PricewaterhouseCoopers

Uncertainty & Risk Analysis

A practical guide from Business Dynamics

PricewaterhouseCoopers, MCS

Business Dynamics

PRICEWATERHOUSECOOPERS

Uncertainty & Risk Analysis

by

Chris Rodger and Jason Petch

BUSINESS DYNAMICS

April 1999

For more information contact Chris Rodger via e-mail or Tel. + 44 (0) 171 213 3519

Table of Contents

1 *Introduction* 1

2 *Phases of risk analysis* 2

3 *Identifying the risks* 4

Introduction 4

Initial risk identification workshops 4

The risk register 8

Using a model to identify risks 9

4 *Quantifying risks* 10

Which risks should be quantified? 10

Theoretical background to quantifying risks 11

Separating risks out 20

A review of the issues that arise when quantifying risks with experts 21

5 *Risk analysis* 24

Introduction to Monte Carlo analysis 24

Identifying the significant risks 28

A practical example of the identification of key risks 31

6 *Presenting the results* 32

Graphical presentation 32

Statistical measures 35

Presenting it in English 37

7 *Beyond presentation* 38

Introduction 38

Valuing risk 38

What does a distribution of NPVs mean? 40

Risk analysis in PFI transactions **41**

Appendix A Categories of generic risks 43

With acknowledgement to; Jonathan Batson, Martin Upcraft and James Angus

1 Introduction

1.1 Uncertainty and risk analysis is not new; however, as a tool in business it has historically been of limited use. This is surprising considering that many business decisions are based on a figure that has been calculated from analysis of some kind. A number on its own is only half the picture; to fully understand the result it is necessary to have an estimate of the uncertainty related to that figure. For example, two projects may both have a value of £10m, but if we know that one has an uncertainty of +/-£1m around the £10m and the other has an uncertainty of +/-£3m, this shows them in a very different light.

1.2 Management will need to consider carefully their attitude to risk before making a decision about whether to accept either or both of the projects. It is frequently the case in project appraisals that large amounts of effort go into generating the expected value (i.e. the '£10m' figure), but very little time is spent understanding the uncertainty around that value.

1.3 This document gives an overview of how to carry out uncertainty and risk analysis modelling projects. In particular, it focuses on the use of Monte Carlo simulation in a spreadsheet model. This is not only because this is a simple but very powerful technique, but also because by covering this technique in the context of carrying out a complete risk project, all the fundamental uncertainty and risk analysis skills will have been covered.

1.4 In general, the word ‘uncertainty’ means that a number of different values can exist for a quantity, and ‘risk’ means the possibility of loss or gain as a result of uncertainties. We have tended to use both terms interchangeably in this document, and indeed this is common practice by most practitioners. However, it is important to understand the difference, as in some situations it may be necessary to apply the absolutely correct term to avoid ambiguity.

1.5 It has been assumed that the reader has a basic knowledge of mathematics and is familiar with spreadsheet modelling. For further help, please contact the Business Dynamics group.

2 Phases of risk analysis

2.1 This section gives a brief overview of our recommended approach to tackling uncertainty and risk analysis projects. The different phases in a typical risk project are shown in Figure 1 below:

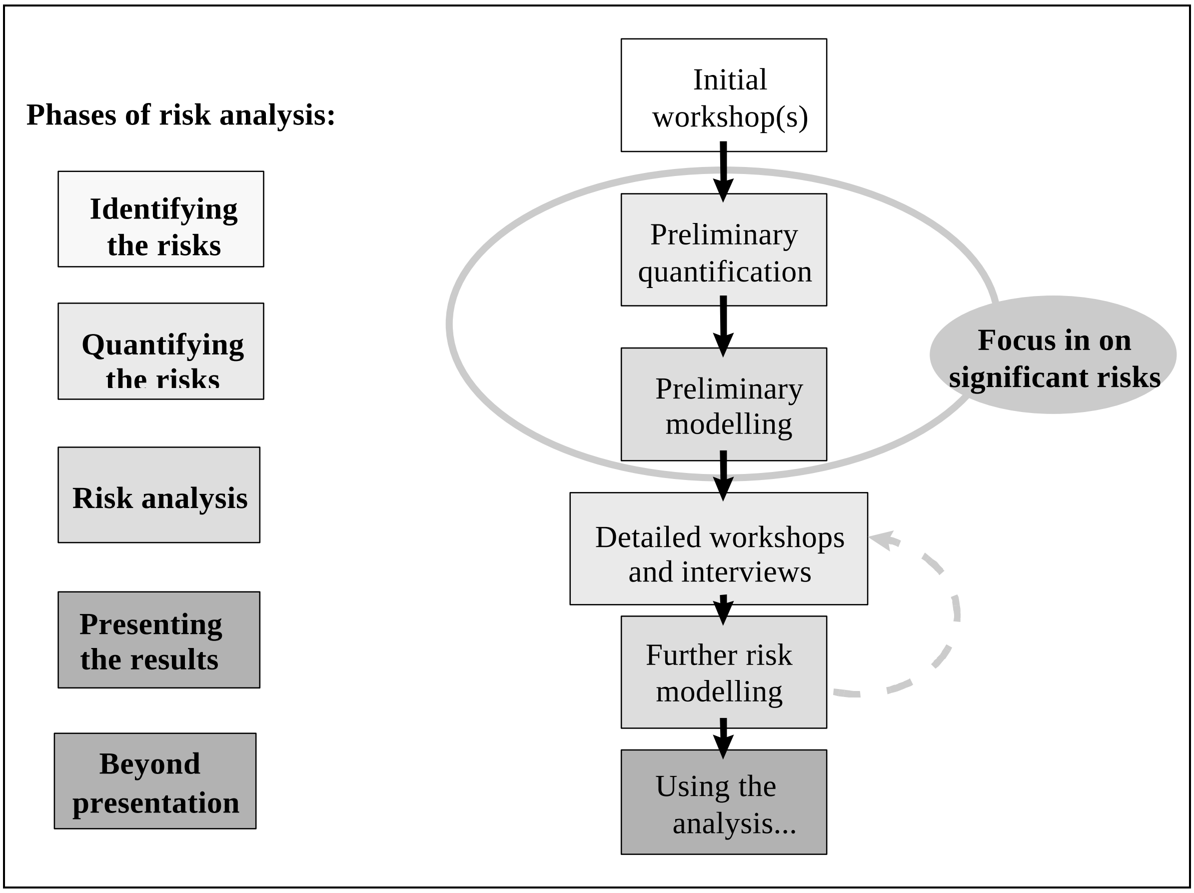

Figure 1: Different phases in a typical risk project

2.2 This flow diagram shows the ideal approach to risk projects; however, it is often the case that risk analysis is included as part of a larger project, and the risk work does not follow such a neat progression. It is important to be pragmatic as to the exact approach adopted, but to bear in mind the above as a conceptual structure that underpins the analysis.

2.3 The structure of the document follows the phases shown above, namely:

- **identifying the risks:** describes techniques used to get a list of possible risks, and how to determine which risks are appropriate for modelling;
- **quantifying the risks:** this section looks at issues that arise when trying to accurately quantify risks, such as which distribution is appropriate for what type of process, what is correlation, etc. It also gives an overview of the issues that can arise when quantifying risks with clients;

- **risk analysis:** this section is devoted to the "how to" of Monte Carlo simulation within a spreadsheet model, from the impact upon model design to the generation of outputs;

- **presenting the results:** describes the different ways of presenting the results of uncertainty and risk analysis, both graphically and in translating the results back into easily understood terms; and

- **beyond presentation:** this last section looks at how to take the outputs from the modelling and interpret them in the context of business decision making.

3 Identifying the risks

Introduction

3.1 This section covers the identification of risks as part of an uncertainty and risk analysis project. We consider the following topics:

- the initial workshop sessions held to identify the risks;
- the output of the identification exercise: the risk register; and
- the use of existing models to identify risks.

Initial risk identification workshops

3.2 The identification of risks is best done by means of a brainstorming exercise in a workshop, or series of workshops. In fact, it is often helpful to split the work along natural boundaries, e.g. hold separate workshops for the construction and operations risks.

3.3 The purpose of the brainstorming should be purely the identification of risks, and a description of them in clear and unambiguous language. There should be no attempt to quantify risks at this point, other than inviting the group to consider the likelihood of occurrence, and the consequences of, each risk. Both of these should be graded on a scale of low, medium or high, and these assessments are used to prioritise the quantification of the risks later in the process.

3.4 We recommend that numerically quantifying the risks is not done at the same time as identifying risks because quantification is a complicated process, and care must be taken to ensure that experts form their own views after some thought. There is a danger that the group dynamics can give rise to a conformist of point of view, and thus a simplification of treatment and underestimation of the level of risk. This is particularly the case if the quantification is done without sufficient preparation and forethought, and too early in the process. The identification of risks is less prone to these problems, and the creative benefits of group work outweigh the dangers.

3.5 In selecting who should attend, it is often a good idea to choose those individuals who will help in quantifying the risks in the next phase. Also, in order to achieve buy-in to the risk analysis, it can also be helpful to include those to whom the final output report is directed, e.g. senior managers, in at least one session.

3.6 Some ideas as to who might be involved are given below:

- construction experts, such as architects, designers, and quantity surveyors;
- project managers;
- operational managers;

- technical consultants, where specific technical issues are relevant;
- financial and legal advisors; and
- the risk analyst.

3.7 Before beginning the risk identification in a workshop, it is useful to explain clearly what the purpose of the exercise is, how the risks will be used and the opportunities that the participants will have to review and modify the output from the session.

3.8 A useful tool to help structure the thinking is a list of typical risks to which the project may be exposed. Appendix A contains such a list, focused mainly towards transactions, although many of the risks shown are generic.

3.9 There are a number of issues that should be borne in mind during the quantification process, namely:

- the nature of the variables;
- the dangers of double counting risks;
- the danger of missing important risks out; and
- the inclusion or not of rare events.

3.10 Each of these is described below in more detail.

Nature of the variables

3.11 A ‘variable’ in the context of risk analysis means a formulation of the uncertainty into a quantity that can be measured or estimated. Not all unknown variables should be treated in a probabilistic manner. In particular there are two types of variable, known as *decision variables* and *value parameters*, that should not be treated in this way. We discuss each in more detail below:

- **decision variables** are quantities over which the decision maker exercises direct control, e.g. whether to construct a new factory in India or Malaysia. Although the decision makers may well be uncertain as to the best location, it is not helpful to carry out risk analysis with the decision variable included as part of the uncertainty. In this example, it would lead to a hybrid result that combines something of the factory in India and something of the factory in Malaysia.

