CAMBRIDGE INTRODUCTIONS TO PHILOSOPHY

An Introduction to Decision Theory

SECOND EDITION

MARTIN PETERSON



2 The Decision Matrix

Before you make a decision you have 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 of 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 formalization 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 visualize 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, because decision rules are only defined relative to a formal representation. For example, it makes no sense to say that the principle of maximizing 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 visualizing 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, we simply adds more lines. At the chance nodes nature decides which line to follow. 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
Take out insurance	No house and $100,000	House and $0
No insurance	No house and $100	House and $100

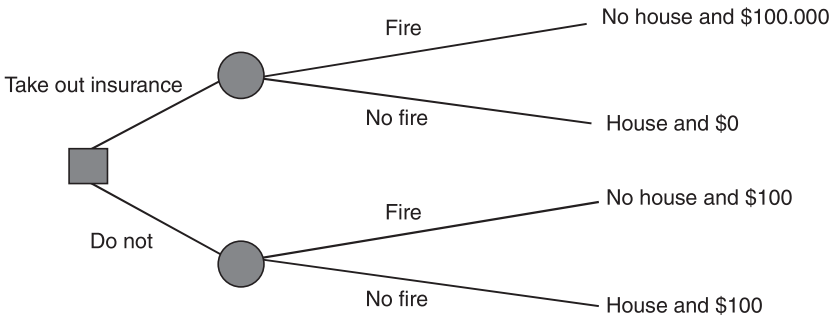


Figure 2.1

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 emphasizing that we are actually dealing with three levels of abstraction:

1. The decision problem
2. A formalization of the decision problem
3. A visualization of the formalization

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 formalization of a decision problem is made up of information about the decision to be made, irrespective of how that information is visualized. Formalizations thus comprise information about acts, states and outcomes, and sometimes also information about probabilities. Of course, one and the same decision problem can be formalized in different ways, not all of which are likely to be equally good. For example, some decision problems *can* be formalized either as decisions under risk or as decisions under ignorance. However, if one knows the probabilities of the relevant states, it is surely preferable to choose the

Table 2.2

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

former type of formalization (since one would otherwise overlook relevant information).

Naturally, any given set of information can be visualized in different ways. We have already demonstrated this by drawing a matrix and a tree visualizing the same formalization. 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 formalize decision problems is arguably more interesting than questions about how to visualize a given formalization. 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 visualized 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; acts

Table 2.3

	I choose the right bet	I do not
Bet on Democrat	$100	$0
Bet on Republican	$200	$0

performed by others can presumably be thought of as states). It is sometimes suggested that we should also think of the outcome brought about by an act as a state of the world. On that view, an act is essentially a device that transforms one state of the world into another state. However, as attractive as this suggestion might be, it will be helpful to keep apart states and outcomes in the present discussion.

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 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 wins the next presidential election (Table 2.3). When offered these bets, it would make little sense to consider the states *I choose the right bet* and *I do not*. That formalization gives the false impression that you will definitively be better off if you bet on the Republican candidate. The reason why this is false, and why the formalization is illegitimate, is that the state *I choose the right bet* is causally dependent on your choice. Whether the state will occur depends on which act you choose.

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 formalization; this is the option that will be

Table 2.4

Bet on Democrat	\$100 (0.6)	\$0 (0.4)
Bet on Republican	\$200 (0.3)	\$0 (0.7)

Table 2.5(a)

	State 1	State 2
Act A	\$50	\$80
Act B	\$30	\$80

Table 2.5(b)

Act A	\$50	\$80
Act B	\$80	\$30

pursued here (with exception for Chapter 9). The second option is to avoid the notion of states altogether. That approach works particularly well if one 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.

From a technical point of view, the formalization 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 the relevant 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. In this representation 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 formalization.

Table 2.6

	You survive	You die
Take first bet	$100	$0
Take second bet	$0	$100

Let us now consider a slightly different kind of problem, having to do with how preferences depend on which state happens to be the true state. Suppose you have been severely injured in a traffic accident and that 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 imagine that both states are equally probable. It then seems 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 anymore. This indicates that the formalization is underspecified. To make the formalization relevant for your decision we would have to add to the outcome the fact that the decision maker will die if the operation fails. If we did so, it would clearly transpire why the first bet is better than the second. To put all this in a more sophisticated vocabulary, to which we will return in Chapter 7, states should be chosen such that the value of the outcomes under all states is causally 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 will examine in more detail in Chapter 5. In the present section we will 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 analyze her attitudes toward 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 the 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 toward 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 will 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.

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

Table 2.7

	Fire	No fire
Take out insurance	1	4
Do not	−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

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 preserve 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

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

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 original 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

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

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.

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 the sake of illustration, 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:

$$F = 1.8 \cdot C + 32 \tag{2}$$

By solving this equation for C, we get:

$$C = (F - 32)/1.8 \tag{3}$$

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

Note that (2) and (3) are straight lines – had the graphs been curved, a difference of, say, one unit on the *x*-axis would 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 mean temperatures in a few well-known cities.

When looking at Table 2.10, a common error is to conclude that it was *twice as warm* in Tokyo as in New York, because 64 units Fahrenheit is twice as 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 lb = 9 kg, and this is twice as much as 10 lb = 4.5 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

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

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 · I. Furthermore, the first four scales, G–J, are equivalent interval scales. For example, scale H can be obtained from G by the formula H = 10 · G + 10, and J = 0.1 · 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.

It is helpful to summarize 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 constants k and m such that:

$$f'(x) = k \cdot f(x) + m \tag{4}$$

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:

$$f'(x) = k \cdot f(x) \tag{5}$$

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 which we use for transforming f' into f.)

2.3 Acts

Imagine that your 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

Box 2.1 Three types of scales

In this chapter we have discussed three different types of scales. Their main characteristics can be summarized 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. 10 kg is twice as much as 5 kg, and 10 lb is also twice as much as 5 lb. But 10 kg is not twice as much as 5 lb.

Table 2.12

	The sixth egg is rotten	The sixth egg is not rotten
Add sixth egg	No omelette	Six-egg omelette
Do not add sixth egg	Five-egg omelette	Five-egg omelette

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.

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 generalized 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, which was 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 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 fulfills 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 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 characterization 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 nonidentical 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 realize 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 $\{a_1 \,\&\, a_3, a_1 \,\&\, a_4\}$, or $\{a_1 \,\&\, a_3 \,\&\, a_5, a_1 \,\&\, a_3 \,\&\, a_6, a_1 \,\&\, a_4\}$? 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 will 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 formalize 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 Formalizations

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

Obviously, rival formalizations are troublesome if an act is judged to be rational in one optimal formalization of a decision problem, but irrational in another optimal formalization 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 formalization are nonrational in another rival formalization of the same decision problem, whereas *all* acts that are rational according to the latter formalization are not rational according to the former.

To give convincing examples of rival formalizations is difficult, mainly because it can always be questioned whether the suggested formalizations are equally reasonable. In what follows we will 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 formalizations.

Imagine that you are a paparazzi photographer and that rumor 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 visualized in Table 2.13.

□

Table 2.13

	P	LA	NY
Stay	$0	$10 k	$10 k
Go to Paris	$5 k	$6 k	$6 k

Table 2.14

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

Table 2.15

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

Table 2.16

	P	LA or NY
Stay	$0	$10 k
Go to Paris	$5 k	$6 k

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.

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.

Table 2.17

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

Which formalization is better, 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 formalizations.

