1

Risk is equal to the expected value

If you throw a die, the outcome will be either 1, 2, 3, 4, 5 or 6. Before you throw the die, the outcome is unknown – to use the terminology of statisticians, it is random. You are not able to specify the outcome, but you are able to express how likely it is that the outcome is 1, 2, 3, 4, 5 or 6. Since the number of possible outcomes is 6 and they are equally probable – the die is fair – the probability that the outcome turns out to be 3 (say), is 1/6. This is simple probability theory, which I hope you are familiar with.

Now suppose that you throw this die 600 times. What would then be the average outcome? If you do this experiment, you will obtain an average about 3.5. We can also deduce this number by some simple arguments: about 100 throws would give an outcome equal to 1, and this gives a total sum of outcomes equal to 100. Also about 100 throws would give an outcome equal to 2, and this would give a sum equal to 2 times 100, and so on. The average outcome would thus be

$$(1 \times 100 + 2 \times 100 + 3 \times 100 + 4 \times 100 + 5 \times 100 + 6 \times 100)/600 = 3.5. \quad (1.1)$$

In probability theory this number is referred to as the expected value. It is obtained by multiplying each possible outcome with the associated probability, and summing over all possible outcomes. In our example this gives

$$1 \times 1/6 + 2 \times 1/6 + 3 \times 1/6 + 4 \times 1/6 + 5 \times 1/6 + 6 \times 1/6 = 3.5. \quad (1.2)$$

We see that formula (1.2) is just a reformulation of (1.1) obtained by dividing 100 by 600 in each sum term of (1.1). Thus the expected value can be interpreted as the average value of the outcome of the experiment if the experiment is repeated

Misconceptions of Risk T. Aven

over and over again. Statisticians would refer to the law of large numbers, which says that the average value converges to the expected value when the number of experiments goes to infinity.

Reflection

For the die example, show that the expected number of throws showing an outcome equal to 2 is 100 when throwing the die 600 times.

In each throw, there are two outcomes: one if the outcome is a 'success' (that is, shows 2), and zero if the outcome is a 'failure' (that is, does not show 2). The corresponding probabilities are 1/6 and 5/6. Hence the expected value for a throw equals $1 \times 1/6 + 0 \times 5/6 = 1/6$, in other words the expected value equals the probability of a success. If you perform 2 throws the expected number of successes equals $2 \times 1/6$, and if you perform 600 throws the expected number of successes equals $600 \times 1/6 = 100$. These conclusions are intuitively correct and are based on a result from probability calculus saying that the expected value of a sum equals the sum of the expected values. Thus the desired result is shown. □

The expected value is a key concept in risk analysis and risk management. It is common to express risk by expected values. Here are some examples:

- For some experts 'risk' equals expected loss of life expectancy (HM Treasury, 2005, p. 33).
- Traditionally, hazmat transport risk is defined as the expected undesirable consequence of the shipment, that is, the probability of a release incident multiplied by its consequence (Verma and Verter, 2007).
- Risk is defined as the expected loss to a given element or a set of elements resulting from the occurrence of a natural phenomenon of a given magnitude (Lirer *et al.*, 2001).
- Risk refers to the expected loss associated with an event. It is measured by combining the magnitudes and probabilities of all of the possible negative consequences of the event (Mandel, 2007).
- Terrorism risk refers to the expected consequences of an existent threat, which, for a given target, attack mode, target vulnerability and damage type, can be expressed as the probability that an attack occurs multiplied by the expected damage, given that an attack occurs (Willis, 2007).
- Flood risk is defined as expected flood damage for a given time period (Floodcite, 2006).

But is an expected value an adequate expression of risk? And should decisions involving risk be based on expected values?

Example. A Russian roulette type of game

Let us look at an example: a Russian roulette type of game where you are offered a play using a six-chambered revolver. A single round is placed in the revolver such that the location of the round is unknown. You take the weapon and shoot, and if it discharges, you lose $24 million. If it does not discharge, you win $6 million.

As the probability of losing $24 million is 1/6, and of winning $6 million is 5/6, the expected gain is given by

$$-24 \times 1/6 + 6 \times 5/6 = -4 + 5 = 1.$$

Thus the expected gain is $1 million. Say that you are not informed about the details of the game, just that the expected value equals $1 million. Would that be sufficient for you to make a decision whether to play or not play? Certainly not – you need to look beyond the expected value. The possible outcomes of the game and the associated probabilities are required to provide the basis for an informed decision. Would it not be more natural to refer to this information as risk, and in particular the probability that you lose $24 million? As we will see in coming chapters, such conceptions of risk are common.

The game has an expected value of $1 million, but that does not mean that you would accept the game as you may lose $24 million. The probability 1/6 of losing may be considered very high as such a loss could have dramatic consequences for you. And how important is it for you to win the $6 million? Perhaps your financial situation is good and an additional $6 million would not change your life very much for the better. The decision to accept the play needs to take into account aspects such as usefulness, desirability and satisfaction. Decision analysts and economists use the term utility to convey these aspects.

