4 Decisions under risk

Decisions under risk differ from decisions under ignorance in that the decision maker knows the probabilities of the outcomes. If you play roulette in Las Vegas you are making a decision under risk, since you then know the probability of winning and thus how much you should expect to lose. However, that you *know* the probability need not mean that you are in a position to immediately determine the probability or expected outcome. It is sufficient that you have enough information for figuring out the answer after having performed a series of calculations, which may be very complex. In this process your 'tacit knowledge' is made explicit. When you play roulette in Las Vegas and bet on a single number, the probability of winning is 1/38: There are 38 equally probable outcomes of the game, viz. 1–36 and 0 and 00, and if the ball lands on the number you have betted on the croupier will pay you 35 times the amount betted, and return your bet. Hence, if you bet \$1, the expected payout is $\$(35+1)\cdot\frac{1}{38}+\$0\cdot\frac{37}{38}=\$\frac{36}{38}\approx\$0.947$. This means that you can expect to lose about $\$1-\$0.947=\$.053$ for every dollar betted.

According to the principle of maximising expected *monetary value* it would obviously be a mistake to play roulette in Las Vegas. However, this does not show that it is *irrational* to play roulette there, all things considered. First, the expected monetary value need not correspond to the overall value of a gamble. Perhaps you are very poor and desperately need to buy some medicine that costs \$35. Then it would make sense to play roulette with your last dollar, since that would entitle you to a chance of winning just enough to buy the medicine. Second, it also seems clear that many people enjoy the sensation of excitement caused by betting. To pay for this is not irrational. (It is just vulgar!) Finally, one may also question the principle of maximising expected value as a general guide to risky choices. Is this really the correct way of evaluating risky acts?

In what follows we shall focus on the last question, i.e. we shall discuss whether it makes sense to think that the principle of maximising expected value is a reasonable decision rule to use in decisions under risk. Somewhat surprisingly, nearly all decision theorists agree that this is the case. There are no serious contenders. This is thus a significant difference compared to decision making under ignorance. As explained in Chapter 3, there is virtually no agreement on how to make decisions under ignorance.

That said, there is significant disagreement among decision theorists about how to *articulate* the principle of maximising expected value. The main idea is simple, but substantial disagreement remains about how to define central concepts such as 'value' and 'probability', and how to account for the causal mechanisms producing the outcomes. In this chapter we shall take a preliminary look at some aspects of these controversies, but it will take several chapters before we have acquired a comprehensive understanding of the debate.

It is worth noting that many situations outside the casino, i.e. in the 'real' world, also involve decision making under risk. For example, if you suffer from a heart disease and your doctor offers you a transplant giving you a 60% chance of survival, you are facing a decision under risk in which it seems utterly important to get the theoretical aspects of the decision right. It would thus be a mistake to think that decision making under risk is essentially linked to gambling. Gambling examples just happen to be a convenient way of illustrating some of the major ideas and arguments.

4.1 Maximising what?

The principles of maximising *expected monetary value* must not be confused with the principle of maximising *expected value*. Money is not all that matters, at least not to all of us. However, in addition to this distinction, we shall also introduce a new distinction between the principle of maximising *expected value* and the principle of maximising *expected utility*. The latter is a more precise version of the former, in which the notion of value is more clearly specified. This gives us three closely related principles, all of which are frequently mentioned in the literature.

1. The principle of maximising expected monetary value
2. The principle of maximising expected value
3. The principle of maximising expected utility

It is helpful to illustrate the difference between these principles in an example. Imagine that you are offered a choice between receiving a million dollars for sure, and receiving a lottery ticket that entitles you to a fifty per cent chance of winning either three million dollars or nothing (Table 4.1).

The expected monetary value (EMV) of these lotteries can be computed by applying the following general formula, in which p_1 is the probability of the first state and m_1 the monetary value of the corresponding outcome:

$$\text{EMV} = p_1 \cdot m_1 + p_2 \cdot m_2 + \cdots + p_n \cdot m_n \tag{1}$$

By applying (1) to our example, we find that the expected monetary values of the two lotteries are:

$$\text{EMV(Lottery A)} = \frac{1}{2} \cdot \$1\text{M} + \frac{1}{2} \cdot \$1\text{M} = \$1\text{M}.$$

$$\text{EMV(Lottery B)} = \frac{1}{2} \cdot \$3\text{M} + \frac{1}{2} \cdot \$0 = \$1.5\text{M}$$

However, even though EMV(Lottery B) > EMV(Lottery A), many of us would prefer a million for sure. The explanation is that the overall value to us of \$3 M is just slightly higher than that of \$1 M, whereas the value of \$1 M is much higher than the value of \$0. Economists say that the *marginal value* of money is decreasing. The graph in Figure 4.1 describes a hypothetical relationship between money and value, for a poor person playing the National Lottery.