Note that the principle of maximizing 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 formalizations 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 formalizations are possible. One may always question the claim that the two formalizations are equally reasonable. Therefore, in order to give a more comprehensive argument for thinking that the formalizations are equally reasonable, we will 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 formalization of a given decision problem. Examples of effective decision rules include the principle of maximizing 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 formalization of a decision problem into another by adding, deleting or modifying information in the initial formalization. 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 formalization 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 nonempty set of (relevant) alternative acts, S is a nonempty 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 formalize a decision problem?" can be restated as: "What sequence of transformative rules $(t \circ u \circ \ldots)$ 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 will 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 will prove that if the technical condition stated below (order-independence) 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 formalizations that are equally reasonable.

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

The left-hand inequality, $(u \circ t)(\pi) \succeq t(\pi)$, states that a transformative rule u should not, metaphorically expressed, throw a wrench in the works carried out by another rule t. Hence, the formalization obtained by first applying u and then t has to be at least as good as the formalization obtained by only applying t. The right-hand inequality, $t(\pi) \succeq (t \circ t)(\pi)$, 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 left-hand side of OI
(2)	$(u \circ t \circ u \circ t)(\pi) \succeq (t \circ u)(\pi)$	From (1) and left-hand side of OI
(3)	$(u \circ t)(\pi) \succeq (t \circ u)(\pi)$	Right-hand side of OI and (2)
(4)	$(t \circ u \circ t)(\pi) \succeq (t \circ u)(\pi)$	Substitute $u \circ t$ for t and u for t in OI
(5)	$(t \circ u \circ t \circ u)(\pi) \succeq (t \circ u)(\pi)$	From (4) and left-hand side of OI
(6)	$(t \circ u)(\pi) \succeq (t \circ u)(\pi)$	Right-hand side of OI and (5)
(7)	$(t \circ u)(\pi) \sim (u \circ t)(\pi)$	From (3) and (6)

□

We will 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 $(\text{ir} \circ \text{ms})(\pi) \sim (\text{ms} \circ \text{ir})(\pi)$. Hence, anyone who thinks that the two transformative rules we use satisfy the OI-condition is logically committed to the view that the two formalizations, 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 dealer will pay you 35 times the amount betted, and return the bet.
   (a) Formalize and visualize the decision problem in a decision matrix.
   (b) Formalize and visualize 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 Formalize 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 a 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 a 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 vacation 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) Formalize and visualize the decision problem in a decision matrix.
   (b) Formalize and visualize 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 $10 million 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 formalization?

	Friend is right	Friend is wrong
Invest	$10 m	$0
Do not	$100,000	$100,000

2.6 Visualize 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 (a) What type of scale is used for ranking universities in *U.S. News & World Report*'s latest ranking of universities? (You may wish to visit their website to learn about the methodology they use.) (b) Is this the appropriate type of scale to use for this purpose?

2.9 Anne has decided to spend $1,000 on a new handbag. There are four brands to choose from: Prada, Gucci, Chanel and Louis Vuitton. Her preferences can be represented on the following interval scale: Prada 100, Gucci 150, Chanel 80 and Louis Vuitton 0. Which of the following scales can be obtained from the original scale by a positive linear transformation?

(a) Prada 200, Gucci 300, Chanel 160 and Louis Vuitton 0.

(b) Prada 210, Gucci 310, Chanel 170 and Louis Vuitton 10.

(c) Prada 10, Gucci 15, Chanel 8 and Louis Vuitton 5.

2.10 Your rich aunt died some time ago. In her will she stipulated that you will receive a painting of your choice from her collection of Impressionist art. The aesthetical values of her four paintings are as follows, as measured on your personal interval scale: Manet 5,000, Monet 8,000, Pissarro 6,000 and Renoir 2,000. Which of the following scales can be obtained from the original scale by a positive linear transformation?

(a) Manet 8, Monet 11, Pissarro 9 and Renoir 5.

(b) Manet −250, Monet 2,750, Pissarro 750 and Renoir −3,250.

(c) Manet 1,000, Monet 3,000, Pissarro 2,950 and Renoir 995.

2.11 Show that scales (a)–(c) in Exercise 2.9 are equivalent ordinal scales.

2.12 Suppose that scale (a) in Exercise 2.9 is a ratio scale. Show that neither (b) nor (c) is equivalent to that ratio scale.

2.13 Show that every ratio scale is an interval scale, and that every interval scale is an ordinal scale.

2.14 Show that some interval scale is not a ratio scale.

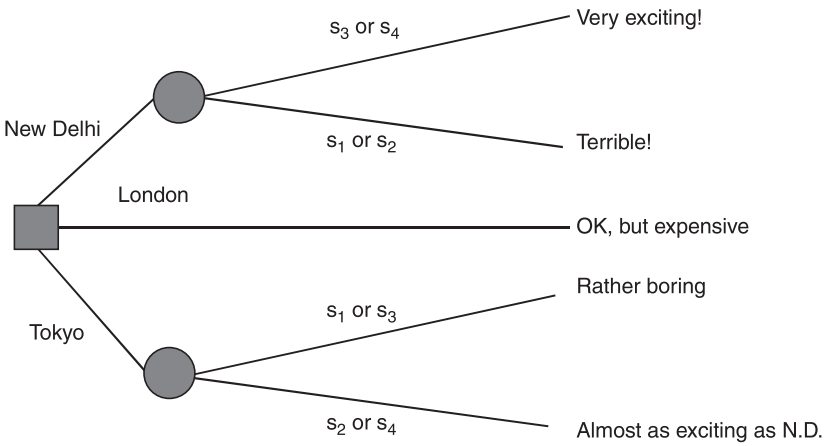
2.15 Suppose f is a ratio scale that can be transformed into f' by multiplying all values of f by k. Show that $f(x) = \frac{1}{k} \cdot f'(x)$.

Solutions

2.1 (a)

	Right number (1/38)	Wrong number (37/38)
Bet	\$35	−\$1
Do not	\$0	\$0

(b)



(c) $\left(\$(35+1)\cdot\frac{1}{38}+\$0\cdot\frac{37}{38}\right)-\$1\approx-\$.053$

2.2

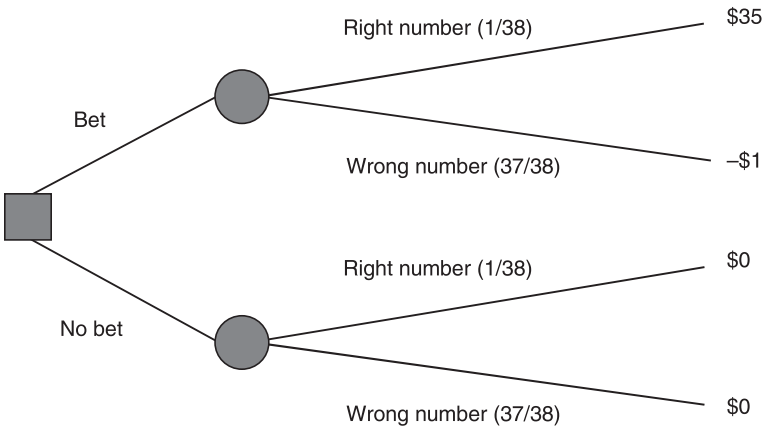
	God exists (1/2)	God does not exist (1/2)
Live as if God exists	Heaven – infinite gain.	You gain nothing and lose nothing.
Live a Godless life	Hell – infinite loss.	You gain nothing and lose nothing.

