MANAGEMENT SCIENCE
Vol. 11, No. 10, August, 1965
Printed in U.S.A.

STOCHASTIC DECISION TREES FOR THE ANALYSIS OF INVESTMENT DECISIONS*

RICHARD F. HESPOS† AND PAUL A. STRASSMANN‡

This paper describes an improved method for investment decision making. The method, which is called the stochastic decision tree method, is particularly applicable to investments characterized by high uncertainty and requiring a sequence of related decisions to be made over a period of time. The stochastic decision tree method builds on concepts used in the risk analysis method and the decision tree method of analyzing investments. It permits the use of subjective probability estimates or empirical frequency distributions for some or all factors affecting the decision. This application makes it practicable to evaluate all or nearly all feasible combinations of decisions in the decision tree, taking account of both expected value of return and aversion to risk, thus arriving at an optimal or near optimal set of decisions. Sensitivity analysis of the model can highlight factors that are critical because of high leverage on the measure of performance, or high uncertainty, or both. The method can be applied relatively easily to a wide variety of investment situations, and is ideally suited for computer simulation.

Investment decisions are probably the most important and most difficult decisions that confront top management, for several reasons. First, they involve enormous amounts of money. Investments of U. S. companies in plant and equipment alone are approaching $50 billion a year. Another $50 billion or so goes into acquisition, development of new products, and other investment expenditures.

Second, investment decisions usually have long-lasting effects. They often represent a "bricks and mortar" permanence. Unlike mistakes in inventory decisions, mistakes in investment decisions cannot be worked off in a short period of time. A major investment decision often commits management to a plan of action extending over several years, and the dollar penalty for reversing the decision can be high. Third, investments are implements of strategy. They are the tools by which top management controls the direction of a corporation.

Finally, and perhaps most important, investment decisions are characterized by a high degree of uncertainty. They are always based on predictions about the future—often the distant future. And they often require judgmental estimates about future events, such as the consumer acceptance of a new product. For all of these reasons, investment decisions absorb large portions of the time and attention of top management.

Investment decision-making has probably benefited more from the development of analytical decision-making methods than any other management area. In the past 10 or 15 years, increasingly sophisticated methods have become available for analyzing investment decisions. Perhaps the most widely known of these new developments are the analytical methods that take into account the time

* Received February 1965.

† McKinsey and Company, Inc., New York
‡ National Dairy Products Corporation, New York

value of money. These include the net present value method, the discounted cash flow method, and variations on these techniques. [4, 13] Complementary to these time-oriented methods, a number of sophisticated accounting techniques have been developed for considering the tax implications of various investment proposals and the effects of investments on cash and capital position. [2, 12, 16] Considerable thought has been given to the proper methods for determining the value of money to a firm, or the cost of capital. [12, 13] The concepts of replacement theory have been applied to investment decisions on machine tools, automobile fleets, and other collections of items that must be replaced from time to time. [16]

In a somewhat different direction, techniques have been developed for the selection of securities for portfolios. These techniques endeavor to select the best set of investments from a number of alternatives, each having a known expected return and a known variability. [11] In this context, the "best" selection of investments is that selection that either minimizes risk or variability for a desired level of return, or maximizes return for a specified acceptable level of risk. (In general, of course, it is not possible to minimize risk and maximize return simultaneously.) The application of these techniques to corporate capital budgeting problems is conceivable but not imminent.

In the evolution of these techniques, each advance has served to overcome certain drawbacks or weaknesses inherent in previous techniques. However, until recently, two troublesome aspects of investment decision making were not adequately treated, in a practical sense, by existing techniques. One of these problems was handling the uncertainty that exists in virtually all investment decisions. The other was analyzing separate but related investment decisions that must be made at different points in time.

Two recent and promising innovations in the methodology for analyzing investment decisions now being widely discussed are directed at these two problems. The first of these techniques is commonly known as risk analysis; [6, 8] the second involves a concept known as decision trees. [9, 10, 15] Each of these techniques has strong merits and advantages. Both are beginning to be used by several major corporations.

It is the purpose of this article to suggest and describe a new technique that combines the advantages of both the risk analysis approach and the decision tree approach. The new technique has all of the power of both antecedent techniques, but is actually simpler to use. The technique is called the stochastic decision tree approach.

To understand the stochastic decision tree approach, it is necessary to understand the two techniques from which it was developed. A review of these two techniques follows.