Daniel Bernoulli: The need to look beyond expected values

The observation that there is a need for seeing beyond the expected values in such decision-making situations goes back to Daniel Bernoulli (1700–1782) more than 250 years ago. In 1738, the *Papers of the Imperial Academy of Sciences in St Petersburg* carried an essay with this central theme: ‘the value of an item must not be based on its price, but rather on the utility that it yields’ (Bernstein, 1996). The author was Daniel Bernoulli, a Swiss mathematician who was then 38 years old. Bernoulli’s St Petersburg paper begins with a paragraph that sets forth the thesis that he aims to attack (Bernstein, 1996):

> Ever since mathematicians first began to study the measurement of risk, there has been general agreement on the following proposition: Expected values are computed by multiplying each possible gain by the number of ways it can occur, and dividing the sum of these products by the total number of cases.

Bernoulli finds this thesis flawed as a description of how people in real life go about making decisions, because it focuses only on gains (prices) and probabilities, and not the utility of the gain. Usefulness and satisfaction need to be taken into account. According to Bernoulli, rational decision-makers will try to maximize expected utility, rather than expected values (see Chapter 6). The attitude to risk and uncertainties varies from person to person. And that is a good thing. Bernstein (1996, p. 105) writes:

> If everyone valued every risk in precisely the same way, many risky opportunities would be passed up. . . . Where one sees sunshine, the other sees a thunderstorm. Without the venturesome, the world would turn a lot more slowly. Think of what life would be like if everyone were phobic about lightning, flying in airplanes, or investing in start-up companies. We are indeed fortunate that human beings differ in their appetite for risk.

Reflection

Bernoulli provides this example in his famous article: two men, each worth 100 ducats (about $4000), decide to play a fair game (i.e. a game where the expectation is the same for both players) based on tossing coins, in which there is a 50–50 probability of winning or losing. Each man bets 50 ducats on the throw, which means that he has an equal probability of ending up worth 150 ducats or of ending up worth only 50 ducats. Would a rational player play such a game?

The expectation is 100 ducats for each player, whether they decide to play or not. But most people would find this play unattractive. Losing 50 ducats hurts more than gaining 50 ducats pleases the winner. There is an asymmetry in the utilities. The best decision for both is to refuse to play the game.

Risk-averse behaviour

Economists and psychologists refer to the players as risk-averse. They dislike the negative outcomes more than the weight given by the expected value. The use of the term ‘risk averse’ is based on a concept of risk that is linked to uncertainties more than expected values (see Chapter 4). Hence this terminology is in conflict with the idea of seeing risk as the expected value.

Let us return to the Russian roulette game described above. Imagine that you were given a choice between a gift of $0.5 million for certain or an opportunity to play the game with uncertain outcomes. The gamble has an expectation equal to $1 million. Risk-averse people will choose the gift over the gamble. As the possible loss is so large ($24 million), they would probably prefer any gift (even a fixed loss) instead of accepting the game. The minimum gift you would require is referred to as the certainty equivalent. A person is risk-averse if the certainty equivalent is less than the expected value. Different people would be risk-averse to different degrees. This degree is expressed by the certainty equivalent. How high (low) would the gift have to go before you would prefer the game to the gift?

For the above examples, most people would show a risk-averse attitude. A risk seeker would have a higher certainty equivalent than the expected value. He or she values the probability of winning to be so great that (s)he would prefer to play the game instead of receiving the gift of say \$1.2 million.

A portfolio perspective

But is it not more rational to be risk-neutral, that is, letting the certainty equivalent be equal to the expected value? Say that you represent an enterprise with many activities and you are offered the Russian roulette type of game. The enterprise is huge, with a turnover of billions of dollars and hundreds of large projects. In such a case the enterprise management would probably accept the game, as the expectation is positive. The argument is that when considering many such games (projects) the expected value would be a good indication of the actual outcome of the total value of the games (projects).

To illustrate this, say that the portfolio of projects comprises $n = 100$ projects and each project is of the Russian roulette type, that is, the probability of losing \$24 million is 1/6, and the probability of winning \$6 million is 5/6. For each project the expected gain equals \$1 million and hence the expected average gain for the 100 projects is \$1 million. Looking at all the projects we would predict \$1 million per project, but the actual gain could be higher or lower. There is a probability that we lose hundreds of million of dollars, but the probability is rather low. In theory, all projects could result in a loss of \$24 million, adding up to loss of \$2400 million. Assuming that all the n projects are independent of each other, the probability of this extreme result is $(1/6)^{100}$, which is an extremely small number; it is negligible. It is, however, quite likely that we end up with a loss, that is, negative average gain. To compute this probability we make use of the central limit theorem, expressing the fact that the probability distribution of the average value can be accurately approximated by the normal (Gaussian) probability curve. As shown in Table 1.1, the probability that the average gain is less than zero

Table 1.1 Probability distribution for the average gain when $n = 100$.