2.3 (a) Pascal's point is that the expected value of believing in God exceeds that of not believing. The secondary literature on this argument is massive, so I leave it to the reader to make up his or her own mind about it. (b) The assumption about a 50% probability is not important. All that matters is that the probability that God exists is > 0. (Since k times infinity equals infinity.)

2.4 We define four states: s_1 = If you are in Delhi you get a stomach infection, and it is cold in Tokyo; s_2 = If you are in Delhi you get a stomach infection, and it is not cold in Tokyo; s_3 = If you are in Delhi you do not get a stomach infection, and it is cold in Tokyo; s_4 = If you are in Delhi you do not get a stomach infection, and it is not cold in Tokyo. Now we draw the following matrix.

	s_1	s_2	s_3	s_4
London	Okay, but expensive	Okay, but expensive	Okay, but expensive	Okay, but expensive
New Delhi	Terrible!	Terrible!	Very exciting!	Very exciting!
Tokyo	Rather boring.	Almost as exciting as N.D.	Rather boring.	Almost as exciting as N.D.

(b) Let the states be defined as in (a).



(c) Here is my suggestion: Very exciting! = 5, Almost as exciting as N.D. = 4, Okay, but expensive = 3, Rather boring = 2, Terrible = 1.

2.5 Your friend can be wrong in many ways. If you get $5 m back your friend was literally speaking wrong, but you would presumably find that outcome quite pleasing. A satisfactory formalization should take this into account, e.g. by listing each possible (monetary) outcome as a separate state, or by distinguishing between only two states: *satisfactory profit* and *unsatisfactory profit*.

2.6

	s_1	s_2
a_1	p	q
a_2	r	s
a_3	t	u

2.7 See Box 2.1.

2.8 (a) The ranking is based on calculating the weighted average of six indicators. Each indicator is measured on a ratio scale. (b) A possible objection is that the weights assigned to the indicators are arbitrary.

2.9 (a) and (b).

2.10 (a) and (b).

2.11 For each scale f it holds that $f(\text{Monet}) > f(\text{Pissarro}) > f(\text{Manet}) > f(\text{Renoir})$.

2.12 (a) and (b) cannot be equivalent, because 8/11 $\neq$ −250/2,750. Furthermore, (a) and (c) cannot be equivalent, because 8/11 $\neq$ 1,000/3,000.

2.13 Let $k = 0$ in (4), and note that (1) will hold true for all positive constants k and all constants m in (4).

2.14 Here is an example: [0, 1, 2] and [1, 2, 3] are equivalent interval scales, but there is no number k such that $k \cdot 1/2 = k \cdot 2/3$.

2.15 By definition, $f'(x) = k \cdot f(x)$; hence, $f(x) = \frac{1}{k} \cdot f'(x)$.

3 Decisions under Ignorance

Jane is having a romantic dinner with her fiancé in a newly opened French bistro in Santa Barbara, California. After enjoying a vegetarian starter, Jane has to choose a main course. There are only two options on the menu, *Hamburger* and *Lotte de mer*. Jane recalls that *Lotte de mer* means monkfish, and she feels that this would be a nice option so long as it is cooked by a first-class chef. However, she has some vague suspicions that this may not be the case at this particular restaurant. The starter was rather poor and cooking monkfish is difficult. She knows that almost every chef can make a decent hamburger.

Jane feels that she cannot assign any probability to the prospect of getting good monkfish. She simply knows too little about this newly opened restaurant. Because of this, she is in effect facing a decision under ignorance. In decision theory *ignorance* is a technical term with a very precise meaning. It refers to cases in which the decision maker (i) knows what her alternatives are and what outcomes they may result in, but (ii) is unable to assign any probabilities to the states corresponding to the outcomes (see Table 3.1). Sometimes the term "decision under uncertainty" is used synonymously.

Jane feels that ordering a hamburger would be a safe option, and a hamburger would also be better than having no main course at all. Furthermore, she feels that good monkfish is better than an edible hamburger, but terrible monkfish is worse than having no main course. This ranking can be represented on an ordinal scale, as explained in Chapter 2. Since ordinal scales only preserve the order between objects, any set of numbers will do as long as better outcomes are represented by higher numbers. An example is given in Table 3.2.

Which alternative should Jane choose? To answer this question the decision maker has either to apply some decision rule or provide some

Table 3.1

	Good chef	Bad chef
Monkfish (lotte de mer)	good monkfish	terrible monkfish
Hamburger	decent hamburger	decent hamburger
No main course	hungry	hungry

Table 3.2

	s_1	s_2
a_1	4	1
a_2	3	3
a_3	2	2

other kind of structured reasoning for choosing one alternative over another. This chapter surveys some of the most influential rules for making decisions under ignorance. Later on, in Chapter 7, we will question the idea that it is sometimes impossible to assess the probability of a state in a meaningful way. If our desires and beliefs about the world are revealed by our choices, and partial beliefs can be represented by a probability function, it seems that we are, in a sense to be explained, always able to assign probabilities to states we know very little about.

3.1 Dominance

When Jane is about to choose a main course, it would clearly be irrational of her to choose a_3. (See Table 3.2.) This is because a_2 will lead to a better outcome than a_3 no matter which state happens to be the true state of the world; decision theorists say that alternative a_2 *dominates* alternative a_3. The widely accepted *dominance principle* prescribes that dominated acts must not be chosen.

Here is another example. Imagine that you wish to travel from London to Las Vegas. This route is serviced by two airlines, Black Jack airlines and Air Mojave, both of which offer equivalent fares and service levels. Just as you are about to book your ticket you read in the *Times* that Air Mojave is on the brink of bankruptcy. You therefore reason as follows: "No matter whether the information about Air Mojave is correct or not, I have nothing

Table 3.3

	Information correct	Information not correct
Black Jack airlines	Valid ticket	Valid ticket
Air Mojave	No valid ticket	Valid ticket

Table 3.4

	s_1	s_2
a_1	1	1
a_2	0	1

to lose by booking my ticket with Black Jack airlines. The fare and service is the same, but I will not face the same risk of losing my money. It is therefore more rational to book with Black Jack instead of Mojave." The decision matrix is given in Table 3.3.

You think that having a valid ticket is better than not having one. This ordinal relation can be represented by two arbitrary numbers, say 1 and 0. This yields the decision matrix in Table 3.4, in which it is easy to see that a_1 dominates a_2.

We are now in a position to state the dominance principle in a slightly more precise way. Briefly put, we wish to formalize the idea that one alternative is better than another if (but not only if) the agent knows for sure that she will be at least as well off if the first alternative is chosen, no matter what state of the world turns out to be the true state. Let $\succ$ be a relation on acts, such that $a_i \succ a_j$ if and only if it is more rational to perform act a_i rather than act a_j. This means that whenever $a_i \succ a_j$ is true, then a_i ought, rationally, to be chosen over a_j. Furthermore, $a_i \succeq a_j$ means that a_i is at least as rational as a_j, and $a_i \sim a_j$ means that the two acts are equally rational. Let v be an ordinal function that assigns values to outcomes, conceived of as ordered pairs of acts and states. For example, $v(a_1, s_1) = 1$ means that the value of doing a_1 in the case that s_1 is the true state of the world is 1.

The dominance principle comes in two versions: Weak and Strong.

Weak Dominance: $a_i \succeq a_j$ if and only if $v(a_i, s) \geq v(a_j, s)$ for every state s.