Note that the graph slopes upwards but with decreasing speed. This means that winning more is always better than winning less. However, the more one wins, the lower is the value of winning yet another million. That said, it is of course not a universal truth that the marginal value of

Table 4.1

	1/2	1/2
Lottery A	\$1 M	\$1 M
Lottery B	\$3 M	\$0

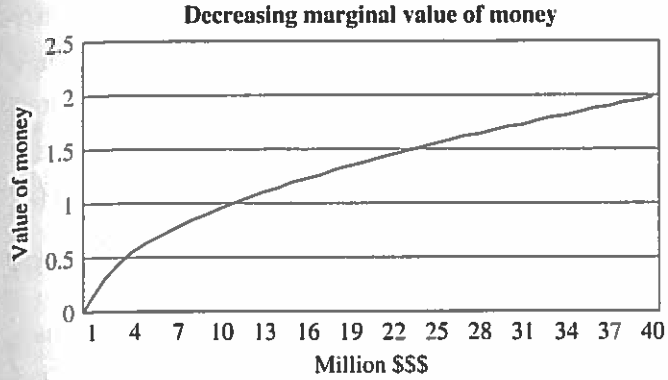

Figure 4.1

money is decreasing. For people with very expensive habits the marginal value may be increasing, and it is even conceivable that it is negative for some people. Imagine, for instance, that you are a multi-billionaire. If you accidentally acquire another billion this will perhaps decrease your well-being, since more money makes it more likely that you will get kidnapped and you cannot stand the thought of being locked up in a dirty basement and having your fingers cut off one by one. If one is so rich that one starts to fear kidnappers, it would perhaps be better to get rid of some of the money.

Clearly, the principle of maximising expected value makes more sense from a normative point of view than the principle of maximising expected monetary value. The former is obtained from the latter by replacing m for v in the formula above, where v denotes value rather than money.

$$\mathrm{EV} = p_1 \cdot v_1 + p_2 \cdot v_2 + \cdots + p_n \cdot v_n \tag{2}$$

Unfortunately, not all concepts of value are reliable guides to rational decision making. Take moral value, for instance. If a billionaire decides to donate his entire fortune to charity, the expected moral value of doing so might be very high. However, this is because many poor people would benefit from the money, not because the billionaire himself would be any happier. (By assumption, this billionaire is very greedy!) The expected moral value of donating a fortune is far higher than the sort of personal value decision theorists are primarily concerned with. In most cases, moral value is not the sort of value that decision theorists think we should base instrumental, ends-means reasoning on. Therefore, in order to single out the kind of value that is the primary object of study in decision theory – the value of an outcome as evaluated from the decision maker's point of view – it is helpful to introduce the concept of *utility*. Utility is an abstract entity that cannot be directly observed. By definition, the utility of an outcome depends on how valuable the outcome is from the decision maker's point of view. The principle of maximising expected utility is obtained from the principle of maximising expected value by replacing v for u in equation (2).

$$EU = p_1 \cdot u_1 + p_2 \cdot u_2 + \cdots + p_n \cdot u_n \tag{3}$$

In the remainder of this chapter we shall focus on the principle of maximising expected utility, rather than any other versions of the expectation thesis. It is worth noticing that the expected utility principle can be applied also in cases in which outcomes are non-monetary. Consider the following example. David and his partner Rose are about to deliver a sailing yacht from La Coruña (in the north of Spain) to English Harbour, Antigua (in the West Indies). Because of the prevailing weather systems, there are only two feasible routes across the Atlantic, either a direct northern route or a slightly longer southern route. Naturally, the couple wish to cross the Atlantic as quickly as possible. The number of days required for the crossing depends on the route they choose and the meteorological situation. Weather-wise, the decisive factor is whether or not a high pressure zone develops over the Azores after they have set off from the coast of Spain. There are reliable meteorological data going back more than a hundred years, and the probability that a high pressure zone will develop is 83%. By studying the meteorological data and the charts, they figure out that the decision problem they are facing is that shown in Table 4.2.

Since David and Rose wish to make the crossing in as few days as possible, the utility of the outcomes is *negatively correlated* with the number of days at sea. Hence, the utility of sailing for 27 days, which we write as $u(27)$, is lower than the utility of sailing for 18 days, $u(18)$. For simplicity, we assume that in this particular case the utility function is linear with respect

Table 4.2

	High pressure zone over the Azores (83%)	No high pressure zone over the Azores (17%)
Northern route	27 days	14 days
Southern route	18 days	21 days

to the number of days spent at sea. It then follows that the expected utilities of the two alternatives are as follows.

$$\text{EU(Northern route)} = 0.83 \cdot u(27) + 0.17 \cdot u(14) = u(24.79)$$

$$\text{EU(Southern route)} = 0.83 \cdot u(18) + 0.17 \cdot u(21) = u(18.51)$$

Clearly, David and Rose ought to choose the southern route, since $u(18.51) > u(24.79)$, according to our assumption about the correlation between utility and the number of days spent at sea.

Box 4.1 A risky decision

Placebo Pharmaceuticals is eager to expand its product portfolio with a drug against cardio-vascular diseases. Three alternative strategies have been identified. The first is to hire a research team of 200 people to develop the new drug. However, to develop a new drug is expensive (about \$50M) and it is also far from certain that the team will manage to successfully develop a drug that meets the regulatory requirements enforced by the Food and Drug Administration; the probability is estimated to be about one in ten. The second alternative is to acquire a small company, Cholesterol Business Inc., that has already developed a drug that is currently undergoing clinical trials. If the trials are successful the Food and Drug Administration will of course license the product rather rapidly, so this alternative is more likely to be successful. According to the executive director of Placebo Pharmaceuticals the probability of success is about 0.8. The downside is that the cost of taking over Cholesterol Inc. is very high, about \$120M, since several other big pharmaceutical companies are also eager to acquire the company. The third alternative is to simply buy a licence for \$170M from a rival company to produce and market an already existing drug. This is the least risky option, since the board of Placebo Pharmaceuticals knows for sure for what it pays. Finally, to complete the list of alternatives also note that there is a fourth alternative: to do nothing and preserve the status quo.

