GAMES OF STRATEGY

THIRD EDITION

Avinash Dixit
Princeton University

Susan Skeath
Wellesley College

David Reiley
University of Arizona and Yahoo! Research

W. W. Norton & Company
New York • London

information, find the payoffs in the two-by-two table of step 4, and find the Nash equilibrium at this step. Work backward in the same way through the rest of the steps to find the Nash equilibrium strategies of the full game.

U13. Describe an example of business competition that is similar in structure to the duel in Exercise U12.

7

Simultaneous-Move Games with Mixed Strategies I: Two-by-Two Games

In our study of simultaneous-move games in Chapter 4, we came across a class of games that the solution methods described there could not solve; in fact, games in that class have no Nash equilibria in pure strategies. To predict outcomes for such games, we need an extension of our concepts of strategies and equilibria. This is to be found in the randomization of moves, which is the focus of this chapter and the next.

Consider the tennis-point game from the end of Chapter 4. This game is zero sum; the interests of the two tennis players are purely in mutual conflict. Evert wants to hit her passing shot to whichever side—down the line (DL) or crosscourt (CC)—is not covered by Navratilova, whereas Navratilova wants to cover the side to which Evert hits her shot. In Chapter 4, we pointed out that in such a situation, any systematic choice by Evert will be exploited by Navratilova to her own advantage and therefore to Evert's disadvantage. Conversely, Evert can exploit any systematic choice by Navratilova. To avoid being thus exploited, each player wants to keep the other guessing, which can be done by acting unsystematically or randomly.

However, randomness doesn't mean choosing each shot half the time, or alternating between the two. The latter would itself be a systematic action open to exploitation, and a 60-40 or 75-25 random mix may be better than 50-50 depending on the situation. In this chapter we develop methods for calculating the best mix and discuss how well this theory helps us understand actual play in such games.

Our method for calculating the best mix can also be applied to non-zero-sum games. However, in such games the players' interests can partially coincide, so when player B exploits A's systematic choice to her own advantage, it is not necessarily to A's disadvantage. Therefore the logic of keeping the other player guessing is weaker or even absent altogether in non-zero-sum games. We will discuss whether and when mixed-strategy equilibria make sense in such games.

Throughout this chapter, we limit the discussion to two-by-two games and to the most direct method for calculating a mixed-strategy equilibrium in order to present the basic ideas in the simplest possible setting. Many of the concepts and methods we develop here continue to be valid in more general games; most important, our discussion in Sections 6 and 7 of how to mix strategies and whether mixing is observed in reality is perfectly general. However, we leave to Chapter 8 the general analysis of best responses in mixed strategies and of mixed-strategy equilibria for games with more than two pure strategies.

1 WHAT IS A MIXED STRATEGY?

When players choose to act unsystematically, they pick from among their pure strategies in some random way. In the tennis-point game, Navratilova and Evert each choose from two initially given pure strategies, DL and CC. We call a random mixture of these two pure strategies a mixed strategy.

Such mixed strategies cover a whole continuous range. At one extreme, DL could be chosen with probability 1 (for sure), meaning that CC is never chosen (probability 0); this "mixture" is just the pure strategy DL. At the other extreme, DL could be chosen with probability 0 and CC with probability 1; this "mixture" is the same as pure CC. In between is the whole set of possibilities: DL chosen with probability 75% (0.75) and CC with probability 25% (0.25); or both chosen with probabilities 50% (0.5) each; or DL with probability 1/3 (33.33 . . . %) and CC with probability 2/3 (66.66 . . . %); and so on.[1]

The payoffs from a mixed strategy are defined as the corresponding probability-weighted averages of the payoffs from its constituent pure

[1]When a chance event has just two possible outcomes, people often speak of the odds in favor of or against one of the outcomes. If the two possible outcomes are labeled A and B, and the probability of A is p so that the probability of B is $(1 - p)$, then the ratio $p/(1 - p)$ gives the odds in favor of A, and the reverse ratio $(1 - p)/p$ gives the odds against A. Thus, when Evert chooses CC with probability 0.25 (25%), the odds against her choosing CC are 3 to 1, and the odds in favor of it are 1 to 3. This terminology is often used in betting contexts, so those of you who misspent your youth in that way will be more familiar with it. However, this usage does not readily extend to situations in which three or more outcomes are possible, so we avoid its use here.

strategies. For example, in the tennis game of Section 4.8, against Navratilova's DL, Evert's payoff from DL is 50 and from CC is 90. Therefore the payoff of Evert's mixture (0.75 DL, 0.25 CC) against Navratilova's DL is $0.75 \times 50 + 0.25 \times 90 = 37.5 + 22.5 = 60$. This is Evert's **expected payoff** from this particular mixed strategy.[2]

You may have noticed that mixed strategies are very similar to the continuous strategies we studied in Chapter 5. In fact, they are just special kinds of continuously variable strategies. With the possibility of mixing made available, each player can now choose from among all conceivable mixtures of the basic or pure strategies initially specified (which, as we just saw, include pure strategies as extreme special cases).