X	Probability that the average gain is less than x
−2	0.004
−1	0.04
0	0.19
1	0.50
2	0.81
3	0.96
4	0.996

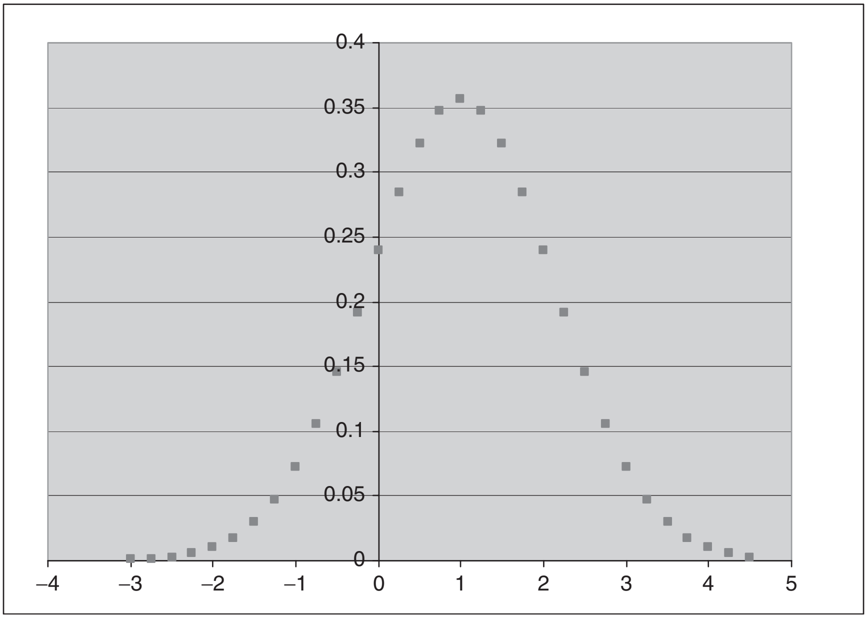

Figure 1.1 The Gaussian curve for the average gain in millions of dollars for the case $n = 100$. The area under the curve from a point a to a point b on the x-axis represents the probability that the gain takes a value in this interval.

equals approximately 0.20, assuming that all projects are independent. Figure 1.1 shows the Gaussian curve for the average gain. The probability that the average gain will take a value lower than any specific number is equal to the area below the curve. The integral of the total curve is 1. Hence, the probability is 0.50 that the average gain is less than 1, and 0.50 that the average gain exceeds 1. Table 1.1 provides a summary of some specific probabilities.

The central limit theorem has an interesting history, as Tijms (2007, p. 162) describes:

The first version of this theorem was postulated by the French-born mathematician Abraham de Moivre, who, in a remarkable article published in 1733, used the normal distribution to approximate the distribution of the number of heads resulting from many tosses of a fair coin. This finding was far ahead of its time, and was nearly forgotten until the famous French mathematician Pierre-Simon Laplace rescued it from obscurity in his monumental work *Théorie Analytique des Probabilités*, which was published in 1812. Laplace expanded De Moivre's finding by approximating the binomial distribution with the normal distribution. But as with De Moivre, Laplace's finding received little attention in his own time. It was not until the nineteenth century was at an end that the importance of the central limit theorem was discerned, when, in 1901, Russian mathematician Aleksandr Lyapunov defined it in general terms and proved precisely how it worked mathematically. Nowadays, the central limit theorem is considered to be the unofficial sovereign of probability theory.

Calculations of the figures in Table 1.1

The calculations are based on the expected value, which equals 1, and the variance, which is a measure of the spread of the distribution relative to the expected value. For one project, the variance equals

$$(-24 - 1)^2 \times 1/6 + (6 - 1)^2 \times 5/6 = 25^2 \times 1/6 + 5^2 \times 5/6 = 750/6 = 125.$$

We see that the variance is computed by squaring the difference between a specific outcome and the expected value, multiplying the result by the probability of this outcome, and then summing over the possible outcomes. If X denotes the outcome we denote by $\mathrm{E}[X]$ the expected value of X, and $\mathrm{Var}[X]$ the variance of X. Formally, we have $\mathrm{Var}[X] = E([X - EX])^2$.

The square root of the variance is called the standard deviation of X, and is denoted $\mathrm{SD}[X]$. For this example we obtain $SD[X] = 11.2$. The variance of a sum of independent quantities equals the sum of the individual variances. Let Y denote the total gain for the 100 projects. Then the variance of Y, $\mathrm{Var}[Y]$, equals 12 500.

The central limit theorem states that

$$\mathrm{P}(Y/n \leq x) \approx \Phi(\sqrt{n}(x - \mathrm{E}X)/\mathrm{SD}[X]),$$

where $\sqrt{n}$ equals the square root of n and Φ is the probability distribution of the standard normal distribution with expectation 0 and variance 1. The approximation $\approx$ produces an accurate result for large n, typically larger than 30. The application of this formula gives

$$\mathrm{P}(Y/n < 0) \approx \Phi(-10/11.2) = \Phi(-0.89) = 0.19,$$

using a statistical table for the Φ function. The standard deviation for Y/n equals $SD[X]]/\sqrt{n} = 1.12$.

