2

Technical Approaches to the Analysis of Risk in Project Economics

Introduction

This part of the Handbook provides a summary of available techniques for dealing with the outcomes of unknown events in project design and economic analysis.

It begins with a conventional definition of the terms “risk” and “uncertainty”. This is then followed by a summary of methods for dealing with the existence of project outcomes which can only be modeled as uncertain, rather than risky. This includes (but is not limited to) techniques already applied as standard practice by ADB. It then summarizes the techniques available for modeling risk on the basis of probability distributions (this is the most widely understood area of risk analysis), and highlights both advantages and limitations of such approaches.

This is followed by consideration of issues involved in modeling what are sometimes called "subjective" attitudes to risk—i.e., risk as perceived (for example) by those investing in or affected by a proposed project. This is an important aspect of risk analysis in project work because it is always the case that choice between alternative risky outcomes involves knowing something about decision-makers' or participants' preferences, and in particular how they are willing to trade off increased rewards against increased risks (i.e., measuring the extent of individuals' or decision-makers' risk aversion).

Risk and Uncertainty: A Definition of Terms

The term "risk and uncertainty" tends to be applied generically to the analysis of situations with unknown outcomes. This document will follow the conventional distinction between risk and uncertainty made in the literature [e.g., following Renbourg (1970) from which the quote below is taken, and also Reutlinger (1970), Pouliquen (1970), etc.].

In essence, risk is a quantity subject to empirical measurement, while uncertainty is of a non-quantifiable type. Thus, in a risk situation it is possible to indicate the likelihood of the realized value of a variable falling within stated limits—typically described by the fluctuations around the average of a probability calculus.

On the other hand, in situations of uncertainty, the fluctuations of a variable are such that they cannot be described by a probability calculus.

Thus risk and uncertainty are best thought of as representing a spectrum of unknown situations with which an analyst may be dealing, ranging from perfect knowledge of the likelihood of all the possible outcomes at one end (i.e., risk) to no knowledge of the likelihood of possible outcomes at the other (i.e., uncertainty).

It is important to realize at the outset that it is not the real-world situation itself which is either "risky" or "uncertain", but merely the information available to planners and analysts which defines it as such. All actual project outcomes are unknown, because they occur in the future and are subject to influence by a number of variables, each of which may take different values. If we have reliable historical or forecast data such that a probability distribution can be constructed for such variables, the situation can be modeled as "risky", if we do not have such data we can only describe the future in terms of "uncertainty".

In dealing with an agricultural project, for example, if historical rainfall patterns or irrigation water supply data exist we may be able to construct a probability distribution such that crop yields can be predicted in terms of expected values with associated levels of variability. If we do not have such data then possibly only "high", "likely" or "low" values for crop yields may be estimated, depending upon whether seasonal rainfall or water flow is above or below certain levels. Similarly, analysis of an energy project may be undertaken in terms of "optimistic" or "pessimistic" assumptions about domestic and commercial power demand levels (and different returns predicted under such different scenarios), or may be modeled on the basis of a distribution of outcomes of future power demand which itself depends upon estimates of economic growth, population growth, etc., and which may be described on the basis of their probabilities of occurrence. In both cases there is nothing inherently different about the circumstances of the projects themselves, only the data available to the analyst which makes modeling of risk more or less possible.

It should be noted that this distinction between risk (unknown but quantified outcomes) and uncertainty (unknown and unquantified outcomes) is not usually so clearly made in typical financial analysis. For example, the UK Treasury Taskforce on Private Initiative (2000) quotes an accounting definition of risk as follows

> "A simple definition of risk as used by the accounting profession is uncertainty as to the amount of benefits. The term includes potential for gain and exposure to loss."

Clearly such a definition blurs the distinction made previously between risk and uncertainty. It is argued that such a distinction is in fact very useful because it helps to separate those situations which may be subject to quantitative analysis from those which are not.

Allowing For Uncertainty

Project economic analysis (and the overwhelming weight of both ADB and World Bank experience) tries to allow for the existence of unknown future outcomes in the most basic sense by modeling the existence of uncertainty rather than dealing with risk *per se*.

Attempts to model the impact of uncertain outcomes and develop decision rules about what choices to make (e.g., between different projects or alternative project designs) derive from the operations research (particularly linear programming models) and game theoretic (e.g., von Neumann, Morgenstern) approaches of the 1960s and 1970s. In situations of different possible project alternatives and uncertain future events ("states of nature"), projects would be chosen on the basis of various proposed criteria, according to decision-makers' preferences. Such proposed criteria included

- "Laplace"—select the project or design alternative which yields the highest return, whatever the "state of nature" obtains
- MAXIMIN—select projects or design alternatives which yield the best returns (e.g., highest NPV) if the situation/"state of nature" turns out as badly as possible, and
- MINIMAX REGRET—select the project which minimizes the maximum opportunity cost of having made a wrong choice by choosing a "state of nature" which does not in fact obtain.

It can be shown that such criteria (and others that were developed around the same time) are in fact all "irrational" in different ways. For example, the "Laplace" criterion effectively ignores uncertainty altogether, the MAXIMIN assumes "nature" to be as malevolent as possible (which is not the case), and the MINIMAX REGRET does away with normal assumptions about decision-makers' preferences (because they are more concerned about minimizing losses *ex post* than about maximizing returns *ex ante*).

