

2 The decision matrix

Before you make a decision you have somehow to determine what to decide about. Or, to put it differently, you have to specify what the relevant *acts*, *states* and *outcomes* are. Suppose, for instance, that you are thinking about taking out fire insurance on your home. Perhaps it costs \$100 to take out insurance on a house worth \$100,000, and you ask: Is it worth it? Before you decide, you have to get the formalisation of the decision problem right. In this case, it seems that you face a decision problem with two acts, two states, and four outcomes. It is helpful to visualise this information in a decision matrix; see Table 2.1.

To model one's decision problem in a formal representation is essential in decision theory, since decision rules are only defined relative to such formalisations. For example, it makes no sense to say that the principle of maximising expected value recommends one act rather than another unless there is a formal listing of the available acts, the possible states of the world and the corresponding outcomes. However, instead of visualising information in a decision matrix it is sometimes more convenient to use a decision tree. The decision tree in Figure 2.1 is equivalent to the matrix in Table 2.1.

The square represents a *choice node*, and the circles represent *chance nodes*. At the choice node the decision maker decides whether to go *up* or *down* in the tree. If there are more than two acts to choose from, one simply adds more lines. At the chance nodes nature decides which line to follow, and the rightmost boxes represent the possible outcomes. Decision trees are often used for representing sequential decisions, i.e. decisions that are divided into several separate steps. (Example: In a restaurant, you can either order all three courses before you start to eat, or divide the decision-making process into three separate decisions taken at three points in time. If you opt for the latter approach, you face a sequential decision problem.) To

Table 2.1

	Fire	No fire
<i>Take out insurance</i>	No house and \$100,000	House and \$0
<i>No insurance</i>	No house and \$100	House and \$100




Figure 2.1

represent a sequential decision problem in a tree, one simply adds new choice and chance nodes to the right of the existing leaves.

Many decision theorists distinguish only between decision problems and a corresponding decision matrix or tree. However, it is worth emphasising that we are actually dealing with three levels of abstraction:

1. The decision problem
2. A formalisation of the decision problem
3. A visualisation of the formalisation

A decision problem is constituted by the entities of the world that prompt the decision maker to make a choice, or are otherwise relevant to that choice. By definition, a formalisation of a decision problem is made up of information about the decision to be made, irrespective of how that information is visualised. Formalisations thus comprise information about acts, states and outcomes, and sometimes also information about probabilities. Of course, one and the same decision problem can be formalised in

Table 2.2

[
$[a_1 = \text{take out insurance},$
$a_2 = \text{do not}];$
$[s_1 = \text{fire},$
$s_2 = \text{no fire}];$
$[(a_1, s_1) = \text{No house and \$100,000},$
$(a_1, s_2) = \text{House and \$0},$
$(a_2, s_1) = \text{No house and \$100},$
$(a_2, s_2) = \text{House and \$100}]$
]

different ways, not all of which are likely to be equally good. For example, some decision problems *can* be formalised either as decisions under risk or as decisions under ignorance, but if probabilities are known it is surely preferable to choose the former type of formalisation (since one would otherwise overlook relevant information).

Naturally, any given set of information can be visualised in different ways. We have already demonstrated this by drawing a matrix and a tree visualising the same formalisation. Table 2.2 is another example of how the same information could be presented, which is more suitable to computers.

In Table 2.2 information is stored in a *vector*, i.e. in an ordered list of mathematical objects. The vector is comprised of three new vectors, the first of which represents acts. The second vector represents states, and the third represents outcomes defined by those acts and states.

From a theoretical perspective, the problem of how to formalise decision problems is arguably more interesting than questions about how to visualise a given formalisation. Once it has been decided what pieces of information ought to be taken into account, it hardly matters for the decision theorist whether this information is visualised in a matrix, a tree or a vector.

2.1 States

The basic building blocks of a decision problem are states, outcomes and acts. Let us discuss each concept in turn, starting with states. What is a state? Intuitively, a state is a part of the world that is not an outcome or an act (that can be performed by the agent in the present decision situation;

Table 2.3

	I choose the right bet	I do not
<i>Bet on Democrat</i>	\$100	\$0
<i>Bet on Republican</i>	\$200	\$0

acts performed by others can presumably be thought of as states). It is difficult to come up with a more precise definition without raising deep metaphysical questions that fall well beyond the scope of this book.

