MasColell ..., MicroTheory

CHAPTER

6

Choice Under Uncertainty

6.A Introduction

In previous chapters, we studied choices that result in perfectly certain outcomes. In reality, however, many important economic decisions involve an element of risk. Although it is formally possible to analyze these situations using the general theory of choice developed in Chapter 1, there is good reason to develop a more specialized theory: Uncertain alternatives have a structure that we can use to restrict the preferences that "rational" individuals may hold. Taking advantage of this structure allows us to derive stronger implications than those based solely on the framework of Chapter 1.

In Section 6.B, we begin our study of choice under uncertainty by considering a setting in which alternatives with uncertain outcomes are describable by means of objectively known probabilities defined on an abstract set of possible outcomes. These representations of risky alternatives are called *lotteries*. In the spirit of Chapter 1, we assume that the decision maker has a rational preference relation over these lotteries. We then proceed to derive the *expected utility theorem*, a result of central importance. This theorem says that under certain conditions, we can represent preferences by an extremely convenient type of utility function, one that possesses what is called the *expected utility form*. The key assumption leading to this result is the *independence axiom*, which we discuss extensively.

In the remaining sections, we focus on the special case in which the outcome of a risky choice is an amount of money (or any other one-dimensional measure of consumption). This case underlies much of finance and portfolio theory, as well as substantial areas of applied economics.

In Section 6.C, we present the concept of *risk aversion* and discuss its measurement. We then study the comparison of risk aversions both across different individuals and across different levels of an individual's wealth.

Section 6.D is concerned with the comparison of alternative distributions of monetary returns. We ask when one distribution of monetary returns can unambiguously be said to be "better" than another, and also when one distribution can be said to be "more risky than" another. These comparisons lead, respectively, ...cepts of *first-order* and *second-order stochastic dominance*.

In Section 6.E, we extend the basic theory by allowing utility to depend on *states of nature* underlying the uncertainty as well as on the monetary payoffs. In the process, we develop a framework for modeling uncertainty in terms of these underlying states. This framework is often of great analytical convenience, and we use it extensively later in this book.

In Section 6.F, we consider briefly the theory of *subjective probability*. The assumption that uncertain prospects are offered to us with known objective probabilities, which we use in Section 6.B to derive the expected utility theorem, is rarely descriptive of reality. The subjective probability framework offers a way of modeling choice under uncertainty in which the probabilities of different risky alternatives are not given to the decision maker in any objective fashion. Yet, as we shall see, the theory of subjective probability offers something of a rescue for our earlier objective probability approach.

For further reading on these topics, see Kreps (1988) and Machina (1987). Diamond and Rothschild (1978) is an excellent sourcebook for original articles.

6.B Expected Utility Theory

We begin this section by developing a formal apparatus for modeling risk. We then apply this framework to the study of preferences over risky alternatives and to establish the important expected utility theorem.

Description of Risky Alternatives

Let us imagine that a decision maker faces a choice among a number of risky alternatives. Each risky alternative may result in one of a number of possible *outcomes*, but which outcome will actually occur is uncertain at the time that he must make his choice.

Formally, we denote the set of all possible outcomes by C.[1] These outcomes could take many forms. They could, for example, be consumption bundles. In this case, $C = X$, the decision maker's consumption set. Alternatively, the outcomes might take the simpler form of monetary payoffs. This case will, in fact, be our leading example later in this chapter. Here, however, we treat C as an abstract set and therefore allow for very general outcomes.

To avoid some technicalities, we assume in this section that the number of possible outcomes in C is finite, and we index these outcomes by $n = 1, \dots, N$.

Throughout this and the next several sections, we assume that the probabilities of the various outcomes arising from any chosen alternative are *objectively known*. For example, the risky alternatives might be monetary gambles on the spin of an unbiased roulette wheel.

The basic building block of the theory is the concept of a *lottery*, a formal device that is used to represent risky alternatives.

Definition 6.B.1: A *simple lottery* L is a list $L = (p_1, \dots, p_N)$ with $p_n \geq 0$ for all n and $\sum_n p_n = 1$, where p_n is interpreted as the probability of outcome n occurring.

1. It is also common, following Savage (1954), to refer to the elements of C as *consequences*.

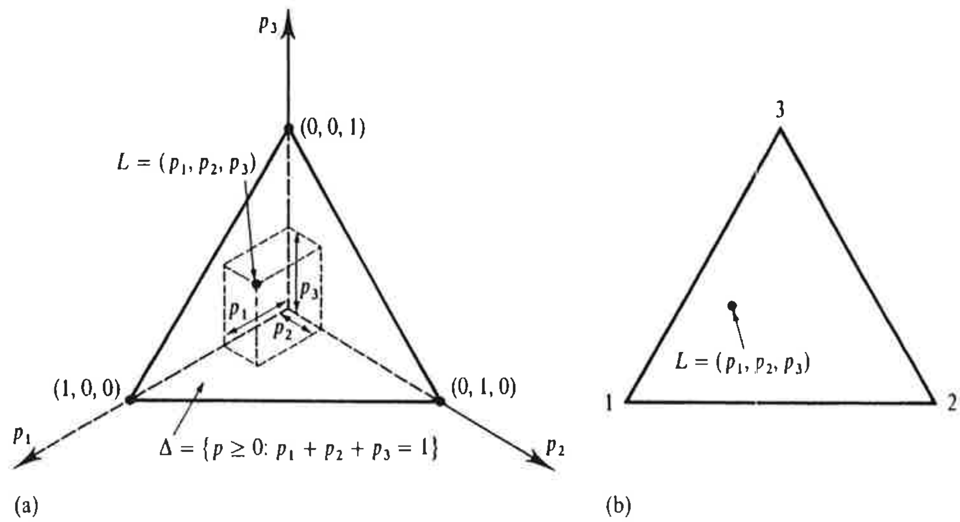

Figure 6.B.1

Representations of the simplex when $N = 3$.
(a) Three-dimensional representation.
(b) Two-dimensional representation.

A simple lottery can be represented geometrically as a point in the $(N - 1)$ *dimensional simplex*, $\Delta = \{p \in \mathbb{R}_+^N: p_1 + \cdots + p_N = 1\}$. Figure 6.B.1(a) depicts this simplex for the case in which $N = 3$. Each vertex of the simplex stands for the degenerate lottery where one outcome is certain and the other two outcomes have probability zero. Each point in the simplex represents a lottery over the three outcomes. When $N = 3$, it is convenient to depict the simplex in two dimensions, as in Figure 6.B.1(b), where it takes the form of an equilateral triangle.[2]

In a simple lottery, the outcomes that may result are certain. A more general variant of a lottery, known as a *compound lottery*, allows the outcomes of a lottery themselves to be simple lotteries.[3]

Definition 6.B.2: Given K simple lotteries $L_k = (p_1^k, \ldots, p_N^k)$, $k = 1, \ldots, K$, and probabilities $\alpha_k \geq 0$ with $\sum_k \alpha_k = 1$, the *compound lottery* $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$ is the risky alternative that yields the simple lottery L_k with probability α_k for $k = 1, \ldots, K$.

For any compound lottery $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$, we can calculate a corresponding *reduced lottery* as the simple lottery $L = (p_1, \ldots, p_N)$ that generates the same ultimate distribution over outcomes. The value of each p_n is obtained by multiplying the probability that each lottery L_k arises, α_k, by the probability p_n^k that outcome n arises in lottery L_k, and then adding over k. That is, the probability of outcome n in the reduced lottery is

$$p_n = \alpha_1 p_n^1 + \cdots + \alpha_K p_n^K$$

2. Recall that equilateral triangles have the property that the sum of the perpendiculars from any point to the three sides is equal to the altitude of the triangle. It is therefore common to depict the simplex when $N = 3$ as an equilateral triangle with altitude equal to 1 because by doing so, we have the convenient geometric property that the probability p_n of outcome n in the lottery associated with some point in this simplex is equal to the length of the perpendicular from this point to the side opposite the vertex labeled n.

3. We could also define compound lotteries with more than two stages. We do not do so because we will not need them in this chapter. The principles involved, however, are the same.

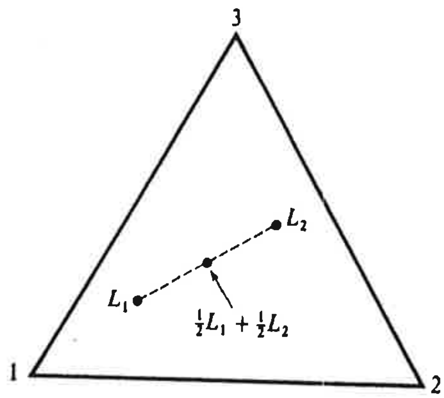

Figure 6.B.2
The reduced lottery of a compound lottery.