We observe that the expected value is a more informative quantity when looking at 100 projects of this form than looking at one in isolation. The prediction is the same, \$1 million per project, but the uncertainties have been reduced. And if we increase the number of projects the uncertainties are further reduced.

Table 1.2 Probability distribution for the average gain when $n = 1000$.

X	Probability that the average gain is less than x
−0.25	0.0002
0	0.002
0.5	0.08
1	0.50
1.5	0.92
2	0.998

Say that we consider $n = 1000$ projects. Then we obtain results as in Table 1.2 and Figure 1.2. We see that the probability of a loss in this case is reduced to 0.2%. The outcome would with high probability be a gain close to $1 million. The uncertainties are small. Increasing n even further would give stronger and stronger concentration of the probability mass around 1. This can be illustrated by the variance or the standard deviation. For the above example the standard deviation of the average gain, SD[Y/n], equals 1.12 in case $n = 100$ and 0.35 when $n = 1000$. As the number of projects increases, the variance and

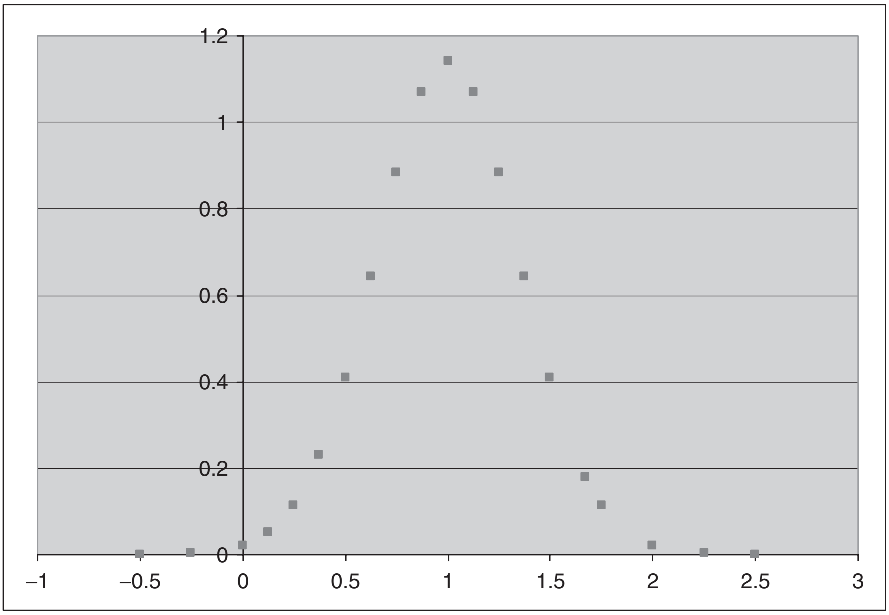

Figure 1.2 The Gaussian curve for the average gain in millions of dollars for the case n = 1000. The area under the curve from a to b represents the probability that the gain takes a value in this interval.

Table 1.3 Probability distribution for the average gain when $n = 10\,000$.

X	Probability that the average gain is less than x
0.60	0.0002
0.90	0.19
1	0.50
1.1	0.81
1.4	0.9998

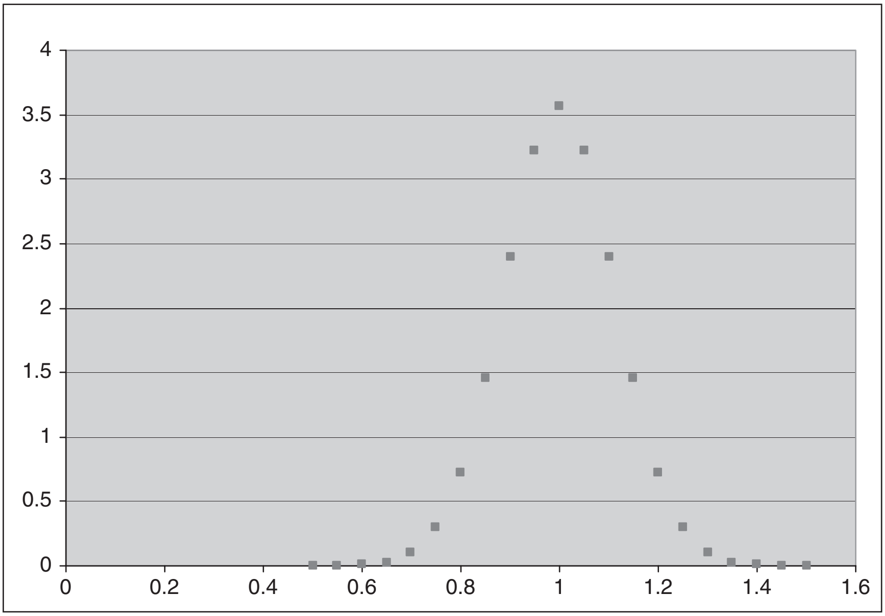

Figure 1.3 The Gaussian curve for the average gain in millions of dollars for the case n = 10 000. The area under the curve from a to b represents the probability that the gain takes a value in this interval.

the standard deviation decrease. When n becomes several thousands, the variance and the standard deviation become negligible and the average gain is close to the expected value 1. See Table 1.3 and Figure 1.3 which present the result for $n = 10\,000$.