That said, not all states are relevant to decision making. For many decisions it is completely irrelevant whether the number of atoms in the universe is odd or even, for instance. Only states that may affect the decision maker's preference among acts need to be taken into account, such as: *The republican candidate wins the election*, or *The interest rate exceeds five per cent next year*, or *My partner loves me*, or *Goldbach's conjecture is true*. For each of these states, we can easily imagine an act whose outcome depends on the state in question. The example of Goldbach's conjecture (a famous mathematical hypothesis) indicates that even necessary truths may be relevant in decision making, e.g. if the decision maker has placed a bet on whether this hypothesis is true or not.

Some states, or at least some ways of *partitioning* states, are clearly illegitimate. In order to understand why, imagine that you are offered a choice between two bets, which pay \$100 and \$200 respectively, depending on whether the Democrat or the Republican candidate will win the next presidential election (Table 2.3). Now, it would make little sense to consider the states *I choose the right bet* and *I do not*.

This formalisation gives the false impression that you will definitively be better off if you choose to bet on the Republican candidate. The reason why this is false, and why the formalisation is illegitimate, is that the state *I choose the right bet* is causally dependent of the act you choose. Whether the state will occur depends on which act is chosen.

The problem of causal dependence can be addressed in two ways. The first is to allow the decision theorist to include only states that are causally independent of the acts in the formalisation; this is the option that we shall pursue here (with exception for Chapter 9). The second option is to avoid the notion of states altogether. That approach works particularly well if one

Table 2.4

<i>Bet on Democrat</i>	\$100 (0.6)	\$0 (0.4)
<i>Bet on Republican</i>	\$200 (0.3)	\$0 (0.7)

Table 2.5(a)

	State 1	State 2
<i>Act A</i>	\$50	\$80
<i>Act B</i>	\$30	\$80

Table 2.5(b)

<i>Act A</i>	\$50	\$80
<i>Act B</i>	\$80	\$30

happens to know the probability of the outcomes. Suppose, for instance, that the probability is 0.6 that you win \$100 if you bet on the Democrat, and 0.3 that you win \$200 if you bet on the Republican.

Arguably, the formalisation in Table 2.4 is impeccable. That said, omitting the states makes little sense in other decision problems, such as decisions under ignorance. If you do not know anything about the probability that your house will burn down and you are offered fire insurance for free, it certainly makes sense to accept the offer. No matter what happens, you will be at least as well off if you accept free insurance than if you do not, and you will be better off if there is a fire. (This line of reasoning has a fancy name: *the dominance principle*. See Chapter 3.) However, this conclusion only follows if we attach outcomes to states. The pair of examples in Table 2.5(a) and 2.5(b) illustrates the difference between including and omitting a set of states.

In Table 2.5(a), act A is clearly better than act B. However, in Table 2.5(b) all states have been omitted and the outcomes have therefore been listed in *arbitrary* order. Here, we fail to see that one option is actually better than the other. This is why it is a bad idea not to include states in a formalisation.

Let us now consider a slightly different kind of problem, having to do with how preferences depend on which state is in fact the true state. Suppose that you have been severely injured in a traffic accident and that

Table 2.6

	You survive	You die
<i>Take first bet</i>	\$100	\$0
<i>Take second bet</i>	\$0	\$100

as a result you will have to undergo a risky and life-threatening operation. For some reason, never mind why, you are offered a choice between a bet in which you win \$100 if the operation is successful and you survive, and nothing if the operation fails and you die, and a bet in which you win nothing if the operation is successful and you survive, and \$100 if the operation fails and you die (Table 2.6).

Let us suppose that both states are equally probable. Then, it is natural to argue that the decision maker should regard both bets as equally attractive. However, most people would prefer the bet in which one wins \$100 if one survives the operation, no matter how improbable that state is. If you die, money does not matter to you any more. This indicates that the formalisation is underspecified. To make the formalisation acceptable we would have to add to the outcome the fact that the decision maker will die if the operation fails. Then, it would clearly transpire why the first bet is better than the second. To put it in a more sophisticated way, to which we shall return in Chapter 7, states should be chosen such that the value of the outcomes under all states is independent of whether the state occurs or not.

2.2 Outcomes

Rational decision makers are not primarily concerned with states or acts. What ultimately matters is the *outcome* of the choice process. Acts are mere instruments for reaching good outcomes, and states are devices needed for applying these instruments. However, in order to figure out which instrument to use (i.e. which act to choose given a set of states), outcomes must be ranked in one way or another, from the worst to the best. Exactly how this should be done is an important topic of debate in decision theory, a topic which we shall examine in more detail in Chapter 5. In the present section we shall merely explain the difference between the various kinds of scales that are used for comparing outcomes.

