Appendix 1: The meaning of risk

In order to understand and apply the RAMP process, it is necessary to understand the basic concepts associated with risk and some of the techniques used in its analysis.

Simple example

Suppose we wish to construct a bird table for the garden in the hope that it will enable us to see many birds from our window.

Baseline plans

First we prepare an outline design, and compile a budget and plan for its construction and installation. Then we consider the risks which could affect the cost and timing of the project, or our enjoyment of the results when the bird table is in use. This will involve identifying and analysing the risks, and then minimising their adverse effects.

Identifying risks

Risks can have outcomes which are more or less favourable than expected (referred to as 'upside' and 'downside' risks respectively). In the case of the bird table, examples of downside risks include:

- suitable timber difficult to obtain
- costs more than expected
- takes too long
- construction proves very difficult
- is struck by lightning after erection
- birds hate it!
- timber soon rots
- tool breaks
- concrete below soil prevents erection
- neighbours complain about the birds
- new legislation prohibits bird tables
- something else goes wrong.

There are also such upside risks as:

- materials cost much less than expected
- takes less time than anticipated
- lasts far longer than thought likely
- attracts more bird varieties than expected
- a neighbour offers to build it
- we find one second hand.

The risks differ considerably in nature. They can arise before and during construction, or after the bird table comes into operation. Some are quite likely, others extremely unlikely. Some are uncertain, in that they are very difficult to assess. For example, the chance of new legislation might seem very unlikely at first sight but could become a real possibility in the event of disease being found to be spread by birds. Some risks are trivial in their effect, while others if they occurred would spell the end of the project. Some risks are independent, but many are dependent on others (e.g. the risk of it taking too long is dependent on, among others, the difficulty of finding suitable timber).

To analyse the risk of the project taking too long, we need to look at all underlying causes of possible delay and their chance of occurrence. Then we can decide what action we can take to minimise the risks and their effects if they arise. For example, we could reduce the chance of construction proving difficult by designing the bird table in accordance with plans in a do-it-yourself magazine. Also we could buy two types of nails in case the first type proves unsuitable for parts of the construction. This would increase the cost slightly, and we would have to consider whether it is worthwhile.

Evaluating risks

We have talked about the 'chance' or 'likelihood' of various things happening. What do we mean by this? Consider the chance of the hammer breaking during construction of our bird table. One approach would be to estimate the proportion of occasions on which the hammer had broken on other jobs. We would not have exact records but might be able to say that, based on past experience, the chance is between one in a hundred and one in two hundred. Or, not having made a bird table before, we might perceive the chance to be somewhat increased, say, to one in 50, because we shall be using longer nails than normal. But suppose the hammer was weakened on the last job without our knowing; then the shaft may break next time it is used, so that the real chance may be as high as, say, one in two.

Faced with the possibility of neighbours objecting, we might visit them to discuss our plans in advance, to ascertain their attitudes and discover how they are likely to react. This, and the previous example, illustrate that additional knowledge can dramatically alter our perception of the risks we are facing. The cost of investigating the true risks can often be worthwhile in terms of giving confidence that the project's objectives can be met.

If we were in the business of constructing and selling bird tables, then we could build up records of risks based on our experience over a large number of units. For example, if we had sold a million and 100 had been struck by lightning over a five-year period, then we could calculate that the chance of this risk event is one in 10 000. We could thus evaluate the likely cost of offering to replace them free of charge. However, if one is constructing only one bird table, does it really matter what the risk is of a lightning strike? It probably will not happen anyway. The answer depends on how seriously we view the loss of our bird table. If we regard this prospect as disastrous, it may be worth paying a small insurance premium to cover the potential loss. If the consequences of the loss are regarded as trivial, we would probably be content to bear the risk ourselves without insurance.

Many more risks will no doubt be considered in practice, even for this simple project, before the first steps are taken. For a major project, there would be many hundreds, or even thousands, of individual risks to be considered.

Unforeseen events

However much effort is put into risk analysis, there is sometimes a totally unexpected event which can scupper our well-laid plans. In our example we completed and installed the bird table successfully. The birds liked it and flocked to it. But after a week we noticed the birds being attacked by sparrow hawks as they fed and this continued. As a result the bird table brought us no pleasure so it has now been taken out of commission.

