Capital investment analysis report

How can business executives make the best investment decisions? Is there a risk analysis method to help managers make wise acquisitions, launch new products, modernize the plant, or avoid overcaze? Risk analysis in capital investment analyzes issues like these and says yes, measuring the multitude of risks involved in each situation. Mathematical formulas that predict a single rate of return or better estimation are not enough. The author's approach emphasizes the nature and processing of the data used and specific combinations of variables such as cash flow, return on investment and risk to estimate the probabilities of each potential outcome. Managers can examine the additional information provided in this way to more accurately assess the potential for substantial profit in their companies. The article, originally presented in 1964, continues to interest HBR readers. In a retrospective commentary, the author discusses the routine use of risk analysis in business and government, emphasizing that the method can and should be used in any situation that requires decisions in our uncertain world. Of all the decisions business executives must make, none is more difficult—and none have received more attention—than choosing between alternative capital investment opportunities. What makes this kind of decision so demanding, of course, is not the problem of projecting the return on investment under a certain set of assumptions. The difficulty is in the assumptions and their impact. Each assumption implies its own degree—often a high degree—of uncertainty; and, taken together, these combined uncertainties can multiply into total uncertainty of critical proportions. This is where the risk element comes in, and it is in the risk assessment that the executive has been able to get little help from the tools and techniques currently available. There is a way to help the executive fine-tune key capital investment decisions by providing him with a realistic measurement of the risks involved. Armed with this indicator, which assesses risk at every possible level of return, then it is able to measure the most skillfully alternative courses of action against corporate objectives. Need for new concept The evaluation of a capital investment project begins with the principle that capital productivity is measured by the rate of return we expect to receive over some future period. A dollar received next year is worth less than a dollar in hand today. Therefore, three-year expenses are less expensive than expenditures of equal magnitude within two years. For this reason we cannot calculate the rate of return realistically unless we take into account (a) when spend the sums involved in an investment and (b) when yields are received. Therefore, the comparison of alternative investments is complicated by the fact that they usually differ not only in size, but also in the period of time during which expenses will have to be made and and Returned. These facts of investment life long ago made clear the shortcomings of approaches that simply wintered expenses and profits, or bulking them up, as in the method of number of years of payment. These deficiencies encouraged decision-making students to explore more accurate methods of determining whether an investment would leave a company better in the long run than another form of action. Unsurprisingly, then, much effort has been made to develop ways to improve our ability to discriminate between investment alternatives. The objective of all these researches has been to refine the definition of the value of capital investments for the company. The controversy and fury that once came out in the business press over the most appropriate way to calculate these values have been largely resolved in favor of the discounted cash flow method as a reasonable means of measuring the rate of return that can be expected in the future of an investment made today. Therefore, we have methods that are more or less elaborate mathematical formulas to compare the results of various investments and the combinations of variables that will affect investments. As these techniques have progressed, the math involved has become increasingly accurate, so we can now calculate discounted returns at a fraction of a percentage. But sophisticated executives know that behind these precise calculations there is data that is not as accurate. At best, the rate-of-return information provided to them is based on an average of different opinions with different flexibilities and different probability ranges. When expected yields on two investments are close, executives are likely to be influenced by intangibles, a precarious pursuit at best. Even when the figures of two investments are quite separate, and the choice seems clear, memories of Edsel and other unfortunate companies lurk. In short, decision makers realize there's something else they should know, something besides the expected rate of return. What is missing has to do with the nature of the data on which the expected rate of return is calculated and how that data is processed. It involves uncertainty, with possibilities and probabilities that extend across a wide range of rewards and risks. (For a summary of the new approach, see Inserting.) After examining current alternative investment comparison methods, the author reports on his company's experience in applying a new approach to the problem. With this approach, management takes the