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, to the concepts 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*.

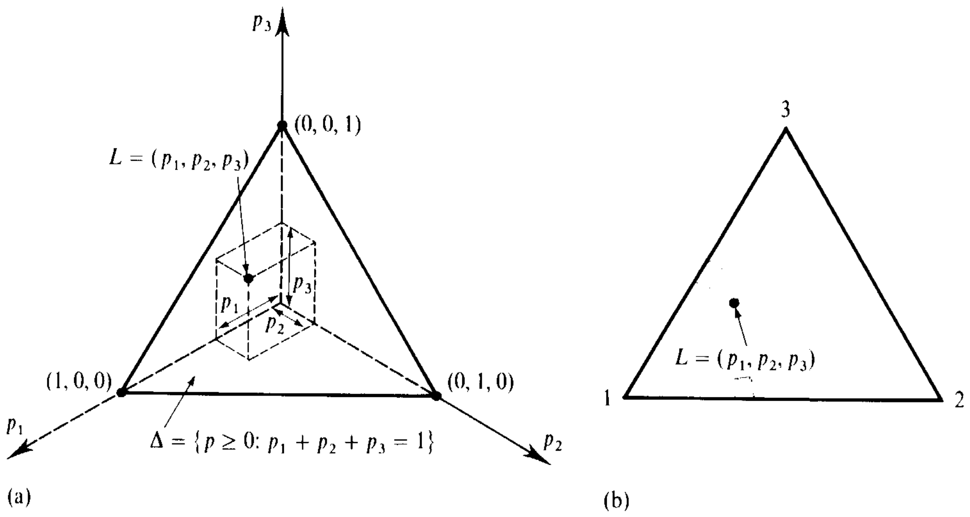

Figure 6.B.1
Representations of the simplex when $N = 3$.
(a) Three-dimensional representation.
(b) Two-dimensional representation.

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.

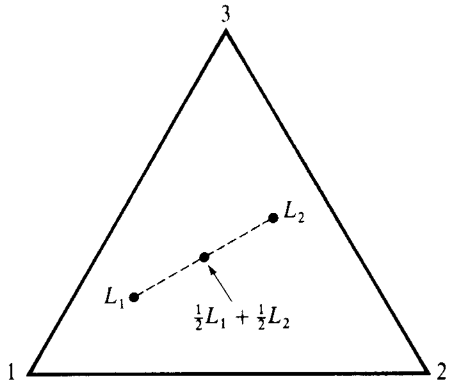

Figure 6.B.2
The reduced lottery of a compound lottery.

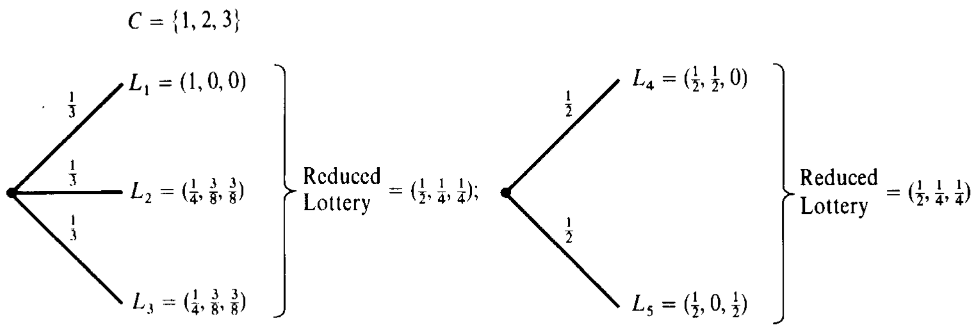

Figure 6.B.3
Two compound lotteries with the same reduced lottery.

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.

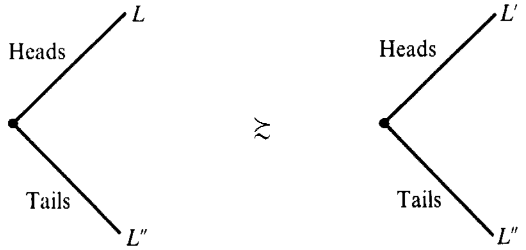

if and only if

$L \succsim L'$

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 $U: \mathcal{L} \to \mathbb{R}$ 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 $L = (p_1, \ldots, p_N) \in \mathcal{L}$ we have

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

A utility function $U: \mathcal{L} \to \mathbb{R}$ 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 $U: \mathcal{L} \to \mathbb{R}$ has an expected utility form if and only if it is *linear*, that is, if and only if it satisfies the property that

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

for any K lotteries $L_k \in \mathcal{L}$, $k = 1, \ldots, K$, 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 $U: \mathcal{L} \to \mathbb{R}$ is a v.N–M expected utility function for the preference relation $\succsim$ on $\mathcal{L}$. Then $\tilde{U}: \mathcal{L} \to \mathbb{R}$ 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 $L \in \mathcal{L}$.

Proof: Begin by choosing two lotteries $\bar{L}$ and $\underline{L}$ with the property that $\bar{L} \succsim L \succsim \underline{L}$ for all $L \in \mathcal{L}$.[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

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

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

Figure 6.B.5 Geometric explanation of the expected utility theorem. (a) $\succsim$ is representable by a utility function with the expected utility form. (b) Contradiction of the independence axiom. (c) Contradiction of the independence axiom.

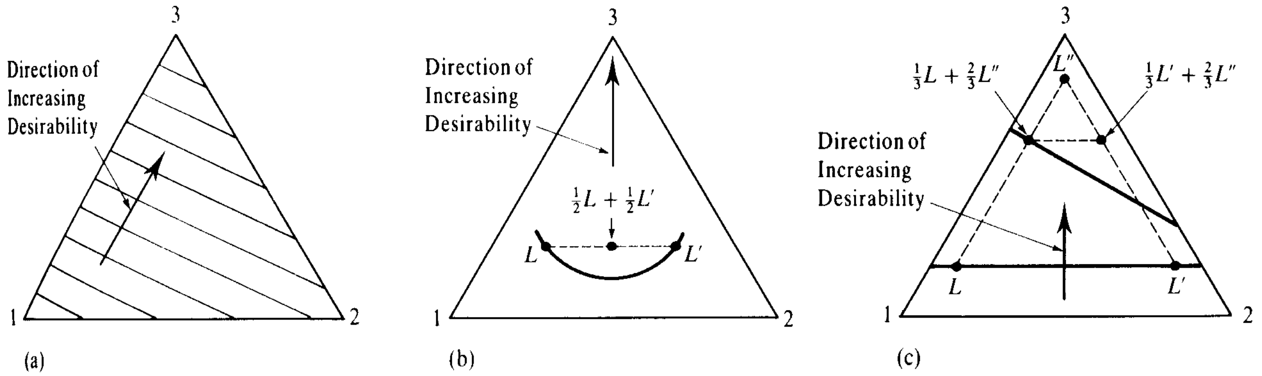

$\frac{1}{2}L' + \frac{1}{2}L \succ L$. This is equivalent to saying that

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

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

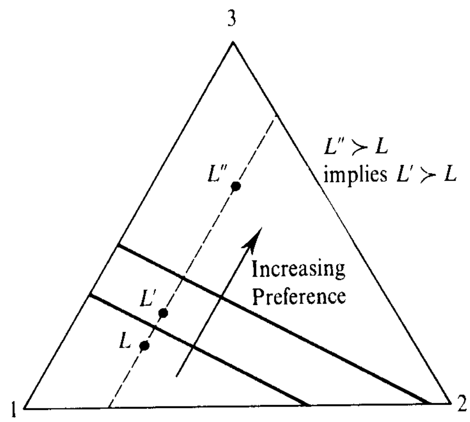

Figure 6.B.6
Expected utility as a guide to introspection.

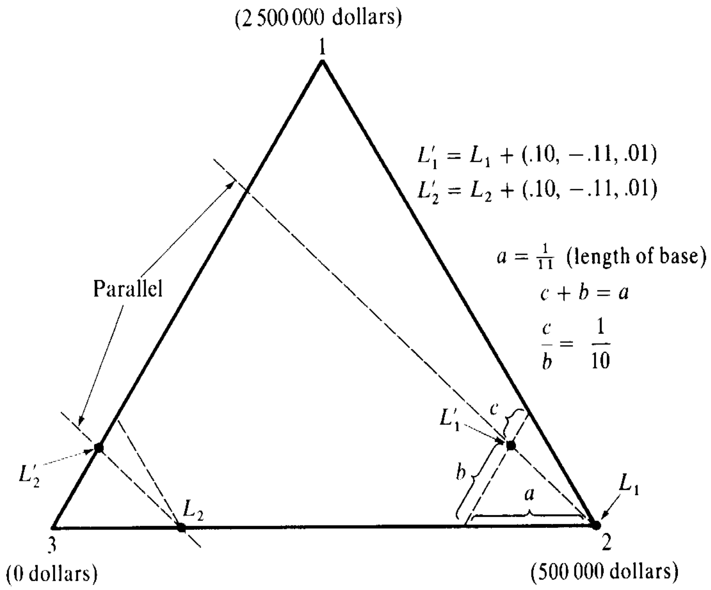

Figure 6.B.7
Depiction of the Allais paradox in the simplex.

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 $B = \{b_1, \ldots, b_N\}$ 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 $L = (p_1, \ldots, p_N)$ 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.

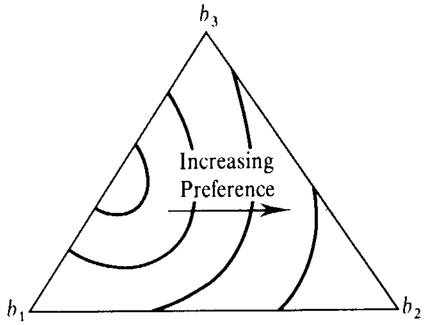

Figure 6.B.8
An indifference map for induced preferences over lotteries on $B = \{b_1, b_2, b_3\}$.

it is nonetheless always *convex*; that is,

$$U(\alpha L + (1 - \alpha)L') \leq \alpha U(L) + (1 - \alpha)U(L').$$

Figure 6.B.8 represents an indifference map for induced preferences in the probability simplex for a case where $N = 3$.

6.C Money Lotteries and Risk Aversion

In many economic settings, individuals seem to display aversion to risk. In this section, we formalize the notion of *risk aversion* and study some of its properties.

From this section through the end of the chapter, we concentrate on risky alternatives whose outcomes are amounts of money. It is convenient, however, when dealing with monetary outcomes, to treat money as a continuous variable. Strictly speaking, the derivation of the expected utility representation given in Section 6.B assumed a finite number of outcomes. However, the theory can be extended, with some minor technical complications, to the case of an infinite domain. We begin by briefly discussing this extension.

Lotteries over Monetary Outcomes and the Expected Utility Framework

Suppose that we denote amounts of money by the continuous variable x. We can describe a monetary lottery by means of a *cumulative distribution function* $F: \mathbb{R} \to [0, 1]$. That is, for any x, $F(x)$ is the probability that the realized payoff is less than or equal to x. Note that if the distribution function of a lottery has a density function $f(\cdot)$ associated with it, then $F(x) = \int_{-\infty}^{x} f(t)\, dt$ for all x. The advantage of a formalism based on distribution functions over one based on density functions, however, is that the first is completely general. It does not exclude a priori the possibility of a discrete set of outcomes. For example, the distribution function of a lottery with only three monetary outcomes receiving positive probability is illustrated in Figure 6.C.1.

Note that distribution functions preserve the linear structure of lotteries (as do density functions). For example, the final distribution of money, $F(\cdot)$, induced by a compound lottery $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$ is just the weighted average of the distributions induced by each of the lotteries that constitute it: $F(x) = \sum_k \alpha_k F_k(x)$, where $F_k(\cdot)$ is the distribution of the payoff under lottery L_k.

From this point on, we shall work with distribution functions to describe lotteries over monetary outcomes. We therefore take the lottery space $\mathscr{L}$ to be the *set of all*

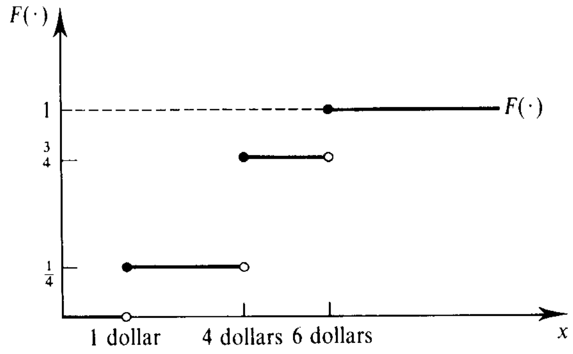

$$\left\{\begin{array}{l}\text{Prob}(1\text{ dollar}) \;\;= \frac{1}{4}\\ \text{Prob}(4\text{ dollars}) = \frac{1}{2}\\ \text{Prob}(6\text{ dollars}) = \frac{1}{4}\end{array}\right\} \rightarrow F(x) = \begin{cases} 0 & \text{if } x < 1 \\ \frac{1}{4} & \text{if } x \in [1, 4) \\ \frac{3}{4} & \text{if } x \in [4, 6) \\ 1 & \text{if } x \geq 6 \end{cases}$$

Figure 6.C.1
A distribution function.

distribution functions over nonnegative amounts of money, or, more generally, over an interval $[a, +\infty)$.