Dependencies

The above analysis is based on the assumption that the projects are independent. The law of large numbers and the central limit theorem require that this assumption is met for the results to hold true. But what does this mean and to what extent is this a realistic assumption? Consider two games (projects), and let X_1 and X_2 denote the gains for games 1 and 2, respectively. These games are independent if the probability distribution of X_2 is not dependent on X_1, that is, the probability that the outcome of X_2 turns out to be -24 or 6 is not dependent on whether the outcome of X_1 is -24 or 6. For the Russian roulette type game independence is a reasonable assumption under appropriate experimental conditions. For real-life projects the situation is, however, more complex. If you know that project 1 has resulted in a loss, this may provide information also about project 2. There could be a common cause influencing both projects negatively or positively. Think about an increase in the oil price or a political event that influences the whole market. Hence, the independence assumption must be used with care.

Reflection

Consider an insurance company that covers the costs associated with work accidents in a country. Is it reasonable to judge the costs as independent?

Yes, a work accident at one moment in one place has a negligible relationship to a work accident at another moment in a different place. □

Returning to the example with n projects, the expected value cannot be the sole basis for the judgement about acceptance or not. The dependencies could give an actual average gain far away from 1. Consider as an example a case where the loss is -24 for all projects if the political event B occurs. The probability of B is set to 10%. If B does not occur, this means that the loss in one game, X, is -24 with probability $4/54 = 0.074$. The projects are assumed independent if B does not occur. We write $P(X = -24|\text{not } B) = 0.074$.

This is seen by using the law of total probability:

$$\begin{aligned}\mathrm{P}(X = -24) &= \mathrm{P}(X = -24|B) \cdot \mathrm{P}(B) + \mathrm{P}(X = -24|\text{not } B) \cdot \mathrm{P}(\text{not } B) \\ &= 1 \cdot 0.1 + \mathrm{P}(X = -24|\text{not } B) \cdot 0.9 = 1/6.\end{aligned}$$

The expected gain and variance given that B does not occur equal

$$\mathrm{E}[X|\text{not } B] = -24 \cdot (4/54) + 6 \cdot (50/54) = 204/54 = 3.78,$$

$$\mathrm{Var}[X|\text{not } B] = (-24 - 3.78)^2 \cdot (4/54) + (6 - 3.78)^2 \cdot (50/54) = 61.7.$$

Hence for one particular project we have the same probability distribution: the possible outcomes are 6 and -24, with probabilities 5/6 and 1/6, respectively. The projects are, however, not independent. The variance of the average gain, $\mathrm{Var}[Y/n]$, does not converge to zero as in the independent case.

To see this, we first note that

$$\begin{aligned}\mathrm{Var}[Y] = \mathrm{E}(Y - \mathrm{E}Y])^2 &= \mathrm{E}[(Y - \mathrm{E}Y])^2|B] \cdot 0.1 + \mathrm{E}[(Y - \mathrm{E}Y])^2|\text{not } B] \cdot 0.9 \\ &\geq (25n)^2 \cdot 0.1 = 62.5n^2.\end{aligned}$$

Consequently $\mathrm{Var}[Y/n] \geq 62.5$ and the desired conclusion is proved.

Hence, for large n the probability distribution of the average gain Y/n takes the following form:

- There is a probability of 0.1 that Y/n equals -24.
- There is a probability close to 0.90 that Y/n is in an interval close to 3.78.

If for example $n = 10\,000$, the interval is [3.6, 3.9]. This interval is computed by using the fact that if B does not occur, the expected value and standard deviation (SD) of Y/n equals 3.78 and $\sqrt{61.7/n}$, respectively. Using the Gaussian approximation, the interval $3.78 \pm 1.64 \cdot \mathrm{SD}$ has a 90% probability.

We can conclude that there is a rather high probability of a large loss even if the number of projects is large. The dependence causes the average gain not to converge to the expected value 1.

Should not risk as a concept explicitly reflect this probability of a loss equal to -24? The expected value 1 is not very informative in this case as the distribution has two peaks, -24 and 3.78. Often in real life we may have many such peaks, but the probabilities could be rather small. Other definitions of risk do, however, incorporate this type of distribution, as we will see in the coming chapters.