This illustrates that in practice it is often the unanticipated risks which can destroy even a major project or render it obsolete or unprofitable. Identification and study of the *catastrophic* risks can often be the most worthwhile aspect to concentrate on, especially near the beginning of the project appraisal process. By catastrophic risks, we mean those risk events which could have very serious or catastrophic consequences, even if the probability of occurrence is thought at first sight to be low. The iceberg risk for the *Titanic* is an obvious example.

RAMP is a process which helps to identify, analyse and respond to the risks and then bring them under proper control. It puts all these different considerations into a logical framework so that they can be treated methodically without so much likelihood of important aspects being missed.

The concepts and issues illustrated by the above example are explained more fully in the following sections.

Defining risk

The word 'risk' can have a number of different meanings. Consider for example, the following statements:

- there is a risk of rain today
- there is an 80% risk of rain today
- there is a risk of getting wet if it rains today
- there is a real risk to motorists from the weather which is forecast for today.

Each of these statements is using 'risk' in a different sense. Risk is commonly used as a synonym for 'hazard', 'danger' or 'threat' – i.e. an undesirable event. It can also refer to the likelihood of an event occurring. A third meaning is the loss, injury or other outcome resulting from an event. Yet another usage is to describe the generality of volatility and uncertainty – the combined effect of all the individual risks in an investment or situation (i.e. the overall 'riskiness').

The 'R' in RAMP stands for the word 'risk' in the latter sense, meaning the potential impact of all the threats (and opportunities) which can affect the achievement of the objectives for an investment. RAMP analyses and responds to the uncertainty relating to the objectives – for both favourable and unfavourable effects. Thus, unlike the usual dictionary definition of 'risk' (which refers only to undesirable events), the RAMP interpretation of risk includes both downside and upside variations in the values involved (e.g. capital and operating costs, revenues, net present values, etc.).

In analysing risks we are contemplating future events, the outcomes of which are therefore uncertain. We cannot generally predict with absolute confidence a particular outcome. Nevertheless, using relevant experience and judgement, we can usually define the range of possible outcomes, and then derive estimates of the likelihood and consequences of each, with a reasonable degree of confidence. This is the basis of risk analysis.

In their 1992 study the Royal Society used the word 'hazard' to mean a situation having the potential for human injury, damage to property, damage to the environment, or economic loss. They reserved the word 'risk' to mean a combination of the probability, or frequency, of occurrence of a defined hazard and the magnitude of the consequences of the occurrence. However, we have found it necessary to use the word 'risk' in the wider senses referred to above, in line with common parlance.

Key elements

There are six main concepts associated with evaluating risk.

- *Overall risk:* the combined effect of all individual risks or sources of uncertainty in a situation. It can be divided into two portions: overall upside risk and overall downside risk.
- *Risk event:* a possible occurrence which could affect (positively or negatively) the achievement of the objectives for the investment.
- *Likelihood:* the chance (or probability) of the risk event occurring within a defined time period.
- *Impact:* the value of the effect of the risk event on one or more objectives if it occurs.
- *Expected value:* a best estimate of the average outcome, i.e. all possible outcomes weighted by their probabilities.
- *Risk efficiency:* a state achieved when the downside risks have been sufficiently mitigated and the upside risks have been optimised.

We shall explain each of these in turn, in the context of their use in RAMP.

Overall risk

Overall risk is the combined effect of all individual risks or sources of uncertainty in a situation, both upside and downside. We commonly say that a project is 'risky' if there is expected to be considerable downside variation or volatility in the eventual possible outcomes, or 'safe' if there is expected to be little significant downside variation. One way of measuring overall risk is to express all the possible outcomes in monetary terms, allowing for time, so as to have a common and meaningful unit of measurement. In this handbook we suggest doing this by calculating the NPV of the various possible outcomes, thus obtaining a distribution of NPVs around the NPV for the base case, showing the likelihood of achieving each NPV. This distribution is a measure of the overall risk and gives an indication of whether the project is 'risky' or 'safe', as well as an indication of the value of the possible upside and downside outcomes. (To keep the work within reasonable bounds we would not usually calculate the NPV for every single possible outcome but instead approximate by using scenario analysis or stochastic modelling.)

Of course we should not necessarily prefer a 'safe' project to a 'risky' project. Even though the 'safe' project may have less downside risk, it may also have less upside risk. Much may depend on how risk tolerant the sponsor is. If the sponsor has a high risk tolerance for a project of this size, and can accept the possibility of a large negative NPV, he may prefer the 'risky' project if it has a relatively high 'expected' NPV or if it has the possibility of a large positive NPV in some scenarios.