As in Section 6.B, we begin with a decision maker who has rational preferences $\succsim$ defined over $\mathscr{L}$. The application of the expected utility theorem to outcomes defined by a continuous variable tells us that under the assumptions of the theorem, there is an assignment of utility values $u(x)$ to nonnegative amounts of money with the property that any $F(\cdot)$ can be evaluated by a utility function $U(\cdot)$ of the form

$$U(F) = \int u(x)\, dF(x). \tag{6.C.1}$$

Expression (6.C.1) is the exact extension of the expected utility form to the current setting. The v.N–M utility function $U(\cdot)$ is the mathematical expectation, over the realizations of x, of the values $u(x)$. The latter takes the place of the values $(u_1, \ldots, u_N)$ used in the discrete treatment of Section 6.B.[11] Note that, as before, $U(\cdot)$ is linear in $F(\cdot)$.

The strength of the expected utility representation is that it preserves the very useful expectation form while making the utility of monetary lotteries sensitive not only to the mean but also to the higher moments of the distribution of the monetary payoffs. (See Exercise 6.C.2 for an illuminating quadratic example.)

It is important to distinguish between the utility function $U(\cdot)$, defined on lotteries, and the utility function $u(\cdot)$ defined on sure amounts of money. For this reason, we call $U(\cdot)$ the *von-Neumann–Morgenstern (v.N–M) expected utility function* and $u(\cdot)$ the *Bernoulli utility function.*[12]

11. Given a distribution function $F(x)$, the expected value of a function $\phi(x)$ is given by $\int \phi(x)\, dF(x)$. When $F(\cdot)$ has an associated density function $f(x)$, this expression is exactly equal to $\int \phi(x) f(x)\, dx$. Note also that for notational simplicity, we do not explicitly write the limits of integration when the integral is over the full range of possible realizations of x.

12. The terminology is not standardized. It is common to call $u(\cdot)$ the v.N–M utility function or the expected utility function. We prefer to have a name that is specific to the $u(\cdot)$ function, and so we call it the Bernoulli function for Daniel Bernoulli, who first used an instance of it.

Although the general axioms of Section 6.B yield the expected utility representation, they place no restrictions whatsoever on the Bernoulli utility function $u(\cdot)$. In large part, the analytical power of the expected utility formulation hinges on specifying the Bernoulli utility function $u(\cdot)$ in such a manner that it captures interesting economic attributes of choice behavior. At the simplest level, it makes sense in the current monetary context to postulate that $u(\cdot)$ is *increasing* and *continuous*.[13] We maintain both of these assumptions from now on.

Another restriction, based on a subtler argument, is the *boundedness* (above and below) of $u(\cdot)$. To argue the plausibility of boundedness above (a similar argument applies for boundedness below), we refer to the famous *St. Petersburg–Menger paradox*. Suppose that $u(\cdot)$ is unbounded, so that for every integer m there is an amount of money x_m with $u(x_m) > 2^m$. Consider the following lottery: we toss a coin repeatedly until tails comes up. If this happens in the mth toss, the lottery gives a monetary payoff of x_m. Since the probability of this outcome is $1/2^m$, the expected utility of this lottery is $\sum_{m=1}^{\infty} u(x_m)(1/2^m) \geq \sum_{m=1}^{\infty} (2^m)(1/2^m) = +\infty$. But this means that an individual should be willing to give up all his wealth for the opportunity to play this lottery, a patently absurd conclusion (how much would you pay?).[14]

The rest of this section concentrates on the important property of *risk aversion*, its formulation in terms of the Bernoulli utility function $u(\cdot)$, and its measurement.[15]

Risk Aversion and Its Measurement

The concept of risk aversion provides one of the central analytical techniques of economic analysis, and it is assumed in this book whenever we handle uncertain situations. We begin our discussion of risk aversion with a general definition that does not presume an expected utility formulation.

Definition 6.C.1: A decision maker is a *risk averter* (or exhibits *risk aversion*) if for any lottery $F(\cdot)$, the degenerate lottery that yields the amount $\int x\, dF(x)$ with certainty is at least as good as the lottery $F(\cdot)$ itself. If the decision maker is always [i.e., for any $F(\cdot)$] indifferent between these two lotteries, we say that he is *risk neutral*. Finally, we say that he is *strictly risk averse* if indifference holds only when the two lotteries are the same [i.e., when $F(\cdot)$ is degenerate].

If preferences admit an expected utility representation with Bernoulli utility function $u(x)$, it follows directly from the definition of risk aversion that the decision maker is risk averse if and only if

$$\int u(x)\, dF(x) \leq u\left(\int x\, dF(x)\right) \quad \text{for all } F(\cdot). \tag{6.C.2}$$

Inequality (6.C.2) is called *Jensen's inequality*, and it is the defining property of a concave function (see Section M.C of the Mathematical Appendix). Hence, in the

13. In applications, an exception to continuity is sometimes made at $x = 0$ by setting $u(0) = -\infty$.

14. In practice, most utility functions commonly used are not bounded. Paradoxes are avoided because the class of distributions allowed by the modeler in each particular application is a limited one. Note also that if we insisted on $u(\cdot)$ being defined on $(-\infty, \infty)$ then any nonconstant $u(\cdot)$ could not be both concave and bounded (above and below).

15. Arrow (1971) and Pratt (1964) are the classical references in this area.

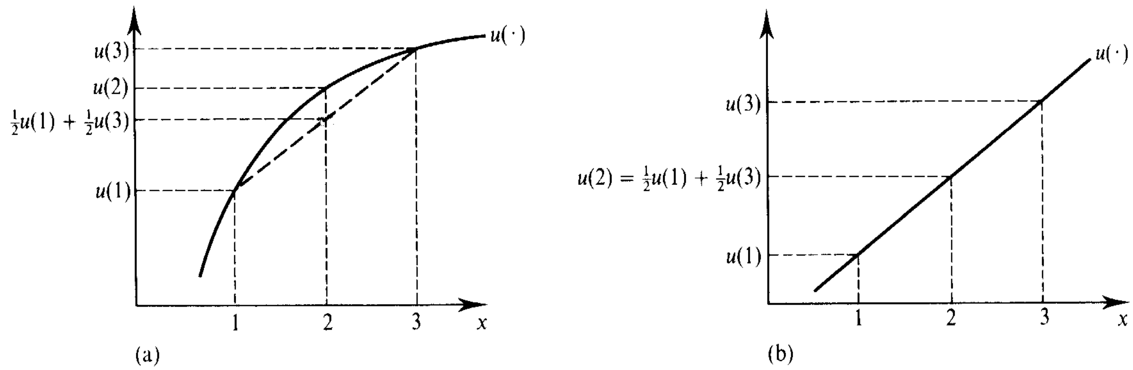

Figure 6.C.2 Risk aversion (a) and risk neutrality (b).

context of expected utility theory, we see that *risk aversion is equivalent to the concavity of* $u(\cdot)$ and that strict risk aversion is equivalent to the strict concavity of $u(\cdot)$. This makes sense. Strict concavity means that the marginal utility of money is decreasing. Hence, at any level of wealth x, the utility gain from an extra dollar is smaller than (the absolute value of) the utility loss of having a dollar less. It follows that a risk of gaining or losing a dollar with even probability is not worth taking. This is illustrated in Figure 6.C.2(a); in the figure we consider a gamble involving the gain or loss of 1 dollar from an initial position of 2 dollars. The (v.N–M)utility of this gamble, $\frac{1}{2}u(1) + \frac{1}{2}u(3)$, is strictly less than that of the initial certain position $u(2)$.

For a risk-neutral expected utility maximizer, (6.C.2) must hold with *equality* for all $F(\cdot)$. Hence, the decision maker is risk neutral if and only if the Bernoulli utility function of money $u(\cdot)$ is linear. Figure 6.C.2(b) depicts the (v.N–M) utility associated with the previous gamble for a risk neutral individual. Here the individual is indifferent between the gambles that yield a mean wealth level of 2 dollars and a certain wealth of 2 dollars. Definition 6.C.2 introduces two useful concepts for the analysis of risk aversion.

Definition 6.C.2: Given a Bernoulli utility function $u(\cdot)$ we define the following concepts:

(i) The *certainty equivalent of* $F(\cdot)$, denoted $c(F, u)$, is the amount of money for which the individual is indifferent between the gamble $F(\cdot)$ and the certain amount $c(F, u)$; that is,

$$u(c(F, u)) = \int u(x)\, dF(x). \qquad (6.C.3)$$

(ii) For any fixed amount of money x and positive number ε, the *probability premium* denoted by $\pi(x, \varepsilon, u)$, is the excess in winning probability over fair odds that makes the individual indifferent between the certain outcome x and a gamble between the two outcomes $x + \varepsilon$ and $x - \varepsilon$. That is

$$u(x) = (\tfrac{1}{2} + \pi(x, \varepsilon, u))u(x + \varepsilon) + (\tfrac{1}{2} - \pi(x, \varepsilon, u))u(x - \varepsilon). \qquad (6.C.4)$$

These two concepts are illustrated in Figure 6.C.3. In Figure 6.C.3(a), we exhibit the geometric construction of $c(F, u)$ for an even probability gamble between 1 and 3 dollars. Note that $c(F, u) < 2$, implying that some expected return is traded for certainty. The satisfaction of the inequality $c(F, u) \le \int x\, dF(x)$ for all $F(\cdot)$ is, in fact,

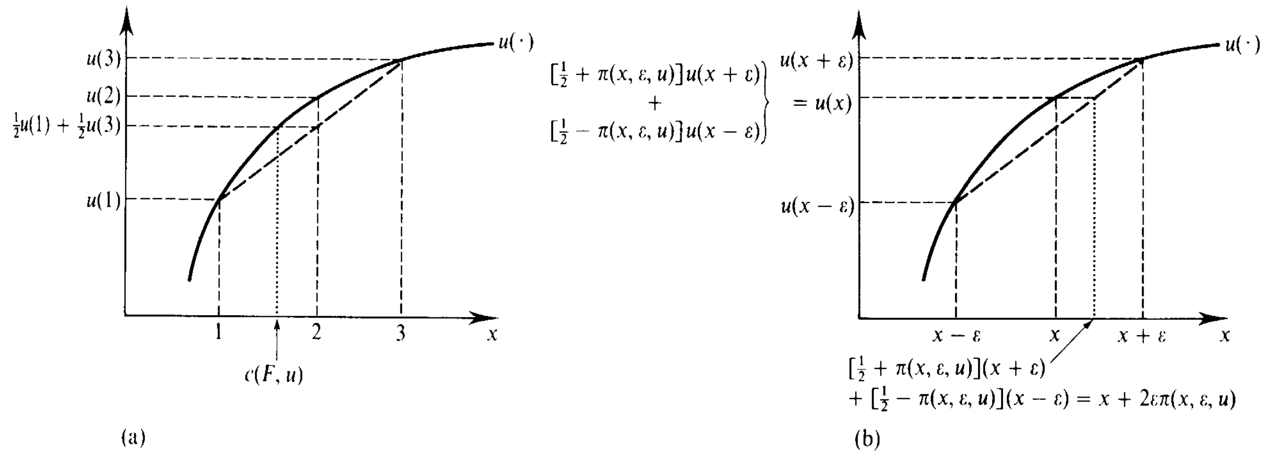

Figure 6.C.3 The certainty equivalent (a) and the probability premium (b).

equivalent to the decision maker being a risk averter. To see this, observe that since $u(\cdot)$ is nondecreasing, we have

$$c(F, u) \le \int x\, dF(x) \Leftrightarrow u(c(F, u)) \le u\left(\int x\, dF(x)\right) \Leftrightarrow \int u(x)\, dF(x) \le u\left(\int x\, dF(x)\right),$$

where the last $\Leftrightarrow$ follows from the definition of $c(F, u)$.

In Figure 6.C.3(b), we exhibit the geometric construction of $\pi(x, \varepsilon, u)$. We see that $\pi(x, \varepsilon, u) > 0$; that is, better than fair odds must be given for the individual to accept the risk. In fact, the satisfaction of the inequality $\pi(x, \varepsilon, u) \ge 0$ for all x and $\varepsilon > 0$ is also equivalent to risk aversion (see Exercise 6.C.3).

These points are formally summarized in Proposition 6.C.1.