Different distributions. Extreme observations

Above we have considered projects that are similar: they have the same distribution. In practice we always have different types of projects and some could be very large. To illustrate this, say that we have one project where the possible outcomes are -2400 and 600 and not -24 and 6. The expectation is thus 100 for this project. Then it is obvious that the outcome of this project dominates the total value of the portfolio. The law of large numbers and the central limit theorem cannot be applied. See Figure 1.4 which shows the case with $n = 100$ standard projects with outcomes -24 and 6 and one project with the extreme outcomes -2400 and 600. The probabilities are the same, 1/6 and 5/6, respectively. We observe that the distribution has two peaks, dominated by the extreme project. There is a probability of 1/6 of a negative outcome. If this occurs, the average gain is reduced to about -23. If the extreme project gives a positive result, the average gain is increased to about 7. The 100 standard projects are not sufficiently many to dominate the total portfolio. The computations of the

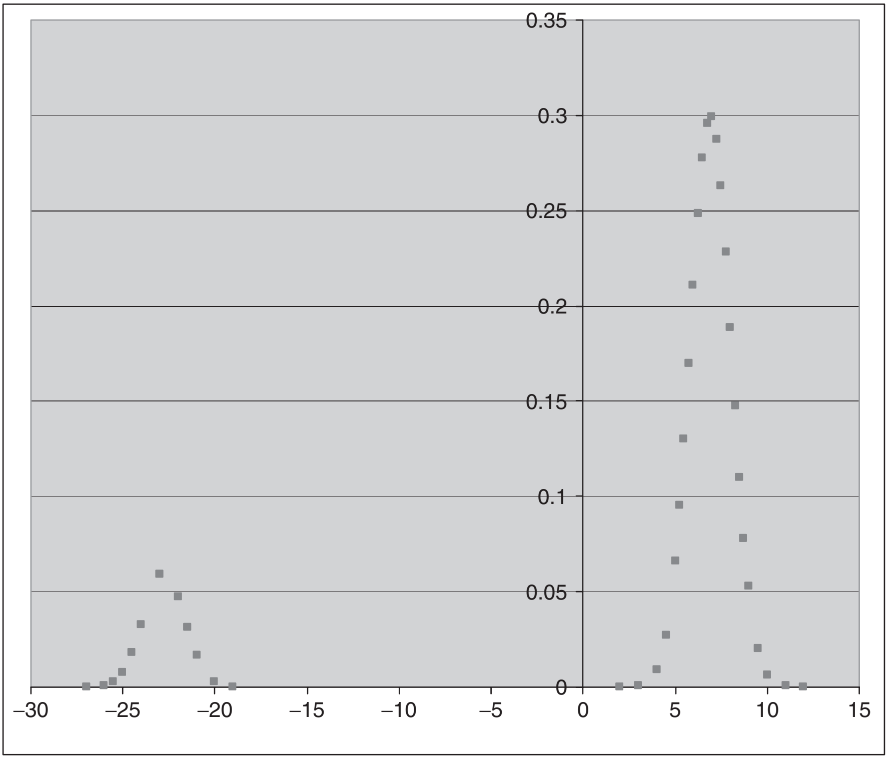

Figure 1.4 The probability distribution for the average gain for n = 100 standard projects and one project with outcomes −2400 and 600. The area under the curve from a to b represents the probability that the gain takes a value in this interval.

numbers in Figure 1.4 are based on the following arguments: If Y denotes the sum of the gains of the $n = 100$ standard projects and Y_L the gain from the extreme project, the task is to compute $P((Y + Y_L)/101 \leq y)$. But this probability can be written as

$$\begin{aligned} P(Y \leq 101y - Y_L) &= P(Y \leq 101y + 2400 | Y_L = -2400) \cdot 1/6 \\ &\quad + P(Y \leq 101y - 600 | Y_L = 600) \cdot 5/6 \\ &= P(Y/100 \leq 1.01y + 24) \cdot 1/6 \\ &\quad + P(Y/100 \leq 1.01y - 6) \cdot 5/6, \end{aligned}$$

and the problem is of the standard form analysed earlier for Y. We have assumed that project gains are independent.

If the number of the standard projects increases, the value of the extreme project is reduced, but it is obvious that some large projects could have a significant influence on the total value of the portfolio. Figures 1.5 and 1.6 are