Let us return to the issue of whether or not one should insure a house worth \$100,000 at a rate of \$100 per annum. Imagine that Jane has made a sincere effort to analyse her attitudes towards safety and money, and that she felt that the four possible outcomes should be ranked as follows, from the best to the worst.

1. House and \$100 *is better than*
2. House and \$0 *is better than*
3. No house and \$100,000 *is better than*
4. No house and \$100.

The first outcome, 'House and \$100', can be thought of as a possible world that is exactly similar to the three others, except for the condition of the house and amount of money in that world. Outcomes are in that sense *comprehensive* – they include much more than we actually need to mention in a decision matrix. Naturally, the ranking of outcomes is to a large extent subjective. Other decision makers may disagree with Jane and feel that the outcomes ought to be ranked differently. For each decision maker, the ranking is acceptable only if it reflects his or her attitudes towards the outcomes. It seems fairly uncontroversial to suppose that people sometimes have different attitudes, but the ranking depicted above does, we assume, accurately reflect Jane's attitudes.

In order to measure the value of an outcome, as it is perceived by the decision maker, it is convenient to assign numbers to outcomes. In decision theory, numbers referring to comparative evaluations of value are commonly called *utilities*. However, the notion of utility has many different technical meanings, which should be kept separate. Therefore, to avoid unnecessary confusion we shall temporarily stick to the rather vague term *value*, until the concept of utility has been properly introduced in Chapter 5.

Value can be measured on two fundamentally different kinds of scales, viz. *ordinal* scales and *cardinal* scales. Consider the set of numbers in Table 2.7, assigned by Jane to the outcomes of her decision problem.

Table 2.7

	Fire	No fire
<i>Take out insurance</i>	1	4
<i>Do not</i>	−100	10

Table 2.8

	Original scale	Scale A	Scale B	Scale C
Best outcome	10	4	100	777
Second best	4	3	98	−378
Third best	1	2	97	−504
Worst outcome	−100	1	92	−777

If Jane assigns a higher number to one outcome than another, she judges the first outcome to be better than the second. However, if the scale she uses is an ordinal scale, nothing more than that follows. In particular, nothing can be concluded about *how much* better one outcome is in relation to another. The numbers merely reflect the qualitative ranking of outcomes. No quantitative information about the ‘distance’ in value is reflected by the scale. Look at Scales A–C in Table 2.8; they could be used for representing exactly the same ordinal ranking.

The transformations of the original scale into scales A, B or C preserves the order between the outcomes. This proves that all four scales are equivalent. Hence, it does not matter which set of numbers one uses. Mathematicians express this point by saying that ordinal scales are *invariant* up to *positive monotone transformations*. That a transformation of a scale is invariant under some sort of change means that the ranking of the objects is preserved after this type of change. In the case of an ordinal scale, the change is describable by some function f such that

$$f(x) \geq f(y) \text{ if and only if } x \geq y. \quad (1)$$

In this expression, x and y are two arbitrary values of some initial scale, e.g. the values corresponding to the best and second best outcomes on the original scale above, and f is some mathematical function. It can be easily verified that the transformation of the initial scale into scale D in Table 2.9 *does not* satisfy condition (1), because if $x = 10$ and $y = 4$, then $f(x) = 8$, and $f(y) = 9$. Of course, it is false that $8 \geq 9$ if and only if $10 \geq 4$. It can also be shown in analogous ways that scales E and F are not permissible ordinal transformations of the original scale.

As pointed out above, ordinal scales are usually contrasted with cardinal scales. Cardinal scales embody more information than ordinal scales. There are two different kinds of cardinal scales, viz. *interval* scales and *ratio* scales.

Table 2.9

	Scale D	Scale E	Scale F
Best outcome	8	−60	100
Second best	9	−50	90
Third best	6	−40	80
Worst outcome	7	0	80

To start with, we focus on interval scales. Unlike ordinal scales, interval scales accurately reflect the difference between the objects being measured. Let us suppose, for illustrative purposes, that scale F in the example above is an interval scale. It would then be correct to conclude that the difference in value between the best and the second best outcome is exactly the same as the distance in value between the second best and the third best outcome. Furthermore, the difference between the best and the worst outcome is twice that between the best and the second best outcome. However, scale E cannot be a permissible transformation of F, because in F the distance between the third best and the worst outcome is zero, whereas the corresponding difference in E is strictly greater than zero.