Proposition 6.C.1: Suppose a decision maker is an expected utility maximizer with a Bernoulli utility function $u(\cdot)$ on amounts of money. Then the following properties are equivalent:

(i) The decision maker is risk averse.
(ii) $u(\cdot)$ is concave.[16]
(iii) $c(F, u) \le \int x\, dF(x)$ for all $F(\cdot)$.
(iv) $\pi(x, \varepsilon, u) \ge 0$ for all x, ε.

Examples 6.C.1 to 6.C.3 illustrate the use of the risk aversion concept.

Example 6.C.1: *Insurance.* Consider a strictly risk-averse decision maker who has an initial wealth of w but who runs a risk of a loss of D dollars. The probability of the loss is π. It is possible, however, for the decision maker to buy insurance. One unit of insurance costs q dollars and pays 1 dollar if the loss occurs. Thus, if α units of insurance are bought, the wealth of the individual will be $w - \alpha q$ if there is no loss and $w - \alpha q - D + \alpha$ if the loss occurs. Note, for purposes of later discussion, that the decision maker's expected wealth is then $w - \pi D + \alpha(\pi - q)$. The decision maker's problem is to choose the optimal level of α. His utility maximization problem is

16. Recall that if $u(\cdot)$ is twice differentiable then concavity is equivalent to $u''(x) \le 0$ for all x.

therefore

$$\underset{\alpha \geq 0}{\text{Max}} \ (1-\pi)u(w-\alpha q) + \pi u(w - \alpha q - D + \alpha).$$

If α^* is an optimum, it must satisfy the first-order condition:

$$-q(1-\pi)u'(w-\alpha^* q) + \pi(1-q)u'(w - D + \alpha^*(1-q)) \leq 0,$$

with equality if $\alpha^* > 0$.

Suppose now that the price q of one unit of insurance is *actuarially fair* in the sense of it being equal to the expected cost of insurance. That is, $q = \pi$. Then the first-order condition requires that

$$u'(w - D + \alpha^*(1-\pi)) - u'(w - \alpha^*\pi) \leq 0,$$

with equality if $\alpha^* > 0$.

Since $u'(w - D) > u'(w)$, we must have $\alpha^* > 0$, and therefore

$$u'(w - D + \alpha^*(1-\pi)) = u'(w - \alpha^*\pi).$$

Because $u'(\cdot)$ is strictly decreasing, this implies

$$w - D + \alpha^*(1-\pi) = w - \alpha^*\pi,$$

or, equivalently,

$$\alpha^* = D.$$

Thus, *if insurance is actuarially fair, the decision maker insures completely.* The individual's final wealth is then $w - \pi D$, regardless of the occurrence of the loss.

This proof of the complete insurance result uses first-order conditions, which is instructive but not really necessary. Note that if $q = \pi$, then the decision maker's expected wealth is $w - \pi D$ for any α. Since setting $\alpha = D$ allows him to reach $w - \pi D$ with certainty, the definition of risk aversion directly implies that this is the optimal level of α. ■

Example 6.C.2: *Demand for a Risky Asset.* An *asset* is a divisible claim to a financial return in the future. Suppose that there are two assets, a safe asset with a return of 1 dollar per dollar invested and a risky asset with a random return of z dollars per dollar invested. The random return z has a distribution function $F(z)$ that we assume satisfies $\int z \, dF(z) > 1$; that is, its mean return exceeds that of the safe asset.

An individual has initial wealth w to invest, which can be divided in any way between the two assets. Let α and β denote the amounts of wealth invested in the risky and the safe asset, respectively. Thus, for any realization z of the random return, the individual's *portfolio* (α, β) pays $\alpha z + \beta$. Of course, we must also have $\alpha + \beta = w$.

The question is how to choose α and β. The answer will depend on $F(\cdot)$, w, and the Bernoulli utility function $u(\cdot)$. The utility maximization problem of the individual is

$$\underset{\alpha, \beta \geq 0}{\text{Max}} \int u(\alpha z + \beta) \, dF(z)$$

$$\text{s.t. } \alpha + \beta = w.$$

Equivalently, we want to maximize $\int u(w + \alpha(z-1)) \, dF(z)$ subject to $0 \leq \alpha \leq w$. If

α^* is optimal, it must satisfy the Kuhn–Tucker first-order conditions:[17]

$$\phi(\alpha^*) = \int u'(w + \alpha^*[z - 1])(z - 1)\, dF(z) \begin{cases} \leq 0 & \text{if } \alpha^* < w, \\ \geq 0 & \text{if } \alpha^* > 0. \end{cases}$$

Note that $\int z\, dF(z) > 1$ implies $\phi(0) > 0$. Hence, $\alpha^* = 0$ cannot satisfy this first-order condition. We conclude that the optimal portfolio has $\alpha^* > 0$. The general principle illustrated in this example, is that *if a risk is actuarially favorable, then a risk averter will always accept at least a small amount of it.*

This same principle emerges in Example 6.C.1 if insurance is not actuarially fair. In Exercise 6.C.1, you are asked to show that if $q > \pi$, then the decision maker will not fully insure (i.e., will accept some risk). ■

Example 6.C.3: *General Asset Problem.* In the previous example, we could define the utility $U(\alpha, \beta)$ of the portfolio (α, β) as $U(\alpha, \beta) = \int u(\alpha z + \beta)\, dF(z)$. Note that $U(\cdot)$ is then an increasing, continuous, and concave utility function. We now discuss an important generalization. We assume that we have N assets (one of which may be the safe asset) with asset n giving a return of z_n per unit of money invested. These returns are jointly distributed according to a distribution function $F(z_1, \ldots, z_N)$. The utility of holding a *portfolio* of assets $(\alpha_1, \ldots, \alpha_N)$ is then

$$U(\alpha_1, \ldots, \alpha_N) = \int u(\alpha_1 z_1 + \cdots + \alpha_N z_N)\, dF(z_1, \ldots, z_N).$$

This utility function for portfolios, defined on $\mathbb{R}^N_+$, is also increasing, continuous, and concave (see Exercise 6.C.4). This means that, formally, we can treat assets as the usual type of commodities and apply to them the demand theory developed in Chapters 2 and 3. Observe, in particular, how risk aversion leads to a convex indifference map for portfolios. ■

Suppose that the lotteries pay in vectors of physical goods rather than in money. Formally, the space of outcomes is then the consumption set $\mathbb{R}^L_+$ (all the previous discussion can be viewed as the special case in which there is a single good). In this more general setting, the concept of risk aversion given by Definition 6.C.1 is perfectly well defined. Furthermore, if there is a Bernoulli utility function $u: \mathbb{R}^L_+ \to \mathbb{R}$, then risk aversion is still equivalent to the concavity of $u(\cdot)$. Hence, we have here another justification for the convexity assumption of Chapter 3: Under the assumptions of the expected utility theorem, the convexity of preferences for perfectly certain amounts of the physical commodities must hold if for any lottery with commodity payoffs the individual always prefers the certainty of the mean commodity bundle to the lottery itself.

In Exercise 6.C.5, you are asked to show that if preferences over lotteries with commodity payoffs exhibit risk aversion, then, at given commodity prices, the induced preferences on money lotteries (where consumption decisions are made after the realization of wealth) are also risk averse. Thus, in principle, it is possible to build the theory of risk aversion on the more primitive notion of lotteries over the final consumption of goods.

---

17. The objective function is concave in α because the concavity of $u(\cdot)$ implies that $\int u''(w + \alpha(z - 1))(z - 1)^2\, dF(x) \leq 0$.

The Measurement of Risk Aversion

Now that we know what it means to be risk averse, we can try to measure the extent of risk aversion. We begin by defining one particularly useful measure and discussing some of its properties.

Definition 6.C.3: Given a (twice-differentiable) Bernoulli utility function $u(\cdot)$ for money, the *Arrow–Pratt coefficient of absolute risk aversion* at x is defined as $r_A(x) = -u''(x)/u'(x)$.

The Arrow–Pratt measure can be motivated as follows: We know that risk neutrality is equivalent to the linearity of $u(\cdot)$, that is, to $u''(x) = 0$ for all x. Therefore, it seems logical that the degree of risk aversion be related to the *curvature* of $u(\cdot)$. In Figure 6.C.4, for example, we represent two Bernoulli utility functions $u_1(\cdot)$ and $u_2(\cdot)$ normalized (by choice of origin and units) to have the same utility and marginal utility values at wealth level x. The certainty equivalent for a small risk with mean x is smaller for $u_2(\cdot)$ than for $u_1(\cdot)$, suggesting that risk aversion increases with the curvature of the Bernoulli utility function at x. One possible measure of curvature of the Bernoulli utility function $u(\cdot)$ at x is $u''(x)$. However, this is not an adequate measure because it is not invariant to positive linear transformations of the utility function. To make it invariant, the simplest modification is to use $u''(x)/u'(x)$. If we change sign so as to have a positive number for an increasing and concave $u(\cdot)$, we get the Arrow–Pratt measure.

A more precise motivation for $r_A(x)$ as a measure of the degree of risk aversion can be obtained by considering a fixed wealth x and studying the behavior of the probability premium $\pi(x, \varepsilon, u)$ as $\varepsilon \to 0$ [for simplicity, we write it as $\pi(\varepsilon)$]. Differentiating the identity (6.C.4) that defines $\pi(\cdot)$ twice with respect to ε (assume that $\pi(\cdot)$ is differentiable), and evaluating at $\varepsilon = 0$, we get $4\pi'(0)u'(x) + u''(x) = 0$. Hence

$$r_A(x) = 4\pi'(0).$$

Thus, $r_A(x)$ measures the rate at which the probability premium increases at certainty with the small risk measured by ε.[18] As we go along, we will find additional related interpretations of the Arrow–Pratt measure.

18. For a similar derivation relating $r_A(\cdot)$ to the rate of change of the certainty equivalent with respect to a small increase in a small risk around certainty, see Exercise 6.C.20.

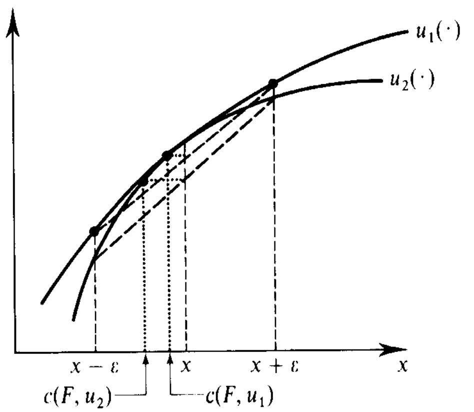

Figure 6.C.4
Differing degrees of risk aversion.

Note that, up to two integration constants, the utility function $u(\cdot)$ can be recovered from $r_A(\cdot)$ by integrating twice. The integration constants are irrelevant because the Bernoulli utility is identified only up to two constants (origin and units). Thus, the Arrow–Pratt risk aversion measure $r_A(\cdot)$ fully characterizes behavior under uncertainty.

Example 6.C.4: Consider the utility function $u(x) = -e^{-ax}$ for $a > 0$. Then $u'(x) = ae^{-ax}$ and $u''(x) = -a^2e^{-ax}$. Therefore, $r_A(x, u) = a$ for all x. It follows from the observation just made that the general form of a Bernoulli utility function with an Arrow–Pratt measure of absolute risk aversion equal to the constant $a > 0$ at all x is $u(x) = -\alpha e^{-ax} + \beta$ for some $\alpha > 0$ and β. ■

Once we are equipped with a measure of risk aversion, we can put it to use in comparative statics exercises. Two common situations are the comparisons of risk attitudes across individuals with different utility functions and the comparison of risk attitudes for one individual at different levels of wealth.

Comparisons across individuals

Given two Bernoulli utility functions $u_1(\cdot)$ and $u_2(\cdot)$, when can we say that $u_2(\cdot)$ is unambiguously *more risk averse than* $u_1(\cdot)$? Several possible approaches to a definition seem plausible:

(i) $r_A(x, u_2) \geq r_A(x, u_1)$ for every x.

(ii) There exists an increasing concave function $\psi(\cdot)$ such that $u_2(x) = \psi(u_1(x))$ at all x; that is, $u_2(\cdot)$ is a concave transformation of $u_1(\cdot)$. [In other words, $u_2(\cdot)$ is "more concave" than $u_1(\cdot)$.]

(iii) $c(F, u_2) \leq c(F, u_1)$ for any $F(\cdot)$.

(iv) $\pi(x, \varepsilon, u_2) \geq \pi(x, \varepsilon, u_1)$ for any x and ε.

(v) Whenever $u_2(\cdot)$ finds a lottery $F(\cdot)$ at least as good as a riskless outcome $\bar{x}$, then $u_1(\cdot)$ also finds $F(\cdot)$ at least as good as $\bar{x}$. That is, $\int u_2(x)\,dF(x) \geq u_2(\bar{x})$ implies $\int u_1(x)\,dF(x) \geq u_1(\bar{x})$ for any $F(\cdot)$ and $\bar{x}$.[19]

In fact, these five definitions are equivalent.