- **value parameters** represent aspects of the preferences of the decision makers or people they represent. Examples of these are the discount rate applied to cashflows, and the value of a human life (i.e. the amount of investment per statistical death averted). Again these parameters have no "true" value, but are rather the reflection of a value that is selected as appropriate for the situation.

3.12 Both decision variables and value parameters can, and probably should, be varied as a part of the risk analysis, but they should be varied parametrically; that is, the analysis is repeated for a range of particular values of the decision variable or value parameter.

3.13 As part of the brainstorming session it is valuable to separate out which of the uncertain variables are decision variables or value parameters. Note that this classification can be dependent upon the point of view that the problem is being considered from. For example: in the valuation of a proposed PFI transaction from the Government's perspective, the performance of a contractor may be a true empirical variable that should be treated probabilistically. From the point of view of the potential service provider, however, this would be classed as a decision variable in that the management of the firm can largely control the performance standards that it will achieve.

The dangers of double counting risks

3.14 A common error in uncertainty and risk analysis is to double count the risks. An example of this can be found in PFI projects: there is a risk to the contractor that they have mis-estimated the resources required to provide the service. There is also a risk that they will incur penalties under a performance mechanism that penalises failure to provide the service. However, these risks are not independent: as mentioned in the example in the previous section, management can generally control the level of performance: they do this by deciding on the level of resources that will be deployed.

3.15 Thus, the risk is either that:

- insufficient resources have been budgeted, and more are required. Hence costs increase, but there are few performance penalties levied; or

- resources are kept at the budgeted level, and performance penalties are incurred.

3.16 The decision as to which facet of this risk is modelled would be taken in discussion with the relevant managers; in practice, it may be that a combination of the two are included. Whatever solution is chosen, it is important that the descriptions of the risks contained in the risk register make clear the links between the risks, and are explicit about the joint treatment that has been used.

The danger of missing important risks out

3.17 A significant danger in uncertainty and risk analysis is that a risk is simply missed out. This can be more of a problem when a generic risk list is used, in that this can make the group feel constrained and channel participants' thinking. Despite this, we feel that the benefits of a generic list normally outweigh this disadvantage.

3.18 After using a generic list it is often useful to run through the project, considering how it might turn out in practice. It is likely that most of the risks identified this way will already have been captured, but it may throw up some that were not. This is a useful 'catch all' to help ensure that nothing has been missed.

3.19 There are no other easy solutions to the problem of missing out risks. It is important to emphasise, therefore, that the normal consultancy skills involved in running brainstorming sessions, asking probing questioning of the participants and challenging their thinking, are all the more vital.

The inclusion or not of rare events

3.20 The treatment of very rare and potentially catastrophic risks needs to be carefully considered. Rare events can have a significant effect on the overall risk, but the usefulness of including them must always be taken into account. For example an oil production company might have a 0.3% chance of failing to complete the installation of a pipeline by the end of a weather ‘window’ (i.e. a period during which conditions are calm enough to allow work to proceed). Should this occur, the loss could be very large, such as £100m; this would give an expected value of the risk of 0.3% * £100m = £0.3m (we explain expected value in section 7.3).

3.21 However, it is important to consider whether or not the inclusion of this risk is useful, bearing in mind the other risks under consideration. If all the other risks in the analysis involve realistic variations in costs, but do not involve such a catastrophic event, the inclusion of this risk may distort the analysis. And in fact, the consequences of missing the weather window are so extreme that valuing it at £0.3m is almost meaningless. Hence it may be better in this situation to omit it from the risk analysis.

3.22 Note that although we suggest that this risk is not included in this case, if the oil company is doing very many such projects, the total risk across all the projects is the sum of the expected values for each project. However, although it may be useful to understand the total corporate risk, it does not help in assessing individual projects.

3.23 As a general rule, we suggest that those risks that are of interest only at corporate level are not included in risk assessments at project level. A common example of such risks is those that threaten the entire viability of a company. For example, if a plastics company is faced with the risk that the plastic that it manufactures may be found to be carcinogenic, then including this risk in the analysis of one particular new product does not add any value.

3.24 It is also helpful when considering whether to include a risk to compare the likelihood that it will occur with the length of the project appraisal. For example, if considering a period of 10 years, a risk with a likelihood of occurrence of once in a thousand years may not be relevant. However, if there is a risk that on average occurs once in a hundred years, this probably should be included. As an example, North Sea oil installations are designed to withstand the ‘once in a hundred years wave’, but not the once in a thousand year one.

3.25 Finally, it is worth noting that some rare events can be covered by insurance; for instance, fire risk is often insured. In cases where insurance has been taken out, the risk, or at least most of it, has been transferred to someone else, and hence should not be included in the analysis of a particular project.

A summary checklist for every risk identified

3.26 The following sets out a checklist of issues to consider for every risk identified, summarising the previous points in this section:

- what is the variable that can be used to quantify the uncertainty? It must be a variable that relates to the output of interest. If this is a Net Present Value (NPV), then the variable must relate to a cost or revenue;

- is this a variable that the client can exercise direct control over (i.e. is it a decision variable?);

- is it a variable that has no true value, but is an expression of a subjective quantity (i.e. discount rate or value of life, value parameters)?;

- is there any overlap between this risk and any that have gone before? If there is, then either one risk or the other should be included, or the definition of the risks must be amended such that there is no overlap; and

- is this risk project specific, or is it a global risk that applies to the entire organisation (in which case, it might be better to omit it from the analysis)?

3.27 Once the session has been completed, briefly consider if there are any rare events that should be included.

The risk register

3.28 The output from the risk identification phase is normally a risk register; this is a document that contains some or all of the following for each risk:

- generic risk - the general heading the risk falls under, e.g. construction risks;

- specific risk - the particular risk being considered, e.g. bad weather delays construction;

- qualitative description of risk - a brief description and perhaps an example scenario;

- assessment of the impact of the risk - this is a subjective assumption as to how sensitive the output is to this particular risk, classified into high (H), medium (M) and low (L). It can be useful to set boundaries for these assessments, for example: high is more than about 2% of total project cost, medium is 1% to 2%, and low is less than 1%. These figures are only guidelines, and should be varied as appropriate for particular projects; and

- assessment of the probability of the risk - this is a subjective indication as to how likely this risk is to occur, again classified into high, medium, and low.

Using a model to identify risks

3.29 In addition to the risk identification approach described above, it is possible that a model constructed for the purpose of appraising a project, and not specifically for risk analysis, may yield valuable information. Sophisticated sensitivity analysis can be carried out on the model (described later) to identify those inputs that have a significant impact on the output. The complicated nature of contractual arrangements and payment mechanisms in complex deals, which should all be captured in the model, can lead to surprising results. Inputs that were not thought of as particularly important can have unexpectedly large effects on the output.

3.30 This technique may be less useful if the model has been purpose-built following a risk elicitation exercise specifically to do uncertainty and risk analysis, rather than as a cashflow model of a deal. In this case, the model will represent the risks in the way that the experts have thought about them.

4 Quantifying risks

4.1 In this section we discuss the next step in the risk analysis methodology: quantifying the risks that have been previously identified. The section is divided into the following topics:

- a brief commentary on which risks should be quantified;
- a discussion of the theoretical background to quantifying risks;
- the benefits of separating risks out; and
- a review of the issues that arise when quantifying risks with experts.

Which risks should be quantified?

4.2 The output from the identification of risks phase should be a risk register that includes an assessment of the likelihood of each risk, and its impact. The two criteria can be shown on a three by three matrix;

	Low	Medium	High
High	3	2	1
Medium	4	3	2
Low	5	4	3

Figure 2: Criteria for risk quantification

4.3 The numbers in the boxes show the level of priority that risks falling in that category should be given in the quantification. The closer to the top right of this matrix, the more significant the risk is likely to be, and the more important it is to focus in upon it.

4.4 Given the constraints of time, it may be that many of the risks rated 4 and 5 in the above matrix will not be quantified. In deciding not to quantify such risks, the effect that the risk has must also be considered: for example, there may be a large number of risks identified that are all fairly low priority, however they all contribute to the same effect, such as extending the length of construction time. In this case these low priority risks can be grouped together into one generic risk.

4.5 Also, it is useful to revisit the preliminary assessment of likelihood and impact, because clients can, and often do, change their assessment of these. Hence, our recommendation is at least to review the category 4 and 5 risks, even if little time is spent quantifying them.

Theoretical background to quantifying risks

4.6 In this section we consider the following questions, and describe our approach to each in some detail:

- which distribution is appropriate? In other words, how do we describe the uncertainty around the value of a variable;
- a discussion of expert opinion versus historical data about probabilities;
- what is correlation? Missing correlation out often results in significantly underestimating the risks. It is important to understand and be able to use this concept; and
- how should the risks be separated out? Separating risks out is one of the best ways of improving the accuracy of the analysis.

Which distribution is appropriate?

4.7 A probability distribution describes the probability that a variable will have a given value or occur within a given range. The following diagram is a hypothetical probability distribution; it shows the probability of any particular cost occurring.

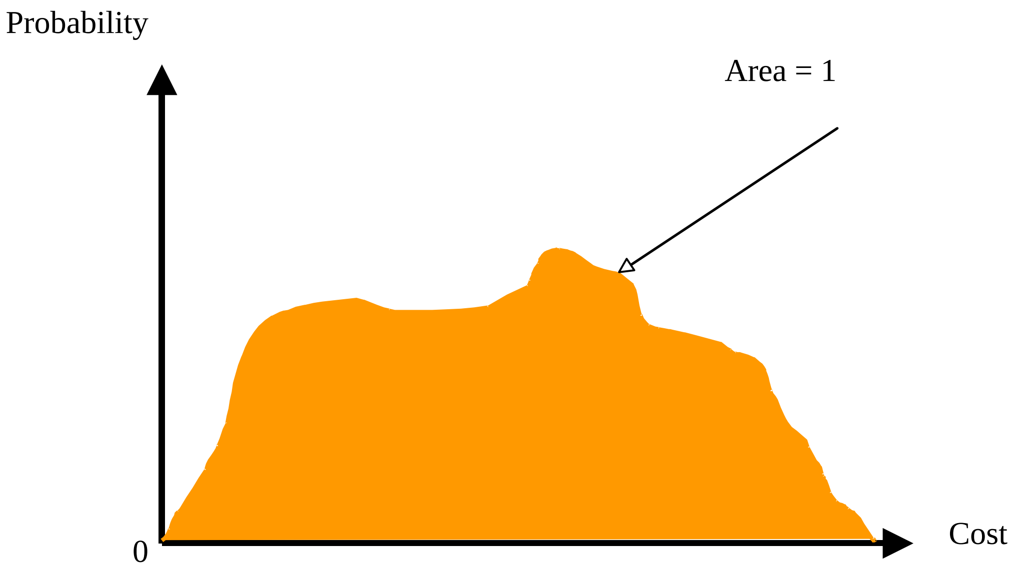

Figure 3: Probability distribution

4.8 The fact that the area of the graph is equal to one means that it is certain that the cost will fall within the range of costs shown on the graph.