The notion of Nash equilibrium also extends easily to include mixed strategies. Nash equilibrium is defined as a list of mixed strategies, one for each player, such that the choice of each is her best choice, in the sense of yielding the highest expected payoff for her, given the mixed strategies of the others. Allowing for mixed strategies in a game solves the problem of possible nonexistence of Nash equilibrium, which we encountered for pure strategies, automatically and almost entirely. Nash's celebrated theorem shows that, under very general circumstances (which are broad enough to cover all the games that we meet in this book and many more besides), a Nash equilibrium in mixed strategies exists.

At this broadest level, therefore, incorporating mixed strategies into our analysis does not entail anything different from the general theory of continuous strategies developed in Chapter 5. However, the special case of mixed strategies does bring with it several special conceptual as well as methodological matters, and therefore deserves separate study.

2 UNCERTAIN ACTIONS: MIXING MOVES TO KEEP THE OPPONENT GUESSING

We begin with the tennis example of Section 4.8, which did not have a Nash equilibrium in pure strategies. We show how the extension to mixed strategies remedies this deficiency, and we interpret the resulting equilibrium as one in which each player keeps the other guessing.

[2]Game theory assumes that players will calculate and try to maximize their expected payoffs when probabilistic mixtures of strategies or outcomes are included. We consider this further in the Appendix to this chapter, but for now we proceed to use it, with just one important note. The word *expected* in "expected payoff" is a technical term from probability and statistics. It merely denotes a probability-weighted average. It does not mean this is the payoff that the player should expect in the sense of regarding it as her right or entitlement.

A. The Benefit of Mixing

We reproduce in Figure 7.1 the payoff matrix of Figure 4.15 with both players' payoffs included. Even though this is a constant-sum game, we now show the payoffs for the two players separately, because it is more intuitive to think of each player as trying to get a higher payoff for herself. In this game, if Evert always chooses DL, Navratilova will then cover DL and hold Evert's payoff down to 50. Similarly, if Evert always chooses CC, Navratilova will choose to cover CC and hold Evert down to 20. If Evert can only choose one of her two basic (pure) strategies, and Navratilova can predict that choice, Evert's better (or less bad) pure strategy will be DL, yielding her a payoff of 50. (This argument follows the minimax reasoning of Chapter 4, Section 5.)

But suppose Evert is not restricted to using only pure strategies and can choose a mixed strategy, perhaps one in which the probability of playing DL on any one occasion is 75%, or 0.75; this makes her probability of playing CC 25%, or 0.25. Using the method outlined in Section 1, we can calculate Navratilova's expected payoff against this mixture as

$$0.75 \times 50 + 0.25 \times 10 = 37.5 + 2.5 = 40 \text{ if she covers DL, and}$$
$$0.75 \times 20 + 0.25 \times 80 = 15 + 20 = 35 \text{ if she covers CC.}$$

If Evert chooses this mixture, the expected payoffs show that Navratilova can best exploit it by covering DL.

When Navratilova chooses DL to best exploit Evert's 75-25 mix, her choice works to Evert's disadvantage because this is a zero-sum game. Evert's expected payoffs are

$$0.75 \times 50 + 0.25 \times 90 = 37.5 + 22.5 = 60 \text{ if Navratilova covers DL, and}$$
$$0.75 \times 80 + 0.25 \times 20 = 60 + 5 = 65 \text{ if Navratilova covers CC.}$$

By choosing DL, Navratilova holds Evert down to 60 rather than 65. But notice that Evert's payoff with the mixture is still better than the 50 she would get by playing purely DL, or the 20 she would get by playing purely CC.[3]

The 75-25 mix does leave Evert's strategy open to some exploitation by Navratilova. By choosing to cover DL she can hold Evert down to a lower expected payoff than when she chooses CC. Ideally, Evert would like to find a mix that would be exploitation proof. To find such a mix requires taking a more general approach to describing Evert's mixed strategy. Specifically, we denote the

[3]Not every mixed strategy will perform better than the pure strategies. For example, if Evert mixes 50:50 between DL and CC, Navratilova can hold Evert's expected payoff down to 50, exactly the same as from pure DL. And a mixture that attaches a probability of less than 30% to DL will be worse for Evert than pure DL. We ask you to verify these statements as a useful exercise to acquire the skill of calculating expected payoffs and comparing strategies.