A Review of Risk Analysis

Risk analysis consists of estimating the probability distribution of each factor affecting an investment decision, and then simulating the possible combinations of the values for each factor to determine the range of possible outcomes and the

FIGURE 1

TWO ANALYSES OF AN INVESTMENT PROPOSAL

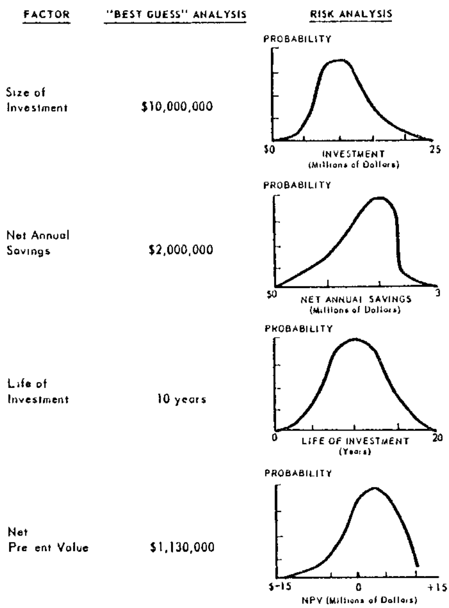

probability associated with each possible outcome. If the evaluation of an investment decision is based only on a single estimate—the "best guess"—of the value of each factor affecting the outcome, the resulting evaluation will be at best incomplete and possibly wrong. This is true especially when the investment is large and neither clearly attractive nor clearly unattractive. Risk analysis is thus an important advance over the conventional techniques. The additional information it provides can be a great aid in investment decision making.

To illustrate the benefit of the risk analysis technique, Figure 1 shows the results of two analyses of an investment proposal. First, the proposal was analyzed by assigning a single, "best guess" value to each factor. The second analysis used an estimate of the probability distribution associated with each factor and a simulation to determine the probability distribution of the possible outcomes.

The best-guess analysis indicates a net present value of $1,130,000, whereas the risk analysis shows that the most likely combination of events gives the project an expected net present value of only $252,000. The conventional technique

FIGURE 2

USE OF SENSITIVITY ANALYSIS TO HIGHLIGHT CRITICAL FACTORS

AN UNFAVORABLE CHANGE OF 10 PERCENTILES FROM THE MEAN VALUE IN THIS FACTOR	WHICH CORRESPONDS TO A PERCENTAGE CHANGE OF	WOULD REDUCE NPV BY
Annual net cash flow		
Sales level	12	17
Selling price	10	21
Manufacturing cost	18	58
Fixed cost	4	6
Amount of investment	5	12
Life of investment	12	30

fails to take into account the skewed distributions of the various factors, the interactions between the factors, and is influenced by the subjective aspects of best guesses Furthermore, the conventional analysis gives no indication that this investment has a 48 percent chance of losing money Knowledge of this fact could greatly affect the decision made on this proposal, particularly if the investor is conservative and has less risky alternatives available

The risk analysis technique can also be used for a sensitivity analysis The purpose of a sensitivity analysis is to determine the influence of each factor on the outcome, and thus to identify the factors most critical in the investment decision because of their high leverage, high uncertainty, or both In a sensitivity analysis, equally likely variations in the values of each factor are made systematically to determine their effect on the outcome, or net present value Figure 2 shows the effect of individually varying each input factor (several of which are components of the net cash inflow)

This analysis indicates that manufacturing cost is a highly critical factor, both in leverage and uncertainty Knowing this, management may concentrate its efforts on reducing manufacturing costs or at least reducing the uncertainty in these costs

Risk analysis is rapidly becoming an established technique in American industry Several large corporations are now using various forms of the technique as a regular part of their investment analysis procedure [1, 3, 7, 17, 18] A backlog of experience is being built up on the use of the technique, and advances in the state of the art are continually being made by users For example, methods have been devised for representing complex interrelationships among factors Improvements are also being made in the methods of gathering subjective probability estimates, and better methods are being devised for performing sensitivity analysis

One aspect of investment decisions still eludes the capabilities of this technique This is the problem of sequential decision making—that is, the analysis of a number of highly interrelated investment decisions occurring at different points