different levels of potential cash flows, return on investment and other results from a and gets an estimate of the probabilities for each potential outcome. Currently, decisions from many facilities are based on discounted cash flow calculations. Management is told, for example, that Investment X has an expected internal rate of return of 9.2%, while for investment Y Y you can expect it. Instead, the new approach would put a calendar in front of executives that gives them the most likely return of X, but also tells them that X has a 1 in 20 chance of being a total loss, 1 in 10 to win from 4% to 5%, 2 in 10 to pay from 8% to 10%, and 1 chance in 50 to reach a return rate of 30%. From another calendar they learn what is Y's most likely rate of return, but also that Y has 1 chance on 10 of what results in a total loss, 1 in 10 to gain from 3% to 5% return, 2 in 10 to pay between 9% and 11%, and 1 chance in 100 of a 30% rate of return. In this case, the estimates of the rates of return provided by the two approaches would not be substantially different. However, for the decision maker with the added information, investment Y no longer looks like the clearly better option, since with X the chances of substantial gain are greater and the risks of loss lower. Two things have made this approach attractive to managers who have used it: 1. It is certainly in all cases a more descriptive statement of the two opportunities. And in some cases you could well reverse the decision, in line with particular corporate goals. 2. This is not a difficult technique to use, as much of the necessary information is already available – or easily accessible – and the validity of the principles involved, for the most part, has already been tested in other applications. The enthusiasm with which managers exposed to this approach have received suggests that it may have a wide application. It has particular relevance, for example, in problems as knotty as investments related to acquisitions or new products and in decisions that may involve overcaly capacity. Return Rate Investment X Return Rate Investment And Achilles Heel The fatal weakness of past approaches has nothing to do with the mathematics of calculating the rate of return. We have pushed down this path so far that the accuracy of our calculation is, if anything, illusory. The fact is that, regardless of the mathematics used, each of the variables that enter the return rate calculation is subject to a high level of uncertainty. For example, the service life of a new piece of capital equipment is rarely known in advance with any degree of certainty. It may be affected by variations in obsolescence or deterioration, and relatively small changes in service life can lead to major changes in return. However, an expected value is entered for the life of the equipment, based on a large amount of data from which a single forecast of the best possible forecast has been developed, is entered in the calculation of the rate of return. The same is done for the other factors that have a significant relationship in the in question. Let's see how this works in a simple case, one in which the odds seem to be in favor of a particular decision. Food company executives must decide whether to launch a new packaged cereal. They have come to the conclusion five factors are the determining variables: advertising and promotion costs, total cereal market, market share for this product, operating costs and new capital investments. Based on the most likely estimate for each of these variables, the image looks very bright-a healthy yield of 30%. This future, however, depends on whether each of these estimates actually becomes a reality. If each of these educated guesses is, for example, 60% likely to be correct, there is only an 8% chance that all five are correct (.60 × to .60 ×.60 ×.60 ×.60). So the expected return actually depends on a rather unlikely coincidence. Decision makers need to know much more about the other values used to make each of the five estimates and what they can earn or lose from various combinations of these values. This simple example illustrates that the rate of return actually depends on a specific combination of values from many different variables. But only the expected levels of ranges (worse, average, better; or pessimistic, more likely, optimistic) of these variables are used in a formal mathematical manner to provide the figures given to management. Therefore, predicting a more likely rate of return gives accurate numbers that don't tell the whole story. The expected rate of return represents only a few points in a continuous cure for possible combinations of future events. It's a bit like trying to predict the result in a dice game by saying that the most likely result is a 7. The description is incomplete because it doesn't tell us about all the other things that might happen. In Exhibit I, for example, we see the probabilities that only two dice have 6 faces. Now suppose each of the eight dice has 100 sides. This is a situation more comparable to business investment, where the company's market share could become any 1 of 100 different sizes and where there are eight factors (prices, promotion, etc.) that can affect the outcome. Test I. Describing uncertainty—a dice release is also not the only problem. Our willingness to bet on a roll of the dice depends not only on odds, but also on bets. Since the probability of rolling a 7 is 1 in 6, we could be very willing to risk a few bucks