that the decision maker has access to a set of preferences over acts before he starts to deliberate. Now, it can be shown that the expected utility principle can be derived from four simple axioms. The presentation given here is informal, but the sceptical reader can rest assured that the argument can be (and has been) formalised.

We use the term utility for referring both to the value of an act and to the value of its outcomes. The aim of the axiomatisation is to show that the utility of an act equals the expected utility of its outcomes. Now, the *first axiom* holds that if all outcomes of an act have utility u, then the utility of the act is u. In Table 4.3 axiom 1 thus entails that the utility of act a_1 is 5, whereas the utility of act a_2 is 7.

The *second axiom* is the dominance principle: If one act is certain to lead to outcomes with higher utilities under all states, then the utility of the former act exceeds that of the latter (and if both acts lead to equal outcomes they have the same utility). Hence, in Table 4.3 the utility of a_2 exceeds that of a_1. Note that this axiom requires that states are causally independent of acts. In Chapter 9 we discuss a type of decision problem for which this assumption does not hold true. The present axiomatisation thus supports the expected utility principle only in a restricted class of decision problems.

The *third axiom* holds that every decision problem can be transformed into a decision problem with equiprobable states, by splitting the original states into parallel ones, without affecting the overall utility of any of the acts in the decision problem; see Table 4.4.

The gist of this axiom is that a_1 and a_2 in the leftmost matrix are exactly as good as a_1 and a_2 in the rightmost matrix, simply because the second

Table 4.3

	s_1	s_2	s_3	s_4	s_5
a_1	5	5	5	5	5
a_2	7	7	7	7	7

Table 4.4

	0.2	0.8		0.2	0.2	0.2	0.2	0.2
a_1	1	3	a_1	1	3	3	3	3
a_2	6	2	a_2	6	2	2	2	2

matrix can be obtained from the first by dividing the set of states corresponding to the outcomes slightly differently.

The *fourth* and last axiom is a trade-off principle. It holds that if two outcomes are equally probable, and if the best outcome is made slightly worse, then this can be compensated for by adding *some* (perhaps very large) amount of utility to the other outcome. Imagine, for instance, that Adam offers you to toss a fair coin. If it lands heads up you will be given 10 units of utility, otherwise you receive 2 units. If you refuse to take part in the gamble you receive 5 units. Before you decide whether to gamble or not, Adam informs you that he is willing to change the rules of the gamble such that instead of giving you 10 units of utility if the coin lands heads up he will give you a little bit less, $10 - \varepsilon_1$, but compensate you for this potential loss by increasing the other prize to $2 + \varepsilon_2$ units (Table 4.5). He adds that you are free to choose the value of ε_2 yourself! The fourth axiom does not say anything about whether you should choose 5 units for sure instead of the gamble yielding either 2 or 10 units of utility, or vice versa. Such choices must be determined by other considerations. The axiom only tells you that there is some number $\delta > 0$, such that for all ε_1, $0 \leq \varepsilon_1 \leq \delta$, there is a number ε_2 such that the trade-off suggested by Adam is unimportant to you, i.e. the utility of the original and the modified acts is the same.

If a sufficiently large value of ε_2 is chosen, even many risk-averse decision makers would accept the suggested trade-off. This means that this axiom can be accepted by more than just decision makers who are neutral to risk-taking. However, this axiom is nevertheless more controversial than the others, because it implies that once ε_1 and ε_2 have been established, these constants can be added over and over again to the utility numbers representing this pair of outcomes. Put in mathematical terms, it is assumed that ε_2 is a function of ε_1, but not of the initial levels of utility. (The axiom can be weakened, however, such that ε_2 becomes a function of more features of the decision problem, but it is beyond the scope of this book to explore this point any further here.)

Table 4.5

	0.5	0.5		0.5	0.5
a_1	5	5	a_1	5	5
a_2	2	10	a_2	$2 + \varepsilon_2$	$10 - \varepsilon_1$

The axioms informally outlined above together entail that the utility of an act equals the expected utility of its outcomes. Or, put in slightly different words, the act that has the highest utility (is most attractive) will also have the highest expected utility, and vice versa. This appears to be a strong reason for letting the expected utility principle guide one's choices in decisions under risk. A more stringent formulation of this claim and a proof is provided in Box 4.2.