FIGURE 3
USE OF DECISION TREE TO ANALYZE INVESTMENT ALTERNATIVES FOR A NEW PRODUCT INTRODUCTION

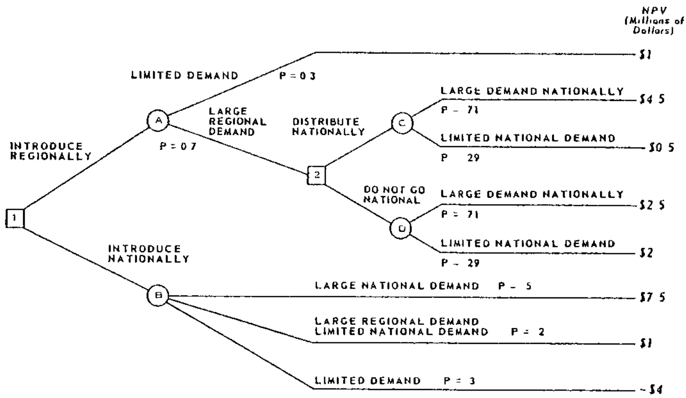

in time Until now no extension of risk analysis has been developed that can handle this problem well

A Review of Decision Trees

The decision tree approach, a technique very similar to dynamic programming, is a convenient method for representing and analyzing a series of investment decisions to be made over time (see Figure 3) Each decision point is represented by a numbered square at a fork or node in the decision tree Each branch extending from a fork represents one of the alternatives that can be chosen at this decision point At the first decision point the two alternatives in the example shown in Figure 3 are "introduce product nationally" and "introduce product regionally " (It is assumed at this point that the decision has already been made to introduce the product in *some* way)

In addition to representing management decision points, decision trees represent chance events The forks in the tree where chance events influence the outcome are indicated by circles The chance event forks or nodes in the example represent the various levels of demand that may appear for the product

A node representing a chance event generally has a probability associated with each of the branches emanating from that node This probability is the likelihood that the chance event will assume the value assigned to the particular branch The total of such probabilities leading from a node must equal 1 In our example, the probability of achieving a large demand in the regional introduction of the product is 0 7, shown at the branch leading from node A Each combination of decisions and chance events has some outcome (in this case, net present value, or NPV) associated with it

FIGURE 4

NET PRESENT VALUE OF INVESTMENT ALTERNATIVES FOR A NEW PRODUCT INTRODUCTION

ALTERNATIVE	CHANCE EVENT	PROBABILITY OF CHANCE EVENT	NET PRESENT VALUE	EXPECTED NPV
Introduce product nationally	Large national demand	.5	$7.5	$2.75
	Large regional, limited national demand	.2	1.0	
	Limited demand	.3	-4.0	
Introduce product regionally *(and distribute nationally if regional demand is large)*	Large national demand	.5	4.5	2.44
	Large regional, limited national demand	.2	-0.5	
	Limited demand	.3	1.0	
Introduce product regionally *(and do not distribute nationally)*	Large national demand	.5	2.5	1.95
	Large regional, limited national demand	.2	2.0	
	Limited demand	.3	1.0	

The optimal sequence of decisions in a decision tree is found by starting at the right-hand side and "rolling backward." At each node, an expected NPV must be calculated. If the node is a chance event node, the expected NPV is calculated for *all* of the branches emanating from that node. If the node is a decision point, the expected NPV is calculated for *each* branch emanating from that node, and the highest is selected. In either case, the expected NPV of that node is carried back to the next chance event or decision point by multiplying it by the probabilities associated with branches that it travels over.

Thus in Figure 3 the *expected* NPV of all branches emanating from chance event node C is $3.05 million ($4.5 × .71 + $-0.5 × .29). Similarly, the expected NPV at node D is $2.355 million. Now "rolling back" to the next node—decision point 2—it can be seen that the alternative with the highest NPV is "distribute nationally," with an NPV of $3.05 million. This means that, if the decision maker is ever confronted with the decision at node 2, he will choose to distribute nationally, and will expect an NPV of $3.05 million. In all further analysis he can ignore the other decision branch emanating from node 2 and all nodes and branches that it may lead to.

To perform further analysis, it is now necessary to carry this NPV backward in the tree. The branches emanating from chance event node A have an overall expected NPV of $2.435 million ($1 × 0.3 + $3.05 × 0.7). Similarly, the expected NPV at node B is 2.75 million. These computations, summarized in Figure 4, show that the alternative that maximizes expected NPV of the entire decision tree is "introduce nationally" at decision point 1. (Note that in this particular case there are *no* subsequent decisions to be made.)