Risk event

Risk events are the specific happenings that can influence the success of the investment, which therefore need to be identified, evaluated and responded to as part of the risk analysis and management process. Examples are:

- delay in tunnelling for a new underground railway due to unforeseen ground conditions
- overspend caused by increased cost of land acquisition
- intervention by regulator to limit price increases
- introduction of new statutory maximum noise levels
- more use than expected.

Each risk event can be triggered by one or more causes and can result in one or several outcomes. For the first three of the above examples, these could be as shown in Table 3.

As illustrated in Fig.6, there is a likelihood that each cause will lead to the event, and a further likelihood that if the event does occur it will result in each of the outcomes. Of course, several causes can arise together, increasing the chance of the event, and several outcomes can result from the event (e.g. a delay is usually associated with an extra cost).

Causes	Risk events	Possible outcomes
Unforeseen geological conditions Man-made obstructions Site flooding	Delay in tunnelling	Late completion Less time for installation of track and equipment Increased capital cost
Higher property prices More land required Unexpected need for decontamination	Increased cost of land	Overspend on capital budget Need to reduce scope
Reduced total investment Recent price rises Appeals from customers	Regulator limits prices	More customers Lower or higher revenue

Table 3. Risk events and possible outcomes

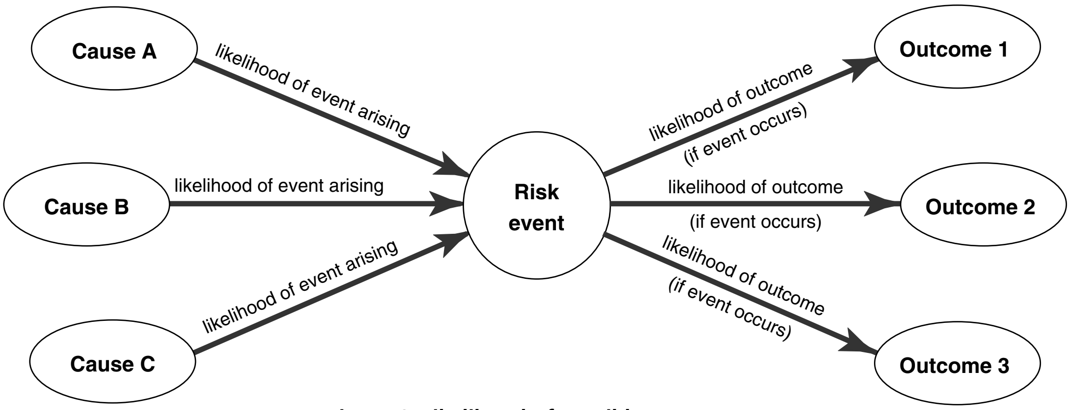

Figure 6. Likelihood of possible outcomes

Likelihood

Likelihood (or probability) is the degree of certainty that a risk event will occur during a specified time period. Typically it is measured on a scale of 0 to 100% (as in 'there is an 80% chance of rain today') or, more usually, as a probability on a scale of 0 to 1, in which 0 represents an impossible event, 1 a certainty, and 0.5 an evens chance of occurring.

The theory of probability was originally developed for games of chance, such as tossing coins, drawing cards and spinning roulette wheels. For such cases, probabilities can be calculated easily by analysis. For example, when an unbiased coin is spun there are only two possible occurrences (ignoring the unlikely event of the coin landing on its edge) – a head or a tail. There is thus an evens chance that a toss will result in a head – i.e. the probability is 1/2. Likewise, the probability of getting a black King or Queen when a card is drawn at random from a pack of 52 is 4/52 or 1/13. In each case, the probability of a particular outcome is calculated as the number of occurrences which give rise to that outcome, divided by the exposure (i.e. number of possible occurrences).

This applies provided all occurrences are equally likely and mutually exclusive (i.e. no two occurrences can happen simultaneously). Such calculations can be confirmed by experience. For example, if a coin is tossed a large number of times, say 500, the proportion of *heads* is likely to be approximately 50% of the total – i.e. about 250. Of course, for real life events the estimation of probabilities is less clear cut, but similar principles can be applied. Generally, the best approach is to use observed frequencies, from experience of similar events, if appropriate information can be obtained. If not, subjective probabilities should be used instead.

