Before addressing any decision problem under uncertainty, it is necessary to build a preference functional that evaluates the level of satisfaction of the decision maker who bears risk. If such a functional is obtained, decisions problems can be solved by looking for the decision that maximizes the decision maker's level of satisfaction. This first chapter provides a way—the now classical way—to evaluate the level of satisfaction under uncertainty. The rest of this book deals with applications of this model to specific decision problems.

The modern theory of risk-bearing has been founded in 1944 by von Neumann and Morgenstern in their famous book entitled *Theory of Games and Economic Behavior*. In the 1930s and early 1940s, economists like John Hicks (1931), Jacob Marschak (1938), and Gerhard Tintner (1941) were debating whether the ordering of risks could be based on a function of their mean and variance alone. A first axiomatic approach to the ordering of probability distributions had been introduced by Frank Ramsey (1931). This approach was revived in successively clearer and simpler terms by von Neumann and Morgenstern (1947) and Marschak (1950). The aim of this chapter is to summarize this axiomatic approach.

1.1 Simple and Compound Lotteries

The description of an uncertain environment contains two different types of information. First, one must enumerate all possible outcomes. An outcome is a list of variables that affect the well-being of the decision maker. The list might give a health status, some meteorological parameters, levels of pollution, and the quantities of different goods consumed. As long as we do not introduce saving (chapter

15), we will assume that the outcome can be measured in a one-dimensional unit, namely money (consumed at a specific date in time). But at this stage, there is no cost to allow for a multidimensional outcome. Let X be the set of possible outcomes. To avoid technicalities, we assume that the number of possible outcomes is finite, so $X \equiv \{x_s\}_{s=1,...,S}$.

The second characteristic of an uncertain environment is the vector of probabilities of the possible outcomes. Let $p_s \ge 0$ be the probability of occurrence of x_s , with $\sum_{s=1}^{S} p_s = 1$. We will assume that these probabilities are objectively known.

A lottery L is described by a vector $(x_1, p_1; x_2, p_2; \dots; x_S, p_S)$. Since the set of potential outcomes is invariant, we will define a lottery L by its vector of probabilities (p_1, \dots, p_S) . The set of all lotteries on outcomes X is denoted $\mathcal{L} \equiv \{(p_1, \dots, p_S) \in R_+^S | p_1 + \dots + p_S = 1\}$.

When S=3, we can represent a lottery by a point (p_1,p_3) in R^2 , since $p_2=1-p_1-p_3$. More precisely, in order to be a lottery, this point must be in the so-called Machina triangle $\{(p_1,p_3)\in R_+^2|1-p_1-p_3\leq 1\}$. If the lottery is on a edge of the triangle, one of the probabilities is zero. If it is at a corner, the lottery is degenerated, by which we mean that it takes one of the values x_1, x_2, x_3 with probability 1. This triangle is represented in figure 1.1.

A compound lottery is a lottery whose outcomes are lotteries. Consider a compound lottery L which yields lottery $L^a = (p_1^a, \dots, p_S^a)$ with probability α and lottery $L^b = (p_1^b, \dots, p_S^b)$ with probability $1 - \alpha$. The probability that the outcome of L be x_1 is $p_1 = \alpha p_1^a + (1 - \alpha)p_1^b$. More generally, we obtain that

L has the same vector of probabilities as
$$\alpha L^a + (1 - \alpha)L^b$$
. (1.1)

A compound lottery is a convex combination of simple lotteries. Such a compound lottery is represented in figure 1.1. From condition (1.1), it is natural to confound L with $\alpha L^a + (1 - \alpha)L^b$. Whether a specific uncertainty comes from a simple lottery or from a complex compound lottery has no significance. Only the probabilities of potential outcomes matter (axiom of reduction).