Strong Dominance: $a_i \succ a_j$ if and only if $v(a_i, s_m) \geq v(a_j, s_m)$ for every state s_m, and there is some state s_n such that $v(a_i, s_n) > v(a_j, s_n)$.

Weak Dominance articulates the plausible idea that one alternative is at least as rational as another if all its outcomes, under every possible state of the world, will be at least as good as that of the other alternative. Strong Dominance stipulates that for an alternative to be more rational than another, two conditions must be fulfilled. First, its outcomes under all states must be at least as good as those of the other alternative. Secondly, there must also be *at least one* state under which the outcome is strictly better than that of the other alternative. For example, booking a ticket with Black Jack airlines strongly dominates booking a ticket with Air Mojave.

The dominance principle can also be applied to decisions under risk. If we know the probabilities of the states, it still makes sense to choose one alternative over another if it is certain to lead to at least as good an outcome no matter which state turns out to be the true state. Hence, it does not really matter if the probabilities of the states involved are known by the decision maker or not. All that matters is that we know that we will be at least as well off if we choose one alternative over another, no matter which state of the world turns out to be the true state.

Unfortunately, the dominance principle is not the whole truth about decision making under ignorance. It can only provide a partial answer. This is because it sometimes fails to single out a set of optimal acts. For example, in Jane's decision at the restaurant in Santa Barbara, act a_1 is not dominated by a_2, nor is a_2 dominated by a_1. All that can be concluded is that it would *not* be rational to choose a_3. This means that some additional criterion is required for reaching a decision.

3.2 Maximin and Leximin

The maximin principle focuses on the worst possible outcome of each alternative. According to this principle, one should *maximize* the *minimal* value obtainable with each act. If the worst possible outcome of one alternative is better than that of another, then the former should be chosen. Consider the decision matrix in Table 3.5.

The worst possible outcome of each alternative is underlined. The best of these worst-case scenarios is underlined and printed in bold face. The maximin rule thus prescribes that one should choose alternative a_3, because the worst possible outcome of that act is 2, which is better than the worst-case scenarios of all other alternatives (0, -5, and -8).

Table 3.5

	s_1	s_2	s_3	s_4
a_1	6	9	3	0
a_2	−5	7	4	12
a_3	6	4	5	**2**
a_4	14	−8	5	7

Table 3.6

	Good chef	Bad chef
Monkfish	Best outcome	Worst outcome
Hamburger	2nd best outcome	2nd best outcome
No main course	3rd best outcome	3rd best outcome

The maximin rule is easy to formalize. Let $\min(a_i)$ be the minimal value obtainable with act a_i. Then,

Maximin: $a_i \succeq a_j$ if and only if $\min(a_i) \geq \min(a_j)$.

The maximin rule does not require that we compare outcomes on an interval scale. All we need is an ordinal scale, since the "distance" between outcomes is irrelevant. In order to illustrate this point, we return to Jane in Santa Barbara. The matrix in Table 3.6 comprises enough information for applying the maximin rule. Clearly, the maximin rule prescribes that Jane should order a hamburger, because the worst possible outcome would then be better than if any other alternative is chosen.

If the worst outcomes of two or more acts are equally good, the maximin rule tells you to be indifferent. Hence, if the outcome of some act will be either \$1 or \$2, while some other act will give you either \$1 or \$10,000, you should be indifferent between those options. Yet this does not ring true. The standard solution to this problem is to invoke the *lexical* maximin rule, which is slightly more sophisticated. It holds that if the worst outcomes are equal, one should choose an alternative in which the *second worst* outcome is as good as possible. If this does not single out a unique act, then the *third worst* outcome should be considered, and so on. Hence, unless all possible outcomes of the alternatives are exactly parallel, the lexical maximin rule (leximin rule) will at some point single out one or several acts that are better

Table 3.7

	s_1	s_2	s_3	s_4
a_1	6	14	8	<u>5</u>
a_2	<u>5</u>	6	100	100
a_3	7	7	7	<u>5</u>

Table 3.8

	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	s_9	s_{10}
a_1	1	100	100	100	100	100	100	100	100	100
a_2	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1

than all the others. Consider the example in Table 3.7, in which the traditional maximin rule ranks all alternatives as equally good, whereas the leximin rule recommends a_3.

To formalize the leximin rule, let $\min^1(a_i)$ be the value of the worst outcome of act a_i, and $\min^2(a_i)$ be the value of its second worst outcome, and $\min^n(a_i)$ be the value of the nth worst outcome of a_i. Then,

Leximin: $a_i \succ a_j$ if and only if there is some positive integer n such that $\min^n(a_i) > \min^n(a_j)$ and $\min^m(a_i) = \min^m(a_j)$ for all $m < n$.

What is the best argument for accepting maximin or leximin? It is tempting to reason as follows: In a certain sense, the maximin and leximin rules allow the decision maker to transform a decision under ignorance into a decision under certainty. Surely, a decision maker accepting the maximin or leximin rule does not know what the *actual* outcome of her choice will be. However, she knows for sure what the *worst possible* outcome will be. This is a kind of partial certainty, which many people seem to find attractive. If one accepts the premise that the worst possible outcome should guide one's decisions, then the maximin and leximin rules allow one to make decisions without fearing the actual outcome – the actual outcome will always be at least as good as the worst possible outcome.

However, in some decision problems it seems rather irrational to focus entirely on the worst-case scenario. Consider the example in Table 3.8.

The maximin and leximin rules prescribe that a_2 should be chosen over a_1. However, from a normative point of view this appears a bit

Table 3.9

	s_1	s_1 or s_2 or s_3 . . . or s_{10}
a_1	1	100
a_2	1.1	1.1

strange. Most people would surely choose a_1, because most people would judge that it is worth sacrificing 0.1 units of value for the possibility of getting 100 units.

What could one say in reply? To start with, note that the decision matrix in Table 3.8 gives the false impression that 100 units of value will be a much more *probable* outcome than 1 if the first alternative is chosen. However, the decision maker is making a decision under ignorance, so nothing can be concluded about the probabilities of the outcomes. The mere fact that there are more states corresponding to 100 rather than 1 does not imply that 100 is a more probable outcome than 1. In fact, by merging the nine rightmost states (s_2–s_{10}) into a single (disjunctive) state, we get Table 3.9, which is an equally good representation of the same decision problem.

Now, one no longer gets the false impression that 100 is a more probable outcome than 1. But despite this maneuver, it may still be reasonable to maintain that the first alternative is better than the second, contrary to the recommendations of the maximin and leximin rules. Most of us would probably agree that it is rational to trade the risk of getting 1 instead of 1.1 for a chance of getting 100.

3.3 Maximax and the Optimism–Pessimism Rule

The maximin rule singles out the worst possible outcome of a decision as its only normatively relevant feature. However, one could of course equally well single out the *best* possible outcome as being particularly relevant. As I write this text, best-selling author J.K. Rowling is just about to release her seventh Harry Potter book. We all know that it will sell millions of copies. Clearly, as Ms. Rowling wrote her first story about an unknown boy called Harry Potter and submitted it to a publisher, she did not think much about the worst-case scenario, i.e. to have the manuscript rejected. What triggered her to actually submit the book to the publisher was no doubt her desire to achieve the *best* possible outcome of this decision. Luckily for her,

the best possible outcome also turned out to be the *actual* outcome. If all authors considering submitting a manuscript always decided to focus on the worst-case scenario, i.e. rejection and disappointment, no books would ever be published.