To illustrate what kind of information is represented in an interval scale, consider the two most frequently used scales for measuring temperature, i.e. the Centigrade (C) and Fahrenheit (F) scales, respectively. Both scales accurately reflect differences in temperature, and any temperature measured on one scale can easily be transformed into a number on the other scale. The formula for transforming Centigrade to Fahrenheit is:

$$F = 1.8 \cdot C + 32 \quad (2)$$

By solving this equation for C, we get:

$$C = (F - 32)/1.8 \quad (3)$$

Note that (2) and (3) are straight lines – had the graphs been curved, a difference of, say, one unit on the x-axis will not always produce the same difference on the y-axis. As an illustration of how the Fahrenheit scale can be transformed into the centigrade scale, consider Table 2.10. It shows the temperatures in a number of cities on a sunny day a few years ago.

When looking at Table 2.10, a common error is to conclude that it was *twice as warm* in Tokyo as in New York, since 64 units Fahrenheit is twice as

Table 2.10

City	Degrees Fahrenheit	Degrees Centigrade
Los Angeles	82	27.8
Tokyo	64	17.8
Paris	62	16.7
Cambridge (UK)	46	7.8
New York	32	0
Stockholm	-4	-20

much as 32 units Fahrenheit. In order to see why that conclusion is incorrect, note that 64°F corresponds to 17.8°C and 32°F to 0°C . Now, 17.8°C is of course not twice as much as 0°C . Had it been twice as warm in Tokyo as in New York, it would certainly have been twice as warm according to every scale. Interval scales accurately reflect differences, but not ratios. Expressed in mathematical terminology, interval scales are invariant up to *positive linear transformations*. This means that any interval scale can be transformed into another by multiplying each entry by a positive number and adding a constant, without losing or gaining any information about the objects being measured. For example, if the value of some outcome is 3 according to scale X , and $Y = 10 \cdot X + 5$, then the value of the same outcome would be 35 if measured on scale Y . Obviously, scale Y is obtained from X by a positive linear transformation.

Unlike interval scales, ratio scales accurately reflect ratios. Mass, length and time are all examples of entities that can be measured on ratio scales. For example, $20\text{ lb} = 9\text{ kg}$, and this is twice as much as $10\text{ lb} = 4.5\text{ kg}$. Furthermore, two weeks is twice as much as one, and 14 days is of course twice as much as 7 days. Formally put, a ratio scale U can be accurately transformed into an equivalent ratio scale V by multiplying U by a positive constant k . Consider, for instance, the series of numbers in Table 2.11, which denote the values of four different outcomes as measured on five different scales, G–K.

Scale G and I are equivalent ratio scales, because $G = 5 \cdot I$. Furthermore, the first four scales, G–J, are equivalent interval scales. For example, scale H can be obtained from G by the formula $G = 10 \cdot H + 10$, and $J = 0.1 \cdot G - 3$. However, there is no equation of the form $V = k \cdot U + m$ that transforms G into K. Hence, since K is not a positive linear transformation of G, it does not reflect the same differences in value, nor the same ratios.

Table 2.11

	Scale G	Scale H	Scale I	Scale J	Scale K
Best outcome	40	410	8	1	5
Second best	30	310	6	0	3
Third best	20	210	4	-1	2
Worst	10	110	2	-2	1

It is helpful to summarise the technical properties of the two kinds of cardinal scales discussed here in two mathematical conditions. To start with, a function f that takes an argument x and returns a real number as its value is an interval scale if and only if condition (1) on page 24 holds *and* for every other function f' that satisfies (1) there are some positive constants k and m such that:

$$f'(x) = k \cdot f(x) + m \quad (4)$$

Condition (4) states what transformations of an interval scale are permissible: As we have shown above, every transformation that can be mapped by an upward sloping straight line is permissible. Furthermore, a function f that takes an argument x and returns a real number as its value is a ratio scale if and only if condition (1) holds *and* for every other function f' that satisfies (1) there is some positive constant k such that:

$$f'(x) = k \cdot f(x) \quad (5)$$

This condition is even more simple than the previous one: A pair of ratio scales are equivalent if and only if each can be transformed into the other by multiplying all values by some positive constant. (Of course, the constant we use for transforming f into f' is not the same as that we use for transforming f' into f .)

2.3 Acts

Imagine that your best friend Leonard is about to cook a large omelette. He has already broken five eggs into the omelette, and plans to add a sixth. However, before breaking the last egg into the omelette, he suddenly starts to worry that it might be rotten. After examining the egg carefully, he decides to take a chance and break the last egg into the omelette.