1.2 Axioms on Preferences under Uncertainty

We assume that the decision maker has a rationale preference relation \succeq over the set of lotteries \mathscr{L} . This means that order \succeq is complete and transitive. For any pair (L^a, L^b) of lotteries in \mathscr{L} , either L^a is

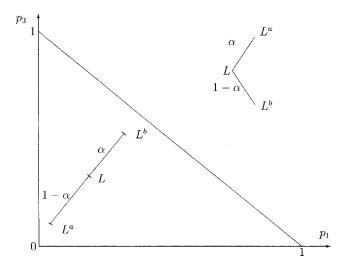


Figure 1.1 Compound lottery, $L = \alpha L^a (1 - \alpha) L^b$, in the Machina triangle

preferred to L^b ($L^a \succeq L^b$), or L^b is preferred to L^a ($L^b \succeq L^a$) (or both). Moreover, if L^a is preferred to L^b which is itself preferred to L^c , then L^a is preferred to L^c . To this preference order \succeq , we associate the indifference relation \sim , with $L^a \sim L^b$ if and only if $L^a \succeq L^b$ and $L^b \succeq L^a$.

A standard hypothesis that is made on preferences is that they are continuous. This means that small changes in probabilities do not change the nature of the ordering between two lotteries. Technically, this axiom is written as follows:

AXIOM 1 (CONTINUITY) The preference relation \succeq on the space of simple lotteries $\mathscr L$ is such that for all $L^a, L^b, L^c \in \mathscr L$ such that $L^a \succeq L^b \succeq L^c$, there exists a scalar $\alpha \in [0,1]$ such that

$$L^b \sim \alpha L^a + (1 - \alpha)L^c$$
.

As is well-known in the theory of consumer choice under certainty, the continuity axiom implies the existence of a functional $U: \mathcal{L} \to R$ such that

$$U(L^a) \ge U(L^b) \Leftrightarrow L^a \succeq L^b. \tag{1.2}$$

The preference functional U is an index that represents the degree of satisfaction of the decision maker. It assigns a numerical value to each lottery, ranking the lotteries in accordance to the individual's

preference \succeq . Notice that U is not unique: take any increasing function $g: R \to R$. Then the functional V that is defined as V(.) = g(U(.)) also represents the same preferences \succeq . The preference functional is ordinal in the sense that it is invariant to any increasing transformation. It is only the ranking of lotteries that matters. A preference functional can be represented in the Machina triangle by a family of continuous indifference curves.

The assumptions above on preferences under uncertainty are minimal. If no other assumption is made on these preferences, the theory of choice under uncertainty would not differ from the standard theory of consumer choice under certainty. The only difference would be on how to define the consumption goods. Most developments in the economics of uncertainty and its applications have been made possible by imposing more structure on preferences under uncertainty. This additional structure originates from the independence axiom.

AXIOM 2 (INDEPENDENCE) The preference relation \succeq on the space of simple lotteries $\mathscr L$ is such that for all $L^a, L^b, L^c \in \mathscr L$ and for all $\alpha \in [0,1]$,

$$L^a \succeq L^b \Leftrightarrow \alpha L^a + (1 - \alpha)L^c \succeq \alpha L^b + (1 - \alpha)L^c$$
.

This means that if we mix each of two lotteries L^a and L^b with a third one L^c , then the preference ordering of the two resulting mixtures is *independent* of the particular third lottery L^c used. The independence axiom is at the heart of the classical theory of uncertainty. Contrary to the other assumptions made above, the independence axiom has no parallel in the consumer theory under certainty. This is because there is no reason to believe that if I prefer a bundle A containing 1 cake and 1 bottle of wine to a bundle B containing 3 cakes and no wine, I also prefer a bundle A' containing 2 cakes and 2 bottles of wine to a bundle B' containing 3 cakes and 1.5 bottles of wine, just because

$$(2,2) = 0.5(1,1) + 0.5(3,3)$$
 and $(3,1.5) = 0.5(3,0) + 0.5(3,3)$.

1.3 The Expected Utility Theorem

The independence axiom implies that the preference functional U must be linear in the probabilities of the possible outcomes. This is the essence of the expected utility theorem, which is due to von

Neumann and Morgenstern (1947). It has been extended later on by Savage (1954) to the case where there is no objective probabilities.