One drawback of the decision tree approach is that computations can quickly become unwieldy. The number of end points on the decision tree increases very rapidly as the number of decision points or chance events increases. To make this approach practical, it is necessary to limit the number of branches emanating from chance event nodes to a very small number. This means that the probability distribution of chance events at each node must be represented by a very few point estimates

As a result, the answers obtained from a decision tree analysis are often inadequate The single answer obtained (say, net present value) is usually close to

FIGURE 5

RANGE OF POSSIBLE OUTCOMES FOR EACH OF THREE ALTERNATIVES

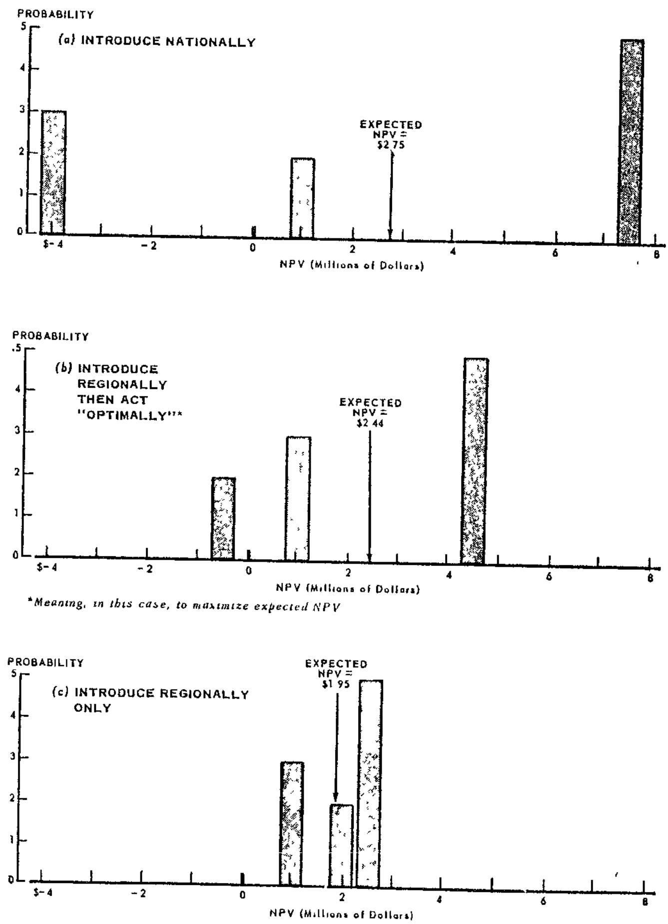

the expectation of the probability distribution of all possible NPVs. However, it may vary somewhat from the expected NPV, depending on how the point estimates were selected from the underlying distributions and on the sensitivity of the NPV to this selection process. Furthermore, the decision tree approach gives *no* information on the range of possible outcomes from the investment or the probabilities associated with those outcomes. This can be a serious drawback.

In the example in Figures 3 and 4, the decision tree approach indicated that introducing the product nationally at once would be the optimal strategy for maximizing expected NPV. However, the NPV of \$2.75 million is simply the mean of three possible values of NPV, which are themselves representative of an entire range of possible values, as shown in Figure 5a. Comparing the range of NPVs possible under each possible set of decisions shows a vastly different view of the outcome. (See Figures 5b and 5c.)

Although the first alternative has the highest expected NPV, a rational manager could easily prefer one of the other two. The choice would depend on the utility function or the aversion to risk of the manager or his organization. A manager with a linear utility function would choose the first alternative, as shown in Figure 6a. However, it is probably true that *most* managers would *not* choose the first alternative because of the high chance of loss, and the higher utility value that they would assign to a loss, as shown in Figure 6b. This conservatism in management is, to a large extent, the result of the system of rewards and punishments that exists in many large corporations today. Whether it is good or bad is a complex question, not discussed here.

In spite of these shortcomings, the decision tree approach is a very useful analytical tool. It is particularly useful for conceptualizing investment planning and for controlling and monitoring an investment that stretches out over time. For these reasons, the decision tree approach has been, and will continue to be an important tool for the analysis of investment decisions.

FIGURE 6
EXAMPLES OF UTILITY FUNCTIONS

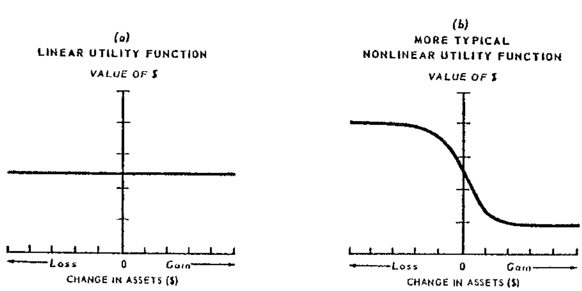

Combining These Approaches: Stochastic Decision Trees

The complementary advantages and disadvantages of risk analysis and decision trees suggest that a new technique might be developed that would combine the good points of each and eliminate the disadvantages The concept of stochastic decision trees, introduced in the remainder of this article, is intended to be such a combination

The stochastic decision tree approach is similar to the conventional decision tree approach, except that it also has the following features.