In order to make a rational decision, Placebo Pharmaceuticals decides to hire a decision analyst. After conducting series of interviews with the board members the decision analyst is able to establish that Placebo Pharmaceuticals' utility of a profit greater than zero is linear and directly proportional to the profit, whereas its disutility of losses L (i.e. a profit equal to or smaller than zero) is determined by the formula $u = 2 \cdot L$. The

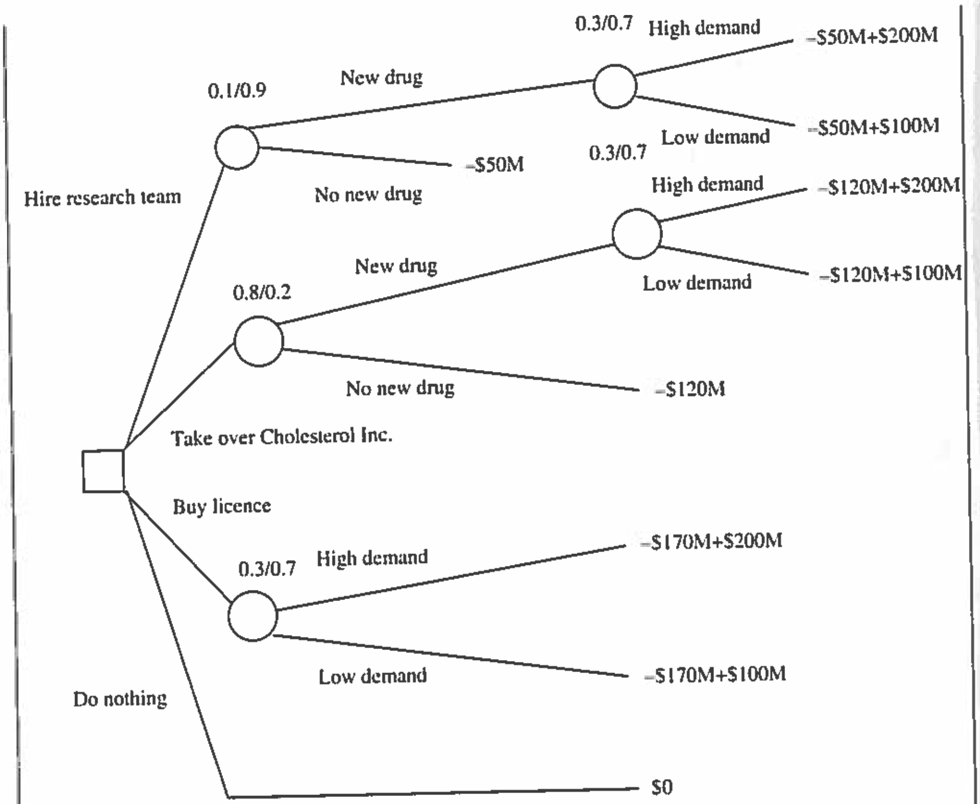

Figure 4.2

decision analyst also concludes that the revenue will depend on future demand of the new drug. The probability that demand will be high is 0.3, in which case the revenue will be $200M. The probability that demand is low is 0.7, and the revenue will then be $100M. To explain her findings to the executive director the decision analyst draws the decision tree shown in Figure 4.2.

It can be easily verified that the four alternatives illustrated in the decision tree may lead to nine different outcomes.

Hire research team:	+$150M with a probability of $0.1 \cdot 0.3$
	+$50M with a probability of $0.1 \cdot 0.7$
	–$50M with a probability of 0.9
Take over Cholesterol Inc.:	+$80M with a probability of $0.8 \cdot 0.3$
	–$20M with a probability of $0.8 \cdot 0.7$
	–$120M with a probability of 0.2
Buy licence:	+$30M with a probability of 0.3
	–$70M with a probability of 0.7
Do nothing:	$0 with a probability of 1

In order to reach a decision, the executive director now applies the principle of maximising expected utility. She recalls that the utility of losses is $u = 2 \cdot L$.

Hire research team:	$+150\text{M} \cdot 0.03 + 50\text{M} \cdot 0.07 - (2 \cdot 50\text{M}) \cdot 0.9 = -82\text{M}$ utility units
Take over Cholesterol Inc.:	$+80\text{M} \cdot 0.24 - (2 \cdot 20\text{M}) \cdot 0.56 - (2 \cdot 120\text{M}) \cdot 0.2 = -51.2\text{M}$ utility units
Buy licence:	$+30\text{M} \cdot 0.3 - (2 \cdot 70\text{M}) \cdot 0.7 = -89\text{M}$ utility units
Do nothing:	0 utility units

Based on the calculations above, the executive director of Placebo Pharmaceuticals decides that the rational thing to do is to do nothing, i.e. to abandon the plan to expand the product portfolio with a drug against cardio-vascular diseases.

4.2 Why is it rational to maximise expected utility?

Decision theorists have proposed two fundamentally different arguments for the expected utility principle. The first argument is based on the law of large numbers; it seeks to show that in *the long run* you will be better off if you maximise expected utility. The second argument aims at deriving the expected utility principle from some more fundamental axioms for rational decision making, which make no reference to what will happen in the long run. We shall return to the axiomatic approach in the next section.

The law of large numbers is a mathematical theorem stating that everyone who maximises expected utility will almost certainly be better off in the long run. In this context the term 'almost certainly' has a very precise meaning. If a random experiment (such as rolling a die or tossing a coin) is repeated n times and each experiment has a probability p of leading to a predetermined outcome, then the probability that the percentage of such outcomes differs from p by more than a very small amount ε converges to 0 as the number of trials n approaches infinity. This holds true for every $\varepsilon > 0$, no matter how small. Hence, by performing the random experiment sufficiently many times, the probability that the average outcome differs from the expected outcome can be rendered *arbitrarily* small.