As an example, suppose we wish to estimate the probability that a new toll bridge spanning a river estuary will have to be closed due to high winds (in excess of 60 mile/h) on the day of the Royal Opening scheduled for 1 November. By consulting meteorological records kept at a weather station near the site of the bridge, we see that during the 20-day period around that date (i.e. from 10 days before to 10 days after mid-day on 1 November), over the last 10 years that records have been kept, there have been 17 days when winds have been stronger than 60 mile/h. Hence the probability of the bridge being closed because of winds is 17/(10 x 20) or 17/200, i.e. about a 1 in 12 chance.

A crucial factor in deriving measures of probability is the validity of the assumptions which need to be made. So these must be carefully checked for realism, noted and monitored. If they change, the probabilities will need to be revised. An illustration of the probabilities and underlying assumptions for the above risk events is set out in Table 4.

Impact

As defined above, the impact of an event is the value of the effect of the risk event, if it occurs, on one or more of the financial parameters of the investment – e.g. on capital cost, revenue or operating cost. Thus impacts are derived by assessing the consequences of the outcomes of an event in terms of their financial consequences for the objectives. The impact can often be expressed as the NPV of the cash flows resulting from the outcomes. Table 5 gives some examples.

Expected value

Any potential risk may or may not occur in a particular investment. In order to get a good idea of what we might reasonably expect to be the overall impact of any category of risks for an investment, we need some simple measure. That is given by the 'expected value', which is calculated by multiplying impacts by the associated probabilities of events. It is equivalent to the average impact of the risk event, which would result if we were to carry out a large number of identical projects.

Thus, for example, if the chance that a cricket ground will become unsuitable for a test match is 10% and the impact on revenue if the game is cancelled is estimated as £2 million, then the expected value of cancellation is 10% of £2 million, that is £200 000. More generally, where there are several alternative possible scenarios, the overall expected value of the variations from the base case is the sum of the corresponding expected values for each separate scenario. Let us illustrate this by returning to our test match. If the match takes place and all available seats are occupied, the gross revenues will be £3 million. There will be fixed expenses of £400 000 irrespective of seat occupancy or whether the match takes place. The base case, or most likely outcome (with a 40% probability), is that the match takes place with two-thirds of the available seats occupied. The forecast net revenue is thus 2/3 x £3 million, less £400 000, i.e. £1 600 000. If, however, the game is cancelled (which has a 10% probability), the net revenue will be £2 million worse than the base case – i.e. minus £400 000. If the match takes place but only one-third of the seats are

Risk event	Probability	Assumptions
Head on toss of coin	1/2	Unbiased coin Properly tossed Does not end on edge
Black King or Queen on draw of card	1/13	Selected at random Normal pack of 52 cards
Toll bridge being closed at Royal Opening (in above example)	17/200 = 0.085	Closure enforced if 60 mile/h winds at any time in day Wind limit does not change Weather system unchanged over last 10 years

Table 4. Probabilities and assumptions for specific risk events

Risk events	Outcomes	Impact (NPV): £ million
Ground subsidence due to tunnelling for underpass	Damage to 6 buildings	Cost of compensation 2.0
New competitor enters market for bus service	Loss of passengers	Reduced revenue 3.7
Major contractor is bankrupt	Extra cost and delay for re-tendering	Increased capital cost 1.0 Loss of early revenue 1.5

Table 5. Examples of impact of risk events (if they occur) expressed as the NPV of the resulting variations in cash flows

occupied (which has a 30% probability), the net revenue will be £1 million worse than the base case – i.e. £600 000. If there is a very favourable outcome and all the seats are occupied (which has a 20% probability), the net revenue will be £1 million better than the base case – i.e. £2 600 000.

Thus, as well as the base case, there are three alternative scenarios being considered, namely that the match is cancelled, or that only one-third of the seats are occupied, or that all the seats are occupied. We can calculate the expected value of the variation from the base case under each scenario as illustrated in Table 6.

Thus the overall expected value of the risk scenarios is minus £300 000. To obtain the expected value of the net revenue from the match, we must deduct £300 000 from the base case forecast of £1 600 000, to leave £1 300 000.

Risk efficiency

This is a key objective for any risk management process. It is achieved when we reach the point, in devising responses to both downside risks and opportunities, beyond which we believe that the marginal cost of introducing an additional response would exceed the utility to the sponsor of the resulting risk reduction or opportunity increase. The process will usually involve trial and error, to find the right set of risk responses.