4.9 There are three ways of classifying probabilistic risk distributions:

- **continuous/discrete:** smooth profiles, in which any value within the limits can occur, are described as continuous, whereas if the variable can only represent discrete items, for example the number of warehouses in use, a discrete distribution is more appropriate;

- **bounded/unbounded:** unbounded distributions extend to minus infinity and/or plus infinity, for example a normal distribution. Although this can appear ridiculous, the actual probability of a value lying a large distance from the mean may be vanishingly small; and

- **parametric/non-parametric:** a parametric distribution is one that has been theoretically derived, for example an exponential distribution, after making assumptions about the nature of the process that is being modelled. Non-parametric distributions are those that have been artificially created, for example triangular distributions.

4.10 Figure 4 gives examples of a number of different distributions and explains when it is appropriate to use them.

Name *defining characteristics*	**Example distributions**	**Use**	**Situations where suitable**	**Examples**
Triangular *Minimum, most probable, maximum*	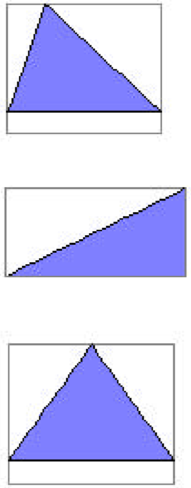	This is the most commonly used distribution. It has no theoretical justification; however, it is a very simple and clear distribution to use. Note that it overestimates the size of the tails at the expense of values close to the mean.	Where the distribution is not known, and it is thought not suitable for a normal distribution, either because it is bounded or because it is not symmetrical. Situations where a simple intuitive understanding is paramount and flexibility is a great advantage.	The operational maintenance costs of a project have been estimated as being a minimum of £40k, with the most probable £60k and a maximum of £100k. The actual cost could be modelled as a triangular distribution. Note that the mean, or expected value, of a triangular distribution is not the most probable value, but is in fact given by: $\text{mean} = \dfrac{\text{minimum} + \text{most probable} + \text{maximum}}{3}$
Normal / Gaussian *Mean, standard deviation*	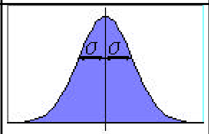	Another frequently used distribution. This is in part due to the result of the central limit theorem which states that the mean of a set of values drawn *independently* from the same distribution will be normally described. Many distributions tend towards normal at their limits (e.g. Poisson and binomial).	Many natural variables fall into a normal distribution, such as human heights (male or female), elephant weight etc. Distribution of errors A situation where the distribution is not known, but it is known to be symmetrical around a mean value, and more likely to be near the centre than the extremes.	The retail price inflation has been assumed to be 3% per annum. However there is a chance that it could be above or below this rate. The mean here is 3%, and the standard deviation (σ) should be estimated bearing in mind that the probability that a value falls within: +/- 1σ of the mean = 68% probable +/ 2σ of the mean = 95% probable +/- 3σ of the mean = 99.7% probable

Name *defining characteristics*	**Example distributions**	**Use**	**Situations where suitable**	**Examples**
Uniform distribution *Minimum, maximum*		Used if the variable is bounded by a known maximum and minimum, and all values in between occur with equal likelihood.	Like the other non-parametric distributions, this has the advantage of being intuitively obvious, and highlights the risk as one where there is very little information about its distribution.	The position of a leak along a pipeline, or the price at any given point in time of a highly market sensitive commodity such as petrol.
Binomial *Number of trials, probability for each trial*	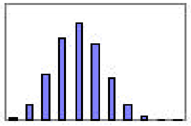	For each trial there are only two outcomes (i.e. pass/fail, heads/tails) The trials are independent: what happens in one trial does not affect the subsequent trials The probability remains the same from trial to trial	This should be used if you require the number of events that will occur given a certain number of trials and a known probability of occurrence.	You want to describe the total number of defective items in a sample of 100 manufactured items, given that the probability of any one item being defective is 7%. The number of defective items will be given by a binomial distribution with n=100, p=0.07.
Poisson distribution *Rate of occurrence*	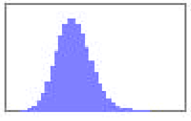	The rate of occurrences remains constant The number of occurrences is not limited The occurrences are independent	This discrete distribution describes the number of events that will occur in a given unit of time, given that the rate is known.	If there is a performance measurement system that deducts payment every time a failure occurs, and it is assumed that the rate of occurrence will be 20 times a year: the number of such events that occur in a given quarter will be described by a Poisson distribution with a rate of 20/4 = 5/quarter.

Name *defining characteristics*	**Example distributions**	**Use**	**Situations where suitable**	**Examples**
Exponential *Rate of occurrence*	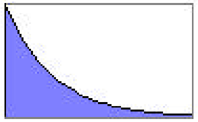	Describes the amount of time between occurrences The rate of occurrence is independent of previous occurrences	Only used for describing the time between (or until) occurrences.	If destructive tests show that a light bulb lasts on average 5200 hours, how long a given light bulb will last will be described by an exponential distribution if we further assume that the rate of failure is constant (i.e. the chances of it failing are the same throughout its life).
Log normal *Mean, standard deviation*	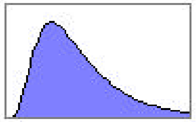	This distribution is also used reasonably frequently. The central limit theorem states that if a quantity is the *product* of two or more independently chosen variables, the distribution will tend to log normal.	Naturally occurring variables that are themselves the product of a number of naturally occurring variables. Any variable that extends from zero to +infinity and is positively skewed. Useful for representing quantities that vary over several orders of magnitude.	The volume of gas in a naturally occurring gas reservoir is often log normally distributed, being a product of its volume, pressure, gas/liquid ratio etc.
Beta distribution $\beta(\alpha_1,\alpha_2)$ *Number of trials, n, number of positive events, r.*	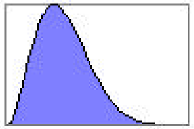	Used to determine the probability of an event given a number of trials *n* have been made with a number of recorded successes *r*. This distribution is primarily used to extrapolate the data taken from a sample to the whole population.	If you only have a limited set of data and have to generate a probability distribution from them. Note that this gives a distribution of the probability of an event or series of events, rather than how many events will occur.	If in 100 (*n*) firings of a gun, it mis-fired 16 (*r*) times, what is the probability that it will misfire? Use Beta(17, 85). This also works for estimating cases where there have been no misfires (i.e. *r* = 0) provided there is some chance of failure.

Name *defining characteristics*	Example distributions	Use	Situations where suitable	Examples
$\sigma_1 = r + 1$ $\sigma_2 = n - r + 1$	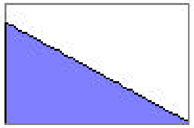 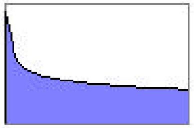			

Figure 4: Examples of different distributions

4.11 In addition to the above distributions, there are others that risk analysts should be aware of. These are very briefly described below, with the situation in which they might be used:

Type	Where used
Gamma	An extension of the exponential distribution: if there is a process with a given rate, the time until *n* events occur.
Hyper-geometric	When you need to know the number of items that there are in a sample of size *n*, out of a total population of *N*, given that you know there are *R* such samples within the total population. For example how many red balls will you take out of a bag that contains 34 red balls, and 57 green, given that you only remove 10.
Log triangular	Similar to the log normal but with its basis a triangular rather than a normal distribution.
Negative binomial	Returns the number of failures there will be before the *n*'th success, assuming the probability of success remains constant.

Figure 5: Additional risks

A guide to when to use which distribution

4.12 The following is a rough guide as to when to use which distribution:

- is it a discrete variable: if yes, try Poisson, or Binomial (if modelling an ‘either/or’ variable);
- is it a symmetric and unbounded variable: if yes, use the normal;
- is it a symmetric and bounded variable: if yes, use the triangular;
- is it asymmetric and bounded variable: if yes, use the triangular; and
- if you have no idea about the shape, use a uniform.

4.13 In practice, the most commonly used distributions are the normal, triangular and uniform.

Historical data versus expert opinion

4.14 If historical empirical data are available for a variable, these can be analysed to determine the correct distribution to represent the uncertainty in the variable. Essentially there are two approaches to using historical data:

- fitting an empirical distribution: in which the histogram of the empirical data is itself used as the probability distribution; and

- fitting a theoretical distribution: a distribution, such as normal, is used to represent the data. The parameters that describe the distribution (for example, the mean and standard deviation for a normal) must then be determined from the data. There are various sophisticated statistical techniques for doing this, which are beyond the scope of this document.

4.15 Historical data about a variable are very useful, and using them would seem to provide a more accurate assessment of the uncertainty than asking for expert opinion. However, caution should be exercised when doing this: the implicit assumption made when using historical data is that the future will continue in the same way that the past has. This may or may not be the case, but it is an issue that should be explicitly addressed and agreed with the sponsor of the project.

4.16 As an example of the dangers of relying on historical data, a computer leasing company was being considered as a potential acquisition by a large conglomerate. The computer leasing company had included in their leases a break option to terminate the lease early. The assumption in the initial analysis of the value of the company was that the number of contracts where the break option was taken would not change materially from the historical level. Hence, there was very little risk associated with the options. Unfortunately, what had been overlooked was that the previous break options had been designed such that it was prohibitively expensive to take the break. When a new option was introduced that made it realistic for customers to take the break, they did in very large numbers, and the value of the operation was reduced to nothing.

4.17 In contrast to the use of data to determine the distribution, we often rely on expert opinion to describe the uncertainty. In these cases, it is normal to use non-parametric distributions: although they are rarely theoretically justified, their simplicity and immediate intuitive nature, together with their flexibility, often make them the most appropriate choice.

What is correlation?

4.18 Some risks are mutually independent: the occurrence of either is independent of the occurrence of the other. Others are correlated: that is, the state of one variable gives us information about the likely occurrence of another. A frequent error in uncertainty and risk analysis is to ignore the correlation between risks. This results in an under-estimation of the overall level of risk. It can also result in scenarios arising that could not in practice occur. For example if a distribution on base rates and another on mortgage rates are treated independently, a nonsense scenario with the base rate higher then the mortgage rate could arise.

4.19 Correlation can result from one variable being directly influenced by another, as in the case of mortgage and base rates, or they could both be dependent upon a third un-modelled variable e.g. the weather could affect the amount of heating required in a building as well as transportation costs.