Imagine, for instance, that you are offered 1 unit of utility for sure or a lottery ticket that will yield either 10 units with a probability of 0.2, or nothing with a probability of 0.8. The expected utility of choosing the lottery ticket is 2, which is more than 1. According to the law of large numbers, you cannot be sure that you will *actually* be better off if you choose the lottery, given that the choice is offered to you only once. However, what you do know for sure is that *if* you face the same decision over and over again, *then* the probability that you will not be better off by choosing the lottery can be made arbitrarily small by repeating the same decision over and over again. Furthermore, if you repeat the decision in question infinitely many times, then the probability that the average utility and the expected utility differ by more than ε units decreases to zero.

Keynes famously objected to the law of large numbers that, "in the long run we are all dead" (1923: 80). This claim is no doubt true, but what should we say about its relevance? Keynes' point was that no real-life decision maker will ever face any decision an infinite number of times; hence, mathematical facts about what would happen after an infinite number of repetitions are therefore of little normative relevance. A different way to express this concern is the following: Why should one care about what *would* happen if some condition were to be fulfilled, given that one knows for sure at the time of making the decision that this condition is certain not to be fulfilled? Personally I think Keynes was right in questioning the relevance of the law of large numbers. But perhaps there is some way in which it could be saved?

It is also worth pointing out that the relevance of the law of large numbers is partially defeated by another mathematical theorem, known as *gambler's ruin*. Imagine that you and I flip a fair coin, and that I pay you $1 every time it lands heads up, and you pay me the same amount when it lands tails up. We both have $1,000 in our pots as we start to play. Now, if we toss the coin *sufficiently* many times each player will at some point encounter a sequence of heads or tails that is longer than he can afford, i.e. longer than the number of one dollar bills in his pot. If you encounter that very long sequence first, you will go bankrupt. Otherwise I will go bankrupt first. It is mathematically impossible that both of us can 'survive' infinitely many rounds of this game, given that each player starts with pots containing finite amounts of money. This means that the law of large numbers guarantees that you will be better off in the long run by maximising expected utility *only if* your initial pot is infinite, which is a rather unrealistic assumption.

An additional worry about the law of large numbers is that it seems perfectly reasonable to question whether decision makers ever face *the same* decision problem several times. Even if you were to play roulette in a Las Vegas casino for weeks, it seems obvious that each time the croupier drops the ball on the roulette wheel she will do it a little bit differently each time, and to some extent it also seems reasonable to claim that the physical constitution of the wheel will change over time, because of dust and wear and tear. Hence, it is not *exactly* the same act you perform every time you play.

That said, for the law of large numbers to work it is strictly speaking not necessary to assume that the agent is facing the same decision problem, in a literal sense, over and over again. All we need to assume is that the *probability* of each outcome is fixed over time. Note, however, that if this is the correct way of understanding the argument, then it will become sensitive to the definition of probability, to be discussed in Chapter 7. According to some theories of probability, the probability that you win when playing roulette is the same over time, but according to other theories this need not necessarily be the case. (For example, if probabilities are defined as relative frequencies, or as subjective degrees of beliefs, then they are very likely to change over time.)

A final worry about the relevance of the law of large numbers is that many decisions under risk are unique in a much stronger sense. It might very well hold true that the probability that John would become happy if he was to marry his partner Joanna is 95%. But so what? He will only marry Joanna once (or at least a very limited number of times). Why pay any attention to the law of large numbers in this decision? The same remark seems relevant in many other unique decisions, i.e. decisions that are made only once, such as a decision to start a war, or appointing a chief executive, or electing a new president. For instance, if the probability is high that the republican candidate will win the next presidential election, the expected utility of investing in the defence industry might be high. However, in this case we cannot justify the expected utility principle by appealing to the law of large numbers, because every presidential election is unique.

4.3 The axiomatic approach

The axiomatic approach to the expected utility principle is not based on the law of large numbers. Instead, this approach seeks to show that the expected utility principle can be derived from axioms that hold independently of what would happen in the long run. If successful, an axiomatic argument can thus overcome the objection that it would not make sense to maximise expected utility in a decision made only once. Here is an extreme example illustrating this point. You are offered to press a green button, and if you do, you will either die or become the happiest person on Earth. If you do not press the button, you will continue to live a rather mediocre life. Let us suppose that the expected utility of pressing the green button exceeds that of not pressing it. Now, what should a rational decision maker do? Axiomatic arguments should entail, if successful, that one should maximise expected utility even in this case, even though the decision is taken only once and the outcome may be disastrous.

Decision theorists have proposed two fundamentally different strategies for axiomatising the expected utility principle. Some axiomatisations are direct, and some are indirect. In the indirect approach, which is the dominant approach, the decision maker does not prefer a risky act to another *because* the expected utility of the former exceeds that of the latter. Instead, the decision maker is asked to state a set of preferences over a set of risky acts. It is irrelevant *how* these preferences are generated. Then, if the set of preferences stated by the decision maker is consistent with a number of structural constraints (axioms), it can be shown that her decisions can be described *as if* she were choosing what to do by assigning numerical probabilities and utilities to outcomes and then maximising expected utility. For an example of such a structural constraint, consider the plausible idea that if act *A* is judged to be better than act *B*, then it must not be the case that *B* is judged to be better than *A*. Given that this constraint is satisfied, as well as a number of more complex and controversial constraints, it is possible to assign numbers representing hypothetical probabilities and utilities to outcomes in such a way that the agent prefers one act over another if and only if the hypothetical expected utility attributed to that alternative is higher than that of all alternatives. A detailed overview of some influential axiomatic constraints on preferences will be given in Chapters 5 and 7.