To be satisfied that risk efficiency has been achieved, we need to be able to answer positively the following four questions:

- Have we exhausted all the opportunities to reduce downside risk or increase upside risk without significantly increasing expected cost?
- Have we sufficiently explored those possible responses to risk which would involve extra cost?
- Have we considered possible responses to risk which would reduce expected cost without increasing downside risk or reducing upside risk to an unacceptable extent?
- Would the marginal cost of an additional risk response, including the time and effort in searching for it, exceed the marginal increase in the sponsor's utility from the resulting risk reduction or opportunity increase, considering each possible additional risk response in turn and allowing for secondary risks?

Achieving risk efficiency will tend to determine which set of risk response actions should be adopted in practice. However, the project may still not proceed if it is not financially viable or if there is too much residual downside risk or uncertainty to meet the sponsor's risk tolerance.

Probability distributions

The expected value of £1 300 000 can be regarded as the average net revenue which would be obtained from each test match if a large number of matches were to be played. If only one match was due to be played, however, it would in addition be useful to look at the probabilities of occurrence of each of the possible outcomes, including the base case, as in Table 7: this shows the probability distribution of the net revenue. It demonstrates that there is a 10% risk of losing £400 000 and this possible outcome needs to be taken into account by the sponsor before deciding whether to proceed or

Scenario	Probability	Impact	Expected value of variation from base case
Game cancelled	10%	–£2 000 000	–£200 000
Only one-third of seats occupied	30%	–£1 000 000	–£300 000
Seats fully occupied	20%	+£1 000 000	+£200 000
	60%		–£300 000

Table 6. A test match – the calculation of expected values of three scenarios different from the base case

Scenario	Net revenue	Probability
Game cancelled	–£400 000	10%
Only one-third of seats occupied	+£600 000	30%
Base case (two-thirds occupied)	+£1 600 000	40%
Seats fully occupied	+£2 600 000	20%
		100%

Table 7. A test match – probability distribution of net revenue

not. Also, the sponsor will want to take account of the 20% chance of gaining as much as £2 600 000.

Let us take as another example a specialist developer of high-quality houses. He has found from experience that, at current prices, he can sell the following proportions of his standard 'classic' houses at each price (to the nearest £100 000):

- 0.3 at £600 000
- 0.5 at £700 000
- 0.2 at £800 000.

This can be plotted as a probability distribution and as a cumulative probability distribution, as shown in Figs 7 and 8. The cumulative probability distribution shows the probability of selling at a particular price or less. Hence the chance of being able to sell at up to £750 000 is 0.3 + 0.5 = 0.8.

Thus, from the cumulative probability distribution, we can estimate that there is approximately a 50% chance that the selling price of a house will reach £690 000. In this example, this is the same as the expected value which is given by: £[(6 × 0.3) + (7 × 0.5) + (8 × 0.2)] × 100 000, i.e. £690 000.

Wherever practicable, recorded data is used to derive probability distributions, with adjustments if necessary to allow for future trends. However, there are many cases where this is not feasible. In such cases, if the risks involved are such as to justify quantifying, then it is necessary to estimate the probability using judgement – ideally exercised by someone with expertise in the particular area concerned.

For example, we may wish to elicit from an art expert the probability distribution for the price likely to be realised at auction for a masterpiece owned by a charity of which we are patrons. We may ask the expert to