The J.K. Rowling example serves as a nice illustration of the maximax rule. Advocates of the maximax rule maintain that rationality requires us to prefer alternatives in which the best possible outcome is as good as possible, i.e. one should *maximize* the *maximal* value obtainable with an act. The maximax rule has surprisingly few adherents among decision theorists, but as far as I can tell, it is about as reasonable as the maximin rule. However, maximax also serves as a helpful introduction to another rule, which has attracted considerably more attention by decision theorists. This is the optimism–pessimism rule, originally proposed by Hurwicz. (A synonymous name for this rule is the alpha-index rule.) The optimism–pessimism rule asks the decision maker to consider both the best and the worst possible outcome of each alternative, and then choose an alternative according to her degree of optimism or pessimism. In order to render this suggestion precise, suppose that a decision maker's degree of optimism can be represented by a real number α between 0 and 1, such that $\alpha = 1$ corresponds to maximal optimism and $\alpha = 0$ to maximal pessimism. Then, if $\max(a_i)$ is the best possible outcome of alternative a_i and $\min(a_i)$ its worst possible outcome, its value is $\alpha \cdot \max(a_i) + (1 - \alpha) \cdot \min(a_i)$. Naturally, α is assumed to be fixed throughout the evaluation of all alternatives a_i.

Consider the example in Table 3.10. Note that we have omitted the states corresponding to the outcomes. It is not necessary to always write them down.

The best possible outcome of a_1 is 100 and the worst possible outcome is 10. That is, $\max(a_1) = 100$ and $\min(a_1) = 10$. Hence, if the decision maker is optimistic to degree 0.7, it follows that the total value of a_1 is $0.7 \cdot 100 + (1 - 0.7) \cdot 10 = 73$. Furthermore, the best possible outcome of alternative a_2 is 90, whereas its worst outcome is 50. Hence, the value of

Table 3.10

a_1	55	18	28	10	36	100
a_2	50	87	55	90	75	70

this alternative is $0.7 \cdot 90 + (1 - 0.7) \cdot 50 = 78$. Since $78 > 73$, it follows that alternative a_2 is better than alternative a_1.

Optimism–Pessimism: $a_i \succ a_j$ if and only if $\alpha \cdot \max(a_i) + (1 - \alpha) \cdot \min(a_i) > \alpha \cdot \max(a_j) + (1 - \alpha) \cdot \min(a_j)$

Note that if $\alpha = 0$, then the optimism–pessimism rule collapses into the maximin rule, since $0 \cdot \max(a_i) + (1 - 0) \cdot \min(a_i) = \min(a_i)$. Furthermore, the optimism–pessimism rule will be identical with the maximax rule if $\alpha = 1$. The optimism–pessimism rule can therefore be conceived of as a generalization of the maximin and maximax rules.

The optimism–pessimism rule clearly requires that value can be measured on an interval scale, because if one were to measure value on an ordinal scale it would make no sense to perform arithmetical operations such as multiplication and addition on the values. However, the maximin and maximax rules only require that value is measured on ordinal scales. The claim that the optimism–pessimism rule is a generalization of maximin and maximax must therefore be taken with a pinch of salt.

The main advantage of the optimism–pessimism rule is that it allows the decision maker to pay attention to the best-case *and* the worst-case scenarios at the same time. The α-index describes the relative importance of both outcomes. The obvious objection is, of course, that it seems irrational to totally ignore the outcomes in between. Should one really ignore all the nonminimal and nonmaximal outcomes? Please take a look at Table 3.11. The max and min values of both alternatives are the same, so for every α it holds that both alternatives are equally good. However, intuitively it seems reasonable to choose the first alternative, because under almost all states it will lead to a significantly better outcome, and there is only one state in which it leads to a slightly worse outcome.

An obvious objection to this argument is that it indirectly invokes a sort of probabilistic reasoning. Instead of assigning numerical probabilities to the states, the argument falsely presupposes that one alternative is more

Table 3.11

a_1	10	10	99	99	...	99	99	100
a_2	10	11	11	11	...	11	11	100

likely than another if there are more states leading to a certain outcome. In a decision under ignorance nothing is known about probabilities.

However, the basic objection remains. Why is it rational to focus exclusively on the best and the worst possible outcomes? What about the other, nonextreme outcomes? A natural response might be to assign α-values to all outcomes, not only to the best and worst ones. If there are n possible outcomes, ordered from the best to the worst, the numbers $\alpha_1, \alpha_2, \ldots, \alpha_n$ should be chosen such that $\alpha_1 + \alpha_2 + \ldots + \alpha_n = 1$. Each number α_i describes the relative weight of outcome i. The value of the act is thereafter calculated as above, i.e. by multiplying the α-index and value of each outcome and thereafter adding all α-weighted outcomes. It should be emphasized that $\alpha_1, \alpha_2, \ldots, \alpha_n$ cannot be interpreted as probabilities, because $\alpha_1, \alpha_2, \ldots, \alpha_n$ are chosen according to how important the decision maker feels that the best and the second-best and the third-best outcomes are, and so on. There is in general no reason to think that the relative importance of an outcome corresponds to its probability.

3.4 Minimax Regret

Imagine that you are invited to invest in a newly launched software company. Some of your friends, who invested in similar companies a few years ago, made a fortune. Of course, you may lose the money, but the investment is rather small, so the fact that the company may go bankrupt is not a genuine issue of concern. However, you know that if you *refrain* from making the investment and the company turns out to be yet another success story, you will feel a lot of regret. Some academics find this intuition important and argue that the concept of regret is relevant to rational decision making. Briefly put, advocates of the *minimax regret* rule maintain that the best alternative is one that minimizes the maximum amount of regret. An example is given in Table 3.12.

We assume that outcomes are compared on an interval scale, not merely an ordinal scale. The numbers in bold face denote the best alternative relative to each state. For example, 20 is highlighted in the rightmost column because a_1 would be the best alternative to choose given that the rightmost state was to occur. Now, for each outcome in the decision matrix, regret values can be calculated by subtracting the value of the best outcome of each state (the number in bold face) from the value of the outcome in

Table 3.12

a_1	12	8	20	**20**
a_2	10	**15**	16	8
a_3	**30**	6	25	14
a_4	20	4	**30**	10

Table 3.13

a_1	**−18**	−7	−10	0
a_2	**−20**	0	−14	−12
a_3	0	**−9**	−5	−6
a_4	−10	**−11**	0	−10

question. For example, the regret value of the lowermost outcome in the rightmost column is −10, because 10 − 20 = −10, and the regret value of the second lowermost outcome in the rightmost column is −6, because 14 − 20 = −6. By doing the analogous calculations for all outcomes in the decision matrix, we obtain a *regret matrix*. See Table 3.13.

In the regret matrix, the numbers in bold face denote the maximum regret value of each alternative act. For example, −18 is worse than −7, −10 and 0. Hence, −18 is the highest amount of regret obtainable with a_1. Note that all regret values will always be negative or zero. There is no such thing as "positive" regret.

The minimax regret rule urges decision makers to choose an alternative under which the maximum regret value is as low as possible. Since −9 corresponds to less regret than −18, −20 and −11, act a_3 is the best alternative in this example. Note that once the regret matrix has been obtained from the original decision matrix, the minimax regret rule is parallel to the maximin rule. In a certain sense, both rules select an alternative under which the worst-case scenario is as good as possible, but the two rules operate on different descriptions of the relevant worst-case scenario (actual outcome versus possible regret). Another important difference is that the minimax regret rule requires that outcomes are compared on a cardinal scale, whereas the maximin rule operates on ordinal values.