4.20 Correlation is one of the most difficult aspects of the quantification of risk. It is quantified through the correlation coefficient r, which can vary between -1 and +1 depending upon the level of correlation. Three important values of r on the scale –1 to +1 are:

- **r = +1:** this signifies that two variables are perfectly positively correlated: in other words the two variables always move together;
- **r = 0:** this signifies that the two variables are completely independent; and
- **r = -1:** this represents perfect negative correlation, where the two variables always move in opposite directions.

4.21 The exact definition of correlation is complicated, and indeed there are many of them. The two most common are the Pearson's and rank order correlation coefficients. The Pearson's correlation coefficient is a measure of the degree of linearity between two variables, or the amount of scatter if one variable was plotted against another. In other words a correlation coefficient of r = +1 means not only that two variables move together, but also that they move together linearly.

4.22 The disadvantage of Pearson's correlation coefficient is that if the relationship is non-linear, it does not work. If one variable is always the square of another, we would expect there to be a correlation (both variables always move together). This problem is addressed by using rank order correlation. In rank order correlation, the two variables are ranked. It is the Pearson's correlation coefficient of the two rankings that are then compared. Rank order correlation is used by most add-in Monte Carlo simulation software (e.g. Crystal Ball and @Risk).

4.23 The correlation coefficient can be determined by comparing two experimental data sets, for example using Excel's CORREL or PEARSON functions, or Crystal Ball. However some common sense must be used in conjunction with this. We are all aware of spurious correlation's, such as that found between the number of babies born in Sweden in a given year with the number of storks sighted in that year.

4.24 In practice correlation is often estimated based upon opinion, and normally it is only possible to quantify it in terms of 'high, medium or low'. We have arbitrarily, but pragmatically, defined these levels of correlation in terms of the correlation coefficient r; the table below shows these graphically.

Level of correlation	r	Sample plot
None	0.00	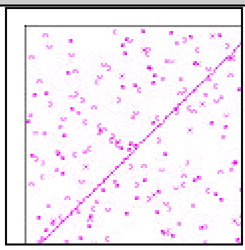

Level of correlation	r	Sample plot
Low	0.71	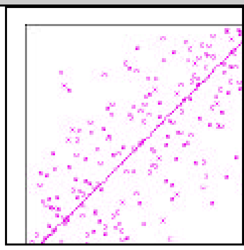
Medium	0.87	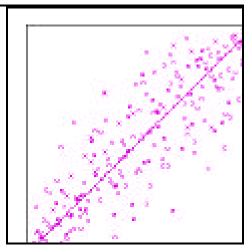
High	0.95	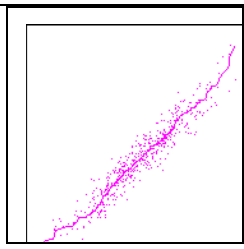
Perfect	1.00	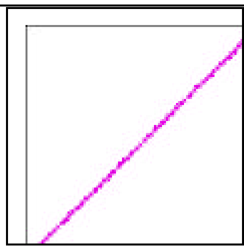

Figure 6: Graphical representation of correlations

4.25 There is an important distinction between correlation and dependency. An example of dependency is where one event can only occur provided another has. This should not be modelled as a 100% positive correlation, as this makes the additional assumption that if the first should occur, the second will *definitely* occur. If possible, dependency should be modelled by using IF statements and/or formulae within the model.

Separating risks out

4.26 A key technique in getting to a more accurate quantification of risk is known as *disaggregation*. This means separating risks out into logical (uncorrelated) components. This has the advantage of making the overall result less dependent upon on the estimate of one critical component. If, after some preliminary risk analysis, the overall level of risk is found to be overwhelmingly dependent on one or two risk elements, where possible these should be separated out into components.

4.27 It is useful to remember when doing this that the total variance, or standard deviation (the dispersion around the mean) squared, of a number of uncorrelated variables is given by the sum of the individual variances, e.g. for two components:

$$\sigma^2_{Total} = \sigma^2_1 + \sigma^2_2$$

4.28 This is useful as a check on the order of magnitude of the combined effect of the disaggregated risks; there is a danger in disaggregation that when the risks are re-combined the size of the total risk can be very different to the original risk that was disaggregated. If this is the case, it is necessary to understand why it has happened, and determine which approach is the most valid representation of the uncertainty.

A review of the issues that arise when quantifying risks with experts

4.29 Given the subjective nature of quantifying risks using expert opinion, it is often better to quantify the risks from empirical data. However, it is frequently the case that there is little, if any, historical data, and expert opinion is the only means available. In this section we discuss some of the problems and issues that arise when experts attempt to quantify uncertainty and risk. We cover the following topics:

- experts' reluctance to provide a distribution;
- factors that can distort experts' opinions; and
- the group dynamics of quantifying risks.

Experts' reluctance to provide a distribution

4.30 Experts are often reluctant to provide probability distributions because they feel that this is more difficult and complicated than the normal approach to estimating, which is to determine a single point value for each input. In fact, providing a distribution for an input instead of one value does not require a greater knowledge of the variable then a single point estimate - quite the reverse. It actually gives the expert a means to express their lack of exact knowledge. This is an important message to convey before any quantification is attempted.

4.31 It is useful to bear in mind when asking experts to quantify uncertainty whether the outcome is truly uncertain, or whether the uncertainty is caused by a lack of information. An obvious example of this is the launch of a new consumer product: undoubtedly there is uncertainty about the level of sales that the product will attain, but this uncertainty might be reduced by carrying out more detailed market research. In a case such as this, it is useful to challenge the experts as to whether they have sufficient information available to them, and discuss the value of gathering more data if this will reduce the uncertainty.

4.32 Some additional suggestions for helping experts through the quantification process are given below:

- reassuring them that the estimation of a probability distribution does not require any great knowledge of probability theory. It is often very helpful to illustrate the most common probability distributions with examples of them, such as those contained earlier in this section; and

- reminding them that there will be an opportunity to revise the estimates at a later stage, particularly if they are found to be significant drivers of the overall risk.

4.33 Considerable reluctance can also been overcome by careful phrasing of the question. For example, if trying to elicit the rates of failure of an average contractor against a service requirement, it makes much more sense to people to be asked "Over the last ten year period, how many failures have you had with your present contractor?" and "How good do you think your contractor is relative to the average?" rather than "What is the rate of failure of an average contractor?".

Factors that can distort experts' opinions

4.34 Investigations into the subjective estimating of risks, in a one-to-one situation, have revealed the following common effects:

Term	Description	Mitigation
Anchoring	Most people start with an initial estimate, and this is then adjusted and used as the basis for evaluating the maximum and minimum points of the distribution. The movement from the initial point tends to be insufficient and the distribution too narrow.	When asking people for a three point estimate (maximum, most probable and minimum), do not start with the most probable, but start instead with both extremes. Ask for justification for those extremes (e.g. what scenario would give rise to this extreme) before moving on to the most probable.
Linguistic imprecision	The quantity under discussion is not well defined, leading to confusion and inaccuracy.	Apply the clairvoyant test: given a description of the event or quantity, could a true clairvoyant state unambiguously whether the event will occur or give the exact numerical value of the quantity? For example the "price of bananas" would not qualify, but the "average price of bananas in the street markets of south east London on the 1/1/99" would.
The tendency to be led	The risk analyst imposes their understanding and structuring of the risks on the expert.	Let the expert explain how they see the logic of the uncertainty in question before going into the quantification.

Figure 7: Subjective estimating of risk

4.35 There are other factors, such as people's tendency to remember significant events more clearly then others, bias, and a systematic tendency in people to underestimate the levels of uncertainty. There are no concrete techniques that will prevent these problems, but a skilled facilitator will constantly challenge participants and question the reasoning underlying their estimates.

The group dynamics of quantifying risk

4.36 There has been considerable analysis done on the psychological factors that give rise to error in the quantification of uncertainty, both in and out of a workshop situation. Some methodologies even go as far as suggesting that experts are presented with sixty page documents on the subject before proceeding. Fortunately there is some agreement on the dangers involved, and those that are relevant to a workshop situation are shown below:

Term	Description
Conformity	When groups attempt to estimate risks, individuals tend to gain unwarranted confidence from each other's estimates and thereby give a narrower range of estimates then when done independently.
Bias	The more senior individual at the session is likely to influence the others merely by their presence. Most people will have a bias in a particular direction, but with a dominant person present their biases will tend to converge.
Personality	The usual issue of the loudest voice dominating, and the quietest not being heard.

Figure 8: Psychological factors

4.37 For these reasons the preferred technique is to carry out small meetings, often on a one-to-one basis with each individual expert. If this process reveals large differences in opinion, the experts are reconvened. This is similar to many respected methodologies (e.g. Delphi technique by the RAND corporation, Stanford/SRI assessment protocol and the Wallsten/EPA protocol).

5 Risk analysis

5.1 This section describes the implementation of risk analysis in a spreadsheet model, and explains how to generate outputs using the Monte Carlo simulation add-in software package Crystal Ball. It has been split up into the following parts:

- introduction to Monte Carlo analysis;
- identifying the significant risks; and
- a practical example of identifying the significant risks (in a large PFI transaction).

Introduction to Monte Carlo analysis

5.2 In this section we consider the following topics:

- what is Monte Carlo simulation?
- the influence of uncertainty and risk analysis on model design;
- impossible scenarios during Monte Carlo simulation;
- setting up distributions in the spreadsheet;
- how many iterations are required?
- Latin Hypercube sampling; and
- what is a random number?

What is Monte Carlo simulation?

5.3 Monte Carlo simulation is a technique that takes the distributions that have been specified on the inputs to the model, and uses them to produce a probability distribution of the output of interest. It does this by running through the following sequence of actions as many times as the user specifies:

- for each input, ‘sample’ (i.e. pick a value from) the distribution for that input (see below for more information about this);
- recalculate the model with the sampled inputs; and
- save the output.

5.4 On each iteration, for each input, the Monte Carlo simulator selects a value from the relevant probability distribution at random such that, over a large number of iterations, the distribution of the selected values reflects the input probability distribution. For example, if the input is modelled as a normal distribution of mean 5 and standard deviation 2, the distribution of the actual inputs selected during the simulation will be approximately mean 5 and approximately standard deviation 2. Obviously, the more iterations that are run, the closer the actual input distribution becomes to that set up (we discuss the number of iterations required below).

5.5 The probability distribution of the output from the model shows the likelihood of occurrence of all modelled values of the output; that is, it makes clear the uncertainty that exists in the value of the output. Note that although we have referred to a single output in the above it is possible to have as many outputs as necessary.