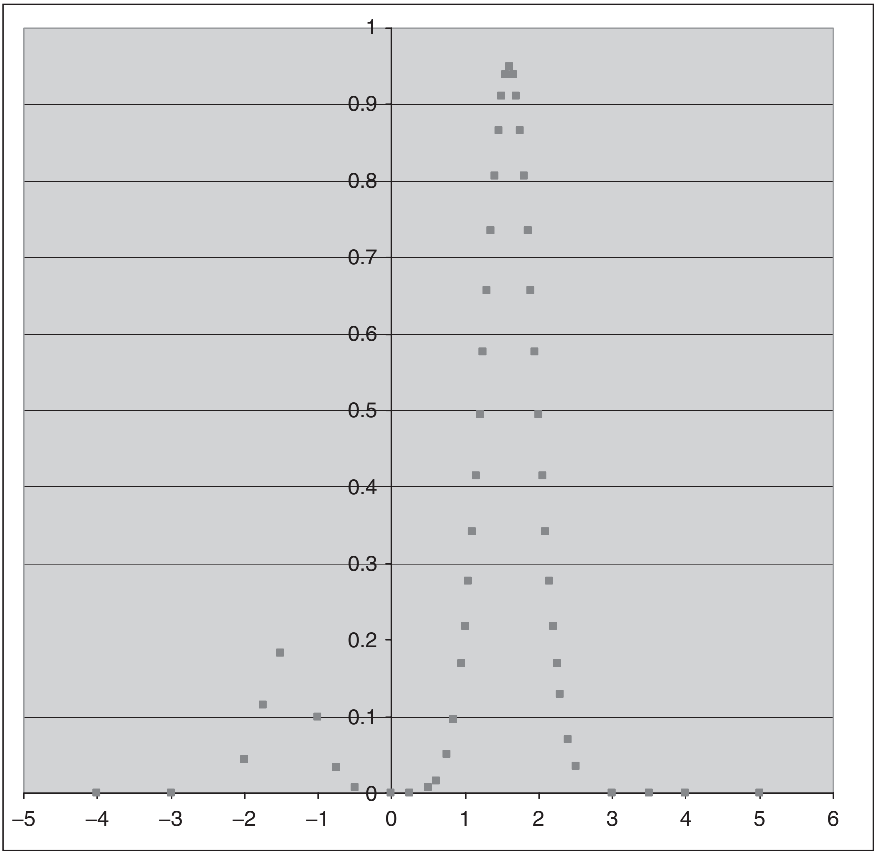

Figure 1.5 The probability distribution for the average gain for n = 1000 and one project with outcomes −2400 and 600. The area under the curve from a to b represents the probability that the gain takes a value in this interval.

similar to Figure 1.4 but $n = 1000$ and 10 000, respectively. The influence of the extreme project is reduced, but for $n = 1000$ the probability of a negative outcome is still about 1/6. However, in the case of $n = 10\,000$, the probability of a negative outcome is negligible. The probability mass is now concentrated around the expected value 1. We see that a very large number of standard projects are required to eliminate the effect of the extreme project.

Difficulties in establishing the probability distribution

In the above analysis there is no discussion about the probability distribution for each project. There is a probability of 1/6 of the negative outcome and a probability of 5/6 of the positive outcome. In practice we seldom have such a distribution available. If you run a business project or you invest your money on

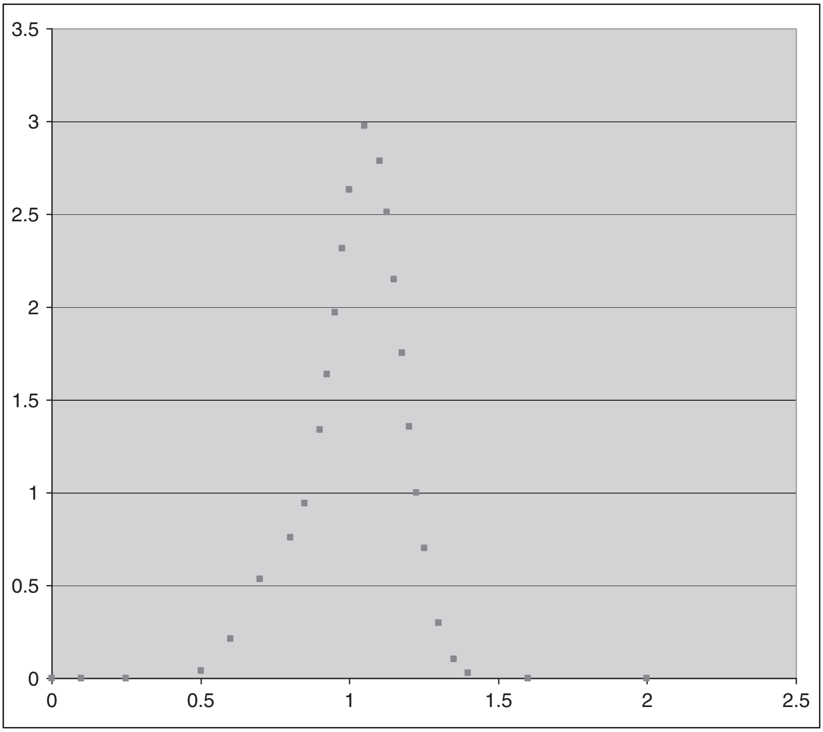

Figure 1.6 The probability distribution for the average gain for n = 10 000 and one project with outcomes −2400 and 600. The area under the curve from a to b represents the probability that the gain takes a value in this interval.

the stock market, it is not obvious how to determine the probability distributions. You may have available some historical data, but to what extent are they relevant for the future? Let us simplify our analysis by considering $n = 10\,000$ future projects and assume that the possible outcomes are −24 and 6 as above, but a probability distribution has not been specified. Historical data are on hand, for similar projects, and these show a distribution between the two possible outcomes which are 1/6 and 5/6 respectively. However, we do not know if this distribution will actually be the distribution for the future projects. We may assume that this will be the case, but we do not know for sure. There are uncertainties – surprises may occur. Under the assumption that the future will be as history indicates, the results are shown by Table 1.3 and Figure 1.3. One may argue that such an assumption is reasonable and that the data are the best information available, but we have to acknowledge that not all uncertainties have been revealed by the numbers shown by Table 1.3 and Figure 1.3. Should not this uncertainty factor be considered an element of risk? We will discuss this in detail in the coming chapter.