THEOREM 1 (EXPECTED UTILITY) Suppose that the rationale preference relation \succeq on the space of simple lotteries $\mathscr L$ satisfies the continuity and independence axioms. Then \succeq can be represented by a preference functional that is linear in probabilities. That is, there exists a scalar u_s associated to each outcome x_s , $s=1,\ldots,S$, in such a manner that for any two lotteries $L^a=(p_1^a,\ldots,p_S^a)$ and $L^b=(p_1^b,\ldots,p_S^b)$, we have

$$L^a \succeq L^b \Leftrightarrow \sum_{s=1}^S p_s^a u_s \geq \sum_{s=1}^S p_s^b u_s.$$

Proof We would be done if we prove that for any compound lottery $L = \beta L^a + (1 - \beta)L^b$, we have

$$U(\beta L^{a} + (1 - \beta)L^{b}) = \beta U(L^{a}) + (1 - \beta)U(L^{b}). \tag{1.3}$$

Applying this property recursively would yield the result. To do this, let us consider the worst and best lotteries in \mathcal{L} , \underline{L} and \overline{L} . They are obtained by solving the problem of minimizing and maximizing U(L) on the compact set \mathcal{L} . By definition, for any $L \in \mathcal{L}$, we have $\overline{L} \succeq L \succeq \underline{L}$. By the continuity axiom, we know that there exist two scalars, α^a and α^b , in [0,1] such that

$$L^a \sim \alpha^a \overline{L} + (1 - \alpha^a) L$$

and

$$L^b \sim \alpha^b \overline{L} + (1 - \alpha^b) \underline{L}.$$

Observe that $L^a \succeq L^b$ if and only if $\alpha^a \ge \alpha^b$. Indeed, suppose that $\alpha^a \ge \alpha^b$ and take $\gamma = (\alpha^a - \alpha^b)/(1 - \alpha^b) \in [0, 1]$. Then we have

$$\begin{split} \alpha^a \overline{L} + (1 - \alpha^a) \underline{L} &\sim \gamma \overline{L} + (1 - \gamma) [\alpha^b \overline{L} + (1 - \alpha^b) \underline{L}] \\ &\succeq \gamma [\alpha^b \overline{L} + (1 - \alpha^b) \underline{L}] + (1 - \gamma) [\alpha^b \overline{L} + (1 - \alpha^b) \underline{L}] \\ &\sim \alpha^b \overline{L} + (1 - \alpha^b) \underline{L}. \end{split}$$

The two equivalences are direct applications of the axiom of reduction. The second line of this sequence of preference orderings is due to the independence axiom together with the fact that $\bar{L} \succeq \alpha^b \bar{L} + (1-\alpha^b)\underline{L}$, by definition of \bar{L} .

We conclude that $U(L) = \alpha$, where α is such that $L \sim \alpha \overline{L} + (1 - \alpha)\underline{L}$, perfectly fits the definition (1.2) of a preference functional associated to \succeq . Thus $U(L^a) = \alpha^a$ and $U(L^b) = \alpha^b$. It remains to prove that

$$U(\beta L^a + (1 - \beta)L^b) = \beta \alpha^a + (1 - \beta)\alpha^b,$$

or equivalently that

$$\beta L^{a} + (1 - \beta)L^{b} \sim (\beta \alpha^{a} + (1 - \beta)\alpha^{b})\overline{L} + (\beta(1 - \alpha^{a}) + (1 - \beta)(1 - \alpha^{b}))L.$$

This is true since, using the independence axiom twice, we get

$$\begin{split} \beta L^a + (1-\beta) L^b &\sim \beta [\alpha^a \overline{L} + (1-\alpha^a) \underline{L}] + (1-\beta) L^b \\ &\sim \beta [\alpha^a \overline{L} + (1-\alpha^a) \underline{L}] + (1-\beta) [\alpha^b \overline{L} + (1-\alpha^b) \underline{L}] \\ &\sim (\beta \alpha^a + (1-\beta) \alpha^b) \overline{L} + (\beta (1-\alpha^a) + (1-\beta) (1-\alpha^b)) \underline{L}. \end{split}$$

This concludes the proof. ■

The consequence of the independence axiom is that the family of indifference curves in the Machina triangle is a set of parallel straight lines. Their slope equals $(u_1 - u_2)/(u_3 - u_2)$. This has the following consequences on the attitude toward risk. Consider the four lotteries, L^a, L^b, M^a , and M^b , depicted in the Machina triangle of figure 1.2. Suppose that the segment L^aL^b is parallel to segment

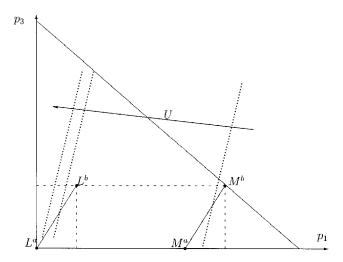


Figure 1.2 Allais paradox in the Machina triangle

 M^aM^b . This implies that L^a is preferred to L^b if and only if M^a is preferred to M^b .