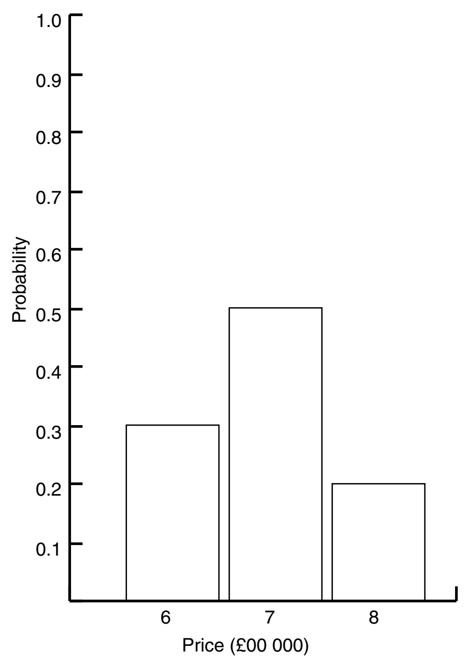

Figure 7. Probability distribution

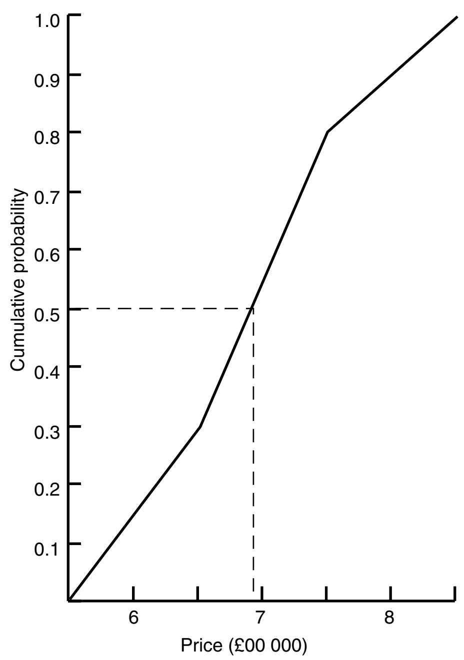

Figure 8. Cumulative probability distribution

proceed as follows (for simplicity, ignoring here the possibility that the painting may remain unsold):

- Estimate the high-impact scenario (H) which has, say, a likelihood of 10% of being reached. For the painting assume this is £400 000.
- Next determine the low-impact scenario (L) for which there is only a 10% chance (say) of it not being exceeded. Let us presume this is set at £200 000 in our example.
- Then choose either one or two intermediate points – either one-half or one-third of the distance between L and H. We assume here one at £300 000.
- Decide on a class interval – usually this can be conveniently set equal to (H – L)/2 or (H – L)/3 depending on whether there are one or two intermediate points between L and H. For the picture, the class interval will be £100 000.
- Then (using expert judgement) assign estimated probabilities (ideally rounded to one significant figure) to each class in the range. These give for the masterpiece the following distribution (to the nearest £100 000):
  - 0.3 at £200 000
  - 0.5 at £300 000
  - 0.2 at £400 000.

The above discussion may seem to imply that any risk event has a single clearly defined impact if it occurs. However, in practice there is often a range of possible outcomes for an individual risk event, each of which has a different probability (i.e. there is a probability distribution for the risk). Furthermore, for each outcome of the risk event, there can be a range of impact values, each with its own likelihood of occurrence.

Chapman and Ward (2003) use a similar 'minimalist' approach. If, however, the expert thinks this approach is too crude because he has more information about the 'shape' of the probability distribution, the method can be modified as necessary.

Combining risks

A primary aim of RAMP is to evaluate and manage risks affecting the overall investment. It is therefore essential to be able to aggregate a number of risks which could potentially affect each of the investment objectives. In order to do this, we need to quantify at least the most significant risks, understand the nature of their relationships, if any, and then combine the probabilities so as to determine their collective effect.

There are two main rules for combining individual risks – probability 'addition' and 'multiplication'. These will be explained by assuming there are two events, A and B respectively, for which the probabilities of their occurring can be denoted as Prob(A) and Prob(B).

The 'addition' rule states

Prob(A or B or both) =
Prob(A) + Prob(B) – Prob(A and B).

The reason for subtracting Prob(A and B) is that Prob(A) includes the possibility that A and B occur together and so does Prob(B), so one Prob(A and B) is subtracted from the sum to eliminate double counting.

If events A and B are mutually exclusive, i.e. cannot occur together, then Prob(A and B) = 0, so Prob(A or B or both) becomes Prob(A) + Prob(B).

The 'multiplication' rule states

Prob(A and B) =
Prob(A) x Prob(B), provided A and B are independent.

By 'independent' we mean that the occurrence of one event does not influence the likelihood of the other. If, as often is the case, the events are not independent, then the multiplication rule becomes:

- If B is dependent on A – i.e. the fact that A has occurred could affect the likelihood of B's occurrence, then

Prob(A and B) =
Prob(A) x Prob(B, given that A has occurred).

Scenario	Selling price of 1st home: £00 000	Selling price of 2nd home: £00 000	Probability	Total revenues: *£00 000
A	6	6	0.3 x 0.3 = 0.09	12
B	6	7	0.3 x 0.5 = 0.15	13
C	6	8	0.3 x 0.2 = 0.06	14
D	7	6	0.5 x 0.3 = 0.15	13
E	7	7	0.5 x 0.5 = 0.25	14
F	7	8	0.5 x 0.2 = 0.10	15
G	8	6	0.2 x 0.3 = 0.06	14
H	8	7	0.2 x 0.5 = 0.10	15
I	8	8	0.2 x 0.2 = 0.04	16
			1.00	* + or – £100 000

Table 8. First estimate of probability distribution (assuming independence)

- If A is dependent on B – i.e. the fact that B has occurred could affect the likelihood of A's occurrence, then

Prob(A and B) =
Prob(B) x Prob(A, given that B has occurred).