on that result with appropriate odds. But would we be equally willing to bet $10,000 or $100,000 on those same odds, or even better odds? In short, the risk is influenced both by the probabilities in the various events that occur and by the magnitude of the rewards or penalties that are involved when they occur. To illustrate again, suppose a company is considering a $1 million investment. The best return is likely to be $200,000 a year. It might well be that this estimate is the average of three possible yields: a 1-in-3 probability of getting no return at all, a 1-in-3 probability of getting $200,000 per year, a 1-in-3 probability of getting $400,000 per year. Suppose getting no return at all would put the company business. Then, by accepting this proposal, management is taking a 1-in-3 chance to go bankrupt. However, if only the best estimate analysis is used, management could move on, not knowing that it is taking a great opportunity. If all available information were examined, management might prefer an alternative proposal with a smaller but safer expectation (i.e. less variable). These considerations have led almost all advocates of the use of modern capital-investment index calculations to advocate recognition of the elements of uncertainty. Perhaps Ross G. Walker summed up the current thought when he spoke of the almost impenetrable mists of any prognosis. 1 How can executives penetrate the fogs of uncertainty surrounding options between alternatives? Limited improvements A number of efforts to address uncertainty have been successful to some extent, but all seem not to be short of the brand in one way or another. 1. More accurate forecasts Reducing error in estimates is a worthy goal. But no matter how many estimates of the future go to a capital investment decision, when everything is said and done, the future remains the future. Therefore, as well as we predict, we still have some knowledge that we cannot eliminate all uncertainty. 2. Empirical adjustments Adjusting the factors that influence the outcome of a decision is subject to serious difficulties. We would like to adjust them to reduce the likelihood that we will make a bad investment, but how can we do so without spoiling at the same time our chances of making a good one? And in any case, what is the basis for the adjustment? We adjust, not because of uncertainty, but because of bias. For example, building estimates are often exceeded. If the company's history of construction costs is that 90% of its estimates have been exceeded by 15%, then in a capital estimate there is all the justification for increasing the value of this factor by 15%. It's about improving the accuracy of the estimate. But suppose that sales estimates for new products have been exceeded by more than 75 percent in a quarter of all historical cases and have not reached 50 percent of the estimate in a sixth of all those cases? The penalties for this overestimation are very real, so the administration is able to reduce the sales estimate to cover the case by six, thus reducing the calculated rate of return. In doing so, you may be missing some of your best opportunities. 3. Review the cut-off rates Select higher cut-off rates to protect against uncertainty you are trying a lot the same. Management would like to have a chance of a return in proportion to risk is assumed. When there is a lot of uncertainty in the various estimates of sales, costs, prices, etc., a high calculated return on investment provides some incentive to take the risk. This is, in fact, a perfectly solid position. The problem is that decision makers still need to know explicitly what risks are taking, and what are the odds of achieving the expected return. 4. Three-level estimates A start at spelling out risks is sometimes made by taking the high, medium and low values of the estimated factors and calculating performance rates based on various combinations of pessimistic, average and optimistic estimates. These calculations give an image of the range of possible results, but do not tell the executive whether the pessimistic result is more likely than the optimistic one, or, in fact, whether the average result is much more likely to occur than either extreme. Therefore, although this is a step in the right direction, it still does not give a clear enough image to compare alternatives. 5. Selected probabilities Several methods have been used to include the probabilities of specific factors in the return calculation. L.C. Grant discussed a program to forecast discounted cash flow rates where service life is subject to obsolescence and deterioration. Calculated the probabilities that the investment will end at any time after it is made depending on the probability distribution of the life factor. Having calculated these factors for each year over the maximum lifespan, he determined a global expected rate of return.2 Edward G. Bennion suggested using game theory to take into account alternative market growth rates, as they would determine the rate of return for various options. It used the estimated probabilities of specific growth rates to develop optimal strategies. Bennion noted: Forecasting can result in a negative contribution to capital budget decisions unless it goes beyond simply providing a single, more likely prediction... [with] an estimated probability coefficient for forecasting, plus knowledge of profits for the company's alternative investments and calculation of probabilities of indifference... the margin of error can be substantially reduced, and the employer can tell how far his prognosis can be before it leads to the wrong decision.