In the remainder of this section we focus on the direct approach. It is easier to understand, although it should be stressed that it is less influential in the contemporary literature. The direct approach seeks to generate preferences over acts from probabilities and utilities *directly* assigned to outcomes. In contrast to the indirect approach, it is not assumed 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

Table 4.7

	Ticket no. 1	Ticket no. 2–11	Ticket no. 12–100
Gamble 1	\$1M	\$1M	\$1M
Gamble 2	\$0	\$5M	\$1M
Gamble 3	\$1M	\$1M	\$0
Gamble 4	\$0	\$5M	\$0

of the axioms frequently used in (indirect) axiomatisation of it. In this section we shall conceive of the paradox as a general argument against the expected utility principle. Consider the gambles in Table 4.7, in which exactly one winning ticket will be drawn at random.

In a choice between Gamble 1 and Gamble 2 it seems reasonable to choose Gamble 1 since it gives the decision maker one million dollars for sure (\$1M), whereas in a choice between Gamble 3 and Gamble 4 many people would feel that it makes sense to trade a ten-in-hundred chance of getting \$5M, against a one-in-hundred risk of getting nothing, and consequently choose Gamble 4. Several empirical studies have confirmed that most people reason in this way. However, no matter what utility one assigns to money, the principle of maximising expected utility recommends that the decision maker prefers Gamble 1 to Gamble 2 if and only if Gamble 3 is preferred to Gamble 4. There is simply no utility function such that the principle of maximising utility is consistent with a preference for Gamble 1 to Gamble 2 *and* a preference for Gamble 4 to Gamble 3. To see why this is so, we calculate the *difference* in expected utility between the two pairs of gambles. Note that the probability that ticket 1 will be drawn is 0.01, and the probability that one of the tickets numbered 2–11 will be drawn is 0.1; hence, the probability that one of the tickets numbered 12–100 will be drawn is 0.89. This gives the following equations:

$$\begin{aligned} u(G1) - u(G2) &= u(1M) - [0.01u(0M) + 0.1u(5M) + 0.89u(1M)] \\ &= 0.11u(1M) - [0.01u(0) + 0.1u(5M)] \end{aligned} \tag{1}$$

$$\begin{aligned} u(G3) - u(G4) &= [0.11u(1M) + 0.89u(0)] - [0.9u(0M) + 0.1u(5M)] \\ &= 0.11u(1M) - [0.01u(0) + 0.1u(5M)] \end{aligned} \tag{2}$$

Equations (1) and (2) show that the difference in expected utility between G1 and G2 is precisely the same as the difference between G3 and G4. Hence, no matter what the decision maker's utility for money is, it is impossible to simultaneously prefer G1 to G2 *and* to prefer G4 to G3 without violating the expected utility principle. However, since many people who have thought very hard about this example still feel it would be rational to stick to the problematic preference pattern described above, there seems to be something wrong with the expected utility principle.

Unsurprisingly, a number of decision theorists have tried to find ways of coping with the paradox. Savage, a pioneer of modern decision theory, made the following point:

> if one of the tickets numbered from 12 through 100 is drawn, it does not matter, in either situation which gamble I choose. I therefore focus on the possibility that one of the tickets numbered from 1 through 11 will be drawn, in which case [the choice between G1 and G2 and between G3 and G4] are exactly parallel ... It seems to me that in reversing my preference between [G3 and G4] I have corrected an error. (Savage 1954: 103)

Savage's point is that it does not matter which alternative is chosen under states that yield the same outcomes, so those states should be ignored. Instead, decision makers should base their decisions entirely on features that differ between alternatives. This idea is often referred to as the sure-thing principle; we will discuss it in more detail in Chapter 7. That said, some people find the sure-thing principle very hard to accept, and argue that this principle is the main target of the paradox. In their view, Savage has failed to explain *why* sure-thing outcomes should be ignored.

Another type of response to Allais' paradox is to question the accuracy of the formalisation of the decision problem. The outcome of getting $0 in G2 is very different from the outcome of getting $0 in G4. The disappointment one would feel if one won nothing instead of a fortune in G2 is likely to be substantial. This is because in the choice between G1 and G2 the first alternative is *certain* to yield a fortune, whereas in the choice between G3 and G4 no alternative is certain to yield a fortune. A more accurate decision matrix would therefore look as in Table 4.8. Note that it no longer holds true that the expected utility principle is inconsistent with the preference pattern people actually entertain.

A drawback of this response is that it seems difficult to tell exactly how fine-grained the description of outcomes ought to be. In principle, it seems that *every* potential violation of the expected utility principle could be

Table 4.8

	Ticket no. 1	Ticket no. 2–11	Ticket no. 12–100
Gamble 1	\$1M	\$1M	\$1M
Gamble 2	\$0 and disappointment	\$5M	\$1M
Gamble 3	\$1M	\$1M	\$0
Gamble 4	\$0	\$5M	\$0

rejected by simply making the individuation of outcomes more fine-grained. However, this would make the principle immune to criticism, unless one has some independent reason for adjusting the individuation of outcomes.

4.5 Ellsberg's paradox

This paradox was discovered by Daniel Ellsberg when he was a Ph.D. student in economics at Harvard in the late 1950s. Suppose the decision maker is presented with an urn containing 90 balls, 30 of which are red. The remaining 60 balls are either black or yellow, but the proportion between black and yellow balls is unknown. The decision maker is then offered a choice between the following gambles:

Gamble 1 Receive \$100 if a red ball is drawn
Gamble 2 Receive \$100 if a black ball is drawn

When confronted with these gambles you may reason in at least two different ways. First, you may argue that it would be better to choose G1 over G2, since the proportion of red balls is known for sure whereas one knows almost nothing about the number of black balls in the urn. Second, you may believe that there are in fact many more black than red balls in the urn, and therefore choose G2. The paradox will arise no matter how you reason, but for the sake of the argument we assume that you prefer G1 to G2. Now, after having made a choice between G1 and G2 you are presented with a second set of gambles.