Proposition 6.C.2: Definitions (i) to (v) of the *more-risk-averse-than* relation are equivalent.

Proof: We will not give a complete proof. (You are asked to establish some of the implications in Exercises 6.C.6 and 6.C.7.) Here we will show the equivalence of (i) and (ii) under differentiability assumptions.

Note, first that we always have $u_2(x) = \psi(u_1(x))$ for some increasing function $\psi(\cdot)$; this is true simply because $u_1(\cdot)$ and $u_2(\cdot)$ are ordinally identical (more money is preferred to less). Differentiating, we get

$$u_2'(x) = \psi'(u_1(x))u_1'(x)$$

and

$$u_2''(x) = \psi'(u_1(x))u_1''(x) + \psi''(u_1(x))(u_1'(x))^2.$$

Dividing both sides of the second expression by $u_2'(x) > 0$, and using the first

19. In other words, any risk that $u_2(\cdot)$ would accept starting from a position of certainty would also be accepted by $u_1(\cdot)$.

expression, we get

$$r_A(x, u_2) = r_A(x, u_1) - \frac{\psi''(u_1(x))}{\psi'(u_1(x))} u_1'(x).$$

Thus, $r_A(x, u_2) \geq r_A(x, u_1)$ for all x if and only if $\psi''(u_1) \leq 0$ for all u_1 in the range of $u_1(\cdot)$. ■

The more-risk-averse-than relation is a *partial ordering* of Bernoulli utility functions; it is transitive but far from complete. Typically, two Bernoulli utility functions $u_1(\cdot)$ and $u_2(\cdot)$ will not be comparable; that is, we will have $r_A(x, u_1) > r_A(x, u_2)$ at some x but $r_A(x', u_1) < r_A(x', u_2)$ at some other $x' \neq x$.

Example 6.C.2 continued: We take up again the asset portfolio problem between a safe and a risky asset discussed in Example 6.C.2. Suppose that we now have two individuals with Bernoulli utility functions $u_1(\cdot)$ and $u_2(\cdot)$, and denote by α_1^* and α_2^* their respective optimal investments in the risky asset. We will show that *if* $u_2(\cdot)$ *is more risk averse than* $u_1(\cdot)$, *then* $\alpha_2^* < \alpha_1^*$; that is, the second decision maker invests less in the risky asset than the first.

To repeat from our earlier discussion, the asset allocation problem for $u_1(\cdot)$ is

$$\underset{0 \leq \alpha \leq w}{\text{Max}} \int u_1(w - \alpha + \alpha z)\, dF(z).$$

Assuming an interior solution, the first-order condition is

$$\int (z - 1) u_1'(w + \alpha_1^*[z - 1])\, dF(z) = 0. \tag{6.C.5}$$

The analogous expression for the utility function $u_2(\cdot)$ is

$$\phi_2(\alpha_2^*) = \int (z - 1) u_2'(w + \alpha_2^*[z - 1])\, dF(z) = 0. \tag{6.C.6}$$

As we know, the concavity of $u_2(\cdot)$ implies that $\phi_2(\cdot)$ is decreasing. Therefore, if we show that $\phi_2(\alpha_1^*) < 0$, it must follow that $\alpha_2^* < \alpha_1^*$, which is the result we want. Now, $u_2(x) = \psi(u_1(x))$ allows us to write

$$\phi_2(\alpha_1^*) = \int (z - 1)\psi'(u_1(w + \alpha_1^*[z - 1])) u_1'(w + \alpha_1^*[z - 1])\, dF(z) < 0. \tag{6.C.7}$$

To understand the final inequality, note that the integrand of expression (6.C.7) is the same as that in (6.C.5) except that it is multiplied by $\psi'(\cdot)$, a positive decreasing function of z [recall that $u_2(\cdot)$ more risk averse than $u_1(\cdot)$ means that the increasing function $\psi(\cdot)$ is concave; that is, $\psi'(\cdot)$ is positive and decreasing]. Hence, the integral (6.C.7) underweights the positive values of $(z - 1)u_1'(w + \alpha_1^*[z - 1])$, which obtain for $z > 1$, relative to the negative values, which obtain for $z < 1$. Since, in (6.C.5), the integral of the positive and the negative parts of the integrand added to zero, they now must add to a negative number. This establishes the desired inequality. ■

Comparisons across wealth levels

It is a common contention that wealthier people are willing to bear more risk than poorer people. Although this might be due to differences in utility functions across people, it is more likely that the source of the difference lies in the possibility that richer people "can afford to take a chance." Hence, we shall explore the implications of the condition stated in Definition 6.C.4.

Definition 6.C.4: The Bernoulli utility function $u(\cdot)$ for money exhibits *decreasing absolute risk aversion* if $r_A(x, u)$ is a decreasing function of x.

Individuals whose preferences satisfy the decreasing absolute risk aversion property take more risk as they become wealthier. Consider two levels of initial wealth $x_1 > x_2$. Denote the increments or decrements to wealth by z. Then the individual evaluates risk at x_1 and x_2 by, respectively, the induced Bernoulli utility functions $u_1(z) = u(x_1 + z)$ and $u_2(z) = u(x_2 + z)$. Comparing an individual's attitudes toward risk as his level of wealth changes is like comparing the utility functions $u_1(\cdot)$ and $u_2(\cdot)$, a problem we have just studied. If $u(\cdot)$ displays decreasing absolute risk aversion, then $r_A(z, u_2) \geq r_A(z, u_1)$ for all z. This is condition (i) of Proposition 6.C.2. Hence, the result in Proposition 6.C.3 follows directly from Proposition 6.C.2.

Proposition 6.C.3: The following properties are equivalent:

(i) The Bernoulli utility function $(\cdot)$ exhibits decreasing absolute risk aversion.

(ii) Whenever $x_2 < x_1$, $u_2(z) = u(x_2 + z)$ is a concave transformation of $u_1(z) = u(x_1 + z)$.

(iii) For any risk $F(z)$, the certainty equivalent of the lottery formed by adding risk z to wealth level x, given by the amount c_x at which $u(c_x) = \int u(x + z)\, dF(z)$, is such that $(x - c_x)$ is decreasing in x. That is, the higher x is, the less is the individual willing to pay to get rid of the risk.

(iv) The probability premium $\pi(x, \varepsilon, u)$ is decreasing in x.

(v) For any $F(z)$, if $\int u(x_2 + z)\, dF(z) \geq u(x_2)$ and $x_2 < x_1$, then $\int u(x_1 + z)\, dF(z) \geq u(x_1)$.

Exercise 6.C.8: Assume that the Bernoulli utility function $u(\cdot)$ exhibits decreasing absolute risk aversion. Show that for the asset demand model of Example 6.C.2 (and Example 6.C.2 continued), the optimal allocation between the safe and the risky assets places an increasing amount of wealth in the risky asset as w rises (i.e., the risky asset is a normal good).

The assumption of decreasing absolute risk aversion yields many other economically reasonable results concerning risk-bearing behavior. However, in applications, it is often too weak and, because of its analytical convenience, it is sometimes complemented by a stronger assumption: *nonincreasing relative risk aversion.*

To understand the concept of relative risk aversion, note that the concept of absolute risk aversion is suited to the comparison of attitudes toward risky projects whose outcomes are *absolute gains or losses* from current wealth. But it is also of interest to evaluate risky projects whose outcomes are *percentage* gains or losses of current wealth. The concept of relative risk aversion does just this.

Let $t > 0$ stand for *proportional* increments or decrements of wealth. Then, an individual with Bernoulli utility function $u(\cdot)$ and initial wealth x can evaluate a random percentage risk by means of the utility function $\tilde{u}(t) = u(tx)$. The initial wealth position corresponds to $t = 1$. We already know that for a small risk around $t = 1$, the degree of risk aversion is well captured by $\tilde{u}''(1)/\tilde{u}'(1)$. Noting that $\tilde{u}''(1)/\tilde{u}'(1) = xu''(x)/u'(x)$, we are led to the concept stated in Definition 6.C.5.

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 $u: \mathbb{R} \to \mathbb{R}$, 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).

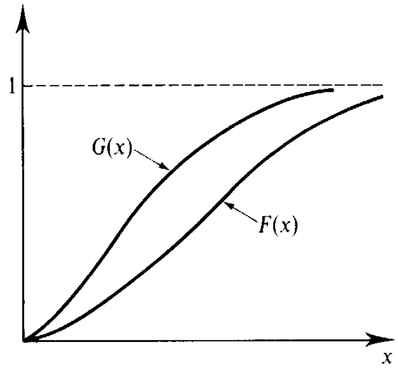

Figure 6.D.1
$F(\cdot)$ first-order stochastically dominates $G(\cdot)$.

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: $\mathbb{R} \to \mathbb{R}$, 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

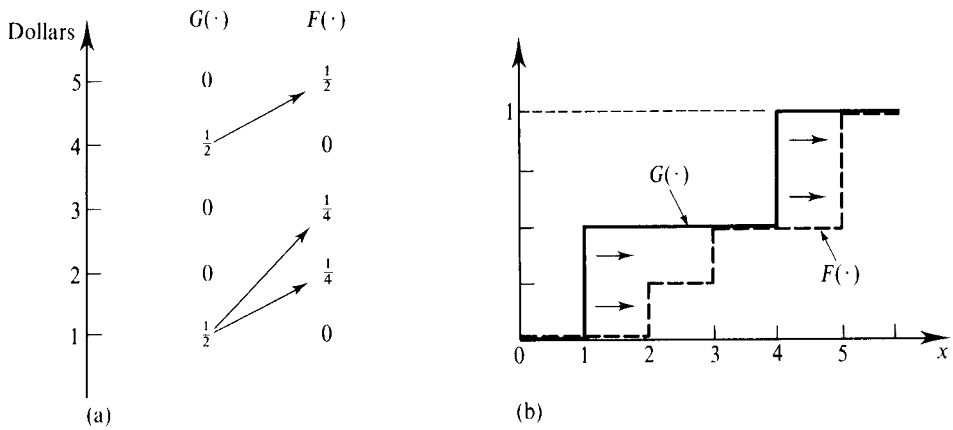

Figure 6.D.2
$F(\cdot)$ first-order stochastically dominates $G(\cdot)$.

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 $u\colon \mathbb{R}_+ \to \mathbb{R}$, 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

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

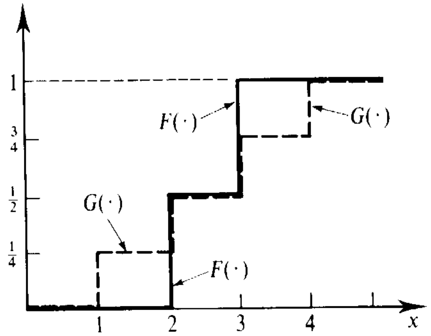

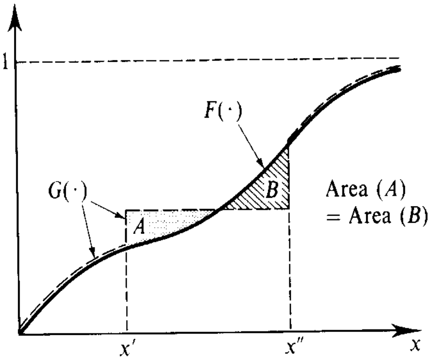

Figure 6.D.3 (left)
$G(\cdot)$ is a mean-preserving spread of $F(\cdot)$.

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).

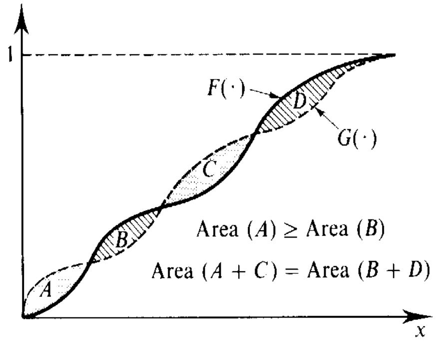

Figure 6.D.5
$F(\cdot)$ second-order stochastically dominates $G(\cdot)$.

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.

State-of-Nature Representations of Uncertainty

In Sections 6.C and 6.D, we modeled a risky alternative by means of a distribution function over monetary outcomes. Often, however, we know that the random outcome is generated by some underlying causes. A more detailed description of uncertain alternatives is then possible. For example, the monetary payoff of an insurance policy might depend on whether or not a certain accident has happened, the payoff on a corporate stock on whether the economy is in a recession, and the payoff of a casino gamble on the number selected by the roulette wheel.

We call these underlying causes *states*, or *states of nature*. We denote the set of states by S and an individual state by $s \in S$. For simplicity, we assume here that the set of states is finite and that each state s has a well-defined, objective probability $\pi_s > 0$ that it occurs. We abuse notation slightly by also denoting the total number of states by S.

An uncertain alternative with (nonnegative) monetary returns can then be described as a function that maps realizations of the underlying state of nature into the set of possible money payoffs $\mathbb{R}_+$. Formally, such a function is known as a *random variable*.

