig. 12–35
CRBR = Clinch River Breeder Reactor
PLBR = Prototype Large Breeder Reactor
CBR–1 = First Commercial Breeder Reactor
Dates are estimated completion times.
Chance nodes:

1. Design succeeds or redesign necessary
2. Combined uncertainties about nuclear moratorium, demand projections, and uranium availability
3. Combined uncertainties about capital cost differentials, availability of other advanced technologies, and more up-to-date estimates of uranium supply
4. Still later estimates of uranium supply
5. Combined uncertainties (at a later date than above) about capital cost differentials and availability of other advanced technologies

are straightforward. Finally, decision analysis also permits us easily to investigate the sensitivity of the conclusions reached to the quantitative assumptions made, a matter of particular interest in the nuclear debate. We turn now to a brief consideration of that aspect of decision analysis.

Sensitivity Analysis

Sometimes experts in a particular field object to the conclusions reached by a decision analysis. If they do, their objections should be rooted not in the methodology but in the assumptions that guided the construction of the tree. Thus an expert on energy might claim that the study of research and development strategies for nuclear power plants discussed in the preceding section neglected certain alternative strategies, such as proceeding full speed ahead to develop solar power. Perhaps some aspects of the payoffs were understated; maybe the future availability of uranium was misestimated. Or perhaps some of the probabilities were miscalculated. The estimate of the likelihood that the first design for the Clinch River breeder reactor would prove successful by 1986 was in fact much too high.

Assume for the moment that at least some of these claims are correct. Should we blame decision analysis? No, it is merely a tool. To discredit the method because it is misapplied would be equivalent to blaming the manufacturer of the architect's table and T-square for the design of an inadequate structure. Indeed, a more forceful defense can be made for decision analysis. It is particularly well suited for examining the effect of changing some of the critical underlying assumptions. This process is called sensitivity analysis.

Usually the probability distributions and payoffs used in a decision analysis are estimates rather than hard numbers. In fact, sometimes they are little more than informed guesses. Nevertheless, if this is the best information the decision maker can come up with, he should go ahead and use it. In such circumstances he naturally would like to know how sensitive his final decision is to the estimates he has used. If he finds it to be very sensitive, he should then spend more time refining the estimates. Decision analysis lends itself very well to a careful calculation of that sensitivity. Here is a very simple example.

A risk neutral decision maker must choose between two construction sites. Site preparation for Location I will cost $85,000. For Location II the cost is iffier; the decision maker thinks there is a 60 percent chance it will cost $100,000, but if he's lucky it will cost only $40,000. The tree is shown in Figure 12–36.

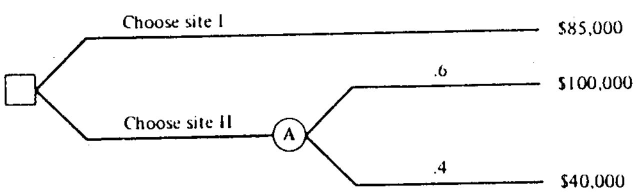

Fig. 12–36

The EMV at chance node *A* is $76,000, and so Site II appears to be the preferred choice. But in fact the probabilities .6 and .4 are based on limited information; the decision maker believes they could be refined significantly with a little effort. How sensitive is his decision to these probabilities? To put it another way, how different would they have to be to change his decision? He can find out by calculating the probabilities for which Location I would be preferred to Location II. As the probability of a cost of $100,000 increases, Site II will become less and less attractive, until for some probability p the decision maker is neutral between the two sites. Thereafter he will prefer Site I. Let's find the p for which he would be indifferent, for which the expected cost of the two alternatives is the same. That p will satisfy the equation

$$p(100{,}000) + (1 - p)(40{,}000) = 85{,}000$$

which yields $p = .75$. In other words, whenever the probability that the cost at Site II is $100,000 exceeds .75, Site I will be preferred. Now the decision maker may feel unsure that p is .6, but at the same time feel reasonably certain that no information he gathers will push it higher than .7. If this is his belief, he should proceed with Location II; no benefit will result from gathering further information even if it is costless. If the sensitivity analysis still leaves him uneasy with the decision, he can always figure out what it would be worth to get more information before proceeding. And perhaps he should reexamine his assumption of risk neutrality.

A similar procedure can be used to test sensitivity to payoff values of other critical parameters. Let's assume that at Location II getting clear title to the land will involve legal expenses. The decision maker might wish to determine just how large those expenses can be before he should choose Location I.