In this book, we will consider the outcome to be a monetary wealth. The ex-post environment of the decision maker is fully described by the amount of money that he can consume. A lottery on wealth can be expressed by a random variable \tilde{w} whose realization is an amount of money w that can be consumed. This random variable can be expressed by a cumulative distribution function F, where F(w) is the probability that \tilde{w} be less or equal than w. This covers the case of continuous, discrete, or mixed random variables. By the expected utility theorem, we know that to each wealth level w, there exists a scalar u(w) such that

$$\tilde{w}_1 \succeq \tilde{w}_2 \Leftrightarrow Eu(\tilde{w}_1) \geq Eu(\tilde{w}_2) \Leftrightarrow \int u(w)dF_1(w) \geq \int u(w)dF_2(w),$$

where F_i is the cumulative distribution function of \tilde{w}_i . It is intuitive in this context to assume that the utility function is constant.

Notice that the utility function is *cardinal*: an increasing *linear* transformation of u, v(.) = au(.) + b, a > 0, will not change the ranking of lotteries: if $\tilde{w}_1 \succeq \tilde{w}_2$, we have

$$Ev(\tilde{w}_1) = E[au(\tilde{w}_1) + b] = aEu(\tilde{w}_1) + b \ge aEu(\tilde{w}_2) + b = Ev(\tilde{w}_2).$$

To sum up, expected utility is ordinal, whereas the utility function is cardinal. Differences in utility have meaning, whereas differences in expected utility have no significance.

1.4 Critics of the Expected Utility Model

The independence axiom is not applied without difficulties. The oldest and most famous challenge was proposed by Allais (1953).

Table 1.1
Outcome as a function of the number of the ball

Lottery	0	1–10	11-99
La	50	50	50
L^b	0	250	50
M^a	50	50	0
M^b	0	250	0

The paradox that he raised generated thousands of papers. Whether the Allais's counterexample knocks out expected utility theory or comes from outwitting the players is still subject of lively discussions among the specialists. We first present the paradox. We then summarize the main arguments that have been proposed by the EU proponents to sustain that the paradox comes from a misconception by the players.

1.4.1 The Allais Paradox

Allais proposed the following experiment: An urn contains 100 balls numbered from 0 to 99. There are four lotteries whose monetary outcomes depend in different ways on the number that is written on the ball that is taken out of the urn. An example of an outcome expressed in thousands of dollars is described in table 1.1.

Decision makers are subjected to two choice tests. In the first test, they are asked to choose between L^a and L^b , and in the second test, they must choose between M^a and M^b . Many people report that they prefer L^a to L^b but subsequently prefer M^b to M^a . Notice that since L^a and L^b have the same outcome when the number of the ball is larger than 10, the independence axiom tells us that most people will prefer $L^{a'}$, which takes value 50 with certainty, to $L^{b'}$, which takes value 0 with probability 1/11 and value 250 with probability 10/11. The paradox is that the same argument can be used with the opposite result when considering the preference of M^b over M^a ! Thus these choices among pairs of lotteries are not compatible with the independence axiom.

This problem can be analyzed using the Machina triangle in figure 1.2. Let us define

```
x_1 = 0,

x_2 = 50,

x_3 = 250.
```

The four lotteries in the Allais experiment are respectively $L^a = (0,1,0), \ L^b = (0.01,0.89,0.10), \ M^a = (0.89,0.11,0)$ and $M^b = (0.90,0,0.10)$. Segment L^aL^b is parallel to segment M^aM^b . Because indifference curves are parallel straight lines under expected utility, it must be the case that $M^a \succeq M^b$ if $L^a \succeq L^b$.