Box 2.1 Three types of scales

In this chapter we have discussed three different types of scales. Their main characteristics can be summarised as follows.

1. Ordinal scale: Qualitative comparison of objects allowed; no information about differences or ratios. Example: The jury of a song contest award points to the participants. On this scale, 10 points is more than 5.
2. Cardinal scales
 - (a) Interval scale Quantitative comparison of objects; accurately reflects differences between objects. Example: The Centigrade and Fahrenheit scales for temperature measurement are the most well-established examples. The difference between 10°C and 5°C equals that between 5°C and 0°C, but the difference between 10°C and 5°C does not equal that between 10°F and 5°F.
 - (b) Ratio scale Quantitative comparison of objects; accurately reflects ratios between objects. Example: Height, mass, time, etc. 10kg is twice as much as 5kg, and 10 lb is also twice as much as 5 lb. But 10kg is not twice as much as 5 lb.

The act of adding the sixth egg can be conceived of as a function that takes either the first state (*The sixth egg is rotten*) or the second (*The sixth egg is not rotten*) as its argument. If the first state happens to be the true state of the world, i.e. if it is inserted into the function, then it will return the outcome *No omelette*, and if the second state happens to be the true state, the value of the function will be *Six egg omelette*. (See Table 2.12.) This definition of acts can be trivially generalised to cover cases with more than two states and outcomes: an act is a function from a set of states to a set of outcomes.

Did you find this definition too abstract? If so, consider some other function that you are more familiar with, say $f(x) = 3x + 8$. For each argument x , the function returns a value $f(x)$. Acts are, according to the suggestion above and originally proposed by Leonard Savage, also functions. However, instead of taking numbers as their arguments they take states, and instead of returning other numbers they return outcomes. From a mathematical

Table 2.12

	The sixth egg is rotten	The sixth egg is not rotten
<i>Add sixth egg</i>	No omelette	Six egg omelette
<i>Do not add sixth egg</i>	Five egg omelette	Five egg omelette

point of view there is nothing odd about this; a function is commonly defined as any device that takes one object as its argument and returns exactly one other object. Savage's definition fulfils this criterion. (Note that it would be equally appropriate to consider states and acts as primitive concepts. Outcomes could be conceived of as ordered pairs of acts and states. For example, the outcome *No omelette* is the ordered pair comprising of the act *Add sixth egg* and the state *The sixth egg is rotten*. States can be defined in similar ways, in terms of acts and outcomes.)

Decision theory is primarily concerned with *particular* acts, rather than generic acts. A *generic* act, such as sailing, walking or swimming can be instantiated by different agents at different time intervals. Hence, Columbus' first voyage to America and James Cook's trip to the southern hemisphere are both instantiations of the same generic act, viz. sailing. Particular acts, on the other hand, are always carried out by specific agents at specific time intervals, and hence Columbus' and Cook's voyages were different particular acts. Savage's definition is a characterisation of particular acts.

It is usually assumed that the acts considered by a decision maker are *alternative* acts. This requirement guarantees that a rational decision maker has to choose only one act. But what does it mean to say that some acts constitute a set of alternatives? According to an influential proposal, the set A is an *alternative-set* if and only if every member of A is a particular act, A has at least two different members, and the members of A are agent-identical, time-identical, performable, incompatible in pairs and jointly exhaustive. At first glance, these conditions may appear as fairly sensible and uncontroversial. However, as pointed out by Bergström (1966), they do not guarantee that every act is a member of only one alternative-set. Some particular acts are members of several non-identical alternative-sets. Suppose, for instance, that I am thinking about going to the cinema (act a_1) or not going to the cinema (a_2), and that $\{a_1, a_2\}$ is an alternative-set. Then I realise that a_1 can be performed in different ways. I can, for instance, buy

popcorn at the cinema (a_3) or buy chocolate (a_4). Now, also $\{a_1 \& a_3, a_1 \& a_4, a_2\}$ is an alternative-set. Of course, $a_1 \& a_3$ and $a_1 \& a_4$ are different particular acts, so both of them cannot be identical to a_1 . Moreover, $a_1 \& a_3$ can also be performed in different ways. I can buy a small basket of popcorn (a_5) or a large basket (a_6), and therefore $\{a_1 \& a_3 \& a_5, a_1 \& a_3 \& a_6, a_1 \& a_4, a_2\}$ also constitutes an alternative-set, and so on and so forth. So what are the alternatives to a_2 ? Is it $\{a_1\}$, or $\{a_1 \& a_3, a_1 \& a_4, a_2\}$, or $\{a_1 \& a_3 \& a_5, a_1 \& a_3 \& a_6, a_1 \& a_4, a_2\}$? Note that nothing excludes that the outcome of a_2 is better than the outcome of $a_1 \& a_3$, while the outcome of $a_1 \& a_3 \& a_6$ might be better than that of a_2 . This obviously causes problems for decision makers seeking to achieve as good outcomes as possible.