Definition 6.E.1: A *random variable* is a function $g: S \to \mathbb{R}_+$ that maps states into monetary outcomes.[23]

Every random variable $g(\cdot)$ gives rise to a money lottery describable by the distribution function $F(\cdot)$ with $F(x) = \sum_{\{s: g(s) \leq x\}} \pi_s$ for all x. Note that there is a loss in information in going from the random variable representation of uncertainty to the lottery representation; we do not keep track of which states give rise to a given monetary outcome, and only the aggregate probability of every monetary outcome is retained.

Because we take S to be finite, we can represent a random variable with monetary payoffs by the vector $(x_1, \ldots, x_S)$, where x_s is the nonnegative monetary payoff in state s. The set of all nonnegative random variables is then $\mathbb{R}^S_+$.

State-Dependent Preferences and the Extended Expected Utility Representation

The primitive datum of our theory is now a rational preference relation $\succsim$ on the set $\mathbb{R}^S_+$ of nonnegative random variables. Note that this formal setting is parallel to the one developed in Chapters 2 to 4 for consumer choice. The similarity is not merely superficial. If we define commodity s as the random variable that pays one dollar if and only if state s occurs (this is called a *contingent commodity* in Chapter 19), then the set of nonnegative random variables $\mathbb{R}^S_+$ is precisely the set of nonnegative bundles of these S contingent commodities.

As we shall see, it is very convenient if, in the spirit of the previous sections of this chapter, we can represent the individual's preferences over monetary outcomes by a utility function that possesses an *extended expected utility form.*

23. For concreteness, we restrict the outcomes to be nonnegative amounts of money. As we did in Section 6.B, we could equally well use an abstract outcome set C instead of $\mathbb{R}_+$.

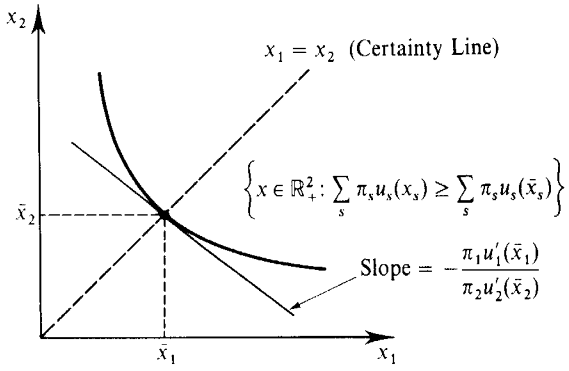

Figure 6.E.1
State-dependent preferences.

Definition 6.E.2: The preference relation $\succsim$ has an ***extended expected utility representation*** if for every $s \in S$, there is a function $u_s: \mathbb{R}_+ \to \mathbb{R}$ such that for any $(x_1, \ldots, x_S) \in \mathbb{R}^S_+$ and $(x_1', \ldots, x_S') \in \mathbb{R}^S_+$,

$$(x_1, \ldots, x_S) \succsim (x_1', \ldots, x_S') \quad \text{if and only if} \quad \sum_s \pi_s u_s(x_s) \geq \sum_s \pi_s u_s(x_s').$$

To understand Definition 6.E.2, recall the analysis in Section 6.B. If only the distribution of money payoffs mattered, and if preferences on money distributions satisfied the expected utility axioms, then the expected utility theorem leads to a *state-independent* (we will also say *state-uniform*) expected utility representation $\sum_s \pi_s u(x_s)$, where $u(\cdot)$ is the Bernoulli utility function on amounts of money.[24] The generalization in Definition 6.E.2 allows for a different function $u_s(\cdot)$ in every state.

Before discussing the conditions under which an extended utility representation exists, we comment on its usefulness as a tool in the analysis of choice under uncertainty. This usefulness is primarily a result of the behavior of the indifference sets around the *money certainty line*, the set of random variables that pay the same amount in every state. Figure 6.E.1 depicts state-dependent preferences in the space $\mathbb{R}^S_+$ for a case where $S = 2$ and the $u_s(\cdot)$ functions are concave (as we shall see later, concavity of these functions follows from risk aversion considerations). The certainty line in Figure 6.E.1 is the set of points with $x_1 = x_2$. The marginal rate of substitution at a point $(\bar{x}, \bar{x})$ is $\pi_1 u_1'(\bar{x})/\pi_2 u_2'(\bar{x})$. Thus, the slope of the indifference curves on the certainty line reflects the nature of state dependence as well as the probabilities of the different states. In contrast, with state-uniform (i.e., identical across states) utility functions, the marginal rate of substitution at any point on the certainty line equals the ratio of the probabilities of the states (implying that this slope is the same at all points on the certainty line).

Example 6.E.1: *Insurance with State-dependent Utility.* One interesting implication of state dependency arises when actuarially fair insurance is available. Suppose there are two states: State 1 is the state where no loss occurs, and state 2 is the state where a loss occurs. (This economic situation parallels that in Example 6.C.1.) The individual's initial situation (i.e., in the absence of any insurance purchase) is a

24. Note that the random variable $(x_1, \ldots, x_S)$ induces a money lottery that pays x_s with probability π_s. Hence, $\sum_s \pi_s u(x_s)$ is its expected utility.

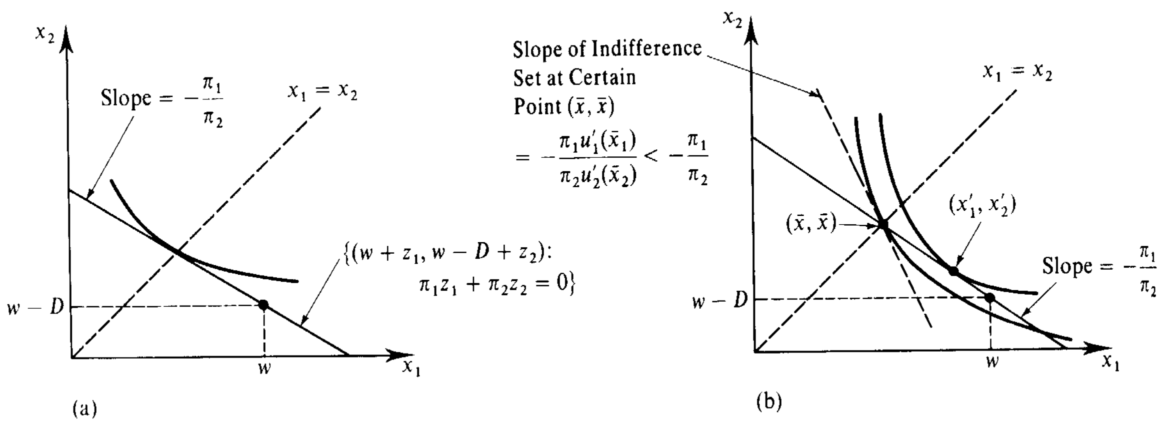

Figure 6.E.2 Insurance purchase with state-dependent utility. (a) State-uniform utility. (b) State-dependent utility.

random variable $(w, w - D)$ that gives the individual's wealth in the two states. This is depicted in Figure 6.E.2(a). We can represent an insurance contract by a random variable $(z_1, z_2) \in \mathbb{R}^2$ specifying the net change in wealth in the two states (the insurance payoff in the state less any premiums paid). Thus, if the individual purchases insurance contract (z_1, z_2), his final wealth position will be $(w + z_1, w - D + z_2)$. The insurance policy (z_1, z_2) is actuarially fair if its expected payoff is zero, that is, if $\pi_1 z_1 + \pi_2 z_2 = 0$.

Figure 6.E.2(a) shows the optimal insurance purchase when a risk-averse expected utility maximizer with state-uniform preferences can purchase any actuarially fair insurance policy he desires. His budget set is the straight line drawn in the figure. We saw in Example 6.C.2 that under these conditions, a decision maker with state-uniform utility would insure completely. This is confirmed here because if there is no state dependency, the budget line is tangent to an indifference curve at the certainty line.

Figure 6.E.2(b) depicts the situation with state-dependent preferences. The decision maker will now prefer a point such as (x_1', x_2') to the certain outcome $(\bar{x}, \bar{x})$. This creates a desire to have a higher payoff in state 1, where $u_1'(\cdot)$ is relatively higher, in exchange for a lower payoff in state 2. ■

Existence of an Extended Expected Utility Representation

We now investigate conditions for the existence of an extended expected utility representation.

Observe first that since $\pi_s > 0$ for every s, we can formally include π_s in the definition of the utility function at state s. That is, to find an extended expected utility representation, it suffices that there be functions $u_s(\cdot)$ such that

$$(x_1, \ldots, x_S) \succsim (x_1', \ldots, x_S') \quad \text{if and only if} \quad \sum_s u_s(x_s) \geq \sum_s u_s(x_s').$$

This is because if such functions $u_s(\cdot)$ exist, then we can define $\tilde{u}_s(\cdot) = (1/\pi_s)u_s(\cdot)$ for each $s \in S$, and we will have $\sum_s u_s(x_s) \geq \sum_s u_s(x_s')$ if and only if $\sum_s \pi_s \tilde{u}_s(x_s) \geq \sum_s \pi_s \tilde{u}_s(x_s')$. Thus, from now on, we focus on the existence of an additively separable form $\sum_s u_s(\cdot)$, and the π_s's cease to play any role in the analysis.

It turns out that the extended expected utility representation can be derived *in exactly the same way* as the expected utility representation of Section 6.B if we appropriately enlarge the domain over which preferences are defined.[25] Accordingly, we now allow for the possibility that within each state s, the monetary payoff is not a certain amount of money x_s but a random amount with distribution function $F_s(\cdot)$. We denote these uncertain alternatives by $L = (F_1, \ldots, F_S)$. Thus, L is a kind of compound lottery that assigns well-defined monetary gambles as prizes contingent on the realization of the state of the world s. We denote by $\mathscr{L}$ the set of all such possible lotteries.

Our starting point is now a rational preference relation $\succsim$ on $\mathscr{L}$. Note that $\alpha L + (1-\alpha)L' = (\alpha F_1 + (1-\alpha)F'_1, \ldots, \alpha F_S + (1-\alpha)F'_S)$ has the usual interpretation as the reduced lottery arising from a randomization between L and L', although here we are dealing with a reduced lottery within each state s. Hence, we can appeal to the same logic as in Section 6.B and impose an independence axiom on preferences.

Definition 6.E.3: The preference relation $\succsim$ on $\mathscr{L}$ satisfies the *extended independence axiom* if for all $L, L', L'' \in \mathscr{L}$ and $\alpha \in (0, 1)$, we have

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

We also make a continuity assumption: Except for the reinterpretation of $\mathscr{L}$, this *continuity axiom* is exactly the same as that in Section 6.B; we refer to Definition 6.B.3 for its statement.

Proposition 6.E.1: (*Extended Expected Utility Theorem*) Suppose that the preference relation $\succsim$ on the space of lotteries $\mathscr{L}$ satisfies the continuity and extended independence axioms. Then we can assign a utility function $u_s(\cdot)$ for money in every state s such that for any $L = (F_1, \ldots, F_S)$ and $L' = (F'_1, \ldots, F'_S)$, we have

$$L \succsim L' \quad \text{if and only if} \quad \sum_s \left(\int u_s(x_s)\, dF_s(x_s)\right) \geq \sum_s \left(\int u_s(x_s)\, dF'_s(x_s)\right).$$

Proof: The proof is identical, almost word for word, to the proof of the expected utility theorem (Proposition 6.B.2).

Suppose, for simplicity, that we restrict ourselves to a finite number $\{x_1, \ldots, x_N\}$ of monetary outcomes. Then we can identify the set $\mathscr{L}$ with Δ^S, where Δ is the $(N-1)$-dimensional simplex. Our aim is to show that $\succsim$ can be represented by a linear utility function $U(L)$ on Δ^S. To see this, note that, up to an additive constant that can be neglected, $U(p_1^1, \ldots, p_N^1, \ldots, p_1^S, \ldots, p_N^S)$ is a linear function of its arguments if it can be written as $U(L) = \sum_{n,s} u_{n,s} p_n^s$ for some values $u_{n,s}$. In this case, we can write $U(L) = \sum_s (\sum_n u_{n,s} p_n^s)$, which, letting $u_s(x_n) = u_{n,s}$, is precisely the form of a utility function on $\mathscr{L}$ that we want.

Choose $\bar{L}$ and $\underline{L}$ such that $\bar{L} \succsim L \succsim \underline{L}$ for all $L \in \mathscr{L}$. As in the proof of Proposition 6.B.2, we can then define $U(L)$ by the condition

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

Applying the extended independence axiom in exactly the same way as we applied the independence axiom in the proof of Proposition 6.B.2 yields the result that $U(L)$ is indeed a linear utility function on $\mathscr{L}$. ■

---

25. By pushing the enlargement further than we do here, it would even be possible to view the existence of an extended utility representation as a corollary of the expected utility theorem.