5.6 We use a package called Crystal Ball to carry out Monte Carlo simulation. This is an easy to use Excel add-in that allows the user to specify distributions on the inputs in the spreadsheet. It does not require any alterations to the spreadsheet itself, so models can be shared with others who do not have Crystal Ball.

The influence of uncertainty and risk analysis on model design

5.7 It is important to consider how the risks being modelled will affect the structure of the model. For example, a model may include an input for ‘total cost’; however, it may be that the risks around the various elements of total cost (e.g. staff, materials, etc.) are quite different, and therefore the components of cost would need to be held separately in the model. This is a matter of judgement, but in large risk projects it may be necessary to construct simple prototypes with different levels of detail in order to judge the impact of the model structure on the outputs.

5.8 In the interests of clarity it is good practice to separate the base case input values from the probability distributions that are superimposed upon those inputs. For example, if there is a base case assumption that inflation is 3%, with a normal distribution of standard deviation 1% superimposed upon it, this should be separated into:

- an input cell, for the base case value of inflation: this contains 3%; and
- an additive factor, for the uncertainty around the base inflation: this consists of a normal distribution, mean 0% and standard deviation 1%.

5.9 This approach makes it possible to change the distributions and the base case inputs independently.

5.10 It may be that the quantification of risks has led to two different probability distributions being applicable for one input. For example the uncertainty around sales of a new product may be very different depending on whether a competitor decides to enter the market. In a case such as this, the uncertainty around sales should be modelled by means of a flag that the model uses to switch between different distributions. Thus, if there is a 25% chance that a competitor enters the market, the flag should be set up with a custom distribution that has a 75% chance of being zero and a 25% chance of being one. The model logic should then use the flag to choose the appropriate sales forecast, namely ‘no competitor entry’ if the flag is 0, and ‘competitor entry’ if the flag is 1. Each sales forecast would then have a suitable distribution applied to it.

Impossible scenarios during Monte Carlo simulations

5.11 A pitfall to beware of in Monte Carlo modelling is the danger of impossible scenarios: that is, during the simulation sets of inputs occur that are mutually exclusive. For example, if the project under analysis is a electricity generating plant, there may be two risks to be modelled:

- a delay in the commencement of operation due to construction over-runs; and
- variability in the amount of energy supplied.

5.12 Obviously, if the plant is not constructed on time, the amount of energy available to be supplied will be less than expected in the first year of operation. Unless the relationship between date of operation and available energy is explicitly contained in the model, it is possible that one of the Monte Carlo simulations will choose a construction delay and a high level of energy supplied in year one. This is clearly nonsense, and thus it is important that the model prevents such ridiculous scenarios occurring.

Setting up distributions in the spreadsheet

5.13 The risk analysis package that we use, Crystal Ball, enables the user to type in the parameters of a distribution; it then applies the distribution to the underlying input cell. However, Crystal Ball itself gives no visible indication in the spreadsheet as to the parameters used, and we therefore recommend that the following approach is adopted to make clear the values that are in use.

5.14 The parameters of the distribution should be entered into a table next to the risk input cell, as shown in Figure 9. For example, the uncertainty around Staff costs has been modelled using a triangular distribution with a minimum of 92%, a most probable of 100% and a maximum of 120%. The risk distribution itself is set up in Crystal Ball on cell E5, using references to cells H5, I5 and J5. Once the risk is set up in Crystal Ball, changing the data in these cells changes the parameters of the distribution that Crystal Ball uses. Note that although the name of the distribution appears in the spreadsheet, this is not actually linked in to Crystal Ball; that is, changing the name in the spreadsheet does not automatically alter the distribution.

	A	B	C	D	E	F	G	H	I	J	K
1	Sensitivities										
2											
3		Operating costs									
4							Type	Minimum	Most prob.	Maximum	Mean
5		Staff costs		%	100%		Triangular	92%	100%	120%	104.0%
6		Maintenance		%	100%		Triangular	67%	85%	140%	97.3%
7		Transport		%	100%		Triangular	98%	100%	103%	100.3%
8											

Figure 9: Spreadsheet abstract

5.15 It is useful to control the application of the risks to the inputs by means of flags that can be set or cleared in a control section of the model. This makes it easy for the user to run the model with different risks switched on and off.

How many iterations are required?

5.16 There is no truly theoretically correct method of estimating the number of iterations required to produce reliable results from a Monte Carlo simulation. In practice, the more the better, and we would normally recommend at least 1,000 iterations.

5.17 An alternative approach is to set Crystal Ball to display the statistics of the output, such as mean and standard deviation, while the simulation is running. As the number of iterations increases the statistics will settle down to a stable value; the simulation can then be stopped at this point. However, it should be noted that this method is not particularly scientific, and may give mis-leading results in some circumstances.

5.18 If it is important to know the level of certainty around the output statistics, it is possible to use various statistical techniques to estimate this. However, these lie outside the scope of this document.

5.19 The need to do thousands of iterations during a Monte Carlo simulation means that run time can become a serious issue. A useful rule of thumb is that it takes Crystal Ball about one second to process 800 different input probability distributions on a 133MHz Pentium PC (this number will vary with the number of correlation’s). Adding in the run-time of the model itself will give the time required for each iteration. For example a model with 400 inputs, and a run-time of approximately one second will take about 1.5s/iteration. Doing one thousand iterations would therefore take about 1500s or 25min. It is usually the run-time of the model that will be the more significant factor rather than the number of probability distributions.

Latin Hypercube sampling

5.20 The conventional sampling approach tends to select values from the areas in the distribution with a higher probability of occurrence. Therefore, many iterations will be required in order to ensure that sufficient samples have been taken to represent extreme, or low probability, values. If a small number of iterations is performed, the output may lead to a ‘clustering’ of values around the high probability values. This problem can be overcome with Latin Hypercube sampling.

5.21 The Latin Hypercube approach involves dividing up the distribution into a number of strata, such that each stratum has the same number of values fall in it as the others. To pick a value from the distribution, the simulator chooses a stratum at random, and then from within that stratum a random value is chosen. All the other strata are sampled from before the original stratum can be sampled again. Thus, Latin Hypercube sampling ensures that the whole distribution is covered much more quickly than in Monte Carlo sampling. Note that this technique works regardless of the shape of the distribution. Crystal Ball allows the user to choose between Monte Carlo or Latin Hypercube sampling, and we recommend that Latin Hypercube is always selected.

What is a random number?

5.22 The selection of values from the input distributions depends on Crystal Ball generating ‘random’ numbers. These numbers are random in the sense that, over a large number of them, there is no discernible pattern or bias in them. However, they are not random in the sense that we cannot predict what will come next. In fact, if we tell the random number generator to use the same starting point, known as the ‘seed’, the sequence of numbers it generates will be identical every time.

5.23 Crystal Ball includes a facility that allows the user to set the seed, and thereby ensure that the same random numbers will be used for every run. This means that if a run must be repeated at any stage later, it will produce exactly the same results (assuming none of the inputs or distributions have been altered). The corollary of this is that if the input distributions have been changed between runs, any differences in the outputs are due to these changes, and not to the use of different random numbers. If the user does not select a seed, Crystal Ball will use the PC's clock to initiate the sequence of random numbers, and hence the numbers will be different on each run. Best practice is to use the same seed value (for example '1') for all runs.

Identifying the significant risks

5.24 Identifying the relative significance of the risks relating to each input is an important step in risk analysis: once the significant risks have been identified more time can be spent on these, at the expense of the less important ones. It may be necessary to revisit the assumptions around the significant risks, and to do more research. Also, if a small number of risks dominate all others they should be disaggregated into smaller components.

5.25 There are three methods of determining the extent to which the risk on a particular input affects the output, although none of them are completely satisfactory. Each is described in turn below, and we also include a summary that sets out when to use which method. They are:

- Crystal Ball sensitivity function;
- simple sensitivity analysis; and
- enhanced sensitivity analysis.

Crystal Ball sensitivity function

5.26 The first technique is to make use of the sensitivity analysis option in Crystal Ball: this works by determining the 'rank order' correlation between all the inputs and the output. We do not describe here the theory behind rank order correlation: it is sufficient to know that the rank order correlation is a measure of the importance of the input to both the level and variance of the output. The advantage of this technique is that only one model run is required to evaluate all the sensitivities, and that all the inputs are varying simultaneously so that the relationships between the different inputs are being taken into account. Unfortunately, it has important drawbacks:

- this method is of limited theoretical validity. If there are any inputs that are not monotonic (i.e. an increase in the input **always** produces either an increase or decrease in the output) then the sensitivity calculation will not work correctly. It will also be inaccurate if there are many correlated assumptions; and

- it may be that some of the inputs have probability distributions that are applied in each year, e.g. the level of inflation. In these cases, the correlation of each year's individual inflation with the output will be very small, and hence Crystal Ball will report that they are insignificant. However, the aggregate effect of inflation over many years within the model may be significant.

5.27 It is possible to overcome problem (b). If the effect of inflation is defined, for example, by its mean level over the contract period then its effect on the output can be calculated outside Crystal Ball. By defining it as an output, Crystal Ball can export every iteration of both the mean inflation and output. Excel can then be used to calculate the rank order correlation, which can then be compared with the rank order correlation's produced within Crystal Ball. One of the advantages of rank order correlation is that there is no necessity for a linear relationship between the two quantities in question, and hence the above technique is valid.

5.28 At lower levels of correlation, say in the -0.25 to 0.25 range, many of the correlation's may be spurious. A good way to determine the cut-off point below which all correlation can be regarded as spurious is to look for correlation's that should be positive, but are in fact negative, and vice versa. For example, if a variable can only increase costs, but it is shown with a positive correlation with the NPV, then this must be spurious. An increase in the costs can only produce a decrease in the NPV, and hence this should be a negative correlation.

Simple sensitivity analysis

5.29 The second method, simple sensitivity analysis, is to keep all inputs constant at their expected values apart from the one of interest which is manually varied over its probability distribution (say from the 2.5^{th} percentile to the 97.5^{th} percentile). Ranking the impacts produces a list showing how the output variable varies with each of the inputs. The disadvantage of this approach is that it ignores non-linear effects that may occur due to combinations of inputs, and it also ignores correlation in the inputs, and is therefore not recommended except with very simple models with no correlation's and no non-linearity's.

Enhanced sensitivity analysis

5.30 The last method, a more sophisticated variant on the above approach, is enhanced sensitivity analysis. Using Crystal Ball, simulations are run keeping all variables constant except the one of interest. The user must repeat this analysis for each risk or category of risks that are thought to be important - Crystal Ball will not do this automatically. The change in the mean from the expected value and the square of the standard deviation in the output gives an indication of the importance of this input. By ranking the difference in the base case output value and the mean value, allowing for variation in the input, the relative importance of the different inputs on the output level is also ranked. Similarly, by ranking the squares of the standard deviations the relative importance of the inputs to the overall variance (or level of risk) of the output.