Despite some historical applications for planning purposes, for such reasons as just given, game theoretic criteria were largely abandoned as models of descriptive or prescriptive behavior, and subsequent practice in ADB and elsewhere has largely concentrated on describing unknown outcomes alone, without attempting to derive decision rules to guide choice under uncertainty.

The most widely-applied technique for describing uncertainty is sensitivity testing, and this is described in detail in the *Guidelines* (page 39, and Appendix 21), and also in the *Financial Guidelines* (section 7.11). A full explanation of the technique and its application is also provided in Belli et al's *Economic Analysis of Investment Operations* (World Bank Institute 2001).

In essence, sensitivity testing involves changing the value of one or more selected variables which affect a project's costs or benefits and calculating the resultant change in the project's NPV or IRR. Although emphases and presentation differ, both ADB and World Bank recommend practices such as:

- testing for the effects of changes in aggregate project costs and benefits
- testing for the effects of changes in individual underlying variables (e.g., areas, yields, crop prices in an agricultural project; prices of cement, operating costs of machinery in a roads project; consumer utilization rates

in a power or water supply project, etc.). This choice of variables will usually be based on previous similar project experience and/or detailed sector knowledge as much as on the particular project in question

- testing variables one at a time, so as to be able to identify the ones with most impact on project NPV
- testing for delays in benefits or implementation (e.g., shift the benefits stream down a year or two)
- testing likely combinations of variables (especially if these may in practice be linked —e.g., project costs go up AND implementation delays simultaneously occur), and
- testing for changes in economic pricing adjustments (e.g., shadow wage rate factor, shadow exchange rate factor, standard conversion factor, etc.) made by the analyst.

Sensitivity testing leads to the calculation of switching values (SVs) and sensitivity indicators (SIs):

- SV identifies the percentage change in a variable for the project NPV to become zero (i.e., for the project decision to switch between "accept" or "reject", average yields would have to fall by 20%). Sometimes SVs are expressed in terms of the absolute value of a variable—e.g., "if passenger traffic volume fell to 15,500 vehicles per day the project would not be viable"
- SI compares the percentage change in a variable with the percentage change in a measure of project worth (e.g., NPV).

The prime utility of sensitivity testing is that it leads to the identification of those variables to which a particular project design is most sensitive, and mitigating action can then be taken (if desired) to minimize the consequences of such outcomes. Likely mitigating actions include undertaking pilot projects, securing long-term supply contracts (for inputs and/or outputs), increasing technical assistance and training levels to support project implementation, publicity campaigns to promote service usage, tax and tariff changes, etc. The technique is extremely easy to apply, as changes to one value in a spreadsheet will reflect instantly in values for NPV, IRR, etc.

However, the technique has a number of limitations:

- most fundamentally, it does not take into account the probability of the occurrence of the events it models. The SV for crop yields may be a fall of 20%, or that for traffic flow may be 15,500 vehicles per day, but how likely is it that either or both will occur in practice?

- where deviations from project "base case" estimates are modeled in sensitivity testing, it is not clear whether the variations in values which are being modeled are changes from "expected" values (i.e., the "base case" estimate of the value of the variable is its average value) or are deviations from "most likely" (or modal) values; depending upon the characteristics of particular distributions (in effect the extent of skewness in the data set), mean and modal values may be very different from one another, and what is being captured in the base case and its variation is not clear;
- the identification of appropriate groups of variables to vary together depends on specialist knowledge, and misunderstanding the nature and extent of correlation between variables can lead to erroneous results; and
- because the distribution characteristics of different variables which determine project outcomes can differ enormously (the variability in commodity prices is less than input prices for example, the variability in power demand is less than in generation, etc.), the use of standard percentages for variations (changes of +/– 10% or 20% are routinely applied for example) in sensitivity testing captures quite differential extents of likely variability. An impression of homogeneous variability is given, which is not warranted by reality.

Overall, sensitivity testing is a highly subjective technique. Its ease of application and familiarity of concept, combined with analysts' understanding of particular sectors and projects (which should lead to reasonably appropriate variable selection and extent of variation being applied) plus its usefulness in leading to development of mitigating measures and project redesign, mean that its use has become widespread. However, its dependence upon judgment rather than empirical evidence and its modeling of uncertainty rather than risk (plus its inability to offer any decision rules following the presentation of its results) mean that its usefulness as a technique is ultimately limited.

One other point may be mentioned in regard to modeling uncertainty in project economic analysis. It is sometimes suggested that uncertainty can be allowed for by either applying a different discount rate in the calculation of NPV or by using a higher cutoff rate (i.e., greater than 10-12%) for investment decisions. While there is a large theoretical literature on this point, in essence there is no justification for this approach—apart from any other consideration (e.g., in determining what an appropriate "risk premium" should be), it assumes that risk always increases with time, which is not necessarily true. The discount rate is a rate of decline in the numeraire of economic value, and has nothing to do with the source of risks facing an investment.

Modeling Risk Quantitatively

Because of the conceptual shortcomings of all approaches to modeling uncertainty, various attempts have been made to properly capture the impacts of unknown outcomes through modeling risk quantitatively in project economic analysis.