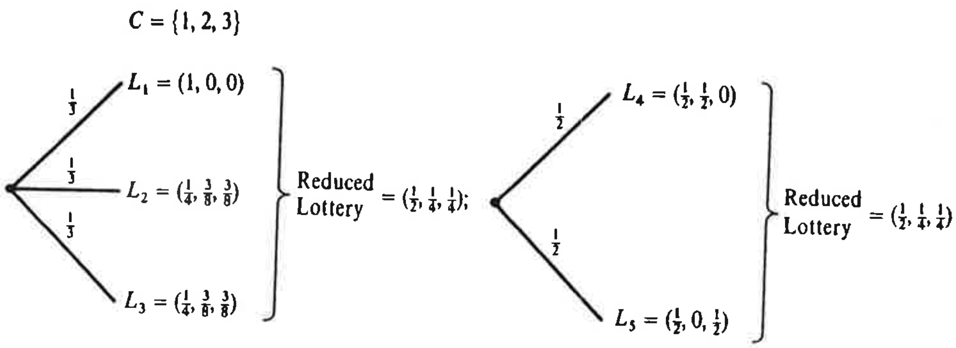

Figure 6.B.3
Two compound lotteries with the same reduced lottery.

for $n = 1, \ldots, N$.[4] Therefore, the reduced lottery L of any compound lottery $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$ can be obtained by vector addition:

$$L = \alpha_1 L_1 + \cdots + \alpha_K L_K \in \Delta.$$

In Figure 6.B.2, two simple lotteries L_1 and L_2 are depicted in the simplex Δ. Also depicted is the reduced lottery $\frac{1}{2}L_1 + \frac{1}{2}L_2$ for the compound lottery $(L_1, L_2; \frac{1}{2}, \frac{1}{2})$ that yields either L_1 or L_2 with a probability of $\frac{1}{2}$ each. This reduced lottery lies at the midpoint of the line segment connecting L_1 and L_2. The linear structure of the space of lotteries is central to the theory of choice under uncertainty, and we exploit it extensively in what follows.

Preferences over Lotteries

Having developed a way to model risky alternatives, we now study the decision maker's preferences over them. The theoretical analysis to follow rest on a basic *consequentialist* premise: We assume that for any risky alternative, only the reduced lottery over final outcomes is of relevance to the decision maker. Whether the probabilities of various outcomes arise as a result of a simple lottery or of a more complex compound lottery has no significance. Figure 6.B.3 exhibits two different compound lotteries that yield the same reduced lottery. Our consequentialist hypothesis requires that the decision maker view these two lotteries as equivalent.

4. Note that $\sum_n p_n = \sum_k \alpha_k (\sum_n p_n^k) = \sum_k \alpha_k = 1$.

We now pose the decision maker's choice problem in the general framework developed in Chapter 1 (see Section 1.B). In accordance with our consequentialist premise, we take the set of alternatives, denoted here by $\mathscr{L}$, to be *the set of all simple lotteries over the set of outcomes C*. We next assume that the decision maker has a rational preference relation $\succsim$ on $\mathscr{L}$, a complete and transitive relation allowing comparison of any pair of simple lotteries. It should be emphasized that, if anything, the rationality assumption is stronger here than in the theory of choice under certainty discussed in Chapter 1. The more complex the alternatives, the heavier the burden carried by the rationality postulates. In fact, their realism in an uncertainty context has been much debated. However, because we want to concentrate on the properties that are specific to uncertainty, we do not question the rationality assumption further here.

We next introduce two additional assumptions about the decision maker's preferences over lotteries. The most important and controversial is the *independence axiom*. The first, however, is a continuity axiom similar to the one discussed in Section 3.C.

Definition 6.B.3: The preference relation $\succsim$ on the space of simple lotteries $\mathscr{L}$ is *continuous* if for any $L, L', L'' \in \mathscr{L}$, the sets

$$\{\alpha \in [0,1]: \alpha L + (1-\alpha)L' \succsim L''\} \subset [0,1]$$

and

$$\{\alpha \in [0,1]: L'' \succsim \alpha L + (1-\alpha)L'\} \subset [0,1]$$

are closed.

In words, continuity means that small changes in probabilities do not change the nature of the ordering between two lotteries. For example, if a "beautiful and uneventful trip by car" is preferred to "staying home," then a mixture of the outcome "beautiful and uneventful trip by car" with a sufficiently small but positive probability of "death by car accident" is still preferred to "staying home." Continuity therefore rules out the case where the decision maker has lexicographic ("safety first") preferences for alternatives with a zero probability of some outcome (in this case, "death by car accident").

As in Chapter 3, the continuity axiom implies the existence of a utility function representing $\succsim$, a function $U: \mathscr{L} \to \mathbb{R}$ such that $L \succsim L'$ if and only if $U(L) \geq U(L')$. Our second assumption, the independence axiom, will allow us to impose considerably more structure on $U(\cdot)$.[5]

Definition 6.B.4: The preference relation $\succsim$ on the space of simple lotteries $\mathscr{L}$ satisfies the *independence axiom* if for all $L, L', L'' \in \mathscr{L}$ and $\alpha \in (0,1)$ we have

$$L \succsim L' \quad \text{if and only if} \quad \alpha L + (1-\alpha)L'' \succsim \alpha L' + (1-\alpha)L''.$$

In other words, if we mix each of two lotteries with a third one, then the preference ordering of the two resulting mixtures does not depend on (is *independent* of) the particular third lottery used.

5. The independence axiom was first proposed by von Neumann and Morgenstern (1944) as an incidental result in the theory of games.

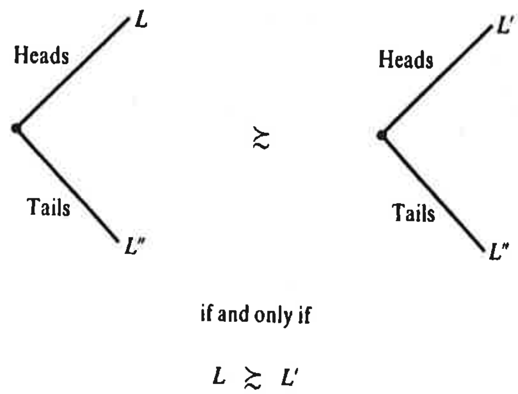

Figure 6.B.4 The independence axiom.

Suppose, for example, that $L \succsim L'$ and $\alpha = \frac{1}{2}$. Then $\frac{1}{2}L + \frac{1}{2}L''$ can be thought of as the compound lottery arising from a coin toss in which the decision maker gets L if heads comes up and L'' if tails does. Similarly, $\frac{1}{2}L' + \frac{1}{2}L''$ would be the coin toss where heads results in L' and tails results in L'' (see Figure 6.B.4). Note that conditional on heads, lottery $\frac{1}{2}L + \frac{1}{2}L''$ is at least as good as lottery $\frac{1}{2}L' + \frac{1}{2}L''$; but conditional on tails, the two compound lotteries give identical results. The independence axiom requires the sensible conclusion that $\frac{1}{2}L + \frac{1}{2}L''$ be at least as good as $\frac{1}{2}L' + \frac{1}{2}L''$.

The independence axiom is at the heart of the theory of choice under uncertainty. It is unlike anything encountered in the formal theory of preference-based choice discussed in Chapter 1 or its applications in Chapters 3 to 5. This is so precisely because it exploits, in a fundamental manner, the structure of uncertainty present in the model. In the theory of consumer demand, for example, there is no reason to believe that a consumer's preferences over various bundles of goods 1 and 2 should be independent of the quantities of the other goods that he will consume. In the present context, however, it *is* natural to think that a decision maker's preference between two lotteries, say L and L', should determine which of the two he prefers to have as part of a compound lottery *regardless* of the other possible outcome of this compound lottery, say L''. This other outcome L'' should be irrelevant to his choice because, in contrast with the consumer context, he does not consume L or L' together with L'' but, rather, only *instead* of it (if L or L' is the realized outcome).

Exercise 6.B.1: Show that if the preferences $\succsim$ over $\mathscr{L}$ satisfy the independence axiom, then for all $\alpha \in (0, 1)$ and $L, L', L'' \in \mathscr{L}$ we have

$$L \succ L' \quad \text{if and only if} \quad \alpha L + (1-\alpha)L'' \succ \alpha L' + (1-\alpha)L''$$

and

$$L \sim L' \quad \text{if and only if} \quad \alpha L + (1-\alpha)L'' \sim \alpha L' + (1-\alpha)L''.$$

Show also that if $L \succ L'$ and $L'' \succ L'''$, then $\alpha L + (1-\alpha)L'' \succ \alpha L' + (1-\alpha)L'''$.

As we will see shortly, the independence axiom is intimately linked to the representability of preferences over lotteries by a utility function that has an *expected utility form*. Before obtaining that result, we define this property and study some of its features.