Unsurprisingly, the minimax regret rule will look a bit complex when formalized. Let Max{ } be a function that returns the maximum value of a sequence of real numbers in the set { }. Furthermore, let max(s_i) be

Table 3.14

a_1	12	8	20	20
a_2	10	**15**	16	8
a_3	**30**	6	25	14
a_4	20	4	**30**	10
a_5	−10	10	10	**39**

Table 3.15

a_1	−18	−7	−10	**−19**
a_2	−20	0	−14	**−31**
a_3	0	−9	−5	**−25**
a_4	−10	−11	0	**−29**
a_5	**−40**	−5	−20	0

a function that returns the best possible outcome that may be achieved with any alternative if s_i is the true state of the world. Then, the minimax regret rule can be written as follows.

Minimax regret: $a_1 \succ a_2$ if and only if $\max\{(v(a_1, s_1) - \max(s_1)), (v(a_1, s_2) - \max(s_2)), \ldots\} < \max\{(v(a_2, s_1) - \max(s_1)), (v(a_2, s_2) - \max(s_2)), \ldots\}$

The most prominent objection to minimax regret is the objection from *irrelevant alternatives*. Consider the decision matrix in Table 3.14, which is obtained from Table 3.12 by adding a new alternative, a_5.

The regret matrix corresponding to Table 3.14 is given in Table 3.15. It can be easily verified that in this matrix, alternative a_1 comes out as the best alternative according to the minimax regret rule – not a_3 as in the previous example. Why is this problematic? The answer is that the ranking of the alternatives has been altered *by adding a nonoptimal alternative*. That is, by simply adding a bad alternative that will not be chosen (a_5), we have altered the order between other, better alternatives. Intuitively, this is strange, because why should one's choice between, say, a_1 and a_3 depend on whether a suboptimal alternative such as a_5 is considered or not? It seems reasonable to maintain that a normatively plausible decision rule must not be sensitive to the addition of irrelevant alternatives. Therefore, since the minimax regret rule does not meet this requirement it has to be ruled out as an illegitimate rule. However, as pointed out in Box 3.1 this objection is not undisputed.

Box 3.1 Minimax and the objection from irrelevant alternatives

The objection from irrelevant alternatives is often thought to be a decisive argument against the minimax regret rule. However, it is far from clear that the objection is as strong as it may at first appear to be. Imagine that Bizarro, the world famous artist, produced fifty paintings in his lifetime. Some of these masterpieces are predominantly red while the others are predominantly blue. For many years no more than two paintings, one red and one blue, have been in private ownership. All remaining paintings are government property and stored in buildings which are closed to the public. However, as luck has it you are rather well off and can afford to buy any one of the two paintings in private ownership. Upon studying both paintings you decide that you much prefer the red to the blue. In an attempt to reduce its budget deficit, the government suddenly decides to sell off one of its Bizarros. The painting that now becomes available for choice is a red one. You find the new painting quite unappealing, and you like both paintings in the initial choice set more than the new one. Furthermore, for some reason, never mind why, the addition of the new painting to the choice set does not affect the prices of the old ones. In fact, *everything* else remains the same. The only difference is that one more painting has been added to the choice set. Nonetheless, you no longer prefer the first red Bizarro to the blue. Since the addition of a second red painting to the choice set makes the blue *unique* (it is the only blue Bizarro in the choice set), you switch your preference. You now prefer the blue over the red, and it is the mere addition of a second red painting to the choice set that makes you switch your preference.

It can hardly be denied that people sometimes reason like this. But is this holistic reasoning rational? I think so. When stating preferences over a set of alternatives, it is rational to compare each alternative with the entire set of alternatives, even if some of those alternatives are certain not to be chosen. By adding a new object to the choice set, the relational properties of the initial objects may change, and sometimes our preferences are based on relational properties, such as uniqueness.

3.5 The Principle of Insufficient Reason

A number of prominent thinkers, including Bernoulli (1654–1705) and Laplace (1749–1827), have defended a rule known as the *principle of insufficient reason*. Imagine that you have booked a weekend trip to a foreign country. Unfortunately, you know nothing about the weather at your destination. What kind of clothes should you take? In your wardrobe you have clothes for three types of weather, viz. hot, moderate and cold weather. The decision matrix is given in Table 3.16.

The principle of insufficient reason prescribes that if one has *no* reason to think that one state of the world is more probable than another, then all states should be assigned *equal* probability. In the example above there are three possible states, so the probability assigned to each state will be 1/3 (Table 3.17). More generally speaking, if there are n states the probability $p(s)$ assigned to each state will be $1/n$.

By applying the principle of insufficient reason, an initial decision problem under ignorance is transformed into a decision problem under risk, that is, a decision problem with known probabilities and values. As will be further explained in Chapter 4, many decision theorists think that one

Table 3.16

	Actual temperature is 100 °F	Actual temperature is 50 °F	Actual temperature is 0 °F
Bring clothes suitable for 100 °F	15	0	−30
Bring clothes suitable for 50 °F	0	15	0
Bring clothes suitable for 0 °F	−15	0	15

Table 3.17

	s_1 (1/3)	s_2 (1/3)	s_3 (1/3)
a_1	15	0	−30
a_2	0	15	0
a_3	−15	0	15

should maximize expected value in decisions under risk. In our example, the expected value of bringing clothes suitable for 100 °F is $1/3 \cdot 15 + 1/3 \cdot 0 + 1/3 \cdot -30 = -5$. The expected value of the second alternative, i.e. bringing clothes suitable for 50 °F, is $1/3 \cdot 0 + 1/3 \cdot 15 + 1/3 \cdot 0 = 5$, and the expected value of the third alternative is $1/3 \cdot -15 + 1/3 \cdot 0 + 1/3 \cdot 15 = 0$. Hence, it is rational to choose the second alternative, i.e. to bring clothes suitable for 50 °F.

Note that the principle of insufficient reason in itself does not recommend any particular act. It merely *transforms* a decision under ignorance into a decision under risk, by assigning equal probabilities to all states. Hence, the principle of insufficient reason could be combined with any decision rule for making decisions under risk. However, nearly everyone advocating the principle of insufficient reason also thinks that decisions under risk should be taken by calculating expected value. This composite version of the principle of insufficient reason can be formalized in the following way.

The principle of insufficient reason:
$a_i \succ a_j$ if and only if $\sum_{x=1}^{n} \frac{1}{n} v(a_i, s_x) > \sum_{x=1}^{n} \frac{1}{n} v(a_j, s_x)$

An obvious problem with the principle of insufficient reason is that it is extremely sensitive to how states are individuated. Suppose that before embarking on a trip you consider whether to bring an umbrella or not. As before, you know nothing about the weather at your destination. If the formalization of the decision problem is taken to include only two states, viz. rain and no rain, the probability of each state will be 1/2. However, it seems that one might just as well go for a formalization that divides the space of possibilities into three states, viz. heavy rain, moderate rain and no rain. If the principle of insufficient reason is applied to the latter set of states, their probabilities will be 1/3. In some cases this difference will affect our decisions. Hence, it seems that anyone advocating the principle of insufficient reason must ascertain the rather implausible hypothesis that there is only one correct way of making up the set of states.