Proposition 6.F.1 gives us a utility representation $\sum_s u_s(x_s)$ for the preferences on state-by-state sure outcomes $(x_1, \ldots, x_S) \in \mathbb{R}^S_+$ that has two properties. First, it is additively separable across states. Second, every $u_s(\cdot)$ is a Bernoulli utility function that can be used to evaluate lotteries over money payoffs in state s by means of expected utility. It is because of the second property that risk aversion (defined in exactly the same manner as in Section 6.C) is equivalent to the concavity of each $u_s(\cdot)$.

There is another approach to the extended expected utility representation that rests with the preferences $\succsim$ defined on $\mathbb{R}^S_+$ and does not appeal to preferences defined on a larger space. It is based on the so-called *sure-thing axiom*.

Definition 6.E.4: The preference relation $\succsim$ satisfies the *sure-thing axiom* if, for any subset of states $E \subset S$ (E is called an *event*), whenever $(x_1, \ldots, x_S)$ and $(x'_1, \ldots, x'_S)$ differ only in the entries corresponding to E (so that $x'_s = x_s$ for $s \notin E$), the preference ordering between $(x_1, \ldots, x_S)$ and $(x'_1, \ldots, x'_S)$ is independent of the particular (common) payoffs for states not in E. Formally, suppose that $(x_1, \ldots, x_S)$, $(x'_1, \ldots, x'_S)$, $(\bar{x}_1, \ldots, \bar{x}_S)$, and $(\bar{x}'_1, \ldots, \bar{x}'_S)$ are such that

$$\text{For all } s \notin E: \quad x_s = x'_s \quad \text{and} \quad \bar{x}_s = \bar{x}'_s.$$
$$\text{For all } s \in E: \quad x_s = \bar{x}_s \quad \text{and} \quad x'_s = \bar{x}'_s.$$

Then $(\bar{x}_1, \ldots, \bar{x}_S) \succsim (\bar{x}'_1, \ldots, \bar{x}'_S)$ if and only if $(x_1, \ldots, x_S) \succsim (x'_1, \ldots, x'_S)$.

The intuitive content of this axiom is similar to that of the independence axiom. It simply says that if two random variables cannot be distinguished in the complement of E, then the ordering among them can depend only on the values they take on E. In other words, tastes conditional on an event should not depend on what the payoffs would have been in states that have not occurred.

If $\succsim$ admits an extended expected utility representation, the sure-thing axiom holds because then $(x_1, \ldots, x_S) \succsim (x'_1, \ldots, x'_S)$ if and only if $\sum_s (u_s(x_s) - u_s(x'_s)) \geq 0$, and any term of the sum with $x_s = x'_s$ will cancel. In the other direction we have Proposition 6.E.2.

Proposition 6.E.2: Suppose that there are at least three states and that the preferences $\succsim$ on $\mathbb{R}^S_+$ are continuous and satisfy the sure-thing axiom. Then $\succsim$ admits an extended expected utility representation.

Idea of Proof: A complete proof is too advanced to be given in any detail. One wants to show that under the assumptions, preferences admit an additively separable utility representation $\sum_s u_s(x_s)$. This is not easy to show, and it is not a result particularly related to uncertainty. The conditions for the existence of an additively separable utility function for continuous preferences on the positive orthant of a Euclidean space (i.e., the context of Chapter 3) are well understood; as it turns out, they are *formally identical* to the sure-thing axiom (see Exercise 3.G.4). ■

Although the sure-thing axiom yields an extended expected utility representation $\sum_s \pi_s u_s(x_s)$, we would emphasize that randomizations over monetary payoffs in a state s have not been considered in this approach, and therefore we cannot bring the idea of risk aversion to bear on the determination of the properties of $u_s(\cdot)$. Thus, the approach via the extended independence axiom assumes a stronger basic framework (preferences are defined on the set $\mathscr{L}$ rather than on the smaller $\mathbb{R}^S_+$), but it also yields stronger conclusions.

6.F Subjective Probability Theory

Up to this point in the development of the theory, we have been assuming that risk, summarized by means of numerical probabilities, is regarded as an objective fact by the decision maker. But this is rarely true in reality. Individuals make judgments about the chances of uncertain events that are not necessarily expressible in quantitative form. Even when probabilities are mentioned, as sometimes happens when a doctor discusses the likelihood of various outcomes of medical treatment, they are often acknowledged as imprecise *subjective* estimates.

It would be very helpful, both theoretically and practically, if we could assert that choices are made *as if* individuals held probabilistic beliefs. Even better, we would like that well-defined probabilistic beliefs be revealed by choice behavior. This is the intent of *subjective probability theory*. The theory argues that even if states of the world are not associated with recognizable, objective probabilities, consistency-like restrictions on preferences among gambles still imply that decision makers behave *as if* utilities were assigned to outcomes, probabilities were attached to states of nature, and decisions were made by taking expected utilities. Moreover, this rationalization of the decision maker's behavior with an expected utility function can be done uniquely (up to a positive linear transformation for the utility functions). The theory is therefore a far-reaching generalization of expected utility theory. The classical reference for subjective probability theory is Savage (1954), which is very readable but also advanced. It is, however, possible to gain considerable insight into the theory if one is willing to let the analysis be aided by the use of lotteries with objective random outcomes. This is the approach suggested by Anscombe and Aumann (1963), and we will follow it here.

We begin, as in Section 6.E, with a set of states $\{1, \dots, S\}$. The probabilities on $\{1, \dots, S\}$ are not given. In effect, we aim to *deduce* them. As before, a random variable with monetary payoffs is a vector $x = (x_1, \dots, x_S) \in \mathbb{R}^S_+$.[26] We also want to allow for the possibility that the monetary payoffs in a state are not certain but are themselves money lotteries with objective distributions F_s. Thus, our set of risky alternatives, denoted $\mathscr{L}$, is the set of all S-tuples $(F_1, \dots, F_S)$ of distribution functions.

Suppose now that we are given a rational preference relation $\succsim$ on $\mathscr{L}$. We assume that $\succsim$ satisfies the continuity and the extended independence axioms introduced in Section 6.E. Then, by Proposition 6.E.1, we conclude that there are $u_s(\cdot)$ such that any $(x_1, \dots, x_S) \in \mathbb{R}^S_+$ can be evaluated by $\sum_s u_s(x_s)$. In addition, $u_s(\cdot)$ is a Bernoulli utility function for money lotteries in state s.

The existence of the $u_s(\cdot)$ functions does not yet allow us to identify subjective probabilities. Indeed, for *any* $(\pi_1, \dots, \pi_S) \gg 0$, we could define $\tilde{u}_s(\cdot) = (1/\pi_s)u_s(\cdot)$, and we could then evaluate $(x_1, \dots, x_S)$ by $\sum_s \pi_s \tilde{u}_s(x_s)$. What is needed is some way to disentangle utilities from probabilities.

Consider an example. Suppose that a gamble that gives one dollar in state 1 and none in state 2 is preferred to a gamble that gives one dollar in state 2 and none in state 1. Provided *there is no reason to think that the labels of the states have any*

26. To be specific, we consider monetary payoffs here. All the subsequent arguments, however, work with arbitrary sets of outcomes.

particular influence on the value of money, it is then natural to conclude that the decision maker regards state 2 as less likely than state 1.

This example suggests an additional postulate. Preferences over money lotteries within state s should be the same as those within any other state s'; that is, risk attitudes towards money gambles should be the same across states. To formulate such a property, we define the state s preferences $\succsim_s$ on state s lotteries by

$$F_s \succsim_s F'_s \quad \text{if} \quad \int u_s(x_s)\, dF_s(x_s) \geq \int u_s(x_s)\, dF'_s(x_s).$$

Definition 6.F.1: The state preferences $(\succsim_1, \ldots, \succsim_S)$ on state lotteries are *state uniform* if $\succsim_s = \succsim_{s'}$ for any s and s'.

With state uniformity, $u_s(\cdot)$ and $u_{s'}(\cdot)$ can differ only by an increasing linear transformation. Therefore, there is $u(\cdot)$ such that, for all $s = 1, \ldots, S$,

$$u_s(\cdot) = \pi_s u(\cdot) + \beta_s$$

for some $\pi_s > 0$ and β_s. Moreover, because we still represent the same preferences if we divide all π_s and β_s by a common constant, we can normalize the π_s so that $\sum_s \pi_s = 1$. These π_s are going to be our subjective probabilities.

Proposition 6.F.1: (*Subjective Expected Utility Theorem*) Suppose that the preference relation $\succsim$ on $\mathscr{L}$ satisfies the continuity and extended independence axioms. Suppose, in addition, that the derived state preferences are state uniform. Then there are probabilities $(\pi_1, \ldots, \pi_S) \gg 0$ and a utility function $u(\cdot)$ on amounts of money such that for any $(x_1, \ldots, x_S)$ and $(x'_1, \ldots, x'_S)$ we have

$$(x_1, \ldots, x_S) \succsim (x'_1, \ldots, x'_S) \quad \text{if and only if} \quad \sum_s \pi_s u(x_s) \geq \sum_s \pi_s u(x'_s).$$

Moreover, the probabilities are uniquely determined, and the utility function is unique up to origin and scale.

Proof: Existence has already been proven. You are asked to establish uniqueness in Exercise 6.F.1. ■

The practical advantages of the subjective expected utility representation are similar to those of the objective version, which we discussed in Section 6.B, and we will not repeat them here. A major virtue of the theory is that it gives a precise, quantifiable, and operational meaning to uncertainty. It is, indeed, most pleasant to be able to remain in the familiar realm of the probability calculus.

But there are also problems. The plausibility of the axioms cannot be completely dissociated from the complexity of the choice situations. The more complex these become, the more strained even seemingly innocent axioms are, For example, is the completeness axiom reasonable for preferences defined on huge sets of random variables? Or consider the implicit axiom (often those are the most treacherous) that the situation can actually be formalized as indicated by the model. This posits the ability to list all conceivable states of the world (or, at least, a sufficiently disaggregated version of this list). In summary, every difficulty so far raised against our model of the rational consumer (i.e., to transitivity, to completeness, to independence) will apply with increased force to the current model.

There are also difficulties specific to the nonobjective nature of probabilities. We devote Example 6.F.1 to this point.

Example 6.F.1: This example is a variation of the *Ellsberg paradox*.[27] There are two urns, denoted R and H. Each urn contains 100 balls. The balls are either white or black. Urn R contains 49 white balls and 51 black balls. Urn H contains an unspecified assortment of balls. A ball has been randomly picked from each urn. Call them the *R-ball* and the *H-ball*, respectively. The color of these balls has not been disclosed. Now we consider two choice situations. In both experiments, the decision maker must choose either the R-ball or the H-ball. After the choices have been made, the color will be disclosed. In the first choice situation, a prize of 1000 dollars is won if the chosen ball is black. In the second choice situation, the same prize is won if the ball is white. With the information given, most people will choose the R-ball in the first experiment. If the decision is made using subjective probabilities, this should mean that the subjective probability that the H-ball is white is larger than .49. Hence, most people should choose the H-ball in the second experiment. However, it turns out that this does not happen overwhelmingly in actual experiments. The decision maker understands that by choosing the R-ball, he has only a 49% chance of winning. However, this chance is "safe" and well understood. The uncertainties incurred are much less clear if he chooses the H-ball. ■

Knight (1921) proposed distinguishing between *risk* and *uncertainty* according to whether the probabilities are given to us objectively or not. In a sense, the theory of subjective probability nullifies this distinction by reducing all uncertainty to risk through the use of beliefs expressible as probabilities. The Example 6.F.1 suggests that there may be something to the distinction. This is an active area of research [e.g., Bewley (1986) and Gilboa and Schmeidler (1989)].

27. From Ellsberg (1961).

REFERENCES

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque, critique des postulats et axiomes de l'école Américaine. *Econometrica* **21**: 503–46.

Anscombe, F., and R. Aumann. (1963). A definition of subjective probability. *Annals of Mathematical Statistics* **34**: 199–205.

Arrow, K. J. (1971). *Essays in the Theory of Risk Bearing.* Chicago: Markham.

Bewley, T. (1986). Knightian Decision Theory: Part 1. New Haven: Cowles Foundation Discussion Paper No. 807.

Dekel, E. (1986). An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom. *Journal of Economic Theory* **40**: 304–18.

Diamond, P., and M. Rothschild. (1978). *Uncertainty in Economics: Readings and Exercises.* New York: Academic Press.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* **75**: 643–69.

Gilboa, I., and D. Schmeidler. (1989). Maximin expected utility with a unique prior. *Journal of Mathematical Economics* **18**: 141–53.

Grether, D., and C. H. Plott. (1979). Economic theory of choice and the preference reversal phenomenon. *American Economic Review* **69**: 623–38.