3 Note that both methods produce an expected yield, each based on a single insecure factor: lifespan in the first case, market growth in the second. Both are useful, and both tend to improve the clarity with which the executive can see investment alternatives. But neither sharpens the range of risk taken or expected return long enough to help much in complex capital planning decisions. Sharpening the image Since each of the many factors that enter the evaluation of a decision is subject to some uncertainty, executives need a useful representation of the effects that uncertainty surrounding each significant factors have in the yields they are likely to achieve. Therefore, I use a method that combines the variability inherent in all the relevant factors considered. The aim is to give a clear picture of the relative risk and chances of coming forward or behind in light of uncertain prior knowledge. A simulation of how these factors can be combined as the future unfolds is the key to extracting maximum information from available forecasts. In fact, the approach is very simple, using a computer to make the necessary arithmetic. To perform the analysis, a company must follow three steps: 1. Estimate the range of values for each of the factors (for example, sales price range and sales growth rate) and within that range the probability of occurrence of each value. 2. Randomly select a value from the value distribution for each factor. Then combine the values of all factors and calculate the rate of return (or present value) from that combination. For example, the lowest in the price range could be combined with the highest in the growth rate range and other factors. (The fact that the elements are dependent must be taken into account, as we will see below.) 3. Do this over and over again to define and evaluate the probabilities of each possible rate of return occurring. Since there are literally millions of possible combinations of securities, we need to test the likelihood of multiple return on investment. This is like figuring out by recording the results of a large number of releases which percentage of 7s or other combinations we can expect in throwing dice. The result will be a list of the rates of return we could achieve, ranging from a loss (if the factors go against us) to any maximum possible gain with the estimates that have been made. For each of these rates we can determine the chances that it may occur. (Note that a specific profitability can typically be achieved through more than one combination of events. The more combinations for a given rate, the greater the chances of achieving it, as with 7s in the roll of dice.) The average expectation is the average of the values of all results weighted by the chances of each occurring. We can also determine the variability of result values from average. This is important because, all other factors are the same, management would presumably prefer less variability for the same return if given the option. This concept has already been applied to investment portfolios. Where the expected return and variability of each of a series of investments has been determined, the same techniques may be used to examine the effectiveness of various combinations of investments in meeting management objectives. Practical test To see how this new approach works in practice, let's take the experience of a management that has already analyzed a specific investment proposal conventional techniques. Taking the same investment calendar and the same expected values actually used, we can find the results that the new method would produce and compare them with the results obtained by conventional methods. As we will see, the new picture of risks and returns is different from the previous one. However, the differences are in no way to changes in basic data, only to the increased sensitivity of the method to administration uncertainties about key factors. Investment proposal In this case, a medium-sized industrial chemical producer is considering an extension of $10 million to its processing plant. The estimated service life of the installation is ten years; engineers expect to use 250,000 tons of processed material worth $510 per tonne at an average processing cost of $435 per tonne. Is this investment a good bet? In fact, what is the profitability the company can expect? What are the risks? We need to make the best and fullest use of all the financial and market research analyses that have been developed, in order to give management a clear picture of this project in an uncertain world. Key entry factors that management has decided to use are market size, sales prices, market growth rate, market share (resulting in physical sales volume), required investment, residual investment value, operating costs, fixed costs, and facility life. These factors are typical of those in many projects of the company that must be analyzed and combined to obtain a measure of the attractiveness of a proposed investment of capital facilities. Obtaining estimates How do we do the recommended type of analysis of this proposal? Our goal is to develop for each of the nine factors listed a frequency distribution or probability curve. The information we need includes the possible range of values for each factor, the average and some idea about the probability that the various