Gamble 3 Receive \$100 if a red or yellow ball is drawn
Gamble 4 Receive \$100 if a black or yellow ball is drawn.

Do you prefer G3 or G4? All four gambles are illustrated in Table 4.9. When confronted with the new pair of gambles, it seems that a person who

Table 4.9

	30	60	
	Red	Black	Yellow
Gamble 1	\$100	\$0	\$0
Gamble 2	\$0	\$100	\$0
Gamble 3	\$100	\$0	\$100
Gamble 4	\$0	\$100	\$100

prefers G1 to G2 is likely to prefer G4 to G3, since G4 is a gamble with *known probabilities*. The probability of winning \$100 in G4 is known for sure to be 60/90.

The point of Ellsberg's example is the following. No matter what the decision maker's utility for money is, and no matter what she believes about the proportion of black and yellow balls in the urn, the principle of maximising expected utility can never recommend G1 over G2 *and* G4 over G3, or vice versa. This is because the expected utility of G1 exceeds that of G2 if and only if the expected utility of G3 exceeds that of G4. To show this, we calculate the difference in expected utility between G1 and G2, as well as the difference between G3 and G4. For simplicity, we assume that the utility of \$100 equals M and that the utility of \$0 equals 0 on your personal utility scale. (Since utility is measured on an interval scale, these assumptions are completely innocent.) Hence, if you believe that there are B black balls in the urn, the difference in expected utility between the gambles is as follows.

$$eu(G1) - eu(G2) = 30/90M - B/90M = 30M - BM$$

$$\begin{aligned} eu(G3) - eu(G4) &= 30/90M + (60 - B)/90M - 60/90M \\ &= 30M + (60 - B)M - 60M = 30M - BM \end{aligned}$$

Note that the paradox cannot be avoided by simply arguing that G2 ought to be preferred over G1, because then it would presumably make sense to also prefer G3 over G4; such preferences indicate that the decision maker seeks to avoid gambles with known probabilities. As shown above, eu(G1) – eu(G2) = eu(G3) – eu(G4), so G2 can be preferred over G1 if and only if G4 is preferred over G3.

The Ellsberg paradox is in many respects similar to the Allais paradox. In both paradoxes, it seems reasonable to violate Savage's sure-thing principle; that is, it does not make sense to ignore entire states just because they have parallel outcomes. However, the *reason* why it seems plausible to take into account outcomes that occur for sure under some states is different. In the Allais paradox, G1 is better than G2 because it guarantees that one gets a million dollars. In the Ellsberg paradox G1 is better than G2 because one knows the exact probability of winning $100 although no alternative is certain to lead to a favourable outcome. Arguably, this shows that the intuitions that get the paradoxes going are fundamentally different. The Ellsberg paradox arises because we wish to avoid epistemic uncertainty about probabilities, whereas the Allais paradox arises because we wish to avoid uncertainty about outcomes.

4.6 The St Petersburg paradox

The St Petersburg paradox is derived from the St Petersburg game, which is played as follows. A fair coin is tossed until it lands heads up. The player then receives a prize worth 2^n units of utility, where n is the number of times the coin was tossed. So if the coin lands heads up in the first toss, the player wins a prize worth 2 units of utility, but if it lands heads up on, say, the fourth toss, the player wins $2^4 = 2 \cdot 2 \cdot 2 \cdot 2 = 16$ units of utility. How much utility should you be willing to 'pay' for the opportunity to play this game? According to the expected utility principle you must be willing to pay any finite amount of utility, because $\frac{1}{2} \cdot 2 + \frac{1}{4} \cdot 4 + \frac{1}{8} \cdot 8 + \cdots = 1 + 1 + 1 + \cdots = \sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n \cdot 2^n = \infty$. But this is absurd. Arguably, most people would not pay even a hundred units. The most likely outcome is that one wins only a very small amount of utility. For instance, the probability that one wins at most 8 units is $0.5 + 0.25 + 0.125 = 0.875$.

The St Petersburg paradox was discovered by the Swiss mathematician Daniel Bernoulli (1700–1782), who was working in St Petersburg for a couple of years at the beginning of the eighteenth century. The St Petersburg paradox is, of course, not a paradox in a strict logical sense. No formal contradiction is deduced. But the recommendation arrived at, that one should sacrifice any finite amount of utility for the privilege of playing the St Petersburg game, appears to be sufficiently bizarre for motivating the use of the term paradox.

In 1745 Buffon argued in response to Bernoulli that sufficiently improbable outcomes should be regarded as "morally impossible", i.e. beyond concern. Hence, in Buffon's view, a rational decision maker should simply disregard the possibility of winning a very high amount, since such an outcome is highly improbable. Buffon's idea is closely related to the principle of *de minimis* risks, which still plays a prominent role in contemporary risk analysis. (Somewhat roughly put, the *de minimis* principle holds that sufficiently improbable outcomes, such as comet strikes, should be ignored.) From a mathematical point of view, it is obvious that if probabilities below a certain threshold are ignored, then the expected utility of the St Petersburg gamble will be finite. That said, this resolution of course seems to be ad hoc. Why on earth would it be rational to ignore highly improbable outcomes?