Green, J. (1987). 'Making book against oneself,' the independence axiom, and nonlinear utility theory. *Quarterly Journal of Economics* **98**: 785–96.

Hey, J. D. and C. Orme. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica* **62**: 1291–326.

Knight, F. (1921). *Risk, Uncertainty and Profit.* Boston, Mass.: Houghton Mifflin. Reprint, London: London School of Economics 1946.

Kreps, D. (1988). *Notes on the Theory of Choice.* Boulder, Colo.: Westview Press.

Machina, M. (1987). Choice under uncertainty: Problems solved and unsolved. *The Journal of Perspectives* **1**: 121–54.

Pratt, J. (1964). Risk aversion in the small and in the large. *Econometrica* **32**: 122–36. Reprinted in Diamond and Rothschild.

Rothschild, M. and J. Stiglitz. (1970). Increasing risk I: A definition. *Journal of Economic Theory* **2**: 225–43. Reprinted in Diamond and Rothschild.

Savage, L. (1954). *The Foundations of Statistics.* New York: Wiley.

Von Neumann, J., and O. Morgenstern. (1944). *Theory of Games and Economic Behavior.* Princeton, N.J.: Princeton University Press.

EXERCISES

6.B.1$^{\mathrm{A}}$ In text.

6.B.2$^{\mathrm{A}}$ In text.

6.B.3$^{\mathrm{B}}$ Show that if the set of outcomes C is finite and the rational preference relation $\succsim$ on the set of lotteries $\mathscr{L}$ satisfies the independence axiom, then there are best and worst lotteries in $\mathscr{L}$. That is, we can find lotteries $\bar{L}$ and $\underline{L}$ such that $\bar{L} \succsim L \succsim \underline{L}$ for all $L \in \mathscr{L}$.

6.B.4$^{\mathrm{B}}$ The purpose of this exercise is to illustrate how expected utility theory allows us to make consistent decisions when dealing with extremely small probabilities by considering relatively large ones. Suppose that a safety agency is thinking of establishing a criterion under which an area prone to flooding should be evacuated. The probability of flooding is 1%. There are four possible outcomes:

(A) No evacuation is necessary, and none is performed.
(B) An evacuation is performed that is unnecessary.
(C) An evacuation is performed that is necessary.
(D) No evacuation is performed, and a flood causes a disaster.

Suppose that the agency is indifferent between the sure outcome B and the lottery of A with probability p and D with probability $1 - p$, and between the sure outcome C and the lottery of B with probability q and D with probability $1 - q$. Suppose also that it prefers A to D and that $p \in (0, 1)$ and $q \in (0, 1)$. Assume that the conditions of the expected utility theorem are satisfied.

(a) Construct a utility function of the expected utility form for the agency.

(b) Consider two different policy criteria:

Criterion 1: This criterion will result in an evacuation in 90% of the cases in which flooding will occur and an unnecessary evacuation in 10% of the cases in which no flooding occurs.

Criterion 2: This criterion is more conservative. It results in an evacuation in 95% of the cases in which flooding will occur and an unnecessary evacuation in 5% of the cases in which no flooding occurs.

First, derive the probability distributions over the four outcomes under these two criteria. Then, by using the utility function in **(a)**, decide which criterion the agency would prefer.

6.B.5$^{\mathrm{B}}$ The purpose of this exercise is to show that the Allais paradox is compatible with a weaker version of the independence axiom. We consider the following axiom, known as the *betweenness axiom* [see Dekel (1986)]:

For all L, L' and $\lambda \in (0,1)$, if $L \sim L'$, then $\lambda L + (1-\lambda)L' \sim L$.

Suppose that there are three possible outcomes.

(a) Show that a preference relation on lotteries satisfying the independence axiom also satisfies the betweenness axiom.

(b) Using a simplex representation for lotteries similar to the one in Figure 6.B.1(b), show that if the continuity and betweenness axioms are satisfied, then the indifference curves of a preference relation on lotteries are straight lines. Conversely, show that if the indifference curves are straight lines, then the betweenness axiom is satisfied. Do these straight lines need to be parallel?

(c) Using **(b)**, show that the betweenness axiom is weaker (less restrictive) than the independence axiom.

(d) Using Figure 6.B.7, show that the choices of the Allais paradox are compatible with the betweeness axiom by exhibiting an indifference map satisfying the betweenness axiom that yields the choices of the Allais paradox.

6.B.6[B] Prove that the induced utility function $U(\cdot)$ defined in the last paragraph of Section 6.B is convex. Give an example of a set of outcomes and a Bernoulli utility function for which the induced utility function is not linear.

6.B.7[A] Consider the following two lotteries:

$$L: \begin{cases} 200 \text{ dollars with probability } .7. \\ 0 \text{ dollars with probability } .3. \end{cases}$$

$$L': \begin{cases} 1200 \text{ dollars with probability } .1. \\ 0 \text{ dollars with probability } .9. \end{cases}$$

Let x_L and $x_{L'}$ be the sure amounts of money that an individual finds indifferent to L and L'. Show that if his preferences are transitive and monotone, the individual must prefer L to L' if and only if $x_L > x_{L'}$. [*Note*: In actual experiments, however, a preference reversal is often observed in which L is preferred to L' but $x_L < x_{L'}$. See Grether and Plott (1979) for details.]

6.C.1[B] Consider the insurance problem studied in Example 6.C.1. Show that if insurance is not actuarially fair (so that $q > \pi$), then the individual will not insure completely.

6.C.2[B]

(a) Show that if an individual has a Bernoulli utility function $u(\cdot)$ with the quadratic form

$$u(x) = \beta x^2 + \gamma x,$$

then his utility from a distribution is determined by the mean and variance of the distribution and, in fact, by these moments alone. [*Note*: The number β should be taken to be negative in order to get the concavity of $u(\cdot)$. Since $u(\cdot)$ is then decreasing at $x > -\gamma/2\beta$, $u(\cdot)$ is useful only when the distribution cannot take values larger than $-\gamma/2\beta$.]

(b) Suppose that a utility function $U(\cdot)$ over distributions is given by

$$U(F) = (\text{mean of } F) - r(\text{variance of } F),$$

where $r > 0$. Argue that unless the set of possible distributions is further restricted (see, e.g., Exercise 6.C.19), $U(\cdot)$ cannot be compatible with any Bernoulli utility function. Give an example of two lotteries L and L' over the same two amounts of money, say x' and $x'' > x'$, such that L gives a higher probability to x'' than does L' and yet according to $U(\cdot)$, L' is preferred to L.

6.C.3[B] Prove that the four conditions of Proposition 6.C.1 are equivalent. [*Hint*: The equivalence of (i), (ii), and (iii) has already been shown. As for (iv), prove that (i) implies (iv) and that (iv) implies $u(\frac{1}{2}x + \frac{1}{2}y) \geq \frac{1}{2}u(x) + \frac{1}{2}u(y)$ for any x and y, which is, in fact, sufficient for (ii).]

6.C.4[B] Suppose that there are N risky assets whose returns z_n $(n = 1, \ldots, N)$ per dollar invested are jointly distributed according to the distribution function $F(z_1, \ldots, z_N)$. Assume also that all the returns are nonnegative with probability one. Consider an individual who has a continuous, increasing, and concave Bernoulli utility function $u(\cdot)$ over $\mathbb{R}_+$. Define the utility function $U(\cdot)$ of this investor over $\mathbb{R}_+^N$, the set of all nonnegative portfolios, by

$$U(\alpha_1, \ldots, \alpha_N) = \int u(\alpha_1 z_1 + \cdots + \alpha_N z_N)\, dF(z_1, \ldots, z_N).$$

Prove that $U(\cdot)$ is (a) increasing, (b) concave, and (c) continuous (this is harder).

6.C.5[A] Consider a decision maker with utility function $u(\cdot)$ defined over $\mathbb{R}_+^L$, just as in Chapter 3.

(a) Argue that concavity of $u(\cdot)$ can be interpreted as the decision maker exhibiting risk aversion with respect to lotteries whose outcomes are bundles of the L commodities.

(b) Suppose now that a Bernoulli utility function $u(\cdot)$ for wealth is derived from the maximization of a utility function defined over bundles of commodities for each given wealth level w, while prices for those commodities are fixed. Show that, if the utility function for the commodities exhibits risk aversion, then so does the derived Bernoulli utility function for wealth. Interpret.

(c) Argue that the converse of part **(b)** does not need to hold: There are nonconcave functions $u: \mathbb{R}_+^L \to \mathbb{R}$ such that for any price vector the derived Bernoulli utility function on wealth exhibits risk aversion.

6.C.6[B] For Proposition 6.C.2:

(a) Prove the equivalence of conditions (ii) and (iii).

(b) Prove the equivalence of conditions (iii) and (v).

6.C.7[A] Prove that, in Proposition 6.C.2, condition (iii) implies condition (iv), and (iv) implies (i).

6.C.8[A] In text.

6.C.9[B] (M. Kimball) The purpose of this problem is to examine the implications of uncertainty and precaution in a simple consumption–savings decision problem.

In a two-period economy, a consumer has first-period initial wealth w. The consumer's utility level is given by

$$u(c_1, c_2) = u(c_1) + v(c_2),$$

where $u(\cdot)$ and $v(\cdot)$ are concave functions and c_1 and c_2 denote consumption levels in the first and the second period, respectively. Denote by x the amount saved by the consumer in the first period (so that $c_1 = w - x$ and $c_2 = x$), and let x_0 be the optimal value of x in this problem.

We now introduce uncertainty in this economy. If the consumer saves an amount x in the first period, his wealth in the second period is given by $x + y$, where y is distributed according to $F(\cdot)$. In what follows, $E[\cdot]$ always denotes the expectation with respect to $F(\cdot)$. Assume that the Bernoulli utility function over realized wealth levels in the two periods (w_1, w_2) is $u(w_1) + v(w_2)$. Hence, the consumer now solves

$$\operatorname*{Max}_{x}\ u(w - x) + E[v(x + y)].$$

Denote the solution to this problem by x^*.

(a) Show that if $E[v'(x_0 + y)] > v'(x_0)$, then $x^* > x_0$.

(b) Define the *coefficient of absolute prudence* of a utility function $v(\cdot)$ at wealth level x to be $-v'''(x)/v''(x)$. Show that if the coefficient of absolute prudence of a utility function $v_1(\cdot)$ is not larger than the coefficient of absolute prudence of utility function $v_2(\cdot)$ for all levels of wealth, then $E[v_1'(x_0 + y)] > v_1'(x_0)$ implies $E[v_2'(x_0 + y)] > v_2'(x_0)$. What are the implications of this fact in the context of part **(a)**?

(c) Show that if $v'''(\cdot) > 0$, and $E[y] = 0$, then $E[v'(x + y)] > v'(x)$ for all values of x.

(d) Show that if the coefficient of absolute risk aversion of $v(\cdot)$ is decreasing with wealth, then $-v'''(x)/v''(x) > -v''(x)/v'(x)$ for all x, and hence $v'''(\cdot) > 0$.

6.C.10^A Prove the equivalence of conditions (i) through (v) in Proposition 6.C.3. [*Hint*: By letting $u_1(z) = u(w_1 + z)$ and $u_2(z) = u(w_2 + z)$, show that each of the five conditions in Proposition 6.C.3 is equivalent to the counterpart in Proposition 6.C.2.]

6.C.11^B For the model in Example 6.C.2, show that if $r_R(x, u)$ is increasing in x then the proportion of wealth invested in the risky asset $\gamma = \alpha/x$ is decreasing with x. Similarly, if $r_R(x, u)$ is decreasing in x, then $\gamma = \alpha/x$ is increasing in x. [*Hint*: Let $u_1(t) = u(tw_1)$ and $u_2(t) = u(tw_2)$, and use the fact, stated in the analysis of Example 6.C.2, that if one Bernoulli utility function is more risk averse than another, then the optimal level of investment in the risky asset for the first function is smaller than that for the second function. You could also attempt a direct proof using first-order conditions.]

6.C.12^B Let $u: \mathbb{R}_+ \to \mathbb{R}$ be a strictly increasing Bernoulli utility function. Show that

(a) $u(\cdot)$ exhibits constant relative risk aversion equal to $\rho \neq 1$ if and only if $u(x) = \beta x^{1-\rho} + \gamma$, where $\beta > 0$ and $\gamma \in \mathbb{R}$.

(b) $u(\cdot)$ exhibits constant relative risk aversion equal to 1 if and only if $u(x) = \beta \ln x + \gamma$, where $\beta > 0$ and $\gamma \in \mathbb{R}$.

(c) $\lim_{\rho \to 1} (x^{1-\rho}/(1 - \rho)) = \ln x$ for all $x > 0$.

6.C.13^B Assume that a firm is risk neutral with respect to profits and that if there is any uncertainty in prices, production decisions are made after the resolution of such uncertainty. Suppose that the firm faces a choice between two alternatives. In the first, prices are uncertain. In the second, prices are nonrandom and equal to the expected price vector in the first alternative. Show that a firm that maximizes expected profits will prefer the first alternative over the second.