As the *Guidelines* state, the purpose of quantitative risk analysis in essence is to

> "provide a means of estimating the probability that the project NPV will fall below zero, or that the project IRR will fall below the opportunity cost of capital." *(Guidelines*, Appendix 21, page 156)

While the *Guidelines* do not themselves provide a methodology for the application of risk techniques, it is suggested that the results of sensitivity testing be used to consider which variable(s) may be appropriate to base a risk analysis upon (i.e., those that have major impacts on project outcomes). Having identified particular variables, a number of possible data points (i.e., values above and below the "base case", upper and lower limits to data values, etc.) are necessary to be specified, together with the frequency (or likelihood) of each of these values occurring. From such data points and associated frequency estimates, a probability distribution can be constructed for the variable(s) in question.

Table 1 shows the numbers of bus passengers who may be expected on a particular route per week (and which may represent, for example, usage of similar services elsewhere in the city for an earlier project). The average number of passenger bus trips is 50,000, the standard deviation (the average difference between all observations and the mean/average) is 27,386 and the coefficient of variation (the ratio of the standard deviation to the mean) is 0.547.

The figure below shows graphically a simple probability distribution (based on frequency of observations falling within particular class intervals) for the number of bus passengers, deriving directly from the data in the table.

Table 1
Number of Bus Trips and Frequency of Occurrence

Number of bus trips ('000)	10	20	30	40	50	60	70	80	90
Probability of occurrence	0.025	0.05	0.1	0.2	0.25	0.2	0.1	0.05	0.025

Figure 1
Passenger Bus Trips Probability Distribution

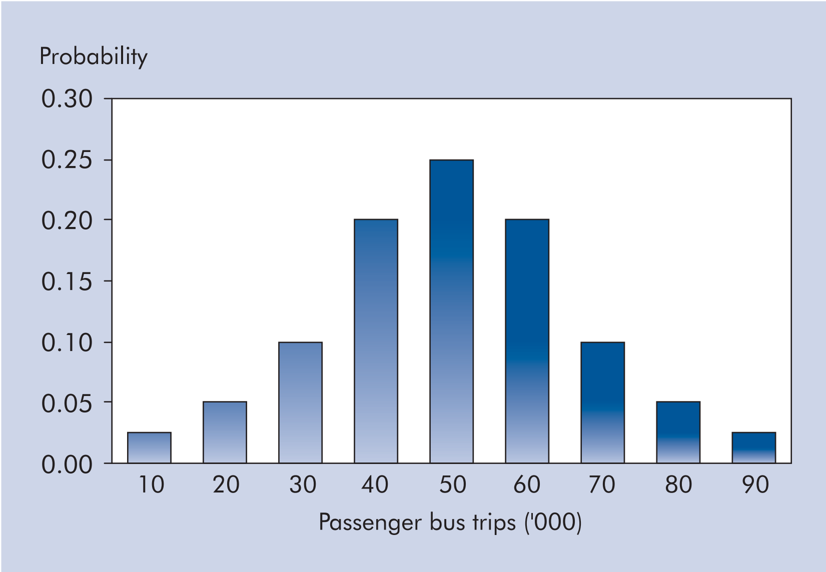

From the data, it is apparent that the expected number of bus trips is 50,000 (and this may be the basis for the estimate of the project "base case" scenario), but there is a 10% chance that the number of passengers may only be 30,000 per week, and a 5% chance that the number of passengers may only be 20,000 passengers per week.

The following figure (Figure 2) shows the same distribution in the form of a cumulative probability distribution function (which has been "smoothed" to allow for the effect of observations of grouped data). The function plots the cumulative probability that the actual outcome will be below a certain level of bus trips. It can be seen that, for example, there is a 27.5% probability that the number of bus passengers will be below 40,000 a week. In this example, actual numbers of passengers on project-financed buses each week may well determine both the financial profitability of the bus company and also the size of the economic NPV of the project. It may be the case, for example, that sensitivity testing has indicated that the SV (switching value—the value of the variable at which the project investment decision is changed) for bus passengers is at or near 40,000 per week—and we can now say what the probability is of this value obtaining. In this case, the probability is relatively high (i.e., over one chance in four), which may well imply that a redesign of the proposed project would be appropriate.

Figure 2
Cumulative Probability Distribution

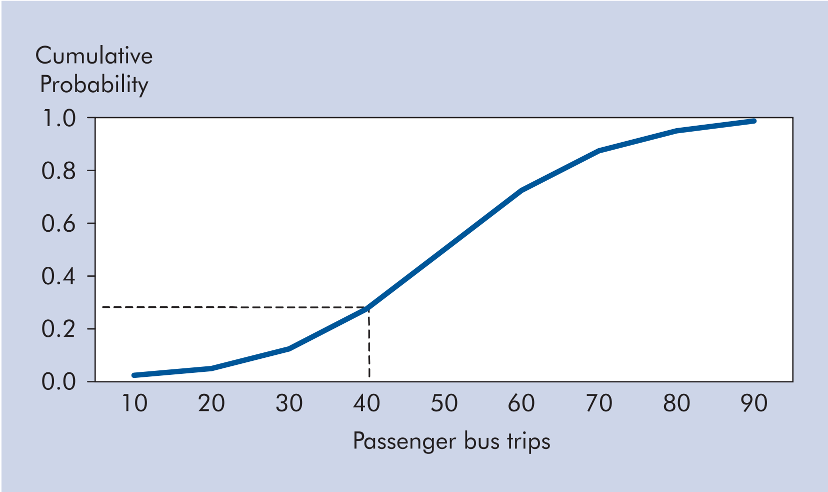

The ultimate purpose of quantitative risk analysis is to generate a similar cumulative distribution probability function for a project's NPV, such that the probability of negative NPV being generated is explicitly identified. The following figure shows such a chart:

Figure 3
Cumulative Distribution Function

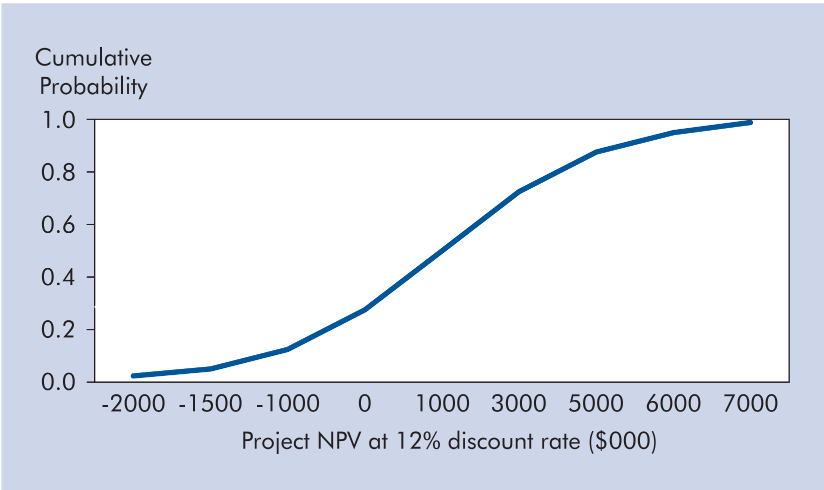

The probability of project NPV being negative as a result of variability in underlying factors is thus able to be identified (i.e., 30% in the figure), and provides further information as to the relative attractiveness (in terms of its riskiness) of the project. The same data could of course be used to model the financial profitability of the bus company in a similar fashion.

With the discussion so far, and most literature describing such techniques, it should be noted that identical procedures to these can be applied to projects where expected ENPV is not typically calculated in ADB practice (e.g., education and health projects, which tend to use measures of cost-effectiveness rather than total net economic worth). The only difference is that instead of expected NPVs would now be absolute measures of expected cost-effectiveness of outputs or impacts (e.g., annual cost per health worker employed, cost per primary school pupil educated) quoted together with distributions for those values (e.g., in 75% of circumstances the cost per health worker would be less than Baht 4,000 per year; the cost per primary school child educated would be Rupees 21,000 in four years out of five, etc.).

Risk analysis typically involves the choice of several variables to be varied simultaneously, as project returns are generally subject to more than one source of risk. Because of the mathematical complexity involved in such calculations, the analysis of risk in this form is invariably undertaken by some kind of computer software. The process which is followed (and which is usually referred to as "Monte Carlo" or simulation analysis) is that values for individual variables are generated randomly according to their respective probability distributions, combined with other randomly-generated values for the other variables, and these figures are used to calculate an estimate of the project NPV. This process is repeated a large number of times (a number which is specified by the analyst—in effect, equivalent to implementing the project again and again in different circumstances—and is usually at least 1000 times, and typically more than this) and an average (or "expected") NPV is produced together with an associated probability distribution. This distribution can usually be viewed in the form of charts like those in Figures 1 to 3.

The early literature on risk modeling (e.g., Reutlinger 1970, Pouliquen 1970) and also standard texts on project appraisal (e.g., Little and Mirrlees 1974, Squire and Van Der Tak 1975, etc.), as well as ADB's 1985 review, all mention the fact that computer time and expertise is likely to be a major constraint to the use of this technique. In recent years this constraint has largely been overcome and more than adequate computational facilities and software are now available to practitioners.

There are even examples of risk modeling using spreadsheets alone (see Clarke and Low 1993, for example). In the model described in this particular article, a standard spreadsheet (e.g., Excel, Lotus) is used to generate random values (through the use of a built-in function—@RAND in the case of Lotus) which are then repeatedly applied to variables identified within the project "base case" scenario so as to generate a distribution of outcomes (counts of observations of variables/expected NPV values are recorded in cells so as to construct a frequency distribution). Expected NPV and IRR for the project are calculated based on the distribution of results, and the base case is then presented in terms of expected NPV/IRR, plus a range in which outcomes will lie 90% of the time (based on all variables being subject to simultaneous change). It is to be noted that the technique is quite simple and practicable, and requires no more information than is already required for traditional sensitivity testing.

Technical Issues in Modeling Risk: Two Considerations

However, despite the overcoming of computational limitations to the application of such techniques, two major practical considerations (and possibly constraints) remain as regards the extent to which such techniques as Monte Carlo simulation can be used in project preparation situations.