Definition 6.B.5: The utility function $U\colon \mathscr{L} \to \mathbb{R}$ has an *expected utility form* if there is an assignment of numbers $(u_1, \ldots, u_N)$ to the N outcomes such that for every simple lottery $L = (p_1, \ldots, p_N) \in \mathscr{L}$ we have

$$U(L) = u_1 p_1 + \cdots + u_N p_N.$$

A utility function $U\colon \mathscr{L} \to \mathbb{R}$ with the expected utility form is called a *von Neumann–Morgenstern (v.N–M) expected utility function*.

Observe that if we let L^n denote the lottery that yields outcome n with probability one, then $U(L^n) = u_n$. Thus, the term *expected utility* is appropriate because with the v.N–M expected utility form, the utility of a lottery can be thought of as the expected value of the utilities u_n of the N outcomes.

The expression $U(L) = \sum_n u_n p_n$ is a general form for a *linear function in the probabilities* $(p_1, \ldots, p_N)$. This linearity property suggests a useful way to think about the expected utility form.

Proposition 6.B.1: A utility function $U\colon \mathscr{L} \to \mathbb{R}$ has an expected utility form if and only if it is *linear*, that is, if and only if it satisfies the property that

$$U\left(\sum_{k=1}^{K} \alpha_K L_k\right) = \sum_{k=1}^{K} \alpha_k U(L_k) \tag{6.B.1}$$

for any K lotteries $L_k \in \mathscr{L}$, $k = 1, \ldots, K$, and probabilities $(\alpha_1, \ldots, \alpha_K) \geq 0$, $\sum_k \alpha_k = 1$.

Proof: Suppose that $U(\cdot)$ satisfies property (6.B.1). We can write any $L = (p_1, \ldots, p_N)$ as a convex combination of the degenerate lotteries $(L^1, \ldots, L^N)$, that is, $L = \sum_n p_n L^n$. We have then $U(L) = U(\sum_n p_n L^n) = \sum_n p_n U(L^n) = \sum_n p_n u_n$. Thus, $U(\cdot)$ has the expected utility form.

In the other direction, suppose that $U(\cdot)$ has the expected utility form, and consider any compound lottery $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$, where $L_k = (p_1^k, \ldots, p_N^k)$. Its reduced lottery is $L' = \sum_k \alpha_k L_k$. Hence,

$$U\left(\sum_k \alpha_k L_k\right) = \sum_n u_n\left(\sum_k \alpha_k p_n^k\right) = \sum_k \alpha_k\left(\sum_n u_n p_n^k\right) = \sum_k \alpha_k U(L_k).$$

Thus, property (6.B.1) is satisfied. ■

The expected utility property is a *cardinal* property of utility functions defined on the space of lotteries. In particular, the result in Proposition 6.B.2 shows that the expected utility form is preserved only by increasing *linear* transformations.

Proposition 6.B.2: Suppose that $U\colon \mathscr{L} \to \mathbb{R}$ is a v.N–M expected utility function for the preference relation $\succsim$ on $\mathscr{L}$. Then $\tilde{U}\colon \mathscr{L} \to \mathbb{R}$ is another v.N–M utility function for $\succsim$ if and only if there are scalars $\beta > 0$ and γ such that $\tilde{U}(L) = \beta U(L) + \gamma$ for every $L \in \mathscr{L}$.

Proof: Begin by choosing two lotteries $\bar{L}$ and $\underline{L}$ with the property that $\bar{L} \succsim L \succsim \underline{L}$ for all $L \in \mathscr{L}$.[6] If $\bar{L} \sim \underline{L}$, then every utility function is a constant and the result follows immediately. Therefore, we assume from now on that $\bar{L} \succ \underline{L}$.

6. These best and worst lotteries can be shown to exist. We could, for example, choose a maximizer and a minimizer of the linear, hence continuous, function $U(\cdot)$ on the simplex of probabilities, a compact set.

Note first that if $U(\cdot)$ is a v.N–M expected utility function and $\tilde{U}(L) = \beta U(L) + \gamma$, then

$$\begin{aligned}\tilde{U}\left(\sum_{k=1}^{K} \alpha_k L_k\right) &= \beta U\left(\sum_{k=1}^{K} \alpha_k L_k\right) + \gamma \\ &= \beta\left[\sum_{k=1}^{K} \alpha_k U(L_k)\right] + \gamma \\ &= \sum_{k=1}^{K} \alpha_k[\beta U(L_k) + \gamma] \\ &= \sum_{k=1}^{K} \alpha_k \tilde{U}(L_k).\end{aligned}$$

Since $\tilde{U}(\cdot)$ satisfies property (6.B.1), it has the expected utility form.

For the reverse direction, we want to show that if both $\tilde{U}(\cdot)$ and $U(\cdot)$ have the expected utility form, then constants $\beta > 0$ and γ exist such that $\tilde{U}(L) = \beta U(L) + \gamma$ for all $L \in \mathscr{L}$. To do so, consider any lottery $L \in \mathscr{L}$, and define $\lambda_L \in [0, 1]$ by

$$U(L) = \lambda_L U(\bar{L}) + (1 - \lambda_L)U(\underline{L}).$$

Thus

$$\lambda_L = \frac{U(L) - U(\underline{L})}{U(\bar{L}) - U(\underline{L})} \qquad (6.B.2)$$

Since $\lambda_L U(\bar{L}) + (1 - \lambda_L)U(\underline{L}) = U(\lambda_L \bar{L} + (1 - \lambda_L)\underline{L})$ and $U(\cdot)$ represents the preferences $\succsim$, it must be that $L \sim \lambda_L \bar{L} + (1 - \lambda_L)\underline{L}$. But if so, then since $\tilde{U}(\cdot)$ is also linear and represents these same preferences, we have

$$\begin{aligned}\tilde{U}(L) &= \tilde{U}(\lambda_L \bar{L} + (1 - \lambda_L)\underline{L}) \\ &= \lambda_L \tilde{U}(\bar{L}) + (1 - \lambda_L)\tilde{U}(\underline{L}) \\ &= \lambda_L(\tilde{U}(\bar{L}) - \tilde{U}(\underline{L})) + \tilde{U}(\underline{L}).\end{aligned}$$

Substituting for λ_L from (6.B.2) and rearranging terms yields the conclusion that $\tilde{U}(L) = \beta U(L) + \gamma$, where

$$\beta = \frac{\tilde{U}(\bar{L}) - \tilde{U}(\underline{L})}{U(\bar{L}) - U(\underline{L})}$$

and

$$\gamma = \tilde{U}(\underline{L}) - U(\underline{L})\frac{\tilde{U}(\bar{L}) - \tilde{U}(\underline{L})}{U(\bar{L}) - U(\underline{L})}.$$

This completes the proof ■

A consequence of Proposition 6.B.2 is that for a utility function with the expected utility form, differences of utilities have meaning. For example, if there are four outcomes, the statement "the difference in utility between outcomes 1 and 2 is greater than the difference between outcomes 3 and 4," $u_1 - u_2 > u_3 - u_4$, is equivalent to

$$\tfrac{1}{2}u_1 + \tfrac{1}{2}u_4 > \tfrac{1}{2}u_2 + \tfrac{1}{2}u_3.$$

Therefore, the statement means that the lottery $L = (\frac{1}{2}, 0, 0, \frac{1}{2})$ is preferred to the lottery $L' = (0, \frac{1}{2}, \frac{1}{2}, 0)$. This ranking of utility differences is preserved by all linear transformations of the v.N–M expected utility function.

Note that if a preference relation $\succsim$ on $\mathscr{L}$ is representable by a utility function $U(\cdot)$ that has the expected utility form, then since a linear utility function is continuous, it follows that $\succsim$ is continuous on $\mathscr{L}$. More importantly, the preference relation $\succsim$ must also satisfy the independence axiom. You are asked to show this in Exercise 6.B.2.

Exercise 6.B.2: Show that if the preference relation $\succsim$ on $\mathscr{L}$ is represented by a utility function $U(\cdot)$ that has the expected utility form, then $\succsim$ satisfies the independence axiom.

The expected utility theorem, the central result of this section, tells us that the converse is also true.

The Expected Utility Theorem

The *expected utility theorem* says that if the decision maker's preferences over lotteries satisfy the continuity and independence axioms, then his preferences are representable by a utility function with the expected utility form. It is the most important result in the theory of choice under uncertainty, and the rest of the book bears witness to its usefulness.

Before stating and proving the result formally, however, it may be helpful to attempt an intuitive understanding of why it is true.

Consider the case where there are only three outcomes. As we have already observed, the continuity axiom insures that preferences on lotteries can be represented by some utility function. Suppose that we represent the indifference map in the simplex, as in Figure 6.B.5. Assume, for simplicity, that we have a conventional map with one-dimensional indifference curves. Because the expected utility form is linear in the probabilities, representability by the expected utility form is equivalent to these indifference curves being straight, parallel lines (you should check this). Figure 6.B.5(a) exhibits an indifference map satisfying these properties. We now argue that these properties are, in fact, consequences of the independence axiom.