possible values will be reached. It has been my experience that for the main capital proposals, administrations often make a significant investment in time and funds to identify information on each of the relevant factors. An objective analysis of the values to be assigned to each can, with little additional effort, produce a subjective probability distribution. Specifically, it is necessary to investigate and question each of the experts involved, to find out, for example, whether the estimated cost of production can actually be said to be exactly a certain value or whether, more likely, it should be estimated to be within a certain range of values. Management generally ignores that range in its analysis. The range is relatively easy to determine; If you have to make a guess, as you often do, it's easier to accurately guess a range instead of a specific value. I have found from experience that a series of meetings with management staff to discuss such distributions are more useful in coming up with realistic answers a priori questions. (The term realistic answers implies that all information management doesn't have as well as everything it has.) The ranges are directly related to the degree of confidence the estimator has in the estimate. Therefore, it can be known that certain estimates are quite accurate. They would be represented by probability distributions indicating, for example, that there are only 1 1 by 10 that the actual value will be different from the best estimate by more than 10%. Others can have up to 100% ranges above and below the best estimate. Therefore, we treat the sale price factor of the finished product by asking executives that they are responsible for the original estimates these questions: Since $510 is the expected selling price, what is the probability that the price will exceed $550? Is there any chance that the price will exceed $650? How likely is it that the price will fall below $475? Administrations should ask similar questions for all other factors until they can build a curve for each. Experience shows that this is not as difficult as it seems. Information about the degree of variation in factors is often easy to obtain. For example, historical information about variations in the price of a good is readily available. Similarly, administrations can estimate the sales variability of industry sales records. Even for factors that have no history, such as the operating costs of a new product, those who make the average estimates should have some idea of how confident they are in their predictions, and are therefore generally too happy to express their feelings. Similarly, the less confidence they have in their estimates, the greater the range of possible values that the variable will assume. This last point is likely to annoy entrepreneurs. Does it really make sense to look for variation estimates? It cannot be stressed too strongly that the less certainty there is in an average estimate, the more important it will be to consider the possible variation of that estimate. In addition, an estimate of the possible variation in a factor, no matter how judicious it may be, is always better than a simple average estimate, as it includes more information about what is known and what is not known. This same lack of knowledge can distinguish one investment possibility from another, so rational decision-making must be taken into account. This lack of knowledge is in itself important information about the proposed investment. Throwing away any information simply because it is highly uncertain is a serious error in the analysis that the new approach is designed to correct. Team executions The next step in the proposed approach is to determine the returns that will result from random combinations of the factors involved. This requires realistic restrictions, such as not allowing the total market to vary more than some reasonable amount from year to year. Of course, any appropriate method of qualification of the return can be used at this point. In the actual case, the address preferred the discounted cash flow for the reasons mentioned above, so this method is followed here. Computer can be used to carry out tests for the simulation method in a very short time and at very little cost. Thus, for one test, 3,600 discounted cash flow calculations were performed, each based on a selection of the nine input factors, in at a cost of $15 per computer time. The resulting return rate probabilities were read immediately and plotted. The process is shown schematically in Exhibit II. Exhibit II. Simulation for investment planning Data comparisons The nine input factors described above are divided into three categories: 1. The market analyses included are market size, market growth rate, company market share and sales prices. For a given combination of these factors, sales revenue can be determined for a particular business. 2. Analysis of investment costs Being linked to the expected types of operating life and cost characteristics, these are subject to various types of error and uncertainty; for example, the progress of automation makes service life uncertain. 