In discussing the operate/don't-operate decision described on page 228, we used conjectural probabilities for recovery and death. It so happens that with these particular probabilities, "test" narrowly edges out "operate without testing." If we perform a crude sensitivity analysis (i.e., take a good hard look at the numbers) we see that this result comes about because the mortality rate when the tumor is benign is 10 percent. Had we hit upon a mortality rate of 0.1 percent, the preferred choice would have been "operate without testing," for in that case an incorrect negative test presents a greater danger than the operation.[10]

In the choice of energy reseach and development strategies, a critical parameter is the discount rate. If the discount rate used is greater than 10 percent, strategies that get substantial R&D programs underway immediately look less desirable. If it turns out that the optimal R&D strategy shifts when the discount rate rises from 10 to, say, 10.2 percent, then much more time should be spent attempting to pin down the current rate. But if the preferred choice does not change until the rate exceeds 18 percent, then the objection that "things came out that way only because they picked a low discount rate" would quickly be shown to be misguided. In such a situation, we say that the conclusion is *robust* with respect to variations in the discount rate. Sensitivity analysis is thus potentially a powerful tool for policy analysis and debate.

In the breeder study, sensitivity analysis was used to check the probability assessments, and all were found to be robust with the exception of the likelihood of a nuclear moratorium. It was found that if concurrent development is ruled out for political reasons, sequential development is the preferred economic alternative provided the probability of a moratorium is less than .63. If that probability is between .63 and .86, the best strategy is to wait. If it is greater than .86, then research and development should be stopped altogether.

[10] Realistically, we should recognize that doctors are strongly influenced by standards of "good medical practice," by the fact that the costs of tests are usually covered by the patient's medical insurance, and by the need to practice defensive medicine in the face of possible malpractice litigation. The type of analysis set forth above enables us to assess the consequences when doctors respond to such influences.

Sometimes we can be confident that we are choosing desirable courses of action because we can formulate relatively precise assumptions about critical variables. At other times we may be in the fortunate position of having policies available that are robust with respect to a realistic range of assumptions. Decision analysis can tell us when our preferred policy choices are robust and when they are sensitive to the numbers assumed in the analysis. If they are sensitive, decision analysis will help us identify which assumptions are critical and which don't matter much. It cannot resolve debates about underlying values, but it is of great assistance when we must engage in a discussion about the implications of alternative sets of values.

The Uses of Decision Analysis

How does a decision tree fit in with the other models we have examined? Many of the analytic techniques studied earlier may be useful for estimating probabilities or payoffs, or even for determining options. The use of a tree is wholly compatible with the other techniques. Indeed, it is more than likely that benefit–cost analysis and discounting will be required in estimating payoffs, and we recognize that other models will be used to formulate probabilities.

This brief discussion of decision analysis has been aimed at convincing you of the wide applicability of the conceptual framework. Sometimes a decision appears so straightforward that it's hardly worth writing down the numbers and running them through the mill, yet defining the problem in the form of a decision tree may uncover issues and perspectives that would otherwise be overlooked. On other occasions a tree may prove helpful in sorting out a decision problem with so many ramifications that only a systematic approach can make it manageable. And the more difficult the task, the more useful decision analysis is likely to prove in comparison to the alternatives of no analysis at all or back-of-the-envelope calculations.

Appendix: Utility Theory

In the foregoing chapter we observed that expected value is frequently an unsatisfactory criterion for choice in risky situations. One example that we explored in detail involved the water supply for an island. In that case, we saw that by attaching dollar values to the water, we could get a much more accurate measure of the true value of alternative lotteries. Utility theory generalizes this procedure to any lottery, including most specifically lotteries where the decision maker is uncomfortable with a decision based on expected monetary value (EMV).

Our discussion focuses on the individual decision maker, who may be acting on his own behalf or on behalf of others; the broad principles we develop are appropriate in either situation. The policy maker who understands the central issues involved in making any risky decision is better equipped to make choices that affect the well-being of his individual constituents.

The Nature of the Problem and a Suggested Solution

Consider the very simple situation set forth in Figure 12A–1. The EMV of each of these two choices is $0, but most of us would have little difficulty in deciding quickly that we prefer *Y*. We say that, faced with a choice between these two courses of action, we would be *risk averse*; we would choose the less risky alternative, and we would even pay something of a

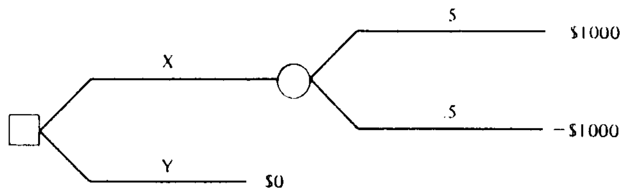

Fig. 12A–1

premium to get it rather than the risky option *X*. If we were to analyze the underlying reasons for our attitudes, we would probably conclude that *Y* appears preferable to *X* because losing $1000 would hurt us far more than winning $1000 would benefit us. This is an example of what economists sometimes loosely refer to as the diminishing marginal utility of money. In laymen's terms, it means that we spend our money on whatever we value most highly. Additional money, if we had it to spend, would go for something less valuable to us. Consequently, the loss of $1000 would hurt more than the gain of $1000 would help.[11]

[11] This principle was recognized at least as early as 1738. Daniel Bernoulli, in analyzing the famous St. Petersburg paradox, noted that expected value does not serve as a guide to action for most people. The paradox arises in the following way: you are offered the opportunity to buy a lottery in which a coin is tossed until it comes up heads. The game then ends and you