Prob(B, given that A has occurred) and Prob(A, given that B has occurred) are called 'conditional probabilities' and will not generally be the same as the unconditional probabilities, Prob(B) and Prob(A).

As a simple example of using the above rules, we can calculate the probability of getting an odd number or a number less than three, or both, when we throw a six-sided die as follows:

Prob('odd' or 'less than 3' or both)
= Prob('odd') + Prob('less than 3') – Prob(both)
= 3/6 + 2/6 – 3/6 x 2/6
= 4/6 = 2/3.

In this case, Prob('odd') and Prob('less than 3') are treated as being independent. The above rules can be extended to cover the combination of any number of individual risks. However, combining risks which have probability distributions for a range of values gets more complicated. To demonstrate how this can be done, let us go back to the example of the house developer above. Suppose he plans to sell two houses over the next two months and wishes to estimate the probability distribution for the total resulting selling price, assuming the two sales are independent. This can be done as shown in Table 8.

Total revenue *£00 000	Scenarios	Combined probability
12	A	0.09
13	B, D	0.30
14	C, E, G	0.37
15	F, H	0.20
16	I	0.04
* + or – £100 000		1.00

Table 9. Result of first estimate of probability distribution (assuming independence)

We now gather together the scenarios giving the same total revenues and hence arrive at the probability distribution shown in Table 9 for the total revenue.

As stated above, this assumes that the price of the second house is independent of that for the first. In practice, risks are rarely perfectly independent.

Two kinds of dependence can be distinguished, representing positive and negative correlation between the risks. Positively correlated risks tend to increase together, whereas negatively correlated risks move in opposite directions.

Let us now suppose the price of the second

Scenario	Selling price of 1st home: £00 000	Selling price of 2nd home: £00 000	Probability	Total revenues: *£00 000
A	6	6	0.30 (0.60 + 0.40 x 0.3) = 0.216	12
B	6	7	0.30 (0.40 x 0.5) = 0.060	13
C	6	8	0.30 (0.40 x 0.2) = 0.024	14
D	7	6	0.50 (0.40 x 0.3) = 0.060	13
E	7	7	0.50 (0.60+0.40 x 0.5) = 0.400	14
F	7	8	0.50 (0.40 x 0.2) = 0.040	15
G	8	6	0.20 (0.40 x 0.3) = 0.024	14
H	8	7	0.20 (0.40 x 0.5) = 0.040	15
I	8	8	0.20 (0.60 + 0.40 x 0.2) = 0.136	16
			1.000	*+ or – £100 000

Table 10. Second estimate of probability distribution (assuming dependence)

Total revenue: *£00 000	Scenarios	Combined probability
12	A	0.216
13	B, D	0.120
14	C, E, G	0.448
15	F, H	0.080
16	I	0.136
* + or – £100 000		1.000

Table 11. Result of second estimate of probability distribution (assuming dependence)

house has a 60% chance of falling within the same price bracket as the first house (due to word getting out) and only a 40% chance of being sold entirely independently. We would then have the situation shown in Table 10 and the resulting probability distribution as is presented in Table 11.

This is *very* different from the preceding result based on an assumption of independence, as can be seen in Figs 9 and 10.

Many users of risk analysis tend to use techniques for aggregating risks which implicitly treat the risks as if they are completely independent. This can give seriously misleading results. Chapman and Ward (2003) point out that a high degree of dependence is typically encountered – up to 80% for costs and typically 50% for time related risks. They propose