1.4.2 The Allais Paradox and Time Consistency

Savage (1954), who was a strong supporter of the expected utility model, participated in the test organized by Allais. It happened that Savage was one of those people who preferred L^a to L^b and M^b to M^a . After realizing that this was not compatible with expected utility, he wanted to revise his choices. He claimed that he had been misled and that a more cautious reading of the problem would have been sufficient to avoid the mistake.

A first argument in favor of the independence axiom is that individuals who violate it would find themselves subject to so-called Dutch book outcomes. To illustrate, suppose that a gambler is offered three lotteries L^a, L^b , and L^c . The gambler ranks them $L^a \succ L^b$ and $L^a \succ L^c$, but contrary to the independence axiom, it turns out that $L^d = 0.5L^b + 0.5L^c$ is preferred to L^a . L^d is a compound lottery in which a fair coin is used to determine which lottery, L^b or L^c , will be played. In making the choice from among L^a, L^b , and L^c , the gambler rationally selects L^a . Since L^d is preferred to L^a , we know that he is willing to pay a fee to replace L^a by the compound lottery L^d . This outcome is thus rationally accepted by the gambler. But, as soon as the coin is tossed, giving the gambler either L^b or L^c , one could get him to pay another fee to trade this lottery for L^a . At this point the gambler has paid two fees but is otherwise back to his original position. This outcome is dynamically inconsistent.

The relation between the independence axiom and time consistency is central to the current debate between the EU and non-EU specialists. Wakker (1999) makes this relation clear by using the Allais's example. We summarize Wakker's analysis in figure 1.3. A simpler version of the Allais paradox is expressed by the comparison of figure 1.3a and e. As is usual, the square in the decision tree represents a decision (up or down), whereas the circle describes the occurrence of a random event. For example, problem a is to choose between a lottery that has a payoff of 250 with probability 10/11, and a sure payoff of 50. Standard choices are to select "down" in figure 1.3a, and to select "up" in figure 1.3e. Any combination of these two choices is not compatible with expected utility.

To see how this is possible, let us examine the following sequence of comparisons:

• Comparison between a and b. Upon being asked to express her preferences between lotteries "up" and "down" of problem a, an agent is

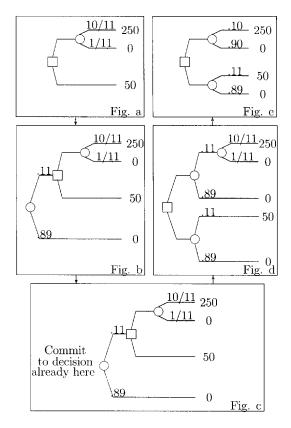


Figure 1.3 Decomposing Allais's paradox

told that this is only by chance that she has to participate. In fact, she has been selected at random with 11 other participants in a group of 100 persons. The decision problem in figure 1.3a is in fact a subtree of a larger problem that is represented by figure 1.3b. Should the agent modify her decision on the basis of this forgone event? Most people agree that forgone events should not affect the decision of the agent. Past events are irrelevant for the current decision. This rationale is called "consequentialism." It tells us that rational agents will make the same decision in both problem a and problem b.

• Comparison between b and c. Let us now take the problem back one step. We have a group of 100 persons. Eleven will be selected at random to participate in game a. The agent under scrutiny does

not yet know whether she will be among the lucky eleven persons. However, she is asked to commit herself to the strategy that she will follow if she is selected. Figure 1.3c describes this problem. If the agent is time consistent, she should commit herself ex ante to do what will be optimal to do ex post. Why, for example, would she commit to the "up" strategy when she knows that "down" is optimal? Time consistency tells us that people will make the same decision in problems b and c.