Indifference curves are straight lines if, for every pair of lotteries L, L', we have that $L \sim L'$ implies $\alpha L + (1-\alpha)L' \sim L$ for all $\alpha \in [0,1]$. Figure 6.B.5(b) depicts a situation where the indifference curve is not a straight line; we have $L' \sim L$ but

3.5 Geometric explanation of the expected utility theorem. (a) $\succsim$ is representable by a utility function with the expected rm. (b) Contradiction of the independence axiom. (c) Contradiction of the independence axiom.

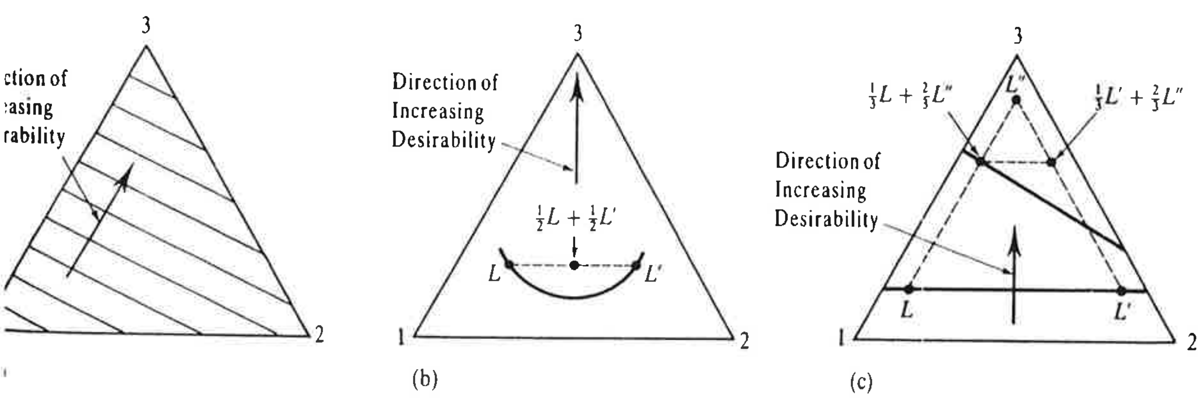

$\frac{1}{2}L' + \frac{1}{2}L \succ L$. This is equivalent to saying that

$$\tfrac{1}{2}L' + \tfrac{1}{2}L \succ \tfrac{1}{2}L + \tfrac{1}{2}L. \tag{6.B.3}$$

But since $L \sim L'$, the independence axiom implies that we must have $\frac{1}{2}L' + \frac{1}{2}L \sim \frac{1}{2}L + \frac{1}{2}L$ (see Exercise 6.B.1). This contradicts (6.B.3), and so we must conclude that indifference curves are straight lines.

Figure 6.B.5(c) depicts two straight but nonparallel indifference lines. A violation of the independence axiom can be constructed in this case, as indicated in the figure. There we have $L \succsim L'$ (in fact, $L \sim L'$), but $\frac{1}{3}L + \frac{2}{3}L'' \succsim \frac{1}{3}L' + \frac{2}{3}L''$ does not hold for the lottery L'' shown in the figure. Thus, indifference curves must be parallel, straight lines if preferences satisfy the independence axiom.

In Proposition 6.B.3, we formally state and prove the expected utility theorem.

Proposition 6.B.3: (*Expected Utility Theorem*) Suppose that the rational preference relation $\succsim$ on the space of lotteries $\mathscr{L}$ satisfies the continuity and independence axioms. Then $\succsim$ admits a utility representation of the expected utility form. That is, we can assign a number u_n to each outcome $n = 1, \ldots, N$ in such a manner that for any two lotteries $L = (p_1, \ldots, p_N)$ and $L' = (p'_1, \ldots, p'_N)$, we have

$$L \succsim L \quad \text{if and only if} \quad \sum_{n=1}^{N} u_n p_n \geq \sum_{n=1}^{N} u_n p'_n. \tag{6.B.4}$$

Proof: We organize the proof in a succession of steps. For simplicity, we assume that there are best and worst lotteries in $\mathscr{L}$, $\bar{L}$ and $\underline{L}$ (so, $\bar{L} \succsim L \succsim \underline{L}$ for any $L \in \mathscr{L}$).[7] If $\bar{L} \sim \underline{L}$, then all lotteries in $\mathscr{L}$ are indifferent and the conclusion of the proposition holds trivially. Hence, from now on, we assume that $\bar{L} \succ \underline{L}$.

Step 1. If $L \succ L'$ and $\alpha \in (0, 1)$, then $L \succ \alpha L + (1 - \alpha)L' \succ L'$.

This claim makes sense. A nondegenerate mixture of two lotteries will hold a preference position strictly intermediate between the positions of the two lotteries. Formally, the claim follows from the independence axiom. In particular, since $L \succ L'$, the independence axiom implies that (recall Exercise 6.B.1)

$$L = \alpha L + (1 - \alpha)L \succ \alpha L + (1 - \alpha)L' \succ \alpha L' + (1 - \alpha)L' = L'.$$

Step 2. Let $\alpha, \beta \in [0, 1]$. Then $\beta\bar{L} + (1 - \beta)\underline{L} \succ \alpha\bar{L} + (1 - \alpha)\underline{L}$ if and only if $\beta > \alpha$.

Suppose that $\beta > \alpha$. Note first that we can write

$$\beta\bar{L} + (1 - \beta)\underline{L} = \gamma\bar{L} + (1 - \gamma)[\alpha\bar{L} + (1 - \alpha)\underline{L}],$$

where $\gamma = [(\beta - \alpha)/(1 - \alpha)] \in (0, 1]$. By Step 1, we know that $\bar{L} \succ \alpha\bar{L} + (1 - \alpha)\underline{L}$. Applying Step 1 again, this implies that $\gamma\bar{L} + (1 - \gamma)(\alpha\bar{L} + (1 - \alpha)\underline{L}) \succ \alpha\bar{L} + (1 - \alpha)\underline{L}$, and so we conclude that $\beta\bar{L} + (1 - \beta)\underline{L} \succ \alpha\bar{L} + (1 - \alpha)\underline{L}$.

For the converse, suppose that $\beta \leq \alpha$. If $\beta = \alpha$, we must have $\beta\bar{L} + (1 - \beta)\underline{L} \sim \alpha\bar{L} + (1 - \alpha)\underline{L}$. So suppose that $\beta < \alpha$. By the argument proved in the previous

7. In fact, with our assumption of a finite set of outcomes, this can be established as a consequence of the independence axiom (see Exercise 6.B.3).

paragraph (reversing the roles of α and β), we must then have $\alpha\bar{L} + (1-\alpha)\underline{L} \succ \beta\bar{L} + (1-\beta)\underline{L}$.

Step 3. For any $L \in \mathscr{L}$, there is a unique α_L such that $[\alpha_L\bar{L} + (1-\alpha_L)\underline{L}] \sim L$.

Existence of such an α_L is implied by the continuity of $\succsim$ and the fact that $\bar{L}$ and $\underline{L}$ are, respectively, the best and the worst lottery. Uniqueness follows from the result of Step 2.

The existence of α_L is established in a manner similar to that used in the proof of Proposition 3.C.1. Specifically, define the sets

$$\{\alpha \in [0,1]: \alpha\bar{L} + (1-\alpha)\underline{L} \succsim L\} \quad \text{and} \quad \{\alpha \in [0,1]: L \succsim \alpha\bar{L} + (1-\alpha)\underline{L}\}.$$

By the continuity and completeness of $\succsim$, both sets are closed, and any $\alpha \in [0,1]$ belongs to at least one of the two sets. Since both sets are nonempty and $[0,1]$ is connected, it follows that there is some α belonging to both. This establishes the existence of an α_L such that $\alpha_L\bar{L} + (1-\alpha_L)\underline{L} \sim L$.

Step 4. The function $U: \mathscr{L} \to \mathbb{R}$ that assigns $U(L) = \alpha_L$ for all $L \in \mathscr{L}$ represents the preference relation $\succsim$.

Observe that, by Step 3, for any two lotteries $L, L' \in \mathscr{L}$, we have

$$L \succsim L' \quad \text{if and only if} \quad \alpha_L\bar{L} + (1-\alpha_L)\underline{L} \succsim \alpha_{L'}\bar{L} + (1-\alpha_{L'})\underline{L}.$$

Thus, by Step 2, $L \succsim L'$ if and only if $\alpha_L \geq \alpha_{L'}$.

Step 5. The utility function $U(\cdot)$ that assigns $U(L) = \alpha_L$ for all $L \in \mathscr{L}$ is linear and therefore has the expected utility form.