5.31 The above method has the advantage of being simple and clear. However it does have three important drawbacks:

- by assuming that all the other input variables stay constant, we are ignoring any effect that they could have on the impact of the particular input on the output. For example, it could be that the importance of rent inflation is very sensitive to the amount of space that is sub-let. By keeping the sub-let area constant, you may over-estimate the impact of rent inflation;

- if there is any correlation between the inputs that are being held constant and those that are not, this method will not pick this up; and

- the model has to be run for the required number of iterations for every input that you require the relative significance of. This can mean a large amount of model run time, particularly as this process will have to be repeated whenever the risk profile changes.

5.32 An improvement over the above approach is to run the Monte Carlo simulation with all the variables free to vary *except* the one of interest. The reduction in the square of the standard deviation due to keeping that variable constant can then be determined for each risk. This information can then be used to order the risks – with those causing the biggest reduction in the square of standard deviation ranked as the most significant. This approach has the advantage of including most correlation's and relationships between the different inputs.

Summary of when to use which

5.33 The table below shows when the different methods should be used:

Situation	Suggested method
A simple model that does not have any risk inputs that apply independently over a number of years (for example inflation)	Crystal Ball sensitivity function

Situation	Suggested method
A model with different risk values in different years, but which is very simple, involving no non-linearity’s in the logic, and no correlation	Simple sensitivity analysis
A model with different risk values in different years, which is also complex (including either correlation’s or non-linearities)	Enhanced sensitivity analysis

Figure 10: Summary of risk methodologies

A practical example of the identification of the significant risks

5.34 The following describes a situation where we used sensitivity analysis to successfully identify the key risks in a very large bid to the UK Government as part of the private finance initiative (PFI). The information was then used to price the bid competitively, and in fact our client was selected on this basis.

5.35 The transaction involved a very complex payment mechanism for the provision of services to Government. We carried out the following process to establish where the risks were, and how to price them:

- the main cost items were identified, and efforts made to reduce them as far as possible, for example by using cheaper contractors and identifying real savings that might be inherent in the proposed deal. This was done without regard as to the levels of extra risk that this might incur;

- a risk analysis was carried out, focussing on identifying how the payment mechanism would work in practice. This enabled us to determine the sensitivity of the financial reward available in the transaction to each of the services contained within the contract; and

- those services that had been identified as critical to success under the contract were re-visited, and extra resources added back in to them to ensure that their performance would be adequate.

5.36 Thus, by using risk analysis to demonstrate the key sensitivities inherent within the project, we were able to assist the client in placing resources in the most appropriate areas. This ensured that the bid was priced as keenly as possible.

6 Presenting the results

6.1 The risk analysis model is of no value unless its result can be communicated. This is necessary not only in presenting the final results, but also in presenting the interim results used to more accurately quantify the significant risks. This section is divided into the following sub-sections:

- graphical presentation;
- statistical measures; and
- presenting it in English.

Graphical presentation

6.2 In this section we describe the following types of graphs:

- histogram;
- cumulative frequency chart;
- tornado chart; and
- scatter diagram.

Histogram

6.3 The histogram is the most commonly used chart in risk analysis. Crystal Ball automatically generates these charts, an example of which is shown below:

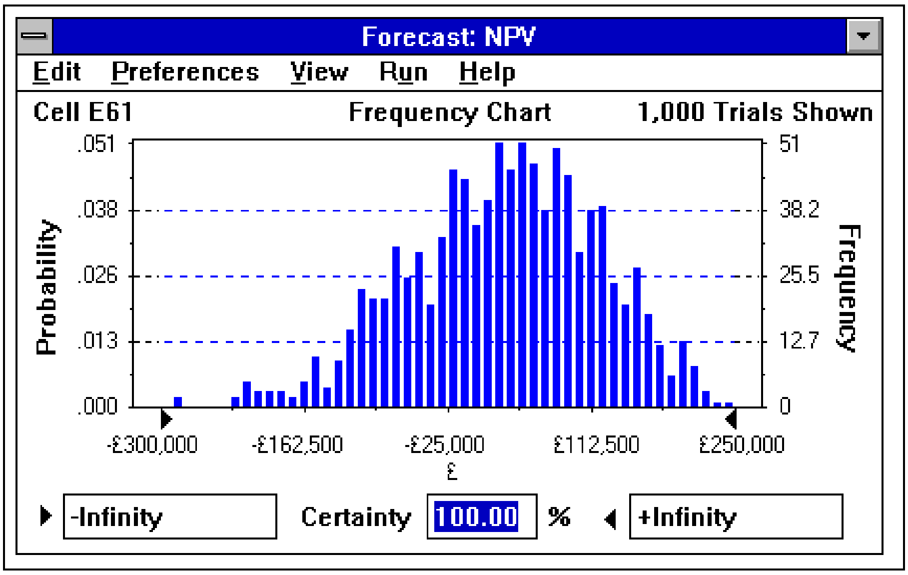

Figure 11: An example of Crystal Ball analysis

6.4 It is important to bear in mind the number of bars in a histogram: too many and the level of random noise dominates, too few and the detail is missed out.

Cumulative frequency chart

6.5 The histogram is very useful for illustrating the degree of uncertainty associated with a variable; however it is not so good for determining quantitative information. The cumulative frequency chart is far better for this purpose: it is much easier to read off a cumulative frequency chart the probability of a value lying within a given range, or above or below a given value.

6.6 Shown below is a typical (ascending) cumulative frequency plot:

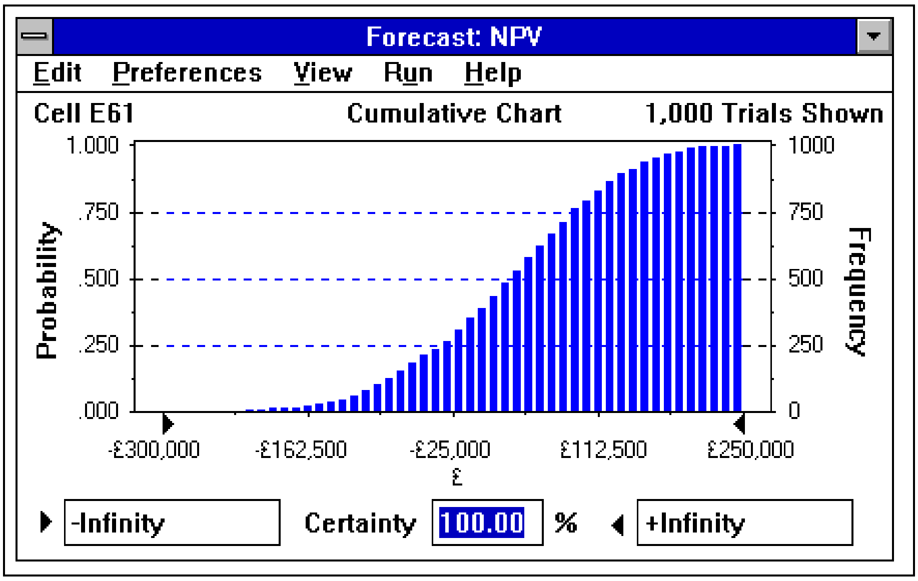

Figure 12: An example of a cumulative frequency plot

6.7 Points along the y axis are known as percentiles, e.g. the '0.25' point is referred to as the 25th percentile. This would be quoted as 'the 25th percentile is -£30,000': this means that 25% of the values are -£30,000 or less.

6.8 This type of plot can be used in determining a measure of the spread or width of the distribution. For example, spread can be quoted as the range of values over which 95% of iterations will lie (for a normal distribution, this is +/- two standard deviations). To determine the spread, we subtract the value at the 97.5th percentile from that at the 2.5th percentile.

Tornado chart

6.9 A Tornado chart is a pictorial representation of a sensitivity analysis of the model, and is very useful for communicating the relative significance of different risks. The sensitivity values for each variable can be calculated as described in section 5, and Excel used to plot them as horizontal bars ranked in descending order, as shown in Figure 13:

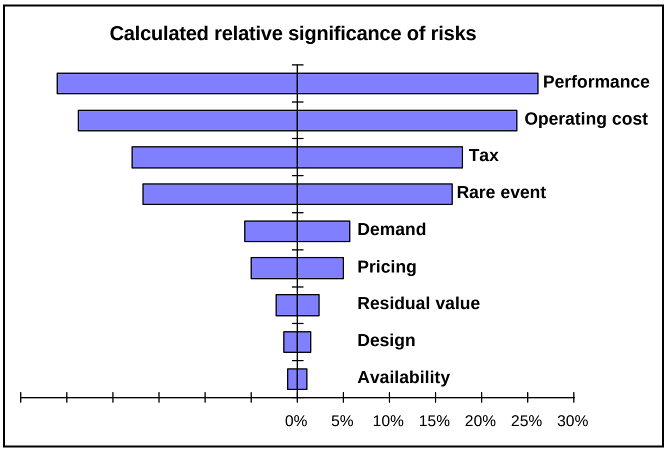

Figure 13: An example of a Tornado chart

6.10 Crystal Ball cannot produce true tornado plots; it does, however, generate a graph that has bars going out from the centre to the right if there is a positive correlation, and to the left if the correlation is negative.

Scatter plots

6.11 Scatter plots are very useful for seeing patterns in the relationships between the inputs and outputs, or indeed between two outputs (for example cost of project and time until completion). The custom is to put the independent variable in the x-axis and the dependent one on the y-axis. Shown below is a typical scatter plot (produced in Excel – Crystal Ball can not generate scatter plots):

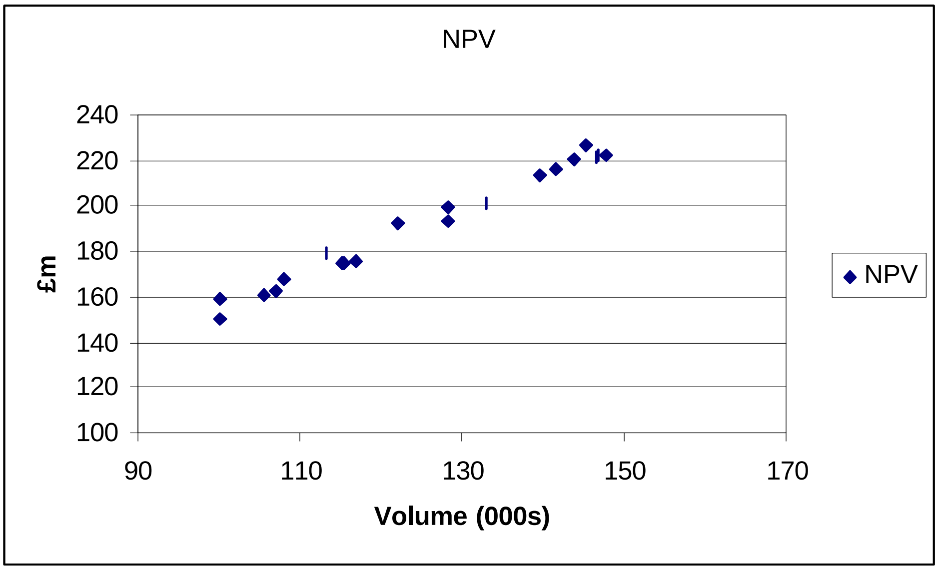

Figure 14: An example of a typical scatter plot

6.12 This shows the relationship between the NPV of a hypothetical project and the volume of products sold. It makes clear the strong relationship between NPV and volume sold.