Traditionally, decision theorists advocating the principle of insufficient reason have tried to make it more plausible by stressing the importance of considering *symmetric* states. For example, if someone gives you a die with n sides and you have no reason to think that any side is more likely than another (i.e. you have no reason to think that the die is biased),

then you should assign a probability of $1/n$ to each side. However, it seems that not all events can be described in symmetric terms, at least not in a way that justifies the conclusion that they are equally probable. Whether Ann's marriage will be a happy one depends on her future emotional attitude toward her husband. According to one description, she could be either in love or not in love with him; then the probability of both states would be 1/2. According to another equally plausible description, she could either be deeply in love, a little bit in love or not at all in love with her husband; then the probability of each state would be 1/3. In both cases the symmetry condition seems to be satisfied, which shows that it is too weak.

Here is another objection to the principle of insufficient reason: If one has *no* reason to think that one state is more probable than another, it seems strange to conclude that they are *equally* probable. If one has *no* reason to think that some state is *twice* as probable as another, we would not conclude that it is twice as probable as the other. More generally speaking, the problem is that it seems *completely arbitrary* to infer that all states are equally probable. Any other distribution of probabilities seems to be equally justified (that is, not at all).

It might be replied that the claim that all states are equally probable is in a certain sense special. Because if one has no reason to think that s_1 is twice as probable as s_2, it may also hold true that one has no reason to think that s_2 is twice as probable as s_1. However, it cannot be the case that s_1 *is* twice as probable as s_2 and that s_2 *is* twice as probable as s_1. Only if we assign equal probabilities to all states can we avoid ascribing several *different* probabilities to one and the same state.

Even if this line of reasoning might have something going for it the basic objection remains. The principle of insufficient reason could equally well be taken to read, "If you have no reason for inferring anything about the probabilities of the states, then assign equal probabilities to all states." Once the principle is formulated in this way, it becomes natural to ask: Why favor any particular way of assigning probabilities to states over another? *Any* assignment of probability seems to be exactly as justifiable as another (that is, not at all). In the case with two states, 1/2 and 1/2 is just one of many options. One could equally well assign probability 0.6 to the first state and 0.4 to the second, and so on.

3.6 Randomized Acts

Instead of selecting one act over another, one may instead toss a coin and, for example, choose the first alternative if it lands heads up and the second otherwise. Consider the decision problem in Table 3.18.

Randomizing between a_1 and a_2 means that a new alternative is introduced, which is the act of randomizing between the other acts, thereby transforming the initial decision matrix into a new matrix with three alternatives. If the randomization is performed by tossing a fair coin, the new decision matrix will be as in Table 3.19.

In the initial, nonrandomized matrix, all acts come out equally well according to all decision rules discussed in this chapter. However, if we add the randomized act a_3, this act will be chosen by the maximin, leximin and minimax regret rules (as well as by the optimism–pessimism rule if $\alpha > 1/2$). Hence, a_3 is strictly better than both other acts. So it seems that something can be gained by allowing randomized acts.

It should be emphasized that the randomized act a_3 will not lead to a *certain* outcome worth 1/2, however. The number 1/2 is the *expected* value of randomizing between a_1 and a_2 with a probability of 1/2. (Of course, any probability could be chosen; 1/2 is just an example.) This means that there are two cases to consider. To start with, consider the case in which the decision maker faces some decision problem only once. Suppose that you have been walking through the desert for days with no food or water. You now come to a crossroads. Should you choose path a_1 or a_2? If you make the right decision you will reach an oasis; otherwise you will die. What would be the point of randomizing between the two alternatives in this case? Depending on which path you choose you will either die or survive.

Table 3.18

a_1	1	0
a_2	0	1

Table 3.19

a_1	1	0
a_2	0	1
a_3	1/2	1/2

To randomize between the two paths does not mean that you will "die to degree 1/2," or get any other outcome between the two extremes. You will either die or reach the oasis. Hence, it is far from clear exactly what is gained by randomizing in this type of case.

In the second type of case, you know that you will face the same type of decision over and over again. Therefore, if you randomize between a_1 and a_2, you will *on average* get 1/2. This conclusion is a bit strange, however. If one faces exactly the same decision problem over and over again, literally speaking, then the truth about the two states would be fixed once and for all. Hence, it would be better to choose either a_1 or a_2 the first time and then *adjust* one's future behaviour to the outcome of the first decision. If one chooses a_2 and gets 0, one should switch to a_1, which would then be certain to lead to 1 in all remaining decisions. Furthermore, if the truth about the states is *not* fixed once and for all, one should of course choose either a_1 or a_2, stick to it, and then adjust one's behavior as one gradually learns more about the relative frequencies of the states.

Box 3.2 Axiomatic analysis of decisions under ignorance

The decision rules discussed in this chapter sometimes yield conflicting prescriptions, so not all of them can be accepted by a rational decision maker. For example, in Table 3.20 maximin recommends a_1, minimax regret and the principle of insufficient reason recommend a_2, whereas the optimism–pessimism rule recommends a_3 (for $\alpha = 1/2$).

Which rule is best? To argue that one rule is better than another because its normative recommendations are more intuitively plausible is unlikely to impress the decision theorist. Some stronger argument is needed. Traditionally, decision theorists seek to show that one rule is better than another by providing an axiomatic justification. Briefly put, the idea is to formulate a set of fundamental principles from which a decision rule can be derived – or, in case one wants to argue against a rule, to find a set of plausible conditions which are inconsistent with the rule in question. Several authors have suggested

Table 3.20

a_1	20	15	1
a_2	19	20	0
a_3	0	30	0

such axiomatic analyses, but the most accessible and elegant one was presented in a classic paper by Milnor.

Milnor assumed that a plausible decision rule must not only single out a set of optimal acts, but also provide an ordinal ranking of all alternatives. Let $\succeq$ be a binary relation on acts such that $a_i \succeq a_j$ if and only if it is at least as rational to choose a_i as a_j. Furthermore, $a_i \succ a_j$ means that $a_i \succeq a_j$ & not-$(a_i \succeq a_j)$, and $a_i \sim a_j$ means that $a_i \succeq a_j$ and $a_j \succeq a_i$. Milnor then asked us to consider a set of ten axioms. (For simplicity, we state the axioms informally.)

1. **Ordering**: $\succeq$ is transitive and complete. (See Chapter 5.)
2. **Symmetry**: The ordering imposed by $\succeq$ is independent of the labeling of acts and states, so any two rows or columns in the decision matrix could be swapped.
3. **Strict dominance**: If the outcome of one act is strictly better than the outcome of another under every state, then the former act is ranked above the latter.
4. **Continuity**: If one act weakly dominates another in a sequence of decision problems under ignorance, then this holds true also in the limit decision problem under ignorance.
5. **Interval scale**: The ordering imposed by $\succeq$ remains unaffected by a positive linear transformation of the values assigned to outcomes.
6. **Irrelevant alternatives**: The ordering between old alternatives does not change if new alternatives are added to the decision problem.
7. **Column linearity**: The ordering imposed by $\succeq$ does not change if a constant is added to a column.
8. **Column duplication**: The ordering imposed by $\succeq$ does not change if an identical state (column) is added.
9. **Randomization**: If two acts are equally valuable, then every randomization between the two acts is also equally valuable.
10. **Special row adjunction**: Adding a weakly dominated act does not change the ordering of old acts.

How do these axioms relate to the rules discussed in this chapter? Consider Table 3.21. The symbol – means that a decision rule and axiom are incompatible, whereas $\times$ indicates that they are compatible.