We want to show that for any $L, L' \in \mathscr{L}$, and $\beta \in [0,1]$, we have $U(\beta L + (1-\beta)L') = \beta U(L) + (1-\beta)U(L')$. By definition, we have

$$L \sim U(L)\bar{L} + (1-U(L))\underline{L}$$

and

$$L' \sim U(L')\bar{L} + (1-U(L'))\underline{L}.$$

Therefore, by the independence axiom (applied twice),

$$\begin{aligned}\beta L + (1-\beta)L' &\sim \beta[U(L)\bar{L} + (1-U(L))\underline{L}] + (1-\beta)L' \\ &\sim \beta[U(L)\bar{L} + (1-U(L))\underline{L}] + (1-\beta)[U(L')\bar{L} + (1-U(L'))\underline{L}].\end{aligned}$$

Rearranging terms, we see that the last lottery is algebraically identical to the lottery

$$[\beta U(L) + (1-\beta)U(L')]\bar{L} + [1 - \beta U(L) - (1-\beta)U(L')]\underline{L}.$$

In other words, the compound lottery that gives lottery $[U(L)\bar{L} + (1-U(L))\underline{L}]$ with probability β and lottery $[U(L')\bar{L} + (1-U(L'))\underline{L}]$ with probability $(1-\beta)$ has the same reduced lottery as the compound lottery that gives lottery $\bar{L}$ with probability $[\beta U(L) + (1-\beta)U(L')]$ and lottery $\underline{L}$ with probability $[1 - \beta U(L) - (1-\beta)U(L')]$. Thus

$$\beta L + (1-\beta)L' \sim [\beta U(L) + (1-\beta)U(L')]\bar{L} + [1 - \beta U(L) - (1-\beta)U(L')]\underline{L}.$$

By the construction of $U(\cdot)$ in Step 4, we therefore have

$$U(\beta L + (1-\beta)L') = \beta U(L) + (1-\beta)U(L'),$$

as we wanted.

Together, Steps 1 to 5 establish the existence of a utility function representing $\succsim$ that has the expected utility form. ■

Discussion of the Theory of Expected Utility

A first advantage of the expected utility theorem is technical: It is extremely convenient analytically. This, more than anything else, probably accounts for its pervasive use in economics. It is very easy to work with expected utility and very difficult to do without it. As we have already noted, the rest of the book attests to the importance of the result. Later in this chapter, we will explore some of the analytical uses of expected utility.

A second advantage of the theorem is normative: Expected utility may provide a valuable guide to action. People often find it hard to think systematically about risky alternatives. But if an individual believes that his choices should satisfy the axioms on which the theorem is based (notably, the independence axiom), then the theorem can be used as a guide in his decision process. This point is illustrated in Example 6.B.1.

Example 6.B.1: *Expected Utility as a Guide to Introspection.* A decision maker may not be able to assess his preference ordering between the lotteries L and L' depicted in Figure 6.B.6. The lotteries are too close together, and the differences in the probabilities involved are too small to be understood. Yet, if the decision maker believes that his preferences should satisfy the assumptions of the expected utility theorem, then he may consider L'' instead, which is on the straight line spanned by L and L' but at a significant distance from L. The lottery L'' may not be a feasible choice, but if he determines that $L'' \succ L$, then he can conclude that $L' \succ L$. Indeed, if $L'' \succ L$, then there is an indifference curve separating these two lotteries, as shown in the figure, and it follows from the fact that indifference curves are a family of parallel straight lines that there is also an indifference curve separating L' and L, so that $L' \succ L$. Note that this type of inference is not possible using only the general

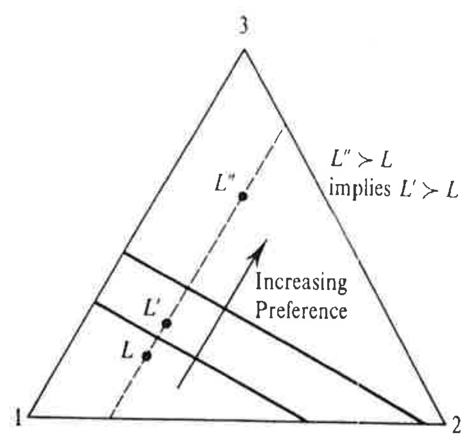

Figure 6.B.6
Expected utility as a guide to introspection.

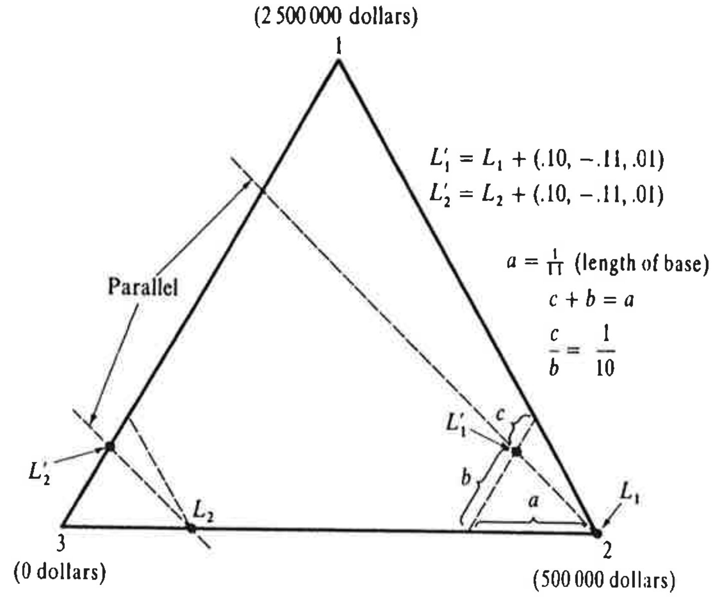

Figure 6.B.7
Depiction of the Allais paradox in the simplex.

choice theory of Chapter 1 because, without the hypotheses of the expected utility theorem, the indifference curves need not be straight lines (with a general indifference map, we could perfectly well have $L'' \succ L$ and $L \succ L'$).

A concrete example of this use of the expected utility theorem is developed in Exercise 6.B.4. ■

As a descriptive theory, however, the expected utility theorem (and, by implication, its central assumption, the independence axiom), is not without difficulties. Examples 6.B.2 and 6.B.3 are designed to test its plausibility.

Example 6.B.2: *The Allais Paradox.* This example, known as the Allais paradox [from Allais (1953)], constitutes the oldest and most famous challenge to the expected utility theorem. It is a thought experiment. There are three possible monetary prizes (so the number of outcomes is $N = 3$):

First Prize	*Second Prize*	*Third Prize*
2 500 000 dollars	500 000 dollars	0 dollars

The decision maker is subjected to two choice tests. The first consists of a choice between the lotteries L_1 and L_1':

$$L_1 = (0, 1, 0) \qquad L_1' = (.10, .89, .01).$$

The second consists of a choice between the lotteries L_2 and L_2':

$$L_2 = (0, .11, .89) \qquad L_2' = (.10, 0, .90).$$

The four lotteries involved are represented in the simplex diagram of Figure 6.B.7.

It is common for individuals to express the preferences $L_1 \succ L_1'$ and $L_2' \succ L_2$.[8]

8. In our classroom experience, roughly half the students choose this way.

The first choice means that one prefers the certainty of receiving 500 000 dollars over a lottery offering a 1/10 probability of getting five times more but bringing with it a tiny risk of getting nothing. The second choice means that, all things considered, a 1/10 probability of getting 2 500 000 dollars is preferred to getting only 500 000 dollars with the slightly better odds of 11/100.

However, these choices are not consistent with expected utility. This can be seen in Figure 6.B.7: The straight lines connecting L_1 to L_1' and L_2 to L_2' are parallel. Therefore, if an individual has a linear indifference curve that lies in such a way that L_1 is preferred to L_1', then a parallel linear indifference curve must make L_2 preferred to L_2', and vice versa. Hence, choosing L_1 and L_2' is inconsistent with preferences satisfying the assumptions of the expected utility theorem.

More formally, suppose that there was a v.N–M expected utility function. Denote by u_{25}, u_{05}, and u_0 the utility values of the three outcomes. Then the choice $L_1 \succ L_1'$ implies

$$u_{05} > (.10)u_{25} + (.89)u_{05} + (.01)u_0.$$

Adding $(.89)u_0 - (.89)u_{05}$ to both sides, we get

$$(.11)u_{05} + (.89)u_0 > (.10)u_{25} + (.90)u_0,$$

and therefore any individual with a v.N–M utility function must have $L_2 \succ L_2'$. ■

There are four common reactions to the Allais paradox. The first, propounded by J. Marshack and L. Savage, goes back to the normative interpretation of the theory. It argues that choosing under uncertainty is a reflective activity in which one should be ready to correct mistakes if they are proven inconsistent with the basic principles of choice embodied in the independence axiom (much as one corrects arithmetic mistakes).

The second reaction maintains that the Allais paradox is of limited significance for economics as a whole because it involves payoffs that are out of the ordinary and probabilities close to 0 and 1.