Firstly is the issue of data availability, and the extent to which the situation can reasonably be defined as risk (as opposed to uncertainty) through the construction of a meaningful probability distribution of outcomes. The actual situation with data availability is likely to vary enormously both between project situations and also across different sources of variability within any one project environment. At one extreme, large volumes of reliable cross-sectional or time series data may be available from historical sources for the variable concerned (e.g., for rainfall, for commodity prices, for traffic flows). At the other extreme may only be the existence of a few data points (e.g., "most likely" values, absolute minimum possible, maximum possible, etc.) which are expectations of experts/analysts involved in preparing the project. Other possibilities lying within these bounds include the forecasting / specification of power-generation theoretical capabilities adjusted for a set of likely different operating conditions, forecasts of trade flows and commodity prices (e.g., based on World Bank publications taking into account world supply and demand factors), etc. Software such as @RISK often has capabilities to fit probability distributions of different types to raw data sets supplied by the analyst. Such routines will fit distributions to data and also provide a measure of the goodness of the particular resulting fit.

It is important to note, therefore, that very large and complete data sets from empirical sources are not always necessary for the undertaking of risk analysis. Simplifying assumptions about variable distributions can be made—as a bare minimum, triangular distribution from three points (i.e., "most likely", "minimum possible", "maximum possible") can be constructed based on "best guesses" of project preparation team members. There is also often considerable expertise within the project preparation environment about the likelihood of variables or outcomes which may not be available from official sources but which can be elicited from potential project participants. Good examples of such knowledge might include rainfall and water flow, crop yields, time taken to collect water or firewood, or travel to market, operating efficiency of agroprocessing machinery, number of family days sick per year, etc. The 'Delphic' method of eliciting opinion from local experts is an example of this type of approach, and has been applied (in a probability-based form) by, for example, World Bank in a risk analysis of institutional reform in the irrigation sector in Pakistan (Dinal et al. 1997).

Well-tried empirical methods exist for developing probability distributions from such subjective sources. These include

- visual impact techniques (e.g., matches or stones piled up to represent frequency of value occurrence),
- structured questions to identify key points in a distribution (e.g., the median, quartiles, etc. – the "judgmental fractile" method), and
- the application of "smoothing" techniques in situations where a few real data points may be available.

Some proponents of probability-based risk analysis (e.g., Clarke and Low 1993) also argue that the shapes of particular distributions of individual variables (e.g., choosing between uniform, normal, triangular distributions for variables such as crop yields, traffic flow, enrolment rates, income differentials, etc.) are less important than the choice of variables themselves which are allowed to be modeled. In the Clarke and Low example (from an agricultural project in East Africa), the random number generation approach necessarily produces a rectangular uniform distribution (i.e., one in which all possible observations within the range defined by the analyst have equal probability of occurrence), although with more complex formulae being written this could be adjusted. Recent experience of preparation of power projects within ADB also suggests that the particular form of distributions also matters less if a large number of simulations are run. Even when considerable effort is made, for example, to replace the quoted discrete distributions with relatively few values by continuous distributions based on large amounts of empirical data, there is little difference in resulting distributions of EIRR/ENPV outcomes (i.e., expected values and variance, minimum and maximum values, etc).

However, this approach may not always be appropriate for all variables, and it still requires judgment on the part of the analyst about what ranges are acceptable for values to fall within. Also, to adjust the spreadsheet model to produce (for example) a normal distribution (as opposed to a uniform one) becomes very much more complex.

It is also noted in the early risk literature (e.g., Pouliquen 1970) that there is no a priori case for the use of normal distributions (even as a default distribution), as all variables are not always subject to relatively large numbers of random influences (which is what typically causes a variable to be normally distributed).

Overall, therefore, the techniques applied to develop definitions and derivations of probability distributions for individual variables in most cases is likely to depend upon some subjective judgment by an appraisal team — and inevitably the extent to which these design assumptions adequately reflect the reality of the project will vary from case to case. The suspicion that what appears as a full-scale risk analysis has in reality only a spurious precision can be ultimately only fully allayed if the data upon which the variables' probability distributions are constructed (either from historical evidence, future computational forecasts, or analysts' 'best guesses') are believable.

The second consideration when applying risk analysis in practice is the extent of covariance between those variables that are to be selected for risk analysis. Projects are rarely subject to only one source of risk, and therefore more than one variable at a time is modeled in the Monte Carlo simulation exercise. However, statistical complexities can arise depending upon the relationships between the selected variables. Where variables are in fact statistically independent of one another there is no problem, as it is appropriate to treat them independently. Where variables may be thought to be related in some way, however, the extent of covariance between them needs to be taken account of when specifying the distribution of individual variables in some type of simulation (again, typical risk analysis software such as @RISK can handle this within the context of specifying correlation matrices).

As an example, project revenues are typically products of both quantities sold and prices obtained. If these underlying variables are correlated in some way (which may well be if project output is large relative to market volume, and negatively so in this case) the expected value of the product of two random variables (i.e., project revenue) is equal to the product of the individual expected values plus the covariance between the two variables. Another typical example of covariance (which would be positive in this case, if it is assumed to be due to improved, project-induced water supply under an irrigation scheme improvement) may be that between area planted and average yield (i.e., with both variables as determinants of farm production volumes). One of the case studies using @RISK includes correlation between estimates of proportions of trainees finding employment, and the numbers of days and months of employment gained following nonformal training in Bangladesh. In practice, the approach to assigning particular levels of covariance between variables is quite pragmatic, and typically simple rank correlation coefficients between pairs of variables are sufficient for most purposes (in the risk software packages, correlation between variables—once specified—can typically be "toggled" on/off).