- Comparison between c and d. Problem d is essentially a rewriting of problem c. If the agent commits to strategy "up," she faces a selection risk and a gamble risk as is represented by the upper part of tree d. If she decides to commit to "down," she only faces the selection risk. Now, let us forget the story of selecting eleven people from a group of 100 persons. Problem d is a decision problem between two compound lotteries. This changes the context of the problem but not the underlying lotteries faced by the agent. Do we think that the context in the way the gamble is processed matters for the decision maker? Most people would say no. The context does not matter. Only the state probabilities and the state payoffs do. This assumption, called "context independence," implies that rational agents will make the same choice in problems c and d.
- Comparison between d and e. Compounded lotteries can be reduced by using the standard law of probabilities. For example, if "up" is selected in problem d, there is an $0.89 + \left(0.11\frac{1}{11}\right) = 0.9$ probability of obtaining 0. This is what we earlier called the reduction axiom. By repeating this calculation, the reader can easily check that problems d and e are the same. Assuming that the agent is able to reduce compounded lotteries in this way, we can conclude that she will make the same choice in both problems d and e.

From this sequence of elementary assumptions—consequentialism, time consistency, context independence, and reduction—about preferences, it should be clear that if an agent prefers "down" to "up" in problem a, he will behave the same way in each of the other problems of figure 1.3, in particular, problem e. The violation of independence is possible only if at least one of the elementary assumptions above is also violated. While it may be possible that some people will violate one of these assumptions in experiments, we believe that there is a strong normative appeal to each assumption.

Lottery	Red	Black	White
L^a	50	0	0
L^b	0	50	0
M^a	50	0	50
M^b	0	50	50

Table 1.2Outcome as a function of the number of the ball

1.5 Concluding Remark

We will learn more in the next chapters about the linearity of the preference functional with respect to probabilities and its central role in finding simple solutions to decision problems under uncertainty. If we relax the independence axiom, most problems presented in this book cannot be solved anymore. Our approach is thus pragmatic: we recognize that the independence axiom may fail to describe real behavior in certain risky environments, especially when there are low-probability events. But the combination of facts showing that the independence axiom, and its related axioms of time consistency and consequentialism, make common sense and make most problems solvable is enough to justify the exploration of its implications.

The search for an alternative model of decision making under uncertainty that does not rely on the independence axiom has been a lively field of research in economics for more than fifteen years. There are several competing models, each with its advantages and faults. Most entail the expected utility model as a particular case. Nevertheless, the use of the expected utility model is pervasive in economics. Probably because of its strong normative basis.

1.6 Exercises and Extensions

Exercise 1 Consider the Allais paradox summarized in table 1.1. Show that there exists no utility function such that the expected utility of L^a is larger than the expected utility of L^b , and the expected utility of M^b is larger than the expected utility of M^a .

^{1.} For more about the history of the economic thought on this field, see Bernstein (1996) and Drèze (1987), for example.

Exercise 2 (Ellsberg's paradox) Savage (1954) extended the expected utility theorem to situations where the probabilities are not objectively known. He introduced the sure thing principle, an axiom that is stronger than the independence axiom. He showed that this stronger axiom could be used to prove the "subjective EU" theorem by noting that

- there exists a subjective probability measure *p*,
- there exists a real-valued utility function *u*

such that the decision maker ranks various distributions of consequences ω by using their subjective expected utility $\sum_s p_s u(\omega_s)$.

The following example is due to Ellsberg (1961). An urn contains 90 colored balls. Thirty balls are red, and the remaining 60 either black or white; the number of black (white) balls is not specified. There are 4 lotteries as described in table 1.2. Given that many people report ranking $L^a > L^b$ and $M^b > M^a$, show that there exists no pair (p,u) that sustains such ranking of lotteries. The agents are averse to ambiguous probabilities, a possibility that is ruled out by the sure thing principle. (For a model with ambiguity aversion, see Gilboa and Schmeidler 1987.)

Exercise 3 (Rank Dependent EU) Consider the lottery $L = (x_1, p_1; x_2, p_2; ...; x_S, p_S)$, with $x_1 < x_2 < \cdots < x_S$. Quiggin (1982) proposed a generalization of the EU model by introducing a function g of transformation of the **cumulative** distribution:

$$U(L) = \sum_{s=1}^{S} u(x_s) \left[g\left(\sum_{t=1}^{s} p_t\right) - g\left(\sum_{t=1}^{s-1} p_t\right) \right]$$
 (1.4)

with g(0) = 0 and g(1) = 1.

- Draw the indifference curves in the Machina triangle for $g(q) = q^2$ and $g(q) = \sqrt{q}$.
- Which restriction on *g* would you consider to solve the Allais paradox?