A third reaction seeks to accommodate the paradox with a theory that defines preferences over somewhat larger and more complex objects than simply the ultimate lottery over outcomes. For example, the decision maker may value not only what he receives but also what he receives compared with what he might have received by choosing differently. This leads to *regret theory*. In the example, we could have $L_1 \succ L_1'$ because the expected regret caused by the possibility of getting zero in lottery L_1', when choosing L_1 would have assured 500 000 dollars, is too great. On the other hand, with the choice between L_2 and L_2', no such clear-cut regret potential exists; the decision maker was very likely to get nothing anyway.

The fourth reaction is to stick with the original choice domain of lotteries but to give up the independence axiom in favor of something weaker. Exercise 6.B.5 develops this point further.

Example 6.B.3: *Machina's paradox.* Consider the following three outcomes: "a trip to Venice," "watching an excellent movie about Venice," and "staying home." Suppose that you prefer the first to the second and the second to the third.

Now you are given the opportunity to choose between two lotteries. The first lottery gives "a trip to Venice" with probability 99.9% and "watching an excellent movie about Venice" with probability 0.1%. The second lottery gives "a trip to

Venice," again with probability 99.9% and "staying home" with probability 0.1%. The independence axiom forces you to prefer the first lottery to the second. Yet, it would be understandable if you did otherwise. Choosing the second lottery is the rational thing to do if you anticipate that in the event of not getting the trip to Venice, your tastes over the other two outcomes will change: You will be severely *disappointed* and will feel miserable watching a movie about Venice.

The idea of disappointment has parallels with the idea of regret that we discussed in connection with the Allais paradox, but it is not quite the same. Both ideas refer to the influence of "what might have been" on the level of well-being experienced, and it is because of this that they are in conflict with the independence axiom. But disappointment is more directly concerned with what might have been if another outcome of a given lottery had come up, whereas regret should be thought of as regret over a choice not made. ■

Because of the phenomena illustrated in the previous two examples, the search for a useful theory of choice under uncertainty that does not rely on the independence axiom has been an active area of research [see Machina (1987) and also Hey and Orme (1994)]. Nevertheless, the use of the expected utility theorem is pervasive in economics.

An argument sometimes made against the practical significance of violations of the independence axiom is that individuals with such preferences would be weeded out of the marketplace because they would be open to the acceptance of so-called "Dutch books," that is, deals leading to a sure loss of money. Suppose, for example, that there are three lotteries such that $L \succ L'$ and $L \succ L''$ but, in violation of the independence axiom, $\alpha L' + (1-\alpha)L'' \succ L$ for some $\alpha \in (0, 1)$. Then, when the decision maker is in the initial position of owning the right to lottery L, he would be willing to pay a small fee to trade L for a compound lottery yielding lottery L' with probability α and lottery L'' with probability $(1-\alpha)$. But as soon as the first stage of this lottery is over, giving him either L' or L'' we could get him to pay a fee to trade this lottery for L. Hence, at that point, he would have paid the two fees but would otherwise be back to his original position.

This may well be a good argument for convexity of the not-better-than sets of $\succsim$, that is, for it to be the case that $L \succsim \alpha L' + (1-\alpha)L''$ whenever $L \succsim L'$ and $L \succsim L''$. This property is implied by the independence axiom but is weaker than it. Dutch book arguments for the full independence axiom are possible, but they are more contrived [see Green (1987)].

Finally, one must use some caution in applying the expected utility theorem because in many practical situations the final outcomes of uncertainty are influenced by actions taken by individuals. Often, these actions should be explicitly modeled but are not. Example 6.B.4 illustrates the difficulty involved.

Example 6.B.4: *Induced preferences.* You are invited to a dinner where you may be offered fish (F) or meat (M). You would like to do the proper thing by showing up with white wine if F is served and red wine if M is served. The action of buying the wine must be taken *before* the uncertainty is resolved.

Suppose now that the cost of the bottle of red or white wine is the same and that you are also indifferent between F and M. If you think of the possible outcomes as F and M, then you are apparently indifferent between the lottery that gives F with certainty and the lottery that gives M with certainty. The independence axiom would then seem to require that you also be indifferent to a lottery that gives F or M with probability $\frac{1}{2}$ each. But you would clearly not be indifferent, since knowing that either F or M will be served with certainty allows you to buy the right wine, whereas, if you are not certain, you will either have to buy both wines or else bring the wrong wine with probability $\frac{1}{2}$.

Yet this example does not contradict the independence axiom. To appeal to the axiom, the decision framework must be set up so that the satisfaction derived from an outcome does not depend on any action taken by the decision maker before the uncertainty is resolved. *Thus, preferences should not be induced or derived from ex ante actions.*[9] Here, the action "acquisition of a bottle of wine" is taken before the uncertainty about the meal is resolved.

To put this situation into the framework required, we must include the ex ante action as part of the description of outcomes. For example, here there would be four outcomes: "bringing red wine when served M," "bringing white wine when served M," "bringing red wine when served F," and "bringing white wine when served F." For any underlying uncertainty about what will be served, you induce a lottery over these outcomes by your choice of action. In this setup, it is quite plausible to be indifferent among "having meat and bringing red wine," "having fish and bringing white wine," or any lottery between these two outcomes, as the independence axiom requires. ■

Although it is not a contradiction to the postulates of expected utility theory, and therefore it is not a serious conceptual difficulty, the induced preferences example nonetheless raises a practical difficulty in the use of the theory. The example illustrates the fact that, in applications, many economic situations do not fit the pure framework of expected utility theory. Preferences are almost always, to some extent, induced.[10]

The expected utility theorem does impose some structure on induced preferences. For example, suppose the complete set of outcomes is $B \times A$, where $B = \{b_1, \dots, b_N\}$ is the set of possible realizations of an exogenous randomness and A is the decision maker's set of possible (ex ante) actions. Under the conditions of the expected utility theorem, for every $a \in A$ and $b_n \in B$, we can assign some utility value $u_n(a)$ to the outcome (b_n, a). Then, for every exogenous lottery $L = (p_1, \dots, p_N)$ on B, we can define a derived utility function by maximizing expected utility:

$$U(L) = \operatorname*{Max}_{a \in A} \sum_n p_n u_n(a).$$

In Exercise 6.B.6, you are asked to show that while $U(L)$, a function on $\mathscr{L}$, need not be linear,

9. Actions taken ex post do not create problems. For example, suppose that $u_n(a_n)$ is the utility derived from outcome n when action a_n is taken after the realization of uncertainty. The decision maker therefore chooses a_n to solve $\text{Max}_{a_n \in A_n} u_n(a_n)$, where A_n is the set of possible actions when outcome n occurs. We can then let $u_n = \text{Max}_{a_n \in A_n} u_n(a_n)$ and evaluate lotteries over the N outcomes as in expected utility theory.

10. Consider, for example, preferences for lotteries over amounts of money available tomorrow. Unless the individual's preferences over consumption today and tomorrow are additively separable, his decision of how much to consume today—a decision that must be made before the resolution of the uncertainty concerning tomorrow's wealth—affects his preferences over these lotteries in a manner that conflicts with the fulfillment of the independence axiom.

Definition 6.C.5: Given a Bernoulli utility function $u(\cdot)$, the *coefficient of relative risk aversion at* x is $r_R(x, u) = -xu''(x)/u'(x)$.

Consider now how this measure varies with wealth. The property of *nonincreasing relative risk aversion* says that the individual becomes less risk averse with regard to gambles that are proportional to his wealth as his wealth increases. This is a stronger assumption than decreasing absolute risk aversion: Since $r_R(x, u) = xr_A(x, u)$, a risk-averse individual with decreasing relative risk aversion will exhibit decreasing absolute risk aversion, but the converse is not necessarily the case.

As before, we can examine various implications of this concept. Proposition 6.C.4 is an abbreviated parallel to Proposition 6.C.3.

Proposition 6.C.4: The following conditions for a Bernoulli utility function $u(\cdot)$ on amounts of money are equivalent:

(i) $r_R(x, u)$ is decreasing in x.

(ii) Whenever $x_2 < x_1$, $\tilde{u}_2(t) = u(tx_2)$ is a concave transformation of $\tilde{u}_1(t) = u(tx_1)$.

(iii) Given any risk $F(t)$ on $t > 0$, the certainty equivalent $\bar{c}_x$ defined by $u(\bar{c}_x) = \int u(tx)\, dF(t)$ is such that $x/\bar{c}_x$ is decreasing in x.