It is specifically recommended that disaggregation of individual variables be limited as much as possible so as to avoid including too much correlation in the analysis. For example, although individual construction cost items (e.g., cement, cost of floor, cost of walls) may each be thought to vary individually, in reality the sources of this variability all arise from one point (e.g., costs of imported cement), and this could be most appropriately captured through some item such as "construction materials" rather than by introducing additional correlation between such items (which would tend to increase unnecessarily the estimate of overall variability). Akin to the nature of some subjective judgment being involved in the allocation of probability distributions, there is therefore a similar judgment to be made about the extent of disaggregation to be applied in individual circumstances.

In sum, the principles to be applied in practical situations to quantitative risk analysis such that the issues just described are dealt with as transparently as possible are summarized in the following table:

Table 2

Principles to Apply in Data Handling for Probabilistic Risk Analysis

	Principles to Apply
1	Identify those variables for which future values are unknown and which are likely to affect project returns (i.e., the 'key' variables)
2	Fully explain the general nature of the data set which is used for modeling those variables' values (its origin—i.e., from objective or subjective sources, whether it is based on historical observations or forecasted projections, the number of observations the data set contains, its extent of completeness/any missing data points, etc.)
3	If the data derives from subjective sources, explain the method by which it was elicited (e.g., from visual techniques, from subjective questioning, from an expert-based 'Delphic' process, etc.)
4	Explain the statistical nature of those variables' assigned probability distributions (i.e., whether these distributions are triangular, uniform, normal, logarithmic, exponential, etc.)
5	Make clear the goodness of fit of the distribution to the data set (if one has been fitted using @RISK or similar software), and quote appropriate statistical measures (e.g., Chi-square, Kolmogorov-Smirnov, Anderson-Darling statistics, etc.)
6	Make explicit any correlation thought to exist between variables used in the risk analysis (i.e., its extent, and the technical, real-world basis for the assumption, etc.), and (based on this)
7	Explain and justify the extent of any variable disaggregation.

The focus of the reporting of the results of risk analysis will be the likelihood that project returns may be negative (e.g., "the project's rate of return is likely to be unacceptable in about 15% of all cases"), but such reporting should not obscure any basic qualifications about the analysis which need also to be included, and especially about the extent of correlation which has been assumed. "Good practice" in this regard is likely to require reporting "with" and "without" correlation being applied in the simulation as the effects of ignoring correlation can be substantial.

Risk Analysis, Decision-Making, and Welfare

The result of risk analysis as just described is therefore to identify projects (or alternative designs of the same project) which now have two essential characteristics—i.e., the expected value of their economic return (as measured by their expected ENPV, EIRR, etc.) and their degree of risk (as measured by their variability in general—captured by the distribution's measures of dispersion, such as variance and coefficient of variation, and also by the probability of the return falling below some unacceptable level in particular). This quantitative measure of risk adds to the information available to the decision-maker, although in itself does not necessarily provide any guide as to whether any individual project is acceptable (or even as to which project among several possible ones actually should be undertaken).

The discussion of the 'acceptability' of particular levels of risk is usually presented in standard economics texts in terms of choice among competing projects. Consider the following three project alternatives (A, B, and C) as shown in Figure 4.

In Figure 4, project B is clearly "inefficient", in the sense that its variability is the same as project A, but its expected value is lower. Project C has a higher expected NPV than project A but it also has a greater variability of returns (including the possibility that its return will be zero). If variance of returns is plotted against expected values for such project situations as A, B, and C, the following is obtained (Figure 5 - in what is usually referred to as E-V space). Projects A and C may lie on the efficient investment frontier—although at different points, while project B lies within the inefficient frontier.

In this case, the expected higher returns of project C have to be weighted against the increased degree of risk of the project. How should a decision be made between the alternative projects? The traditional view taken towards public sector investment was that governments with large project portfolios could afford to ignore the riskiness of investments as long as the expected values were acceptable—i.e.,