Statistical measures

6.13 There are many statistics that can be calculated based upon a distribution; however, most are esoteric and are unlikely to contribute much to business people's understanding of risk. The table below lists the most common statistics, and explains when they might be useful:

Statistic	Definition	Use	Dangers
Mean (expected value)	The average of all the generated outputs	Very useful, this is one of the two most common statistics reported. For example the average NPV of a transaction. It also has the useful property that if two (or more) variables are independent, then: mean(a+b)=mean(a)+mean(b), & mean(a*b)=mean(a)*mean(b).	Confusing the mean with the most probable (mode)
Standard deviation (σ)	The square root of the variance (see later under variance)	Another very useful statistic, it gives a measure to the dispersion around the mean of a distribution. It is frequently used in conjunction with normal distributions to give the level of certainty that a value lies within a certain distance from the mean: +/- σ of the mean = 68% +/- 2σ of the mean = 95% +/- 3σ of the mean = 99.7% So, for example, a normally distributed variable with a mean of 1.0 and a σ = 0.05 can be said to have a 95% certainty of lying between 1.1 and 0.9.	Assuming that the standard deviation of the sum of independent components is the sum of the separate standard deviations! In fact it is the square root of the sum of the squares: $\sigma^2_{TOT}=\sigma^2_1 + \sigma^2_2$ The relationship given above is only valid if the distribution is symmetrical. It becomes more of an approximation the more skewed the distributions are.
Variance (V)	The variance is calculated by determining the mean of a set of values, and then summing the square of the difference between the value and the mean for each value: $V = \dfrac{\sum_{i=1}^{n} (x_i\text{-mean}(x))^2}{(n\text{-}1)}$	This is also a measure of the dispersion around the mean, however it is in the units of a quantity squared. Thus the variance of a distribution in NPV (in £s) will be given in £^{2}. The reason it is used is that it is useful for estimating the widths of a sum or multiple of several independent variables: V(a + b) = V(a) + V(b), & V(a * b) = V(a) * V(b). The risks in a model have been successfully disaggregated if the variances of the different significant risks are similar.	As with the standard deviation, the relationships shown to the left are only valid if the distribution is symmetrical. It should be noted that the variance (and thus the standard deviation) is much more sensitive to the values at the tails of the distribution than those close to the mean.
Median	The median is the value at which there is an equal percentage chance of a being above it as below it. In other words, it is the 50th percentile.	Rarely used as it gives no indication as to the range of the values above it or below it. If the mean is not equal to the medium, then the distribution is skewed.	Confusing the median with the mean or mode.
Percentiles	The nth percentile of a variable is that value for which there is a n% chance of the variable lying at or below that value.	A useful concept, used in measuring the range of a variable. For example the range of a distribution might be defined as the difference between the 5th and 95th percentile; this means the width of a distribution if the top 5% and bottom 5% of all values are ignored. It can also be used to answer questions like "What are the chances that the IRR is below 9%?". The answer would be the percentile for which the value was 9%.	Can be confusing as to exactly what is meant, so it is usually a good idea to explain the concept of percentiles in layman's terms.

Statistic	Definition	Use	Dangers
Mode	The most likely value. For a discrete distribution this is the value with the greatest observed frequency, and for a continuous distribution the point of maximum probability.	Sometimes used to describe a Poisson like distribution: the mode is the most likely number of event to occur in the given time period (and is approximately given by the reciprocal of the rate). Also used in describing triangular distributions (the minimum, the mode and the maximum). In general it has little value in uncertainty and risk analysis.	It is difficult to determine precisely, particularly is a distribution is unusually shaped.
Skewness (*S*)	$S = \dfrac{\sum_{i=1}^{n} (x_i\text{-mean}(x))^3}{\sigma^3}$	This is a measure of the 'lopsidedness' of a distribution. It is positive if a distribution has a longer right tail (and negative if a more prominent left tail). A zero skewness means the distribution is symmetric. Apart from a general measure it is used to determine how 'normal' a distribution, the closer a distribution is to having a skewness of zero, the more normal it is. Examples of skewness: the skewness of normal distribution is 0, triangular distributions vary between 0 and 0.56, and an exponential distribution has a skewness of 2.	The skewness is even more sensitive to the points in the tail of the distribution than the variance. It therefore requires many iterations to be run before it reaches a stable value.
Kurtosis (*K*)	$K = \dfrac{\sum_{i=1}^{n} (x_i\text{-mean}(x))^4}{\sigma^4}$	The kurtosis is a measure of the 'peakedness' of a distribution. Examples of kurtosis: uniform distribution has a kurtosis of 1.8, a triangular distribution 2.4, a normal 3, and an exponential has a kurtosis of 9. If a distribution is approximately bell shaped, and has a skewness of around 0 together with a kurtosis of close to 3, then it can be considered normal.	Stable values of the kurtosis often require even more iterations to be run then skewness. For example a randomly sampled normal distribution required approximately 1500 iterations to be within 2% of 3.
Coefficient of variability (normalised standard deviation) (σ_n)	This is defined as the standard deviation divided by the mean: $\sigma_n = \sigma$ / mean	This is a dimensionless quantity that allows you to compare, for example, the large standard deviation of a large variable with the small standard deviation of a small variable. An example would be comparing the level of fluctuation with time between different currencies.	This is not a meaningful statistic to compare if the mean and standard deviation are unlikely to bear any relation with each other. An example would be the NPV of a project. Here the spread need not be related to the mean value, which could be close to zero. An extreme would be the coefficient of variability of a normal distribution that is centred on zero.
Mean standard error.	This is included purely as it is one of the statistics provided by Crystal Ball.	Crystal Ball calculates this as a measure of the accuracy of the simulation, and whether enough iterations have been run. Specifically it tells you the likely difference between the estimated mean and the actual mean, to a certainty level of 68%.	This is precise **only** for the accuracy of the mean. The accuracy's of the other statistics such as the standard deviation or any percentile value are likely to be considerably less then this figure implies. It should be used only as indicative in rough and ready simulations. For more detail on this subject, see the sub-section on 'how many iterations are necessary' under risk analysis.

Figure15: Table of common statistics

6.14 Two of the statistics described above – skewness and kurtosis – appear at first sight to have little use in communicating the results of a risk analysis. While this is generally true, by looking at the skewness and kurtosis figures for an output distribution the analyst can determine how ‘normal’ the output is. If the output is reasonably close to normal, and thus can be considered normal, the interpretation is more straightforward because we can assume the usual normal characteristics; for example, we know that for a normal distribution approximately 95% of the values are within +/- 2 standard deviations of the mean.

6.15 In general it is more helpful to keep the number of statistics quoted in a report down to a minimum (e.g. the mean and the standard deviation/spread between two percentiles), and not quote them to a large number of significant figures.

Presenting it in English

6.16 The analysis of risk is a complex subject, and often it is very valuable to express the treatment of risks in 'laymen's' terms. Figure 16 below is a table that gives examples of how to phrase descriptions of risks and the output results in simple terms:

Mathematical description	Layman's terms
The base case value of x= £100,000 has been multiplied by a factor that is normally distributed, with a mean of 95% and a standard deviation of 2.5%	x has been assumed to be on average £95,000, with a 95% chance of being between £100,000 and £90,000
The number of events in a quarter has been described by a Poisson distribution, with a rate of 0.1events/quarter	It has been assumed that 4 events are likely to occur over ten years.
Variables x and y are correlated with a correlation coefficient of +0.87 (or +0.95 or +0.71)	There is a reasonable (or high or low) chance that x will move together with y
The NPV varies between £92,483 and £766,003, with a mean of £460,311, a skewness of +1.34, a kurtosis of 3.2, and a standard deviation of £96,314	The average NPV is £460k with a 2 in 3 chance of being between £340k and £580k. The shape is approximately normal, slightly skewed to the right, with the maximum £770k and the minimum of £90k **or** use a graph!

Figure 16: Laymen’s’ terms of risk

7 Beyond presentation

Introduction

7.1 In this section we discuss the interpretation and use of the output from a risk analysis project. We cover the following topics:

- valuing risk;
- what does a distribution of NPVs mean? and
- risk analysis in PFI transactions.

Valuing risk

7.2 In attempting to value risk there are two measures that can be considered:

- the expected value of the risk; and
- the range of uncertainty.

Expected value

7.3 The expected value of a variable is the mean of the distribution that describes the risk on that variable. For example, consider the following risk on construction costs:

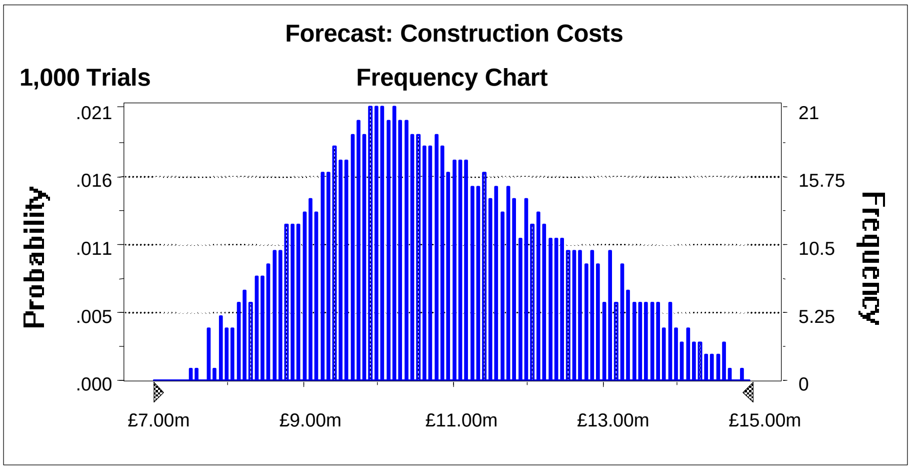

Figure 17: An example of a risk variable

7.4 The expected value of construction cost is £10.8m in this example. If we assume that the most likely construction cost (£10m) is the base case value, in theory we ought to be prepared to pay up to £10.8m - £10m = £0.8m to mitigate, transfer or insure the risk on this cost. In some cases it is sufficient to do this calculation, and no further analysis is required, e.g. in most PFI transactions (see later section).