Proof: Here we show only that (i) implies (iii). To this effect, fix a distribution $F(t)$ on $t > 0$, and, for any x, define $u_x(t) = u(tx)$. Let $c(x)$ be the usual certainty equivalent (from Definition 6.C.2): $u_x(c(x)) = \int u_x(t)\, dF(t)$. Note that $-u_x''(t)/u_x'(t) = -(1/t)tx[u''(tx)/u'(tx)]$ for any x. Hence if (i) holds, then $u_{x'}(\cdot)$ is less risk averse than $u_x(\cdot)$ whenever $x' > x$. Therefore, by Proposition 6.C.2, $c(x') > c(x)$ and we conclude that $c(\cdot)$ is increasing. Now, by the definition of $u_x(\cdot)$, $u_x(c(x)) = u(xc(x))$. Also

$$u_x(c(x)) = \int u_x(t)\, dF(t) = \int u(tx)\, dF(t) = u(\bar{c}_x).$$

Hence, $\bar{c}_x/x = c(x)$, and so $x/\bar{c}_x$ is decreasing. This concludes the proof. ∎

Example 6.C.2 continued: In Exercise 6.C.11, you are asked to show that if $r_R(x, u)$ is decreasing in x, then the proportion of wealth invested in the risky asset $\gamma = \alpha/w$ is increasing with the individual's wealth level w. The opposite conclusion holds if $r_R(x, u)$ is increasing in x. If $r_R(x, u)$ is a constant independent of x, then the fraction of wealth invested in the risky asset is independent of w [see Exercise 6.C.12 for the specific analytical form that $u(\cdot)$ must have]. Models with constant relative risk aversion are encountered often in finance theory, where they lead to considerable analytical simplicity. Under this assumption, no matter how the wealth of the economy and its distribution across individuals evolves over time, the portfolio decisions of individuals in terms of budget shares do not vary (as long as the safe return and the distribution of random returns remain unchanged). ∎

6.D Comparison of Payoff Distributions in Terms of Return and Risk

In this section, we continue our study of lotteries with monetary payoffs. In contrast with Section 6.C, where we compared utility functions, our aim here is to compare payoff distributions. There are two natural ways that random outcomes can be compared: according to the level of returns and according to the dispersion of returns. We will therefore attempt to give meaning to two ideas: that of a distribution $F(\cdot)$ yielding unambiguously higher returns than $G(\cdot)$ and that of $F(\cdot)$ being unambiguously less risky than $G(\cdot)$. These ideas are known, respectively, by the technical terms of *first-order stochastic dominance* and *second-order stochastic dominance*.[20]

In all subsequent developments, we restrict ourselves to distributions $F(\cdot)$ such that $F(0) = 0$ and $F(x) = 1$ for some x.

First-Order Stochastic Dominance

We want to attach meaning to the expression: "The distribution $F(\cdot)$ yields unambiguously higher returns than the distribution $G(\cdot)$." At least two sensible criteria suggest themselves. First, we could test whether every expected utility maximizer who values more over less prefers $F(\cdot)$ to $G(\cdot)$. Alternatively, we could verify whether, for every amount of money x, the probability of getting at least x is higher under $F(\cdot)$ than under $G(\cdot)$. Fortunately, these two criteria lead to the same concept.

Definition 6.D.1: The distribution $F(\cdot)$ *first-order stochastically dominates* $G(\cdot)$ if, for every nondecreasing function $u\colon \mathbb{R} \to \mathbb{R}$, we have

$$\int u(x)\, dF(x) \geq \int u(x)\, dG(x).$$

Proposition 6.D.1: The distribution of monetary payoffs $F(\cdot)$ first-order stochastically dominates the distribution $G(\cdot)$ if and only if $F(x) \leq G(x)$ for every x.

Proof: Given $F(\cdot)$ and $G(\cdot)$ denote $H(x) = F(x) - G(x)$. Suppose that $H(\bar{x}) > 0$ for some $\bar{x}$. Then we can define a nondecreasing function $u(\cdot)$ by $u(x) = 1$ for $x > \bar{x}$ and $u(x) = 0$ for $x \leq \bar{x}$. This function has the property that $\int u(x)\, dH(x) = -H(\bar{x}) < 0$, and so the "only if" part of the proposition follows.

For the "if" part of the proposition we first put on record, without proof, that it suffices to establish the equivalence for differentiable utility functions $u(\cdot)$. Given $F(\cdot)$ and $G(\cdot)$, denote $H(x) = F(x) - G(x)$. Integrating by parts, we have

$$\int u(x)\, dH(x) = [u(x)H(x)]_0^\infty - \int u'(x)H(x)\, dx.$$

Since $H(0) = 0$ and $H(x) = 0$ for large x, the first term of this expression is zero. It follows that $\int u(x)\, dH(x) \geq 0$ [or, equivalently, $\int u(x)\, dF(x) - \int u(x)\, dG(x) \geq 0$] if and only if $\int u'(x)H(x)\, dx \leq 0$. Thus, if $H(x) \leq 0$ for all x and $u(\cdot)$ is increasing, then $\int u'(x)H(x)\, dx \leq 0$ and the "if" part of the proposition follows. ■

In Exercise 6.D.1 you are asked to verify Proposition 6.D.1 for the case of lotteries over three possible outcomes. In Figure 6.D.1, we represent two distributions $F(\cdot)$ and $G(\cdot)$. Distribution $F(\cdot)$ first-order stochastically dominates $G(\cdot)$ because the graph of $F(\cdot)$ is uniformly below the graph of $G(\cdot)$. Note two important points: First, first-order stochastic dominance does *not* imply that every possible return of the

20. They were introduced into economics in Rothschild and Stiglitz (1970).

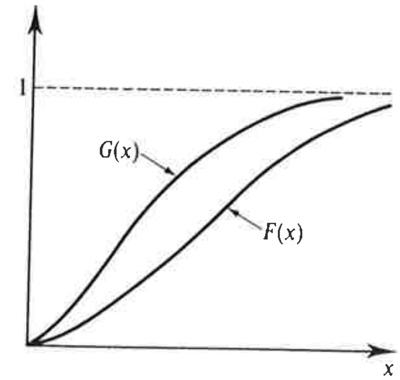

Figure 6.D.1
$F(\cdot)$ first-order stochastically dominates $G(\cdot)$.

superior distribution is larger than every possible return of the inferior one. In the figure, the set of possible outcomes is the same for the two distributions. Second, although $F(\cdot)$ first-order stochastically dominating $G(\cdot)$ implies that the mean of x under $F(\cdot)$, $\int x\, dF(x)$, is greater than its mean under $G(\cdot)$, a ranking of the means of two distributions does *not* imply that one first-order stochastically dominates the other; rather, the entire distribution matters (see Exercise 6.D.3).

Example 6.D.1: Consider a compound lottery that has as its first stage a realization of x distributed according to $G(\cdot)$ and in its second stage applies to the outcome x of the first stage an "upward probabilistic shift." That is, if outcome x is realized in the first stage, then the second stage pays a final amount of money $x + z$, where z is distributed according to a distribution $H_x(z)$ with $H_x(0) = 0$. Thus, $H_x(\cdot)$ generates a *final* return of at least x with probability one. (Note that the distributions applied to different x's may differ.)

Denote the resulting reduced distribution by $F(\cdot)$. Then for any nondecreasing function $u\colon \mathbb{R} \to \mathbb{R}$, we have

$$\int u(x)\, dF(x) = \int \left[\int u(x + z)\, dH_x(z) \right] dG(x) \geq \int u(x)\, dG(x).$$

So $F(\cdot)$ first-order stochastically dominates $G(\cdot)$.

A specific example is illustrated in Figure 6.D.2. As Figure 6.D.2(a) shows, $G(\cdot)$ is an even randomization between 1 and 4 dollars. The outcome "1 dollar" is then

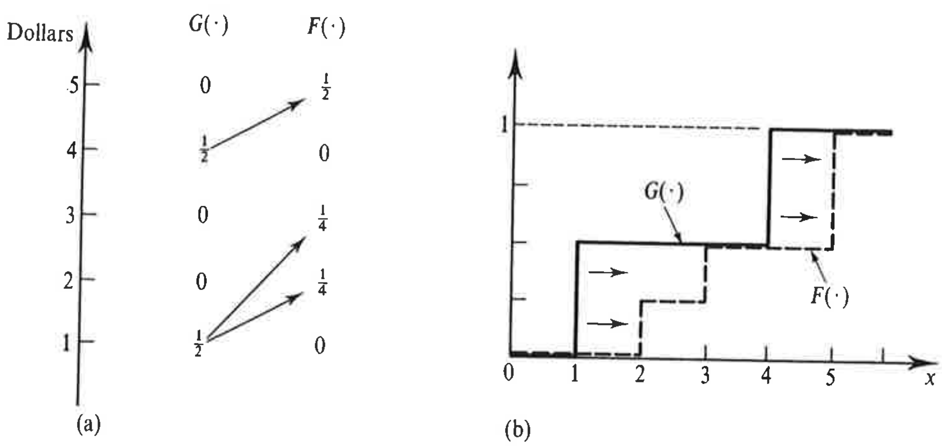

Figure 6.D.2
$F(\cdot)$ first-order stochastically dominates $G(\cdot)$.

shifted up to an even probability between 2 and 3 dollars, and the outcome "4 dollars" is shifted up to 5 dollars with probability one. Figure 6.D.2(b) shows that $F(x) \leq G(x)$ at all x.