Table 3.21

	Maximin	Optimism–pessimism	Minimax regret	Insufficient reason
1. Ordering	⊗	⊗	⊗	⊗
2. Symmetry	⊗	⊗	⊗	⊗
3. Strict dominance	⊗	⊗	⊗	⊗
4. Continuity	⊗	⊗	⊗	×
5. Interval scale	×	⊗	×	×
6. Irrelevant alternatives	⊗	⊗	–	⊗
7. Column linearity	–	–	⊗	⊗
8. Column duplication	⊗	⊗	⊗	–
9. Randomization	⊗	–	⊗	×
10. Special row adjunction	×	×	⊗	×

Table 3.22

a_1	1	2	a_1	11	2
a_2	0	3	a_2	10	3

Furthermore, a decision rule is characterized by the axioms marked with the symbol ⊗, that is, those axioms constitute necessary and sufficient conditions for the decision rule in question.

The table summarizes 44 separate theorems. We shall not prove all of them here, but for illustrative purposes we prove one theorem of each type. Consider the following examples.

Theorem 3.1 Column linearity is incompatible with maximin.

Proof We construct a counterexample to the claim that column linearity and maximin are compatible. Consider the pair of decision problems in Table 3.22.

The rightmost matrix can be obtained from the leftmost by adding 10 to its first column. However, in the leftmost matrix, the maximin rule prescribes a_1, whereas it prescribes a_2 in the rightmost one. Hence, contrary to what is required by column linearity, maximin entails that it is false that the ordering of acts remains unchanged if a constant is added to a column. □

Theorem 3.2 Randomization is compatible with the principle of insufficient reason.

Proof According to the principle of insufficient reason, the value of an arbitrary act a_i equals $\sum_{m=1}^{n} = \frac{1}{n} v(a_i, s_m)$. Call this sum z. Now, it follows that if two acts are to be equally good, then their values z and z' must be the same. Therefore, if we randomize between two acts with identical values z, then the value of the randomized act will also be z. □

Theorem 3.3 The principle of insufficient reason is characterized by Axioms 1, 2, 3, 6 and 7.

Proof We wish to prove that Axioms 1, 2, 3, 6 and 7 are necessary and sufficient conditions for the principle of insufficient reason. I leave it to the reader to verify that they are necessary, i.e. that the principle of insufficient reason entails the axioms.

In order to prove that the axioms are sufficient we first prove the following lemma: The axioms guarantee that two acts which differ only in the order of their outcomes are ranked as equally good. In order to see why this is true, take two arbitrary acts and add a sequence of intermediate acts (by applying Axiom 6) such that any two acts differ only by a permutation of two outcomes. Our lemma now follows by applying Axiom 2 to each pair of successive rows.

Next, we wish to show that if the average outcomes of two acts a_i and a_j are the same, then the axioms guarantee that $a_i \sim a_j$. It is helpful to proceed in two steps. (i) We first use our lemma to permute the outcomes of a_i and a_j in order of increasing value, i.e. such that the worst outcome is listed first and the best last. (ii) In the second step we apply Axiom 7 and subtract from each column of a_i and a_j the number corresponding to its worst outcome. By repeating steps (i) and (ii) a finite number of times, we obtain a decision matrix with two rows in which all numbers are zero. (This is because we stipulated that the two acts have the same average outcome.) Hence, since the new versions have exactly the same outcomes, Axioms 1 and 3 imply that $a_i \sim a_j$. □

Exercises

3.1 Which act(s) should be chosen in the decision problem below according to the following decision rule: (a) the maximin rule, (b) maximax, (c) minimax regret, (d) the optimism–pessimism rule (for $\alpha > ¼$), (e) the principle of insufficient reason.

a_1	20	20	0	10
a_2	10	10	10	10
a_3	0	40	0	0
a_4	10	30	0	0

3.2 Which act(s) should be chosen in the decision matrix below according to the following decision rules: (a) the maximin rule (b) the maximax rule (c) the optimism–pessimism rule (for $\alpha > 1/5$) d) the minimax regret rule (e) the principle of insufficient reason.

a_1	5	45	0	0
a_2	5	35	35	15
a_3	15	15	15	15
a_4	5	35	15	5

3.3 What is the main advantage of leximin as compared to maximin?

3.4 What is so special with an act's *best* and *worst* possible outcomes?

3.5 Construct a decision matrix that corresponds to the following regret table.

a_1	−5	−45	0
a_2	0	−30	−20
a_3	−10	0	−25

3.6 According to the minimax regret principle, regret is defined as the difference between the best outcome obtainable under a given state and the outcome obtained under that state with the chosen act. Now consider an alternative version of this principle, according to which regret is defined (relative to each act) as the difference between the best outcome and the worst outcome obtainable with the act in question. This modification of the minimax regret rule would give us a highly implausible decision rule. Why?

3.7 Prove that Axiom 8 (Column duplication) is incompatible with the principle of insufficient reason.

3.8 Prove that Axiom 7 (Column linearity) is incompatible with the optimism–pessimism rule.

3.9 Explain informally why Axiom 5 (Interval scale) is compatible with all of the following rules: (a) maximin, (b) optimism–pessimism, (c) minimax regret, (d) the principle of insufficient reason.

3.10 Rawls famously argued that the maximin rule could be used for formulating a theory of social justice. Very briefly put, he argued that a distribution A of social goods is better than an alternative distribution B if the worst-off person is better off in A than in B. Suppose that Rawls' theory of social justice is based on the optimism–pessimism rule instead of maximin. What would be the implications for social justice? (Assume that $1 > \alpha > 0$.)

Solutions

3.1 (a) a_2, (b) a_3, (c) a_4, (d) a_3, (e) a_1

3.2 (a) a_3, (b) a_1, (c) a_3, (d) a_2, (e) a_2

3.3 An explicit answer is given in the text.

3.4 This question has no universally accepted answer.

3.5 There are many solutions to this problem; here is one.

a_1	5	0	25
a_2	10	15	5
a_3	0	45	0

3.6 According to the modified rule an act that gives you, say, either 1 or 0 would be ranked higher than an act that yields either 100 or 98. This can't be right.

3.7 Here is a counterexample: The matrix to the right can be obtained from the left by duplicating the second column. However, in the leftmost matrix the principle of insufficient reason is indifferent between a_1 and a_2, whereas in the rightmost matrix it recommends a_2.

a_1	1	0
a_2	0	1

a_1	1	0	0
a_2	0	1	1

3.8 Here is a counterexample: The matrix to the right can be obtained from the left by adding 1 to the first column. However, in the leftmost matrix the optimism–pessimism rule is indifferent between a_1 and a_2 (for all α), whereas in the rightmost matrix it recommends a_1 if $\alpha > 1/2$, and a_2 if $\alpha < 1/2$ and is indifferent between them if and only if $\alpha = 1/2$.

a_1	1	0	a_1	2	0
a_2	0	1	a_2	1	1

3.9 None of these rules require any more information than what is contained in an interval scale. In particular, none of the rules require that we calculate *ratios* between outcomes.

3.10 The optimism–pessimism rule only considers the worst and the best possible outcomes of an act. All intermediate outcomes are irrelevant. If we translate this to a social context in which each possible outcome corresponds to a single individual, we find that it does not matter if 99 out of 100 people are badly off and 1 person is well off, or if 1 person is well off and 99 persons are badly off. Furthermore, since the best *possible* outcomes matter, it follows that we could make society better by making things much better for the best-off individual. I leave it to the reader to judge whether a theory of social justice thus modified is an improvement over Rawls' original theory.