Figure 4
Probability Distribution: Three Projects

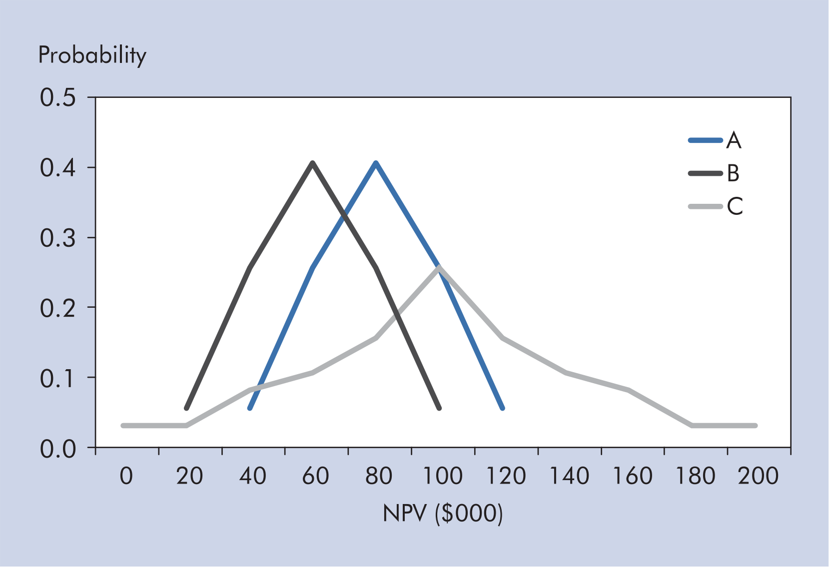

Figure 5
Expected Returns-Variance Frontier

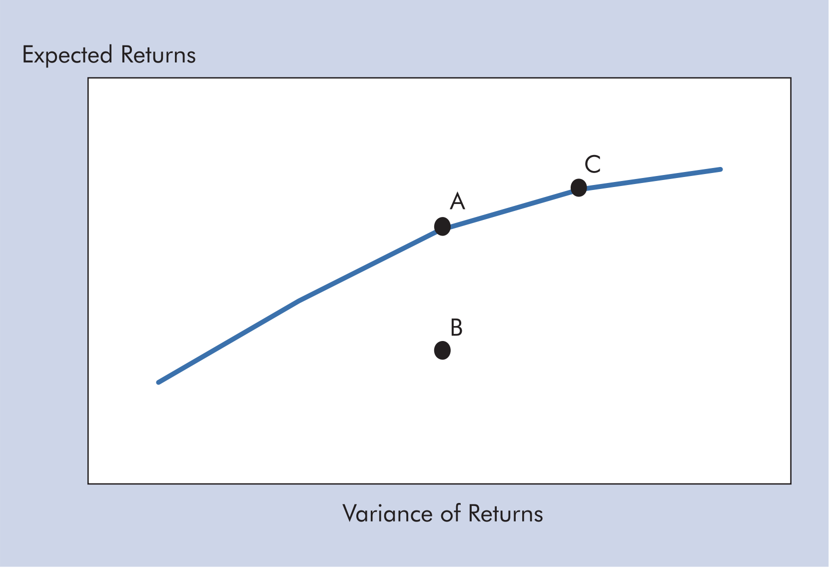

they could afford to be "risk-neutral". This is because, with a large number of investments spread across all of society, the costs of any individual project failures could be absorbed within the portfolio as a whole. Exceptions to this view were projects which were either very large, or were somehow correlated with national circumstances (such that good performance of the project in bad years for the economy as a whole was worth more in terms of a disproportionate contribution to national income), or affected particular groups (e.g., in one region, one type of student, etc.) such that the impact on those particular individuals could not be ignored.

In fact, there is no answer to the question of project choice in such circumstances without reference to knowledge about the extent of decision-makers' risk-aversion (i.e., the rate at which they are prepared to trade-off levels of expected returns—or in effect, income—against levels of variability of returns). Figure 6 shows how many individuals are thought to be risk averse; the line joins points of equal utility (i.e., welfare or satisfaction – and is in fact an indifference curve) from combinations of expected income (measured, say, in $000) and variability of income (measured by annual income variance). What it suggests, in general, is that for people to accept greater variability in their incomes they need to receive increasingly larger expected incomes.

Figure 6
Risk Aversion: Line of Equal Utility

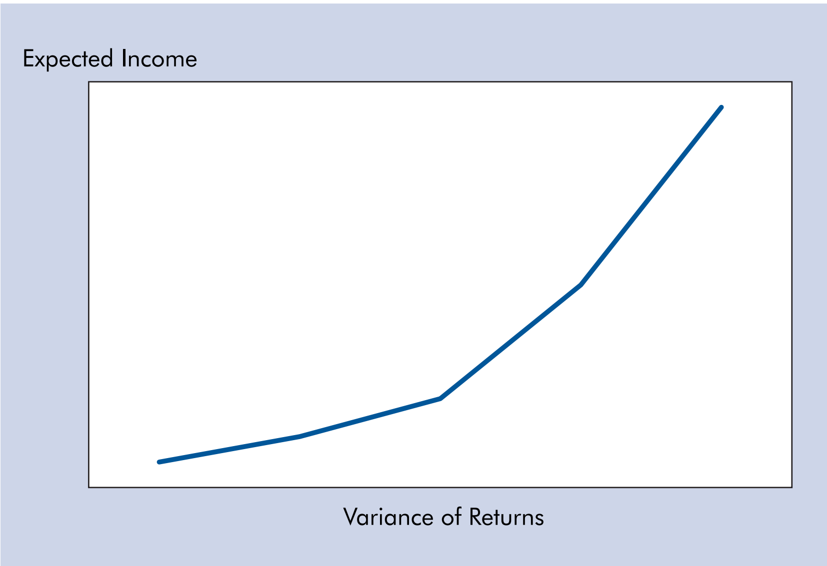

The actual nature of decision-makers' utility functions as regards their risk aversion occupied a large literature (much of it deriving from agricultural economics and work with small farmers in developing countries) around the same time as the classic texts on project appraisal and risk analysis were written, where it was hoped that insights about preferences regarding risk would lead to making optimal project choices. The results of many classic empirical studies of risk perceptions from the 1970s (e.g., Anderson 1974, Swalm 1966, etc.) suggested that most individuals were risk averse, although the specific income levels at which this was the case and also the extent of risk aversion could vary enormously even within apparently similar circumstances (small farmers in Africa, Australia, and the USA, workers in large corporations, professionals and business executives were all studied).