We end this chapter with two general reflections on the use of expected values in risk management.

Reflection

In his paper from 1738 referred to above, Daniel Bernoulli presented and discussed the so-called St Petersburg paradox. This problem was first suggested by his cousin Nicolaus Bernoulli. It is based on a casino game where a fair coin is tossed successively until the moment that heads appears for the first time. The casino payoff is 2 ducats if heads comes up on the first toss, 4 ducats if heads turns up for the first time in the second toss, and so on. In general the payoff is 2^n ducats. Thus, with each additional toss, the payoff is doubled. How much should the casino require the player to stake such that, over the long run, the game will not be a losing endeavour for the casino? (Tijms, 2007; Bernstein, 1996).

The answer is that the casino owners should not allow this game to be played, whatever amount a player is willing to stake, as the expected value of the casino payoff for one player is an infinitely large number of ducats. To see that the expected value exceeds every conceivable large value, note first that the probability of getting heads in the first toss is $1/2$, the probability of getting heads for the first time at the second toss equals $1/2 \times 1/2$, and so on and the probability of getting heads for the first time at the nth toss equals $(1/2)^n$. The expected value of the game for one player is thus

$$1/2 \times 2 + (1/2)^2 \times 4 + \ldots + (1/2)^n \times 2^n + \ldots.$$

In this infinite series, a figure of 1 ducat is added to the sum each time, and consequently the sum exceeds any fixed number.

In Bernoulli's day a heated discussion grew up around this problem. Some indicated that the probabilistic analysis was wrong. But no, the mathematics is correct. However, the model is not a good description of the real world. The model tacitly assumes that the casino is always in a position to pay out, whatever happens, even in the case that heads shows up for the first time after say 30 tosses where the payoff is larger than \$1000 million. So a more realistic model is developed if the casino can only pay out up to a limited amount. Then the expected payoff is limited and can be calculated. See Tijms (2007). □

Reflection

Willis (2007) defines terrorism risk as expected damage, as mentioned above. In view of the discussion in this chapter, why is such a perspective problematic?

For terrorism risk, the possible consequences could be extreme with millions of possible fatalities, and consequently the expectation, even on a national and international level, would produce poor predictions of the actual damages and losses. In addition, the uncertainties in underlying phenomena and processes are large. It is extremely difficult to predict when the next attack will occur and what form it will take. The historical data obviously provide limited information. Any assigned probability distribution is likely to deviate strongly from the observed loss distribution.

Summary

These are the two main issues that we should reflect on:

- Is the concept of risk captured by the expected value?
- Should decisions involving risk be based on expected values?

For both questions, the answer is in general *no*. The problem is that the expected value could deviate strongly from the actual outcomes. There are two main reasons for this:

1. The consequences or outcomes could be so extreme that the average of a large population of activities is dominated by these extreme outcomes.
2. The probability distribution could deviate strongly from the future observed outcome distribution.

In gamble-like situations of repeated experiments, the expected value would provide good predictions of the actual future quantities studied, but not often in other situations.

References

Bernstein, P.L. (1996) *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons, Inc., New York.

Floodcite (2006) Guidelines for Socio-economic Flood Damage Evaluation, Report T9-06-01. February. Floodcite.

HM Treasury (2005) Managing Risk to the Public, Appraisal guidance.

Lirer, L., Petrosino, P. and Alberico, I. (2001) Hazard assessment at volcanic fields: the Campi Flegrei case history. *Journal of Volcanology and Geothermal Research*, **112**, 53–73.

Mandel, D. (2007) Toward a Concept of Risk for Effective Military Decision Making, Defence R&D Canada, Toronto, Technical Report 2007-124. DRDC Toronto.

Tijms, H. (2007) *Understanding Probability: Chance Rules in Everyday Life*, 2nd edn, Cambridge University Press, Cambridge.

Verma, M. and Verter, V. (2007) Railroad transportation of dangerous goods: population exposure to airborne toxins. *Computers and Operations Research*, **34**, 1287–1303.

Willis, H.H. (2007) Guiding resource allocations based on terrorism risk. *Risk Analysis*, **27** (3), 597–606.

Further reading

Bernstein, P.L. (1996) Against the gods: *The remarkable story of risk*, Wiley, New York.

Tijms, H. (2007) *Understanding Probability: Chance Rules in Everyday Life*, 2nd ed. Cambridge University Press, Cambridge.