The problem of defining an alternative-set has been extensively discussed in the literature. Bergström proposed a somewhat complicated solution of the problem, which has been contested by others. We shall not explain it here. However, an interesting implication of Bergström's proposal is that the problem of finding an alternative-set is partly a normative problem. This is because we cannot formalise a decision problem until we know which normative principle to apply to the resolution of the problem. What your alternatives are depends partly on what your normative principle tells you to seek to achieve.

2.4 Rival formalisations

In the preceding sections, we have briefly noted that one cannot take for granted that there exists just one *unique* best formalisation of each decision problem. The decision maker may sometimes be confronted with *rival* formalisations of one and the same decision problem. Rival formalisations arise if two or more formalisations are equally reasonable and strictly better than all alternative formalisations.

Obviously, rival formalisations are troublesome if an act is judged to be rational in one optimal formalisation of a decision problem, but non-rational in another optimal formalisation of the same decision problem. In such cases one may legitimately ask whether the act in question should be performed or not. What should a rational decision maker do? The scope of this problem is illustrated by the fact that, theoretically, there might be cases in which *all* acts that are rational in one optimal formalisation are non-rational in another rival formalisation of the same decision problem,

whereas *all* acts that are rational according to the latter formalisation are not rational according to the former.

To give convincing examples of rival formalisations is difficult, mainly because it can always be questioned whether the suggested formalisations are equally reasonable. In what follows we shall outline a hypothetical example that some people may find convincing, although others may disagree. Therefore, in Box 2.2 we also offer a more stringent and technical argument that our example actually is an instance of two equally reasonable but different formalisations.

Imagine that you are a paparazzi photographer and that rumour has it that actress Julia Roberts will show up in either New York (NY), Los Angeles (LA) or Paris (P). Nothing is known about the probability of these states of the world. You have to decide if you should stay in America or catch a plane to Paris. If you stay and actress Julia Roberts shows up in Paris you get \$0; otherwise you get your photos, which you will be able to sell for \$10,000. If you catch a plane to Paris and Julia Roberts shows up in Paris your net gain after having paid for the ticket is \$5,000, and if she shows up in America you for some reason, never mind why, get \$6,000. Your initial representation of the decision problem is visualised in Table 2.13.

Since nothing is known about the probabilities of the states in Table 2.13, you decide it makes sense to regard them as equally probable, i.e. you decide to assign probability $1/3$ to each state. Consider the decision matrix in Table 2.14.

Table 2.13

	P	LA	NY
<i>Stay</i>	\$0	\$10k	\$10k
<i>Go to Paris</i>	\$5k	\$6k	\$6k

Table 2.14

	P ($1/3$)	LA ($1/3$)	NY ($1/3$)
<i>Stay</i>	\$0	\$10k	\$10k
<i>Go to Paris</i>	\$5k	\$6k	\$6k

Table 2.15

	P (1/3)	LA or NY (2/3)
<i>Stay</i>	\$0	\$10k
<i>Go to Paris</i>	\$5k	\$6k

Table 2.16

	P	LA or NY
<i>Stay</i>	\$0	\$10k
<i>Go to Paris</i>	\$5k	\$6k

Table 2.17

	P (1/2)	LA or NY (1/2)
<i>Stay</i>	\$0	\$10k
<i>Go to Paris</i>	\$5k	\$6k

The two rightmost columns are exactly parallel. Therefore, they can be merged into a single (disjunctive) column, by adding the probabilities of the two rightmost columns together (Table 2.15).

However, now suppose that you instead start with Table 2.13 and *first* merge the two repetitious states into a single state. You would then obtain the decision matrix in Table 2.16.

Now, since you know nothing about the probabilities of the two states, you decide to regard them as equally probable, i.e. you assign a probability of 1/2 to each state. This yields the formal representation in Table 2.17, which is clearly different from the one suggested above in Table 2.15.

Which formalisation is best, 2.15 or 2.17? It seems question begging to claim that one of them must be better than the other – so perhaps they are equally reasonable? If they are, we have an example of rival formalisations.