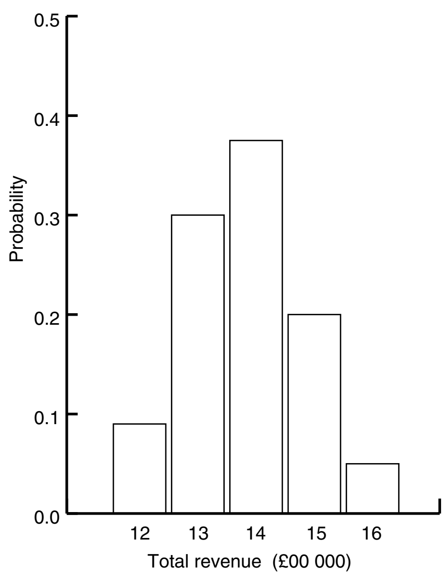

Figure 9. First probability distribution: all sold independently

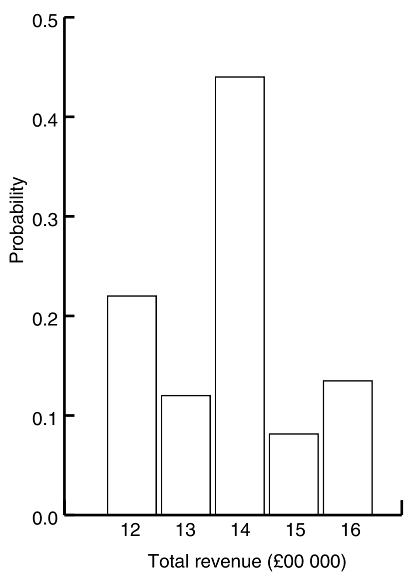

Figure 10. Second probability distribution: only 40% sold independently

Perceived risk distribution		Real risk distribution	
Value of outcome	Probability	Value of outcome	Probability
6	0.2	7	0.1
7	0.4	8	0.2
8	0.3	9	0.4
9	0.1	10	0.2
		11	0.1

Table 12. Perceived and real risk distributions

calculating the cumulative probability distributions for each of the risks being combined and then deriving a weighted average cumulative distribution to represent the combination of the risks, using as weights, *r* and *r – 1* to multiply the dependent and independent distributions, where *r* is the correlation coefficient. The coefficient can be based on calculation or judgement.

Real and perceived risks

Perhaps the most tantalising thing about risk analysis is that, in most situations, we do not, and cannot, know the real risks that we face. Even after exposure to the risk has taken place (whether the risk event occurred or not), we still do not know what the true risks really were. Specifically, we do not know what the real risk distribution is and what will be the specific outcome for a particular risk event if it occurs. Clearly, there is a difference between perceived risk and real risk. The former is what we estimate, the latter is the true risk which currently exists. See Table 12 for an example of possible differences between a perceived and a real risk distribution.

The average values of these risk distributions are 7.3 and 9 respectively, and let us assume that the actual outcome is a value of 7. The risk distribution histograms and the actual outcome, on common axes, are shown in Fig. 11.

When and if the risk occurs we know, of course, that it lies within the real risk distribution, but not where nor what the shape of the real distribution is. In this example, the perceived risk distribution is significantly different from the real one. Yet, when the actual result turns out to be 7 – approximately equal to the estimated average value – it may be presumed, quite wrongly, that the perceived distribution was close to the truth.

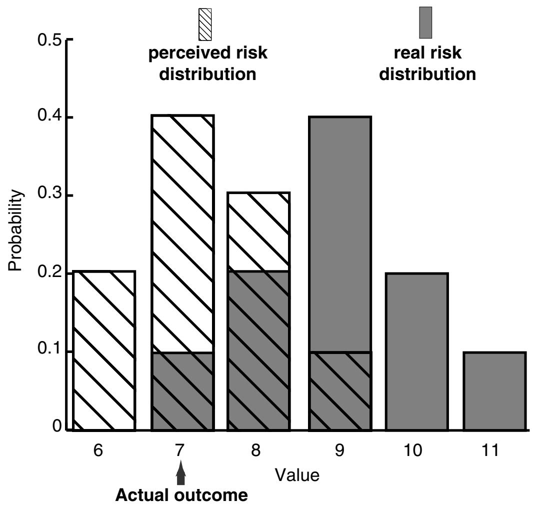

Figure 11. The outcome of the risk distributions shown in Table 12

Even if we had correctly estimated the risk distribution initially, it is distinctly possible that the distribution may have changed significantly before the risk exposure ends. That is why we need to be diligent in recording and monitoring the assumptions which underlie our estimates of risks, and why we need to reassess risks at regular intervals during the life of an investment.

The degree of difference between perceived and real risks depends on our level of knowledge about the risks and the situation in which they exist. We can sometimes improve our knowledge of the risks by carrying out further research in an endeavour to make our perception of the risks more closely match reality.