¶ All quantities and factors, including chance events, can be represented by continuous, empirical probability distributions

¶ The information about the results from any or all possible combinations of decisions made at sequential points in time can be obtained in a probabilistic form

¶ The probability distribution of possible results from any particular combination of decisions can be analyzed using the concepts of utility and risk

A discussion of each of these features follows

Replacement of Chance Event Nodes by Probability Distributions

The inclusion of probability distributions for the values associated with chance events is analogous to adding an arbitrarily large number of branches at each chance event node In a conventional decision tree, the addition of a large number of branches can serve to represent any empirical probability distribution Thus in the previous example, chance event node B can be made to approximate more closely the desired continuous probability distribution by increasing the number of branches, as shown in Figure 7a and 7b However, this approach makes the tree very complex, and computation very quickly becomes burdensome or impractical Therefore, two or three branches are usually used as a coarse approximation of the actual continuous probability distribution

Since the stochastic decision tree is to be based on simulation, it is not necessary to add a great many branches at the chance event nodes In fact, it is possible to reduce the number of branches at the chance event nodes to *one* (See Figure 7c) Thus, in effect, the chance event node can be *eliminated* Instead, at the point where the chance event node occurred, a random selection is made on each iteration from the appropriate probabilistic economic model such as the break-even chart shown in Figure 8 and the value selected is used to calculate the NPV for that particular iteration The single branch emanating from this simplified node then extends onward to the next management decision point, or to the end of the tree This results in a drastic streamlining of the decision tree as illustrated in Figure 9

Replacement of All Specific Values by Probability Distributions

In a conventional decision tree, factors such as the size of the investment in a new plant facility are often assigned specific values Usually these values are expressed as single numbers, even though these numbers are often not known with certainty