Note that the principle of maximising expected value recommends different acts in the two matrices. According to Table 2.15 you should stay, but 2.17 suggests you should go to Paris. Arguably, this example shows that rival formalisations must be taken seriously by decision theorists, although there is at present no agreement in the literature on how this phenomenon ought to be dealt with.

Box 2.2 Why rival representations are possible

The examples illustrated in Tables 2.15 and 2.17 do not *prove* that rival formalisations are possible. One may always question the claim that the two formalisations are equally reasonable. Therefore, in order to give a more comprehensive argument for thinking that the formalisations are equally reasonable, we shall introduce some technical concepts. To begin with, we need to distinguish between two classes of decision rules, viz. *transformative* and *effective* decision rules. A decision rule is effective if and only if it singles out some set of recommended acts, whereas it is transformative if and only if it modifies the formalisation of a given decision problem. Examples of effective decision rules include the principle of maximising expected utility and the dominance principle, mentioned in Chapter 1. Transformative decision rules do not directly recommend any particular act or set of acts. Instead, they transform a given formalisation of a decision problem into another by adding, deleting or modifying information in the initial formalisation. More precisely, transformative decision rules can alter the set of alternatives or the set of states of the world taken into consideration, modify the probabilities assigned to the states of the world, or modify the values assigned to the corresponding outcomes. For an example of a transformative decision rule, consider the rule saying that if there is no reason to believe that one state of the world is more probable than another then the decision maker should transform the initial formalisation of the decision problem into one in which every state is assigned equal probability. This transformative rule is called the principle of insufficient reason.

We assume that all significant aspects of a decision problem can be represented in a triplet $\pi = \langle A, S, O \rangle$, where A is a non-empty set of (relevant) alternative acts, S is a non-empty set of states of the world, and O is a set of outcomes. Let us call such a triplet a *formal decision problem*. A transformative decision rule is defined as a function t that transforms one formal decision problem π into another π' , i.e. t is a transformative decision rule in a set of formal decision problems Π if and only if t is a function such that for all $\pi \in \Pi$, it holds that $t(\pi) \in \Pi$. If t and u form a pair of transformative decision rules, we can construct a new composite rule $(t \circ u)$ such that $(t \circ u)(\pi) = u(t(\pi))$. In this framework the question, “How should the decision maker formalise a decision problem?” can be restated as: “What sequence of transformative rules $(t \circ u \circ \dots)$ should a rational decision maker apply to an initial formal decision problem π ?”.

Let $\succeq$ be a relation on Π such that $\pi \succeq \pi'$ if and only if the formal representation π is at least as reasonable as π' . (If π and π' are equally reasonable we write $\pi \sim \pi'$.) We shall not go into detail here about what makes one formal representation more reasonable than another, but it should be obvious that some representations are better than others. Now, we shall prove that if the technical condition stated below holds for $\succeq$, then two different sequences of a given set of transformative decision rules, $(t \circ u)$ and $(u \circ t)$, will always yield formalisations that are equally reasonable.

Order-independence (OI): $(u \circ t)(\pi) \succeq t(\pi) \succeq (t \circ u)(\pi)$

The left-hand inequality, $(u \circ t)(\pi) \succeq t(\pi)$, states that a transformative rule u should not, metaphorically expressed, throw a spanner in the works carried out by another rule t . Hence, the formalisation obtained by first applying u and then t has to be at least as good as the formalisation obtained by only applying t . The right-hand inequality, $t(\pi) \succeq (t \circ u)(\pi)$, says that nothing can be gained by immediately repeating a rule. This puts a substantial constraint on transformative rules; only ‘maximally efficient’ rules, that directly improve the formal representation as much as possible, are allowed by the OI-condition. Now consider the following theorem.

Theorem 2.1 Let the OI-condition hold for all $\pi \in \Pi$. Then, $(u \circ t)(\pi) \sim (t \circ u)(\pi)$ for all u and t .

Proof We prove Theorem 2.1 by making a series of substitutions:

- (1) $(u \circ t \circ u)(\pi) \succeq (t \circ u)(\pi)$ Substitute $t \circ u$ for t in OI
- (2) $(u \circ t \circ u \circ t)(\pi) \succeq (t \circ u)(\pi)$ From (1) and OI, substitute $t(\pi)$ for π
- (3) $(u \circ t)(\pi) \succeq (t \circ u)(\pi)$ Right-hand side of OI
- (4) $(t \circ u \circ t)(\pi) \succeq (u \circ t)(\pi)$ Substitute $u \circ t$ for t and u for t in OI
- (5) $(t \circ u \circ t \circ u)(\pi) \succeq (u \circ t)(\pi)$ From (4) and OI, substitute t for u and u for t in OI, then substitute $u(\pi)$ for π
- (6) $(t \circ u)(\pi) \succeq (u \circ t)(\pi)$ Right-hand side of OI
- (7) $(t \circ u)(\pi) \sim (u \circ t)(\pi)$ From (3) and (6) □

We shall now illustrate how this technical result can be applied to the paparazzi example. We use the following pair of transformative decision

rules, and we assume without further ado that both rules satisfy the OI-condition.