Box 4.2 A direct axiomatisation of the expected utility principle

Consider the following four axioms.

EU 1 If all outcomes of an act have utility u, then the utility of the act is u.

EU 2 If one act is certain to lead to better outcomes under all states than another, then the utility of the first act exceeds that of the latter; and if both acts lead to equal outcomes they have the same utility.

EU 3 Every decision problem can be transformed into a decision problem with equally probable states, in which the utility of all acts is preserved.

EU 4 If two outcomes are equally probable, and if the better outcome is made slightly worse, then this can be compensated for by adding some amount of utility to the other outcome, such that the overall utility of the act is preserved.

Theorem 4.1 Let axioms EU 1–4 hold for all decision problems under risk. Then, the utility of an act equals its expected utility.

Proof The proof of Theorem 4.1 consists of two parts. We first show that $\varepsilon_1 = \varepsilon_2$ (see page 76) whenever EU 4 is applied. Consider the three decision problems in Table 4.6, in which u_1 and u_2 are some utility levels such that u_1 is higher than u_2, while their difference is less than ε_1. (That is, $u_1 - u_2 < \varepsilon_1$.)

Table 4.6

	s	s'		s	s'		s	s'
a_1	u_1	u_2	a_1	u_1	u_2	a_1	u_1	u_2
a_2	u_1	u_2	a_2	$u_1 - \varepsilon_1$	$u_2 + \varepsilon_2$	a_2	$u_1 - \varepsilon_1 + \varepsilon_2$	$u_2 + \varepsilon_2 - \varepsilon_1$

In the leftmost decision problem a_1 has the same utility as a_2, because of EU 2. The decision problem in the middle is obtained by applying EU 4 to act a_2. Note that the utility of both acts remains the same. Finally, the rightmost decision problem is obtained from the one in the middle by applying EU 4 to a_2 again. The reason why ε_1 is subtracted from $u_2 + \varepsilon_2$ is that the utility of the rightmost outcome of a_2 now exceeds that of the leftmost, since the difference between u_1 and u_2 was assumed to be less than ε_1. By assumption, the utility of both acts has to remain the same, which can only be the case if $\varepsilon_1 = \varepsilon_2$. To see why, assume that it is not the case that $\varepsilon_1 = \varepsilon_2$. EU 2 then entails that either a_2 dominates a_1, or a_1 dominates a_2, since $-\varepsilon_1 + \varepsilon_2 = \varepsilon_2 - \varepsilon_1$.

In the second step of the proof we make use of the fact that $\varepsilon_1 = \varepsilon_2$ whenever EU 4 is applied. Let D be an arbitrary decision problem. By applying EU 3 a finite number of times, D can be transformed into a decision problem D* in which all states are equally probable. The utilities of all acts in D* are equal to the utility of the corresponding acts in D. Then, by adding a small amount of utility ε_1 to the lowest utility of a given act and at the same time subtracting the same amount from its highest utility (as we now know we are allowed to do), and repeating this operation a finite number of times, we can ensure that all utilities of each act over all the equally probable states will be equalised. Since all states are equally probable, and we always withdraw and add the same amounts of utilities, the expected utility of each act in the modified decision problem will be exactly equal to that in the original decision problem. Finally, since all outcomes of the acts in the modified decision problem have the same utility, say u, then the utility of the act is u, according to EU 1. It immediately follows that the utility of each act equals its expected utility. □

4.4 Allais' paradox

The expected utility principle is by no means uncontroversial. Naturally, some objections are more sophisticated than others, and the most sophisticated ones are referred to as paradoxes. In the following sections we shall discuss a selection of the most thought-provoking paradoxes. We start with Allais' paradox, which was discovered by the Nobel Prize winning economist Maurice Allais. In the contemporary literature, this paradox is directed both against the expected utility principle in general, as well as against one