Other scholars have tried to resolve the St Petersburg paradox by imposing an upper limit on the decision maker's utility scale. From a historical perspective, this is probably the most prominent resolution of the paradox. The eighteenth-century mathematician Cramer, who discussed Bernoulli's original formulation of the paradox in which the prizes consisted of ducats instead of utility, suggested that "any amount above 10 millions, or (for the sake of simplicity) above 2^{24} =16677216 ducats [should] be deemed … equal in value to 2^{24} ducats" (in Bernoulli 1738/1954: 33). More recently, Nobel Prize winner Kenneth Arrow has maintained that the utility of wealth should be "taken to be a bounded function.... since such an assumption is needed to avoid [the St Petersburg] paradox" (Arrow 1970: 92). In order to see how the introduction of an upper limit affects the paradox, let L be the finite upper limit of utility. Then the expected utility of the gamble would be finite, because:

$$\begin{aligned}\frac{1}{2}\cdot 2+\frac{1}{4}\cdot 4+\frac{1}{8}\cdot 8+\cdots+\frac{1}{2^k}\cdot L+\frac{1}{2^{k+1}}\cdot L\cdots &= \sum_{i=1}^{k-1}(1/2)^i\cdot 2^i+\sum_{i=k}^{\infty}(1/2)^i\cdot L\\ &= \sum_{i=1}^{k-1}(1/2)^i\cdot 2^i+\left(1-\sum_{j=1}^{k-1}(1/2)^j\right)\cdot L\end{aligned}$$

A common reaction to Arrow's proposal is that the introduction of an upper limit is also ad hoc. Furthermore, even if one could overcome this objection, the introduction of a bounded utility scale may not resolve the paradox anyway. This is because the paradox has little to do with *infinite* utility. Arguably, the paradox arises whenever the expected utility of a gamble is unreasonably high *in comparison to* what we feel would be reasonable to pay for entering the gamble. To see this, note that a slightly modified version of the St Petersburg paradox arises even if only small finite amounts of utility are at stake, as in the following gamble.

> **Gamble 1** A fair coin is tossed until it lands heads up. The player thereafter receives a prize worth min $\{2^n \cdot 10^{-100}, L\}$ units of utility, where n is the number of times the coin was tossed.

Suppose L equals 1. Now, the expected utility of Gamble 1 has to be greater than $332 \cdot 10^{-100}$ units of utility (since 2^{332} is approximately equal to 10^{100}). However, on average in one out of two times the gamble is played, you win only $2 \cdot 10^{-100}$ units, and in about nine times out of ten you win no more than $8 \cdot 10^{-100}$ units. This indicates that even though the expected utility of Gamble 1 is finite, and indeed very small, it is nevertheless paradoxically high *in comparison to* the amount of utility the player actually wins.

Another resolution of the St Petersburg paradox was suggested by Richard C. Jeffrey. He claimed that, "anyone who offers to let the agent play the St. Petersburg game is a liar, for he is pretending to have an indefinitely large bank" (1983: 154). This is because no casino or bank can possibly fulfil its commitments towards the player in the case that a very large number of tails precedes the first head; hence, the premises of the gamble can never be valid. A possible reply to Jeffrey's argument is to point out that all sorts of prizes should, of course, be allowed. Suppose, for instance, that after having played the St Petersburg gamble you will be connected to Robert Nozick's experience machine. By definition, the experience machine can create any experience in you, e.g. intense happiness or sexual pleasure. The fact that there is a limited amount of money in the world is therefore no longer a problem.

There is also another response to Jeffrey's proposal. As before, the main idea is to show that a slightly modified version of the St Petersburg paradox arises even if we accept Jeffrey's restriction, i.e. if we assume that the amount of utility in the bank is finite. Consider Gamble 1 again, and let L equal the total amount of utility available in the bank. Now, Jeffrey's requirement of a finite amount of utility in the bank is obviously satisfied. Arguably, the most important point is that if a new paradoxical conclusion can be obtained just by making some minor alterations to the original problem, then the old paradox will simply be replaced by a new one, and nothing is gained. The following gamble is yet another illustration of this point, which raises a more general issue that goes beyond Jeffrey's proposal.

Gamble 2 A manipulated coin, which lands heads up with probability 0.4, is tossed until it lands heads up. The player thereafter receives a prize worth 2^n units of utility, where n is the number of times the coin was tossed.

Common sense tells us that Gamble 2 should be preferred to the original St Petersburg gamble, since it is more likely to yield a long sequence of tosses and consequently better prizes. However, the expected utility of both gambles is infinite, because $\sum_{n=1}^{\infty} (0.4)^n \cdot 2^n = \infty$. Hence, the principle of maximising expected utility recommends us to judge both gambles as equally valuable. This is also absurd. Any satisfactory account of rationality must entail that Gamble 2 is better than the original St Petersburg gamble.

4.7 The two-envelope paradox

The two-envelope paradox arises from a choice between two envelopes, each of which contains some money. A trustworthy informant tells you that one of the envelopes contains exactly twice as much as the other, but the informant does not tell you which is which. Since this is all you know you decide to pick an envelope at random. Let us say you pick envelope A. Just before you open envelope A you are offered to swap and take envelope B instead. The following argument indicates that you ought to swap. Let x denote the amount in A. Then envelope B has to contain either $2x$ or $x/2$ dollars. Given what you know, both possibilities are equally likely. Hence, the expected monetary value of swapping to B is $\frac{1}{2} \cdot 2x + \frac{1}{2} \cdot \frac{x}{2} = \frac{5}{4}x$. Since $\frac{5}{4}x$ is more than x, it is rational to take B instead of A.