FIGURE 7

PROBABILITY DISTRIBUTIONS AT CHANCE EVENT NODES

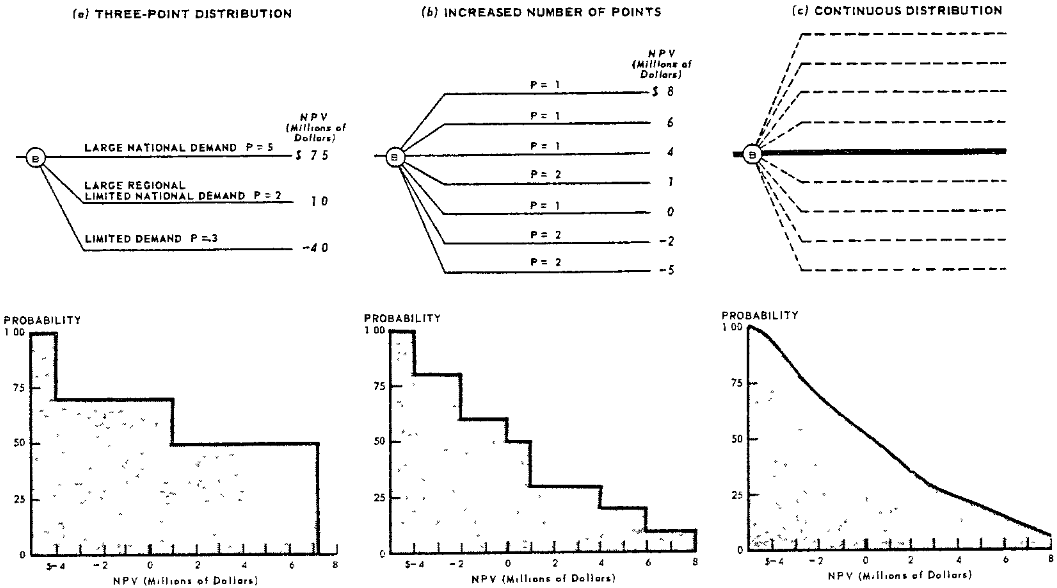

FIGURE 8
TYPICAL PROBABILISTIC ECONOMIC MODEL
USED TO SELECT VALUES OF FACTORS AT CHANCE EVENT NODES

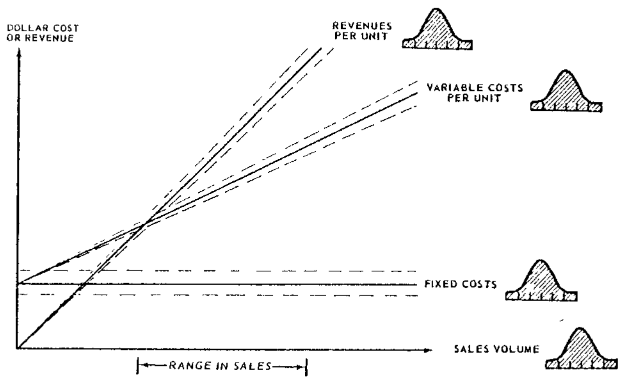

If the values of these factors could be represented instead by probability distributions, the degree of uncertainty characterizing each value could be expressed The stochastic decision tree approach makes it possible to do this Since the approach is basically a simulation, any or all specific values in the investment analysis can be represented by probability distributions On each iteration in the simulation, a value for each factor is randomly selected from the appropriate frequency distribution and used in the computation Thus, in the example, NPV can be calculated from not only empirical distributions of demand, but also probabilistic estimates of investment, cost, price, and other factors

Evaluating all Possible Combinations of Decisions

Since this stochastic decision tree approach greatly simplifies the structure of the decision tree, it is often possible to evaluate by complete enumeration all of the possible paths through the tree For example, if there are five sequential decisions in an analysis and each decision offers two alternatives, there are at most 32 possible paths through the decision tree This number of paths is quite manageable computationally And since most decision points are two-sided ("build" or "don't build," for example), or at worst have a very small number of alternatives, it is often feasible and convenient to evaluate all possible paths through a decision tree when the stochastic decision tree approach is used

Why is it sometimes desirable to evaluate all possible paths through a decision tree? As the inquiry into the risk analysis approach showed, decisions cannot always be made correctly solely on the basis of a single expected value for each factor The roll-back technique of the conventional decision tree necessarily deals

FIGURE 9

SIMPLIFIED DECISION TREE

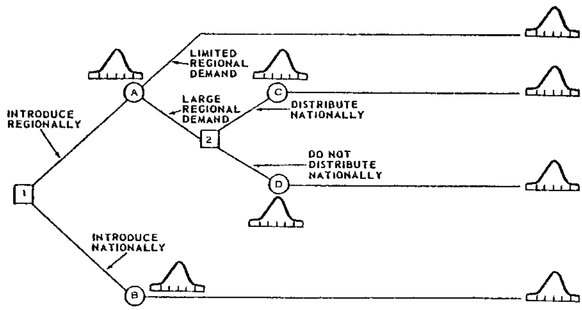

only with expected values It evaluates decisions (more exactly, sets of decisions) by comparing their expectations and selects the largest as the best, in all cases

However, the stochastic decision tree approach produces *probabilistic* results for each possible set of decisions These probability distributions, associated with each possible path through the decision tree, can be compared on the basis of their expectations alone, if this is considered to be sufficient But alternative sets of decisions can *also* be evaluated by comparing the probability distributions associated with each set of decisions, in a manner exactly analogous to risk analysis (The details of this technique are discussed in the next section) Thus, the stochastic decision tree approach makes it possible to evaluate a series of interrelated decisions spread over time by the same kinds of risk and uncertainty criteria that one would use in a conventional risk analysis

In a large decision tree problem, even with the simplifications afforded by the stochastic decision tree approach, complete enumeration of all possible paths through the tree could become computationally impractical, or the comparison of the probability distributions associated with all possible paths might be too laborious and costly

In such a case, two simplifications are possible First, a *modified* version of the roll-back technique might be used This modified roll-back would take account of the probabilistic nature of the information being handled Branches of the tree would be eliminated on the basis of dominance rather than simply expected value [7] For example, a branch could be eliminated if it had both a lower expected return and a higher variance than an alternative branch A number of possible sets of decisions could be eliminated this way without being completely evaluated, leaving an efficient set of decision sequences to evaluate in more detail

Computation could also be reduced by making decision rules before the simulation, such that if, on any iteration, the value of a chance event exceeds some criterion, the resulting decision would not be considered at all This has been done