The principle of insufficient reason (ir): *If π is a formal decision problem in which the probabilities of the states are unknown, then it may be transformed into a formal decision problem π' in which equal probabilities are assigned to all states.*

Merger of states (ms): *If two or more states yield identical outcomes under all acts, then these repetitious states should be collapsed into one, and if the probabilities of the two states are known they should be added.*

It can now be easily verified that Table 2.15 can be obtained from 2.13 by first applying the ir rule and then the ms rule. Furthermore, 2.17 can be obtained from 2.13 by first applying the ms rule and then the ir rule. Because of Theorem 2.1 we know that $(ir \circ ms)(\pi) \sim (ms \circ ir)(\pi)$. Hence, anyone who thinks that the two transformative rules we use satisfy the OI-condition is logically committed to the view that the two formalisations, 2.15 and 2.17, are equally reasonable (and at least as good as 2.13). However, although equally reasonable, 2.15 and 2.17 are of course very different – the expected utility principle even recommends different acts!

Exercises

2.1 If 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, 0 and 00. If the ball lands on the number you have chosen the croupier will pay you 35 times the amount betted, and return the bet.

- (a) Formalise and visualise the decision problem in a decision matrix.
- (b) Formalise and visualise the decision problem in a decision tree.
- (c) How much money can you expect to lose, on average, for every dollar you bet?

2.2 Formalise the following decision problem, known as Pascal's wager:

God either exists or He doesn't. It is abundantly fair to conceive, that there is at least 50% chance that the Christian Creator God does in fact exist. Therefore, since we stand to gain eternity, and thus infinity, the wise and safe choice is to live as though God does exist. If we are right, we gain everything, and lose nothing. If we are wrong, we gain nothing and lose

nothing. Therefore, based on simple mathematics, only a fool would choose to live a Godless life. Since you must choose, let us see which interests you least. You have nothing to lose. Let us estimate these two chances. If you gain, you gain all; if you lose, you lose nothing. Wager, then, without hesitation that He is. (Pascal 1660: §233)

- 2.3 (a) Is Pascal's argument convincing? (b) Is it really necessary for Pascal to assume that, "there is at least 50% chance that the Christian Creator God does in fact exist"? What if the probability is much lower?
- 2.4 Congratulations! You have won a free holiday in a city of your choice: London, New Delhi or Tokyo. You have been to London before, and you know that the city is okay, but expensive. New Delhi would be very exciting, unless you get a stomach infection; then it would be terrible. Tokyo would be almost as exciting, given that it is not too cold; then the trip would be rather boring.
- (a) Formalise and visualise the decision problem in a decision matrix.
 (b) Formalise and visualise the decision problem in a decision tree.
 (c) Represent the five possible outcomes in an ordinal scale.
- 2.5 A friend offers you to invest all your savings, \$100,000, in his dot com company. You find it very hard to understand the business plan he presents to you, but your friend tells you that your \$100,000 will 'certainly' be worth at least \$10M within two years. Naturally, your friend may be right, but he may also be wrong – you feel that you cannot estimate the probabilities for this. Consider the decision matrix below. What is wrong with this formalisation?

	Friend is right	Friend is wrong
<i>Invest</i>	\$10M	\$0
<i>Do not</i>	\$100,000	\$100,000

- 2.6 Visualise the following vector (which is written on a single line, to save space) in a decision matrix: $[[a_1, a_2, a_3]; [s_1, s_2]; [(a_1, s_1)=p, (a_1, s_2)=q, (a_2, s_1)=r, (a_2, s_2)=s, (a_3, s_1)=t, (a_3, s_2)=u]]$.
- 2.7 Explain the difference between (a) ordinal and cardinal scales, and (b) interval scales and ratio scales.
- 2.8 Your rich aunt died some time ago. In her will she stipulated that you shall receive a painting of your choice from her collection of impressionist art. The aesthetical values of her four paintings are, as