6.C.14^B Consider two risk-averse decision makers (i.e., two decision makers with concave Bernoulli utility functions) choosing among monetary lotteries. Define the utility function $u^*(\cdot)$ to be strongly more risk averse than $u(\cdot)$ if and only if there is a positive constant k and a nonincreasing and concave function $v(\cdot)$ such that $u^*(x) = ku(x) + v(x)$ for all x. The monetary amounts are restricted to lie in the interval $[0, r]$.

(a) Show that if $u^*(\cdot)$ is strongly more risk averse than $u(\cdot)$, then $u^*(\cdot)$ is more risk averse than $u(\cdot)$ in the usual Arrow–Pratt sense.

(b) Show that if $u(\cdot)$ is bounded, then there is no $u^*(\cdot)$ other than $u^*(\cdot) = ku(\cdot) + c$, where c is a constant, that is strongly more risk averse than $u(\cdot)$ on the entire interval $[0, +\infty]$. [*Hint*: in this part, disregard the assumption that the monetary amounts are restricted to lie in the interval $[0, r]$.]

(c) Using **(b)**, argue that the concept of a strongly more risk-averse utility function is stronger (i.e., more restrictive) than the Arrow–Pratt concept of a more risk-averse utility function.

6.C.15[A] Assume that, in a world with uncertainty, there are two assets. The first is a riskless asset that pays 1 dollar. The second pays amounts a and b with probabilities of π and $1-\pi$, respectively. Denote the demand for the two assets by (x_1, x_2).

Suppose that a decision maker's preferences satisfy the axioms of expected utility theory and that he is a risk averter. The decision maker's wealth is 1, and so are the prices of the assets. Therefore, the decision maker's budget constraint is given by

$$x_1 + x_2 = 1, \quad x_1, x_2 \in [0, 1].$$

(a) Give a simple *necessary* condition (involving a and b only) for the demand for the riskless asset to be strictly positive.

(b) Give a simple *necessary* condition (involving a, b, and π only) for the demand for the risky asset to be strictly positive.

In the next three parts, assume that the conditions obtained in **(a)** and **(b)** are satisfied.

(c) Write down the first-order conditions for utility maximization in this asset demand problem.

(d) Assume that $a < 1$. Show by analyzing the first-order conditions that $dx_1/da \le 0$.

(e) Which sign do you conjecture for $dx_1/d\pi$? Give an economic interpretation.

(f) Can you prove your conjecture in **(e)** by analyzing the first-order conditions?

6.C.16[A] An individual has Bernoulli utility function $u(\cdot)$ and initial wealth w. Let lottery L offer a payoff of G with probability p and a payoff of B with probability $1-p$.

(a) If the individual owns the lottery, what is the minimum price he would sell it for?

(b) If he does not own it, what is the maximum price he would be willing to pay for it?

(c) Are buying and selling prices equal? Give an economic interpretation for your answer. Find conditions on the parameters of the problem under which buying and selling prices are equal.

(d) Let $G = 10$, $B = 5$, $w = 10$, and $u(x) = \sqrt{x}$. Compute the buying and selling prices for this lottery and this utility function.

6.C.17[B] Assume that an individual faces a two-period portfolio allocation problem. In period $t = 0, 1$, his wealth w_t is to be divided between a safe asset with return R and a risky asset with return x. The initial wealth at period 0 is w_0. Wealth at period $t = 1, 2$ depends on the portfolio α_{t-1} chosen at period $t-1$ and on the return x_t realized at period t, according to

$$w_t = ((1 - \alpha_{t-1})R + \alpha_{t-1}x_t)w_{t-1}.$$

The objective of this individual is to maximize the expected utility of terminal wealth w_2. Assume that x_1 and x_2 are independently and identically distributed. Prove that the individual optimally sets $\alpha_0 = \alpha_1$ if his utility function exhibits constant relative risk aversion. Show also that this fails to hold if his utility function exhibits constant absolute risk aversion.

6.C.18[B] Suppose that an individual has a Bernoulli utility function $u(x) = \sqrt{x}$.

(a) Calculate the Arrow–Pratt coefficients of absolute and relative risk aversion at the level of wealth $w = 5$.

(b) Calculate the certainty equivalent and the probability premium for a gamble $(16, 4; \frac{1}{2}, \frac{1}{2})$.

(c) Calculate the certainty equivalent and the probability premium for a gamble $(36, 16; \frac{1}{2}, \frac{1}{2})$. Compare this result with the one in **(b)** and interpret.

6.C.19[C] Suppose that an individual has a Bernoulli utility function $u(x) = -e^{-\alpha x}$ where $\alpha > 0$. His (nonstochastic) initial wealth is given by w. There is one riskless asset and there are N

risky assets. The return per unit invested on the riskless asset is r. The returns of the risky assets are jointly normally distributed random variables with means $\mu = (\mu_1, \ldots, \mu_N)$ and variance covariance matrix V. Assume that there is no redundancy in the risky assets, so that V is of full rank. Derive the demand function for these $N + 1$ assets.

6.C.20[A] Consider a lottery over monetary outcomes that pays $x + \varepsilon$ with probability $\frac{1}{2}$ and $x - \varepsilon$ with probability $\frac{1}{2}$. Compute the second derivative of this lottery's certainty equivalent with respect to ε. Show that the limit of this derivative as $\varepsilon \to 0$ is exactly $-r_A(x)$.

6.D.1[A] The purpose of this exercise is to prove Proposition 6.D.1 in a two-dimensional probability simplex. Suppose that there are three monetary outcomes: 1 dollar, 2 dollars, and 3 dollars. Consider the probability simplex of Figure 6.B.1(b).

(a) For a given lottery L over these outcomes, determine the region of the probability simplex in which lie the lotteries whose distributions first-order stochastically dominate the distribution of L.

(b) Given a lottery L, determine the region of the probability simplex in which lie the lotteries L' such that $F(x) \leq G(x)$ for every x, where $F(\cdot)$ is the distribution of L' and $G(\cdot)$ is the distribution of L. [Notice that we get the same region as in **(a)**.]

6.D.2[A] Prove that if $F(\cdot)$ first-order stochastically dominates $G(\cdot)$, then the mean of x under $F(\cdot)$, $\int x\, dF(x)$, exceeds that under $G(\cdot)$, $\int x\, dG(x)$. Also provide an example where $\int x\, dF(x) > \int x\, dG(x)$ but $F(\cdot)$ does not first-order stochastically dominate $G(\cdot)$.

6.D.3[A] Verify that if a distribution $G(\cdot)$ is an elementary increase in risk from a distribution $F(\cdot)$, then $F(\cdot)$ second-order stochastically dominates $G(\cdot)$.

6.D.4[A] The purpose of this exercise is to verify the equivalence of the three statements of Proposition 6.D.2 in a two-dimensional probability simplex. Suppose that there are three monetary outcomes: 1, 2, and 3 dollars. Consider the probability simplex in Figure 6.B.1(b).

(a) If two lotteries have the same mean, what are their positions relative to each other in the probability simplex.

(b) Given a lottery L, determine the region of the simplex in which lie the lotteries L' whose distributions are second-order stochastically dominated by the distribution of L.

(c) Given a lottery L, determine the region of the simplex in which lie the lotteries L' whose distributions are mean preserving spreads of L.

(d) Given a lottery L, determine the region of the simplex in which lie the lotteries L' for which condition (6.D.2) holds, where $F(\cdot)$ and $G(\cdot)$ are, respectively, the distributions of L and L'.

Notice that in **(b)**, **(c)**, and **(d)**, you always have the same region.

6.E.1[B] The purpose of this exercise is to show that preferences may not be transitive in the presence of regret. Let there be S states of the world, indexed by $s = 1, \ldots, S$. Assume that state s occurs with probability π_s. Define the expected regret associated with lottery $x = (x_1, \ldots, x_S)$ relative to lottery $x' = (x'_1, \ldots, x'_S)$ by

$$\sum_{s=1}^{S} \pi_s h(\text{Max}\,\{0, x'_s - x_s\}),$$

where $h(\cdot)$ is a given increasing function. [We call $h(\cdot)$ the *regret valuation function*; it measures the regret the individual has after the state of nature is known.] We define x to be at least as good as x' in the presence of regret if and only if the expected regret associated with x relative to x' is not greater than the expected regret associated with x' relative to x.

Suppose that $S = 3$, $\pi_1 = \pi_2 = \pi_3 = \frac{1}{3}$, and $h(x) = \sqrt{x}$. Consider the following three lotteries:

$$x = (0, -2, \quad 1),$$

$$x' = (0, \quad 2, -2),$$

$$x'' = (2, -3, -1).$$

Show that the preference ordering over these three lotteries is not transitive.

6.E.2[A] Assume that in a world with uncertainty there are two possible states of nature ($s = 1, 2$) and a single consumption good. There is a single decision maker whose preferences over lotteries satisfy the axioms of expected utility theory and who is a risk averter. For simplicity, we assume that utility is state-independent.

Two contingent commodities are available to the decision maker. The first (respectively, the second) pays one unit of the consumption good in state $s = 1$ (respectively $s = 2$) and zero otherwise. Denote the vector quantities of the two contingent commodities by (x_1, x_2).

(a) Show that the preference relation of the decision maker on (x_1, x_2) is convex.

(b) Argue that the decision maker is also a risk averter when choosing between lotteries whose outcomes are vectors (x_1, x_2).

(c) Show that the Walrasian demand functions for x_1 and x_2 are normal.

6.E.3[B] Let $g: S \to \mathbb{R}_+$ be a random variable with mean $E(g) = 1$. For $\alpha \in (0, 1)$, define a new random variable $g^*: S \to \mathbb{R}_+$ by $g^*(s) = \alpha g(s) + (1 - \alpha)$. Note that $E(g^*) = 1$. Denote by $G(\cdot)$ and $G^*(\cdot)$ the distribution functions of $g(\cdot)$ and $g^*(\cdot)$, respectively. Show that $G^*(\cdot)$ second-order stochastically dominates $G(\cdot)$. Interpret.

6.F.1[B] Prove that in the subjective expected utility theorem (Proposition 6.F.2), the obtained utility function $u(\cdot)$ on money is uniquely determined up to origin and scale. That is, if both $u(\cdot)$ and $\hat{u}(\cdot)$ satisfy the condition of the theorem, then there exist $\beta > 0$ and $\gamma \in \mathbb{R}$ such that $\hat{u}(x) = \beta u(x) + \gamma$ for all x. Prove also that the subjective probabilities are uniquely determined.

6.F.2[A] The purpose of this exercise is to explain the outcomes of the experiments described in Example 6.F.1 by means of the theory of *nonunique prior beliefs* of Gilboa and Schmeidler (1989).

We consider a decision maker with a Bernoulli utility function $u(\cdot)$ defined on $\{0, 1000\}$. We normalize $u(\cdot)$ so that $u(0) = 0$ and $u(1000) = 1$.

The probabilistic belief that the decision maker might have on the color of the H-ball being white is a number $\pi \in [0, 1]$. We assume that the decision maker has, not a single belief but a *set* of beliefs given by a subset P of $[0, 1]$. The actions that he may take are denoted R or H with R meaning that he chooses the R-ball and H meaning that he chooses the H-ball.

As in Example 6.F.1, the decision maker is faced with two different choice situations. In choice situation W, he receives 1000 dollars if the ball chosen is white and 0 dollars otherwise. In choice situation B, he receives 1000 dollars if the ball chosen is black and 0 dollars otherwise.

For each of the two choice situations, define his utility function over the actions R and H in the following way:

For situation W, $U_W: \{\text{R}, \text{H}\} \to \mathbb{R}$ is defined by

$$U_W(\text{R}) = .49 \quad \text{and} \quad U_W(\text{H}) = \text{Min}\,\{\pi : \pi \in P\}.$$

For situation B, $U_B: \{\text{R}, \text{H}\} \to \mathbb{R}$ is defined by

$$U_B(\text{R}) = .51 \quad \text{and} \quad U_B(\text{H}) = \text{Min}\,\{(1 - \pi) : \pi \in P\}.$$

Namely, his utility from choice R is the expected utility of 1000 dollars with the (objective) probability calculated from the number of white and black balls in urn R. However, his utility from choice H is the expected utility of 1000 dollars with the probability associated with the most pessimistic belief in P.

(a) Prove that if P consists of only one belief, then U_W and U_B are derived from a von Neumann Morgenstern utility function and that $U_W(\mathrm{R}) > U_W(\mathrm{H})$ if and only if $U_B(\mathrm{R}) < U_B(\mathrm{H})$.

(b) Find a set P for which $U_W(\mathrm{R}) > U_W(\mathrm{H})$ and $U_B(\mathrm{R}) > U_B(\mathrm{H})$.