FIGURE 10

THE GPSS CONCEPT OF DECISION TREES WITH RISK SIMULATION

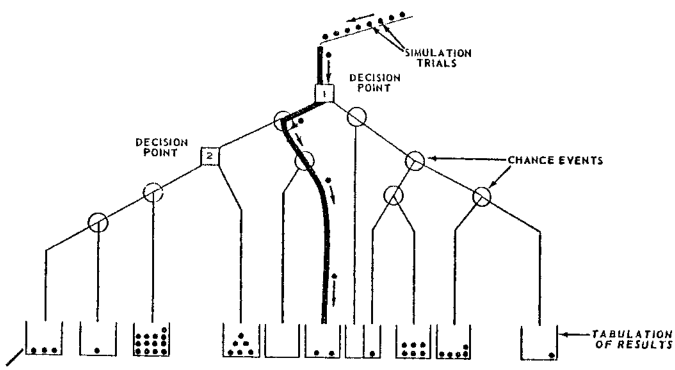

in the example shown in Figure 3 If a limited demand appears at node A, national introduction of the product will not be evaluated In the simulation, if demand were below some specified value, the simulation would not proceed to the decision point 2 This technique only saves computation effort—it does not simplify the structure of the tree, and if the criterion is chosen properly, it will not affect the final outcome

Recording Results in the Form of Probability Distributions

It has already been shown that probability distributions are more useful than single numbers as measures of the value of a particular set of decisions The simulation approach to the analysis permits one to get these probability distributions relatively easily It is true that the method smacks of brute force However, the brute force required is entirely on the part of the computer and not at all on the part of the analyst

The technique is simply this On each iteration or path through the decision tree, when the computer encounters a binary decision point node, it is instructed to "split itself in two" and perform the appropriate calculations along *both* branches of the tree emanating from the decision node (The same logic applies to a node with three or more branches emanating from it) Thus, when the computer completes a single iteration, an NPV will have been calculated for each possible path through the decision tree These NPVs are accumulated in separate probability distributions This simulation concept is illustrated in Figure 10

At the completion of a suitable number of iterations, there will be a probability distribution of the NPV associated with each set of decisions that it is possible to make in passing through the tree These different sets of decisions can then be compared, one against the other, in the usual risk analysis matter, as if they were alternative investment decisions (which in fact they are) That is, they can be compared by taking into account not only the expected return, but also the shape of each probability distribution and the effects of utility and risk On the basis of this, one can select the single best set of decisions, or a small number of possibly

FIGURE 11
RESULTS OF
STOCHASTIC DECISION-TREE ANALYSIS

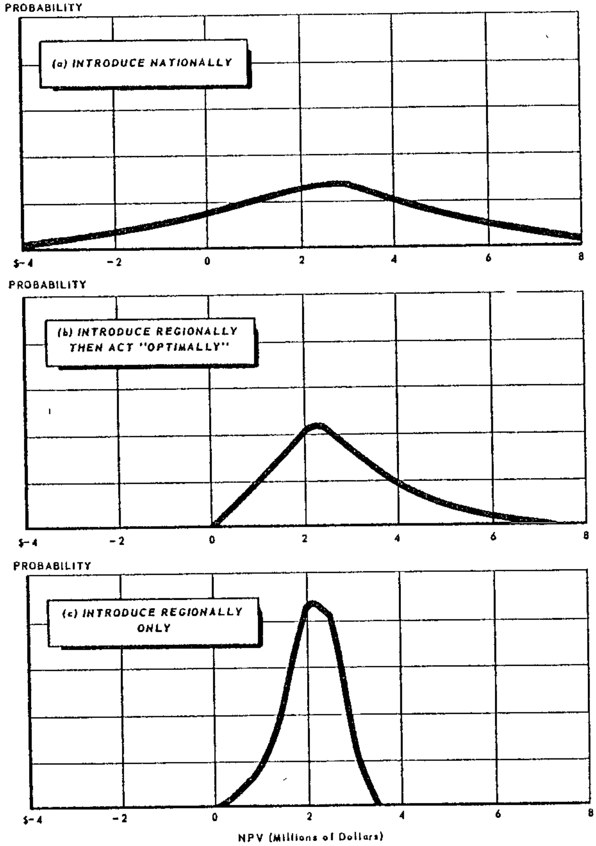

acceptable sets. These sets of sequential decisions can then be evaluated and a decision whether or not to undertake the investment can be made by comparing it to alternative investments elsewhere in the corporation or against alternative uses for the money.