However, just as you are about to open envelope B, you are offered to swap back. The following argument indicates that you ought to take envelope A. Let y denote the amount in envelope B. It then follows that envelope A contains either $2y$ or $y/2$ dollars. As before, both possibilities are equally likely, so the expected monetary value of swapping is $\frac{1}{2} \cdot 2y + \frac{1}{2} \cdot \frac{y}{2} = \frac{5}{4}y$. Since $\frac{5}{4}y$ is more than y you ought to swap.

Table 4.10

	1/2	1/2	Expected value
Envelope A	x	x	x
Envelope B	$2x$	$1/2x$	$5x/4$

Table 4.11

	1/2	1/2	Expected value
Envelope A	$2y$	$1/2y$	$5y/4$
Envelope B	y	y	y

Clearly, there must be something wrong with the reasoning outlined here. It simply cannot hold true that the expected monetary value of choosing A exceeds that of choosing B, at the same time as the expected monetary value of choosing B exceeds that of choosing A. Consider Table 4.10 and Table 4.11.

The present formulation of the paradox presupposes that there is no upper limit to how much money there is in the world. To see this, suppose that there indeed is some upper limit L to how much money there is in the world. It then follows that no envelope can contain more than $(2/3)L$, in which case the other envelope would be certain to contain $(1/3)L$. (If, say, envelope A contains $(2/3)L$, then it would clearly be false that envelope B contains either $2 \cdot (2/3)L$ or $1/2 \cdot (2/3)L$. Only the latter alternative would be a genuine possibility.) Hence, for the paradox to be viable one has to assume that the amount of money in the world is infinite, which is implausible. That said, the paradox can easily be restated without referring to monetary outcomes; if we assume the existence of infinite utilities the paradox will come alive again.

The two-envelope paradox can also be generated by starting from the St Petersburg paradox: A fair coin is flipped n times until it lands heads up. Then a prize worth 2^n units of utility is put in one of the envelopes and either half or twice that amount in the other envelope. It follows that, for every finite n, if the first envelope contains 2^n units of utility, one always has reason to swap to the other envelope, since its expected utility is higher. However, as we know from the discussion of the St Petersburg paradox, the *expected utility* of the contents in each envelope is infinite.

At present there is no consensus on how to diagnose the two-envelope paradox. A large number of papers have been published in philosophical journals. Most attempt to show that the probability assignments are illegitimate, for one reason or another. However, it has also been argued that the way the outcomes are described does not accurately represent the real decision problem. I leave it to the reader to make up her own mind about this surprisingly deep problem.

Exercises

4.1 Consider the decision problem illustrated below.

	1/2	1/4	1/4
a_1	\$49	\$25	\$25
a_2	\$36	\$100	\$0
a_3	\$81	\$0	\$0

(a) The decision maker's utility u of money is linear. Which act should be chosen according to the principle of maximising expected monetary value?

(b) The decision maker's utility u of money x is given by the formula $u(x) = \sqrt{x}$. Which act should be chosen according to the principle of maximising expected utility?

4.2 I am in my office in Cambridge, but I have to catch a flight from Heathrow this afternoon. I must decide whether to go to Heathrow by coach, which comes relatively cheap at £40, or buy a train ticket for £70. If I take the coach I might get stuck in an intense traffic jam and miss my flight. I would then have to buy a new ticket for £100. According to the latest statistics, the traffic jam on the M25 to Heathrow is intense one day in three. (a) Should I travel by train or coach? (b) This description of my decision problem overlooks a number of features that might be relevant. Which?

4.3 (a) You are in Las Vegas. The probability of winning a jackpot of \$350,000 is one in a million. How much should you, who find no reason to reject the principle of maximising expected utility, be prepared to pay to enter this gamble? Your utility of money is $u(x) = \ln(x + 1)$.

(b) This time the probability of winning the jackpot of \$350,000 is one in a thousand. How much should you, who find no reason to reject the principle of maximising expected utility, be prepared to pay to enter this gamble? Your utility of money is $u(x) = \ln(x + 1)$.

(c) Why is the difference between the amount you are willing to pay in (a) and (b) so small?

(d) Why did we assume that your utility function is $u(x) = \ln(x + 1)$, rather than just $u(x) = \ln(x)$?

4.4 (a) What is the law of large numbers?

(b) How is the law of large number related to the theorem known as gambler's ruin?

4.5 (a) Explain why Allais' and Ellsberg's paradoxes pose difficulties for the principle of maximising expected utility. (b) Explain the difference between the two paradoxes – they arise for two different reasons.

4.6 Suppose that you prefer Gamble 1 to Gamble 2, and Gamble 4 to Gamble 3. Show that your preferences are incompatible with the principle of maximising expected utility, no matter what your utility of money happens to be.

	1/3	1/3	1/3
Gamble 1	\$50	\$50	\$50
Gamble 2	\$100	\$50	\$0
Gamble 3	\$50	\$0	\$50
Gamble 4	\$100	\$0	\$0

4.7 (a) Explain why the St Petersburg paradox poses a difficulty for the principle of maximising expected utility.

(b) Construct a new version of the St Petersburg paradox, in which the player rolls a six-sided die instead of tossing a coin.

4.8 There is an interesting connection between the St Petersburg and the two-envelope paradox – explain!

Solutions

4.1 (a) a_2 (b) a_1

4.2 (a) If money is all that matters, and my utility of money is linear, I should buy a train ticket. It will cost me £70, but the expected monetary cost of going by coach is £73.33. (b) Arguably, money is not all that matters. For instance, travelling by train is more comfortable and if