It can be shown that the reverse direction also holds. Whenever $F(\cdot)$ first-order stochastically dominates $G(\cdot)$, it is possible to generate $F(\cdot)$ from $G(\cdot)$ in the manner suggested in this example. Thus, this provides yet another approach to the characterization of the first-order stochastic dominance relation. ■

Second-Order Stochastic Dominance

First-order stochastic dominance involves the idea of "higher/better" vs. "lower/worse." We want next to introduce a comparison based on relative *riskiness* or *dispersion*. To avoid confusing this issue with the trade-off between returns and risk, *we will restrict ourselves for the rest of this section to comparing distributions with the same mean.*

Once again, a definition suggests itself: Given two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean [that is, with $\int x\, dF(x) = \int x\, dG(x)$], we say that $G(\cdot)$ is riskier than $F(\cdot)$ if every risk averter prefers $F(\cdot)$ and $G(\cdot)$. This is stated formally in Definition 6.D.2.

Definition 6.D.2: For any two distributions $F(x)$ and $G(\cdot)$ with the same mean, $F(\cdot)$ *second-order stochastically dominates* (or *is less risky than*) $G(\cdot)$ if for every nondecreasing concave function $u: \mathbb{R}_+ \to \mathbb{R}$, we have

$$\int u(x)\, dF(x) \geq \int u(x)\, dG(x).$$

Example 6.D.2 introduces an alternative way to characterize the second-order stochastic dominance relation.

Example 6.D.2: *Mean-Preserving Spreads.* Consider the following compound lottery: In the first stage, we have a lottery over x distributed according to $F(\cdot)$. In the second stage, we randomize each possible outcome x further so that the final payoff is $x + z$, where z has a distribution function $H_x(z)$ with a mean of zero [i.e., $\int z\, dH_x(z) = 0$]. Thus, the mean of $x + z$ is x. Let the resulting reduced lottery be denoted by $G(\cdot)$. When lottery $G(\cdot)$ can be obtained from lottery $F(\cdot)$ in this manner for some distribution $H_x(\cdot)$, we say that $G(\cdot)$ is a *mean-preserving spread* of $F(\cdot)$.

For example, $F(\cdot)$ may be an even probability distribution between 2 and 3 dollars. In the second step we may spread the 2 dollars outcome to an even probability between 1 and 3 dollars, and the 3 dollars outcome to an even probability between 2 and 4 dollars. Then $G(\cdot)$ is the distribution that assigns probability $\frac{1}{4}$ to the four outcomes: 1, 2, 3, 4 dollars. These two distributions $F(\cdot)$ and $G(\cdot)$ are depicted in Figure 6.D.3.

The type of two-stage operation just described keeps the mean of $G(\cdot)$ equal to that of $F(\cdot)$. In addition, if $u(\cdot)$ is concave, we can conclude that

$$\int u(x)\, dG(x) = \int \left(\int u(x+z)\, dH_x(z) \right) dF(x) \leq \int u\left(\int (x+z)\, dH_x(z) \right) dF(x)$$
$$= \int u(x)\, dF(x),$$

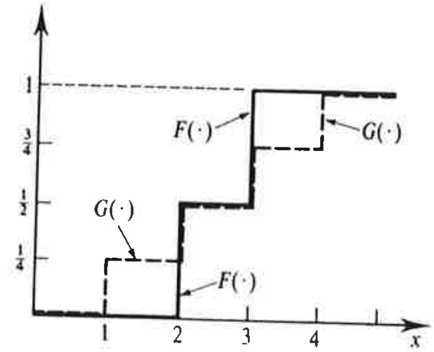

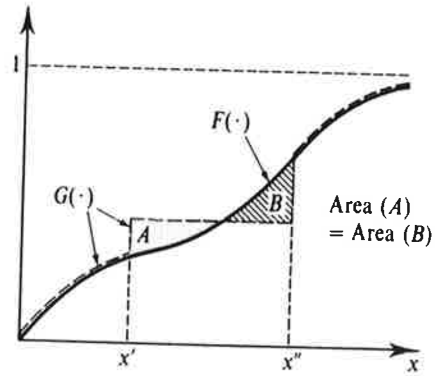

Figure 6.D.3 (left) $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$.

Figure 6.D.4 (right) $G(\cdot)$ is an elementary increase in risk from $F(\cdot)$.

and so $F(\cdot)$ second-order stochastically dominates $G(\cdot)$. It turns out that the converse is also true: If $F(\cdot)$ second-order stochastically dominates $G(\cdot)$, then $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$. Hence, *saying that $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$ is equivalent to saying that $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.* ■

Example 6.D.3 provides another illustration of a mean-preserving spread.

Example 6.D.3: *An Elementary Increase in Risk.* We say that $G(\cdot)$ constitutes an *elementary increase in risk* from $F(\cdot)$ if $G(\cdot)$ is generated from $F(\cdot)$ by taking all the mass that $F(\cdot)$ assigns to an interval $[x', x'']$ and transferring it to the endpoints x' and x'' in such a manner that the mean is preserved. This is illustrated in Figure 6.D.4. An elementary increase in risk is a mean-preserving spread. [In Exercise 6.D.3, you are asked to verify directly that if $G(\cdot)$ is an elementary increase in risk from $F(\cdot)$, then $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.] ■

We can develop still another way to capture the second-order stochastic dominance idea. Suppose that we have two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean. Recall that, for simplicity, we assume that $F(\bar{x}) = G(\bar{x}) = 1$ for some $\bar{x}$. Integrating by parts (and recalling the equality of the means) yields

$$\int_0^{\bar{x}} (F(x) - G(x))\, dx = -\int_0^{\bar{x}} x\, d(F(x) - G(x)) + (F(\bar{x}) - G(\bar{x}))\bar{x} = 0. \qquad (6.D.1)$$

That is, the areas below the two distribution functions are the same over the interval $[0, \bar{x}]$. Because of this fact, the regions marked A and B in Figure 6.D.4 must have the same area. Note that for the two distributions in the figure, this implies that

$$\int_0^x G(t)\, dt \geq \int_0^x F(t)\, dt \quad \text{for all } x. \qquad (6.D.2)$$

It turns out that property (6.D.2) is equivalent to $F(\cdot)$ second-order stochastically dominating $G(\cdot)$.[21] As an application, suppose that $F(\cdot)$ and $G(\cdot)$ have the same mean and that the graph of $G(\cdot)$ is initially above the graph of $F(\cdot)$ and then moves

21. We will not prove this. The claim can be established along the same lines used to prove Proposition 6.D.1 except that we must integrate by parts twice and take into account expression (6.D.1).

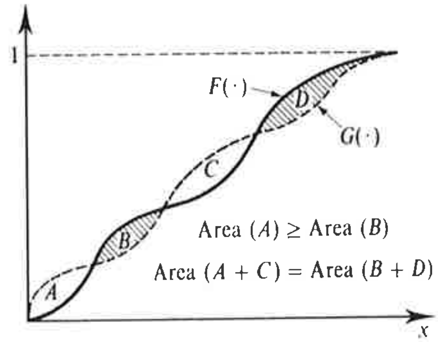

Figure 6.D.5
$F(\cdot)$ second-order stochastically dominates $G(\cdot)$.

permanently below it (as in Figures 6.D.3 and 6.D.4). Then because of (6.D.1), condition (6.D.2) must be satisfied, and we can conclude that $G(\cdot)$ is riskier than $F(\cdot)$. As a more elaborate example, consider Figure 6.D.5, which shows two distributions having the same mean and satisfying (6.D.2). To verify that (6.D.2) is satisfied, note that area A has been drawn to be at least as large as area B and that the equality of the means [i.e., (6.D.1)] implies that the areas $B + D$ and $A + C$ must be equal.

We state Proposition 6.D.2 without proof.

Proposition 6.D.2: Consider two distributions $F(\cdot)$ and $G(\cdot)$ with the same mean. Then the following statements are equivalent:

(i) $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.
(ii) $G(\cdot)$ is a mean-preserving spread of $F(\cdot)$.
(iii) Property (6.D.2) holds.

In Exercise 6.D.4, you are asked to verify the equivalence of these three properties in the probability simplex diagram.

6.E State-dependent Utility

In this section, we consider an extension of the analysis presented in the preceding two sections. In Sections 6.C and 6.D, we assumed that the decision maker cares solely about the distribution of monetary payoffs he receives. This says, in essence, that the underlying cause of the payoff is of no importance. If the cause is one's state of health, however, this assumption is unlikely to be fulfilled.[22] The distribution function of monetary payoffs is then not the appropriate object of individual choice. Here we consider the possibility that the decision maker may care not only about his monetary returns but also about the underlying events, or *states of nature*, that cause them.

We begin by discussing a convenient framework for modeling uncertain alternatives that, in contrast to the lottery apparatus, recognizes underlying states of nature. (We will encounter it repeatedly throughout the book, especially in Chapter 19.)

22. On the other hand, if it is an event such as the price of some security in a portfolio, the assumption is more likely to be a good representation of reality.