3. Operating and fixed costs These are also subject to uncertainty, but are perhaps the easiest to estimate. These categories are not independent, and for realistic results, my approach allows the various factors to be linked together. Therefore, if the price determines the total market, we first select from a probability distribution the price for the specific execution of the equipment and then use for the total market a probability distribution that is logically related to the selected price. We are now ready to compare the values obtained under the new approach with those obtained by the old one. This comparison is shown in Exhibit III. Exhibit III. Comparison of expected values in old and new approaches Note: The range figures in the right column represent approximately 1% to 99% probabilities. That is, there is only a 1 in 100 chance that the value actually reached is respectively greater than or less than the range. Valuable results How do the results of new and old approaches compare? In this case, management had been informed, on the basis of a better estimate approach, that the expected return was 25.2% before taxes. When we run the new data set through the computer program, however, we get an expected return of only 14.6% before taxes. This surprising difference results not only from the range of values under the new approach, but also from weighing each value in the range by the possibilities of its occurrence. Our new analysis can help management avoid reckless investment. In fact, the general result of carefully weighing information and lack of information in the way I have suggested is to indicate the true nature of seemingly satisfactory investment proposals. If this practice were followed, management could avoid much capacity. The computer program developed to carry out the simulation allows easy insertion of new variables. But most programs do not allow dependency relationships between the various input factors. In addition, the program used here allows the choice of a value for the price of a distribution, the value of which determines a given probability distribution (from several) to be used to determine the values for the sales volume. The next scenario scenario how this important technique works. Suppose we have a wheel, as in roulette, with numbers from 0 to 15 representing a price for the product or material, numbers 16 to 30 representing a second price, numbers 31 to 45 a third price, and so on. For each of these segments we would have a different range of expected market volumes, for example, $150,000–$200,000 for the first, $100,000–$150,000 for the second, $75,000–$100,000 for the third. Now let's say we turn the wheel and the ball falls by 37. This means that we chose a sales volume in the range of $75,000–$100,000. If the ball goes at 11, we have a different price, and we move to the range of $150,000–$200,000 for a sales volume. The most significant thing, perhaps, is the fact that the program allows management to determine the sensitivity of the results to each or all of the input factors. By simply running the program with changes in the distribution of an input factor, it is possible to determine the effect of the aggregated or modified information (or the lack of information). It may turn out that fairly large changes in some factors do not significantly affect the results. In this case, in fact, management was particularly concerned about the difficulty in estimating market growth. The implementation of the program with variations in this factor quickly showed that for average annual growth rates of 3% to 5% there was no significant difference in the expected result. In addition, let's see what the implications of the detailed knowledge that the simulation method gives us. Under the method that uses expected unique values, management reaches only an expected expectation of 25.2% after taxes (which, as we have seen, is wrong unless there is no variability in the many input factors, a highly unlikely event). However, with the proposed method, uncertainties are clearly shown, as shown in Exhibit IV. Observe the contrast with the profile obtained under the conventional approach. This concept has also been used for the evaluation of product introduction, business acquisition and plant modernization. Exhibit IV. Expected rates of return under old and new approaches Comparing Opportunities from a decision-making point of view one of the most significant advantages of the new return rate determination method is that it allows management to discriminate between measures of (1) expected return based on the weighted probabilities of all possible returns , (2) return variability and (3) risks. To visualize this advantage, let's take an example based on another real but simplified case for explanation purposes. The example includes two investments considered, A and B. With investment analysis, we get the tabulated data and In Exhibit V. We see that: Investment B has a higher expected return than investment A. Investment B also has substantially more variability than investment A. There is a good chance that Investment B will get a very different return than expected 6.8 per cent, possibly as high as 15 per cent or as low as a 5 per cent loss. Investment A is unlikely to vary widely from the expected return of 5%. Investment B involves much more risk than investment A. There is virtually no possibility of incurring a loss in investment A. However, there is 1 in 10 chance of losing money on investment B. If such a loss occurs, its expected size is approximately $200,000. Exhibit V. Comparison of two investment opportunities Clearly, the new investment assessment method provides management with much more information on which to base a decision. Investment decisions made solely on the basis of the maximum expected return are not