7.5 However, it may be that just calculating the expected value is not sufficient. This is because the implicit assumption underlying the expected value of the construction costs is that it is the value that the costs would come to if they were averaged across many identical projects. Expected value does not provide any information about how much this risk might cost us in **this** project.

The range of uncertainty

7.6 To return to our construction cost example, it may be that we are willing to spend more than £0.8m to mitigate this risk. In particular, if a cost of £15m would be disastrous for the organisation for some reason, we might be prepared to pay up to £2.5m to transfer the risk. The concept of utility attempts to help value the true cost of financial risks.

7.7 Utility is defined as the subjective value that a certain financial value has. The definition of utility means if two options have the same utility, there would be a genuine indifference as to which option to select. If we can determine the utility function of the organisation, or the decision maker (acting as a proxy for the organisation), we can convert the distribution in financial reward into a distribution in utility. The diagram below shows an example utility function for an organisation:

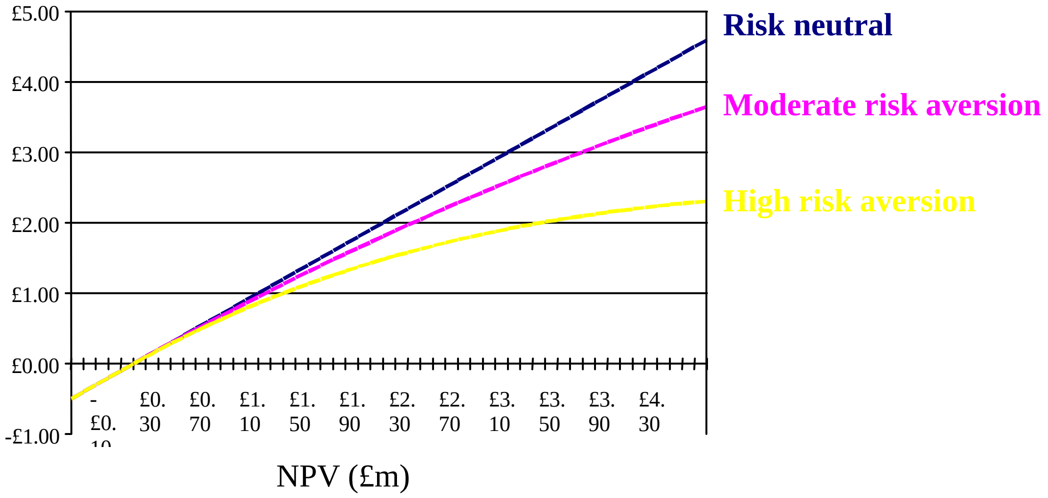

Figure 18: An example utility function

7.8 If we know the utility function for an organisation, the expected value of the utility of two different scenarios can then be compared. The one with the greatest expected utility is the preferable one. Thus, in our construction costs example we could compare the 'do nothing' case with the case where we spend, for example, £2.5 to transfer the risk. Whichever has the greatest utility, we carry out.

7.9 In practice, drawing out a full utility curve from managers in an organisation is very difficult. However, it may be possible to identify some key points on the curve that could be used to calculate utility at those points. At the very least, having a discussion with managers about the consequences of the range of outcomes demonstrated by the uncertainty and risk analysis is useful, and can help to inform decision making in a qualitative sense.

What does a distribution of NPVs mean?

7.10 In practice, it is often the case that we are not interested in the risk around one cost, or even total cost. Usually, it is some other measure of project return, such as net present value (NPV), that is important. It might seem obvious that a distribution of NPVs would provide useful additional information about the risks in a project. However, this is a controversial view: for example, Brealey and Myers, in the ‘Principles of Corporate Finance’, describe these as ‘bastard NPVs’.

7.11 Their view is that NPV analysis should be based on discounting the expected values of the project cashflow by a discount rate that reflects the cost of capital of the project. If this gives a positive NPV, the project should be undertaken. This approach captures all the risks, in the following manner:

- the expected value of the cashflow captures the risk unique to the project; and
- the cost of capital for the project reflects the market risk of the project.

7.12 The argument is that if all the risk has been captured in this way, looking at distributions of NPVs is a pointless activity. While this is true, it ignores an important point: that the calculation of NPV in this regard is adequate from the point of view of investors, but ignores the view-point of individual managers, or groups, or divisions, within an organisation. To these constituencies, it is not particularly helpful to know that, although their project has gone badly, investors with fully diversified portfolios are insulated because another project that they have invested in has done better than expected. This is the critical difference between managers and investors – investors can invest in diversified portfolios, managers work for a single company and are not diversified.

7.13 Thus, we would suggest that it is useful to produce distributions of NPVs, and use them to understand more the nature of the project. In contrast to the basic NPV project appraisal rule (“go ahead if the NPV is bigger than zero”), there are no hard and fast rules for interpreting an NPV distribution. However, common sense would suggest that a project that has a positive expected NPV, and only a 5% chance of a negative NPV, is probably an attractive investment. If, however, a project has a positive expected NPV and a 50% chance of a negative NPV, then this may give cause for concern, an example is given in Figure 19:

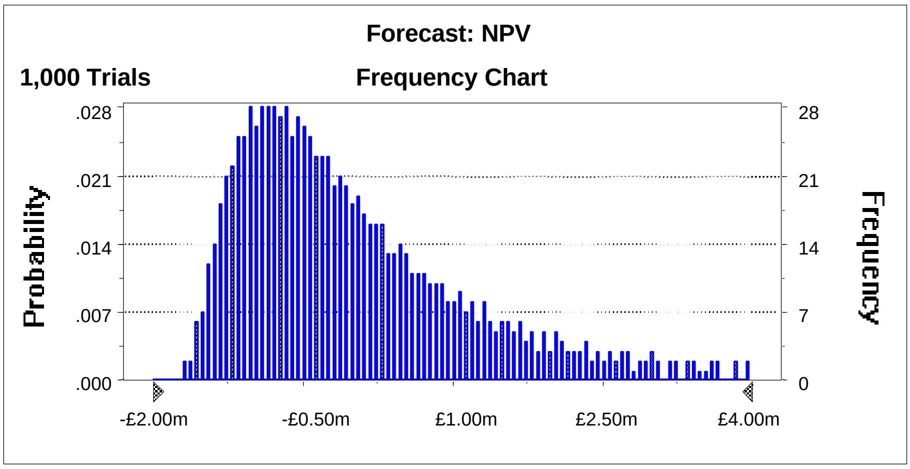

Figure 19: An example of NPV distribution

7.14 This project has an NPV of about £0.2m, but an almost 60% chance of producing a negative NPV. At the very least, the reason for the large bulge of negative NPVs shown here should be investigated. Risk mitigation strategies can be devised, their effects modelled, and the change in the distribution of NPVs reviewed. In general, a narrowing of the range of possible NPVs is to be welcomed.

7.15 There may be some situations where reducing the spread of NPVs is done at the expense of the expected value of the NPV. In cases such as this, managers need to trade off the increase in certainty for the reduction in expected value for the project. The utility theory described earlier in this section can be applied in helping an organisation to make explicit its attitude to uncertainty.

Risk analysis in PFI transactions

7.16 Risk analysis is an integral part of the PFI procurement process, and we outline below the specific requirements for including risk within PFI transactions, from the procurer's (i.e. the public sector's) viewpoint.

7.17 Risk analysis is required at the following stages:

- **Strategic Outline Case**: develops the strategic context. Project risks are assessed at a high level;

- **Outline Business Case**: in which the options for development of the facility by the public sector are set out. Risk analysis is carried on the preferred option to determine its affordability;

- **Full Business Case**: in which the development of the Reference Project for the public sector procurement and a description of the proposed PFI project are set out. The public sector's costs, contained in the public sector comparator (PSC), are risk adjusted in order to compare them with the private sector's bid; and

- **accounting treatment**: determining the correct accounting treatment for the asset within the transaction, and, in particular, whether it is on the Government's balance sheet or not.

7.18 In stages (a) to (c) the risk analysis is mainly concerned with determining the expected value of the risks involved. In stage (d) it is aimed at understanding the relative risk held by the public and private sectors, expressed as potential variations in the financial flows that arise from the asset to each party.

Appendix A Categories of generic risks

We list below some typical risks to which projects are exposed, under the following headings.

Design and construction risk (both to cost and to time);

Commissioning and operating risks (including maintenance);

Demand (or volume/ usage) and pricing risks;

Residual value risks;

Obsolescence risks; and

Regulation and similar risks.

Design and construction risk

Availability of finance

Site availability

Site costs

Existing assets:

- condition; and
- sale value.

Design risks:

- inadequate basis for design;
- new design standards;
- inadequate design;
- design errors; and
- design delays.

Implementation risks:

- site access problems;

- unforeseen ground conditions;
- archaeological discoveries;
- weather;
- strikes;
- interference from third parties;
- interactions with utilities and statutory undertakers;
- noise restrictions;
- buildability;
- delays with procurement of materials;
- availability of plant;
- unproven technology;
- unforeseen incompatibility with existing systems and services;
- delays with approvals;
- late design changes;
- insolvency of subcontractors or suppliers;
- commissioning difficulties;
- contractual claims;
- abandonment of contract;
- corruption;
- construction defects;
- third-party liability;
- fire;
- flood; and
- health and safety.

Commissioning and operating risks

- cost;
- availability of spares;
- downtime;
- availability of skilled maintenance staff;
- availability and cost of consumables;
- new regulations: environmental;
- inflation;
- fire;
- theft;
- accidental damage;
- vandalism;
- health and safety of staff and customers;
- prosecution for breach of environmental regulations;
- mis-estimation of operating costs (maintenance; staff; materials; transport; rent...)
- training costs;
- availability of staff;
- exchange rates;
- supply failures; and
- competition from new and existing facilities.

Demand (or volume/ usage) and pricing risks

Pricing and demand risk are obviously linked

- demand is greater or less than expected;
- price is greater or less than expected;
- timing of demand;
- competition.

Residual value risks

- residual value of assets; and
- market attractiveness.

Obsolescence risks

- technological; and
- environmental.

Regulation and similar risks

Taxation:

- VAT; and
- capital allowances.

Legislation