In order to fully understand how people actually deal with risk in project and other situations in terms of the decisions they make, academic risk literature has also concentrated on the extent to which

- the consequences of particular risks are catastrophic or not,
- the risks are controllable at the micro level or not,
- the consequences are reversible or not, and
- the risks are insurable or not.

These considerations are returned to in Part IV in relation to the attitudes of the poor.

To complete the above analysis and understand what decision actually *should be* taken in any particular circumstance, it is necessary to combine the notions of efficient projects located in E-V space with decision-maker's utility functions expressed in the same space. Thus in Figure 7 below, one decision-maker on the basis of an individual utility function (U1) will choose project A (lower expected NPV, lower risk) while another will choose project C (higher expected return, higher risk) on the basis of a utility function such as (U2).

As well as the E-V approach to considering the risk-reward relationship, other models have attempted to capture the characteristics of this situation within the context of financial investment portfolio analysis. Probably the most well-known of these approaches is the Capital Asset Pricing Model (CAPM, developed by Sharpe and Lintner in the 1980s) which measures an individual stock (or project) risk relative to the volatility of returns relative to a market (or sector) index. Again, however, ADB practice as regards project economic and financial analysis is not primarily concerned with a number of projects as they may comprise an investment portfolio (or even with one project within the context of all projects in the portfolio), but with analyzing risk as faced by individual projects one at a time.

Figure 7
Project Choice and Risk Aversion

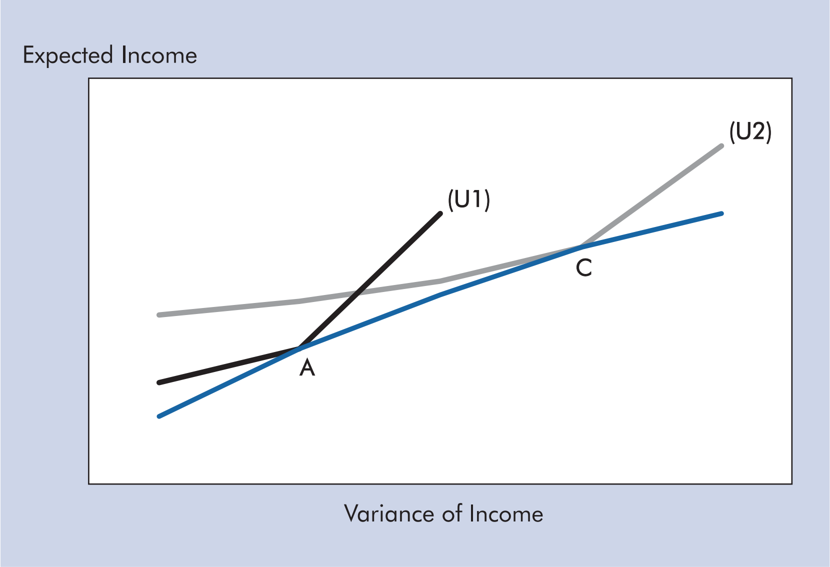

Summary of Approaches to Risk Analysis

This Part of the Handbook has summarized approaches to dealing with unknown outcomes in project analysis through description of various techniques attempting to capture the content and consequences of uncertainty and risk. It is clear that the quantitative modeling of risk is in principle preferable to the simple depiction of uncertainty, although it is also obvious that data and time considerations (and the difficulties in properly identifying and specifying covariance among variables) have often limited the extent of actual risk analysis practice. However, it is also the case that the present availability of computer software to easily process Monte Carlo-type simulations from data already available to the analyst within spreadsheets, plus the existence of statistical routines and computational processes to fit probability distributions to many (even quite sparse) data sets, greatly increase the possibilities for application of quantitative risk analysis as a more commonplace part of project design.

The results of quantitative risk analysis can greatly inform the process of project design, so that mitigation measures can be put in place before projects are implemented. They can also increase the information available to decision-makers about individual projects (i.e., so that their likelihood of failure is identified). While the results of risk analysis are quite easy to understand in conceptual terms for planners and project staff, however, these results in themselves do little to aid project investment decision-making unless and until they are combined with some consideration of planners/policymakers and/or project participants' risk aversion. What level of risk is it appropriate

- for ADB to accept?
- for ADB to suggest that borrowing DMC governments or institutions accept?
- to impose upon groups of project beneficiaries?

are questions which have to be answered outside the scope of those quantitative techniques which provide the measures of project risk. The question is akin to the difficulty of suggesting *a priori* cut-off points for the Poverty Impact Ratio (PIR) as it is inherently associated with social welfare functions as perceived by decision-makers.

Following a review of agency experience with the application of the techniques just outlined (in Part III), it will be suggested (in Parts IV and V) that strengthening risk analysis in ADB operations involves both greater application of the probability-based techniques for modeling sources of risk in project design (located within a framework to guide such application) and also greater awareness of attitude towards risk amongst those being planned for in ADB projects.