unequivocally the best decisions. Final Note The question facing capital investment management is first and foremost: What information is needed to clarify the key differences between the different alternatives? There is agreement on the basic factors to consider: markets, prices, costs, etc. And the way in which future return on investment should be calculated, if not agreed, is limited to at least a few methods, any of which can be consistently used in a given company. If the input variables turn out as estimated, any of the methods commonly used to qualify investments must provide satisfactory (if not necessarily maximum) yields. In real practice, however, conventional methods do not work satisfactorily. Why? The reason, as we have seen earlier in this article and as every executive and economist knows, is that the estimates used to make the early calculations are just that: estimates. More accurate estimates would be useful, but at best residual uncertainty can easily make a mockery of corporate hopes. However, there is a solution. Collecting realistic estimates of key factors means finding out a lot about them. Therefore, the type of uncertainty involved in each estimate can be assessed in advance. Using this knowledge of uncertainty, executives can maximize the value of information for decision-making. The value of software in developing clear representations of the uncertainty and risk surrounding alternative investments has been demonstrated. These programs can produce valuable information on the sensitivity of possible outcomes to the variability of input factors and the likelihood of achieving various possible rates of performance. This information can be extremely important in support of the management trial. Having quota calculations on all possible outcomes gives some certain certainties to decision-makers available information has been used with maximum efficiency. This simulation approach has the inherent advantage of simplicity. It only requires an extension of the input estimates (to the best of our capacity) in terms of probabilities. No projection should be identified unless we are sure of it. The discipline of thinking through the uncertainties of the problem will in itself help ensure that investment decision-making. Understanding uncertainty and risk is understanding the key business problem and key business opportunity. Since the new approach can be applied continuously to each capital alternative as it is taken into account and moves towards fruiting, gradual progress can be expected in improving the estimation of the probabilities of variation. Finally, the courage to act courageously in the face of apparent uncertainty can be greatly reinforced by the clarity of the representation of risks and possible rewards. Achieving these lasting results requires only a slight effort beyond what most companies already exercise in the study of capital investments. Retrospective Comment When this article was published 15 years ago, there were two recurring topics in the management community's responses to it: (1) how the uncertainties surrounding each key element of an investment decision would be determined, and (2) what criteria should be used to decide to proceed with an investment once quantified and uncertainties were shown. I answered this last question in a sequel to HBR, Investment Policies That Pay, describing risk and betting relationships with longer-term investment criteria. This article, published in 1968, showed how risk analyses can provide bases for developing policies to choose from a variety of investment alternatives. Similar approaches to investment fund portfolio management were subsequently developed. Analyzing uncertainty in describing complex decision-making situations is now an integral part of business and government. The elements of an investment decision, private or public, are subject to all uncertainties of an unknown future. As shown in the 1964 article, an estimated probability distribution paints the clearest picture of all possible outcomes. This description contains considerably more information than simplistic combinations of better subjective estimates of input factors. The best estimates are point-in-time estimates (there may be more than one: high, medium, low) of the value of an element of the investment analysis used to determine an outcome decision criterion, such as the internal rate of return or the current value of the investment. Therefore, even when the conventional approach was used for best estimation in a single-point determination for statistically estimated expected values from an element distribution, the single-point approach was shown to be extremely misleading. In Exhibit III, an analysis of the best single-point estimate gave an internal rate of return of 25.2%. And a risk analysis that used distributions of estimated elements showed that an average of possible outcomes, weighted by the relative frequency of their occurrences at 14.6%, was more realistic and significantly different. It presented a more real picture of the actual average expectation of the result of this investment (if it could be repeated over and over again). The case was and the point of this result—that risk and uncertainty were more precisely defined by a simulation of input variables—was little questioned thereafter. Management began to adopt some or all of this procedure to examine some, if not all, significant investments in which there were doubts about the levels of risk