An Example

To illustrate the kinds of results that can be expected from a stochastic decision tree analysis, the new product introduction problem described earlier has been solved using this method. The results are shown in Figure 11.

The differences in the expected values of the outcomes can now be seen in proper perspective, since the results show the relationship of the expected values to the entire distribution of possible outcomes. Moreover, the expected values of these distributions will not necessarily be identical with expectations resulting from the conventional decision tree approach, because:

1. The interdependencies among the variables were not accounted for by the conventional approach.
2. The small number of point estimates used to approximate an entire distribution under the conventional approach did not utilize all the available information.

With the three alternatives presented in this form, it is easier to understand why a rational manager might choose an alternative other than the one with the highest expected value. Presented with the full range of possible outcomes related to each alternative, he can select that alternative most consistent with his personal utility and willingness to take risk.

Using the Stochastic Decision Tree Approach

Stochastic decision trees described here combine the best features of both risk analysis and conventional decision trees and are actually simpler to construct and use than either of these. The steps for collecting data and conceptualizing the problem are the same for the stochastic decision tree approach as they are for the risk analysis approach. These steps are:

1. Gather subjective probability estimates of the appropriate factors affecting the investment.
2. Define and describe any significant interdependencies among factors.
3. Specify the probable timing of future sequential investment decisions to be made.
4. Specify the model to be used to evaluate the investment.

The stochastic decision tree approach is ideally suited to the computer language known as General Purpose Systems Simulator (GPSS). [5, 14] Although this language is not now capable of handling very complex interdependencies without certain modifications, it permits the solution of a very wide range of investment problems.

The structuring and solving of several sample problems have indicated that the stochastic decision tree approach is both easy to use and useful. The example in Figures 4, 5 and 6 shows emphatically how the stochastic decision tree approach can detect and display the probable outcomes of an investment strategy that would be deemed optimal by the conventional decision tree approach, but that many managements would definitely regard as undesirable. Other work is being done on both sample problems and real world problems, and on the development and standardization (to a limited extent) of the computer programs for performing this analysis.

Summary

The stochastic decision tree approach to analyzing investment decisions is an evolutionary improvement over previous methods of analyzing investments. It combines the advantages of several earlier approaches, eliminates several disadvantages, and is easier to apply.

References

1. Anderson, S. L. and Haight, H. G., "A Two-by-Two Decision Problem," *Chemical Engineering Progress*, Vol. 57, No. 5, May 1961.
2. Anthony, Robert N. (Editor), *Papers on Return on Investment*, Harvard Business School, Boston, 1959.
3. "Chance Factors Meaning and Use," Atlantic Refining Company, Producing Department, July 1962.
4. Dean, Joel, *Capital Budgeting*, New York, Columbia University Press, 1951.
5. Gordon, G., "A General Purpose Systems Simulator," *IBM Systems Journal*, Vol. I., September 1962.
6. Hertz, David B., "Risk Analysis in Capital Investment," *Harvard Business Review*, January–February, 1964.
7. Hess, Sidney W. and Quigley, Harry A., "Analysis of Risk in Investments Using Monte Carlo Technique," Chemical Engineering Progress Symposium Series No. 42, Vol. 59.
8. Hillier, Frederick, S., Stanford University, "The Derivation of Probabilistic Information for the Evaluation of Risky Investments," *Management Science*, April, 1963.
9. Magee, John F., "Decision Trees for Decision Making," *Harvard Business Review*, July–August, 1964.
10. Magee, John F., "How to Use Decision Trees in Capital Investment," *Harvard Business Review*, Setpember-October, 1964.
11. Markowitz, Harry, *Portfolio Selection, Efficient Diversification of Investments*, New York, John Wiley and Sons, 1959.
12. Masse, Pierre, *Optimal Investment Decisions*, Prentice Hall, 1962.
13. McLean, John G., "How to Evaluate New Capital Investments," *Harvard Business Review*, November-December, 1958.
14. Reference Manual General Purpose Systems Simulator II, IBM, 1963.
15. Schlaifer, Robert, *Probability and Statistics for Business Decisions*, McGraw-Hill, 1959.
16. Terborgh, George, *Business Investment Policy*, Machinery and Allied Products Institute, Washington, D.C., 1958.
17. Thorne, H. C. and Wise, D. C., American Oil Company, "Computers in Economic Evaluation," *Chemical Engineering*, April 29, 1963.
18. "Venture Analysis," Chemical Engineering Progress Technical Manual, American Institute of Chemical Engineers.