involved. My article in the sequel sought to demonstrate that if sufficient investments were chosen consistently on the basis of criteria related to such risk representations, the overall results would be stabilized around the desired expected value or the best estimate of the criterion. All this now seems simple and straightforward. It was previously falsely thought that risk analysis was aimed at eliminating uncertainty, which was not worth doing at all, as the future is so desperately uncertain. Thus, in 1970 the Financial Times (London) published an article aimed at showing the futility of risk analysis. He was referring to a baker of geriatric biscuits who made an investment only to go bankrupt when his nursing home market disappeared hastily with the death of its founder. The author quoted as a moral, Don't put all your dough in a cookie. The points took some time to spread through executive circles that (1) exactly such an analysis would have been just as bad, or worse, done through subjective single-point estimates, and (2) no analytical technique could handle future events, even with sensitive inputs and tracking control requirements to improve the probabilities projected by the original risk analyses. But in the end, judgment would be necessary in both the estimation of inputs and in the decision. It did not mean that the article was a methodology argument, but rather a cautionary note for examining the data surrounding an investment proposal in the light of all the widespread uncertainties in the world, of which the business is simply a part. The years since 1964 have made it clear to me that this message should have been amplified and emphasized more strongly in the article. If this point had been clearer, the question of whether taking the risk and proceeding with an investment might have been less problematic. If I had been able to look more closely, I could have seen that the area of risk analysis would become routine in business and be adopted virtually universally in matters of public cost-benefit. Of course, cost-benefit analysis for public decisions is just a special form of investment analysis. Government issues that require decisions involving significant uncertainties are too numerous to fully catalogue: energy from both fossil and nuclear sources; chemical risks, and food carcinogens; Manipulation of DNA and its progeny from genetic splicing. The Three Mile Island nuclear accident brought home the fallibility of presenting a conclusion of risk analysis in simplistic terms. The well-known Rasmussen report on the safety of nuclear reactors, commissioned by the Nuclear Regulatory Commission, undertook what amounted to a risk analysis that was intended to provide a base investment decisions regarding future nuclear energy production. The Nuclear Regulatory Commission, in January 1979, disalted the risk estimates of that report; new studies to estimate risk are already underway. But there is also a school of thought saying that we face too many risks every day to worry about one more. A commonly indicated estimate of the risk of a major nuclear power plant accident is 1 probability in 1,000,000 years. In the 1964 article, I delayed the risk image with a table of two-dice releases that would be required to give several results, from two 1s to two 6s, each of which has a probability of 1 in 36. There should be no problem visualizing or testing the meaning and possibilities of any of the events imagined by these dice. And, although 1 in 1,000,000 is presented in some way as mind-boggling compared to 1 in 36, and it's so unlikely to happen that it's beyond our ken, I suggest it's just visualized. We just need to use eight dice at once. If we plot all possible results for eight dice, as we did for both of us, we find that the sum of 8 (or 48) can occur in one way, through the 1s (or 6s). The odds of this happening are about 1 in 1,680,000. Therefore, the visualization of such probabilities, and more importantly, the lesson we must learn about risk—what incidents like Three Mile Island should teach us—is that what can happen will happen if we stay on it long enough. Any of us can simulate a statistical picture of the estimated risks or even the complexities of The Rasmussen analysis with enough patience and sufficient dice (or a computer). By the way, to make the eight dice act more like the odds of 1 in 1,000,000, simply mark two sides not 1 with a felt pen and count them as 1s if they appear; the odds of getting the 1s become a little less than 1 in 100,000. And the possibilities of human error can be included by similarly marking others given on the set. The difficulty is not in building such a simulation to portray probabilities, but in determining events that can lead to these probabilities and estimate the frequencies of their occurrence. Risk analysis has become one with public policy. Without it, any important choice that leads to uncertain results is uninformed; with it, properly implemented and understood, the decision-maker—business executive, government administrator, scientist, legislator—is better able to decide why one course of action might be more desirable than another. 1. The judgment factor in investment decisions, HBR March-April 1961, p. 99. 2. Monitoring Investments, Financial Executive, Apr. 1963, 19. 3. Capital Budgeting and Game Theory, HBR November-December 1956, 123. A version of this article appeared in the September 1979 issue of Harvard Business Review. Review. Review.