		NAVRATILOVA	
		DL	CC
EVERT	DL	50, 50	80, 20
	CC	90, 10	20, 80

FIGURE 7.1 Payoffs of Both Players in the Tennis Point

probability of Evert choosing DL by the algebraic symbol p (so the probability of choosing CC is $1 - p$) and find the condition for this mixture to be exploitation proof (in terms of p). Solving the condition for p provides the answer we seek.

So suppose Evert chooses DL with probability p and CC with probability $(1 - p)$. We will refer to this mixture as Evert's p-mix for short. Against the p-mix, Navratilova's expected payoffs are

$$50p + 10(1 - p) \text{ if she covers DL, and}$$
$$20p + 80(1 - p) \text{ if she covers CC.}$$

For Evert's strategy to be exploitation proof, these two expected payoffs should be equal. That implies $50p + 10(1 - p) = 20p + 80(1 - p)$; or $30p = 70(1 - p)$; or $100p = 70$; or $p = 0.7$. Thus Evert's exploitation-proof mix uses DL with probability 70% and CC with probability 30%. Evert's expected payoff from this mixed strategy is

$$50 \times 0.7 + 90 \times 0.3 = 35 + 27 = 62 \text{ if Navratilova covers DL, and also}$$
$$80 \times 0.7 + 20 \times 0.3 = 56 + 6 = 62 \text{ if Navratilova covers CC.}$$

This expected payoff is better than the 50 that Evert would get if she used the pure strategy DL, and better than the 60 from the 75-25 mixture. We know this mixture is exploitation proof, but does it ensure Evert the best possible expected payoff?

To answer this question, we first consider how Navratilova will respond to various choices of p by Evert. To see this explicitly, we construct a diagram, Figure 7.2, showing Navratilova's expected payoffs from covering DL and from covering CC for all of Evert's possible p-mixes. Evert's choice of p for her p-mix is shown on the horizontal axis in Figure 7.2; Navratilova's expected payoffs are on the vertical axis. In the figure are two lines, one corresponding to Navratilova's choice of pure DL and the other to her choice of pure CC. The equations of the lines are given by the two expressions above—that is, Navratilova's expected payoff from covering DL $= 50p + 10(1 - p) = 10 + 40p$ and her expected payoff from covering CC $= 20p + 80(1 - p) = 80 - 60p$.

The upward sloping straight line in Figure 7.2 shows Navratilova's expected payoff from her DL against Evert's p-mix. The intuition for this shape is that

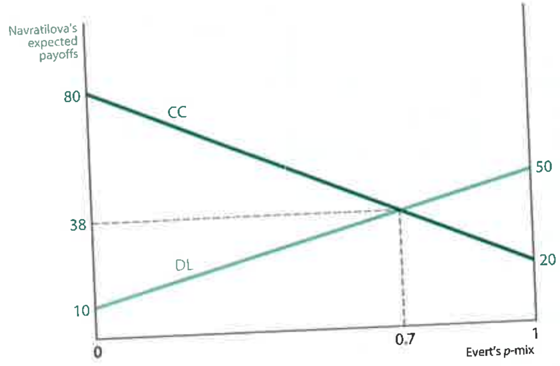

FIGURE 7.2 Navratilova's Expected Payoffs Against Evert's p-mixes

Navratilova does better by covering DL (so her expected payoff is higher) when Evert is more likely to use DL (or when p is higher). The falling straight line in the figure depicts Navratilova's expected payoff from her CC against Evert's p-mix. Navratilova fares worse as the likelihood that Evert uses DL rises.

If Navratilova were to use a mixture of DL and CC instead of a pure strategy, her expected payoff would be a weighted average of the payoffs from the two pure strategies. Thus her expected payoff from choosing a mixed strategy would lie somewhere between the two expected payoff lines in Figure 7.2. (As we did for Evert, we will use an algebraic symbol, this time a q, to represent the probability with which Navratilova covers DL. Then we will refer to the mixture putting probability q on DL and probability $1 - q$ on CC as Navratilova's q-mix, for short.)

It is useful to calculate the value of p at the point of intersection of the two lines in Figure 7.2. That p satisfies $10 + 40p = 80 - 60p$, or $p = 0.7$. The expected payoff for Navratilova at this value of p is the same on both lines: it is $50 \times 0.7 + 10 \times 0.3 = 35 + 3 = 38$ when calculated using the line for DL, and $20 \times 0.7 + 80 \times 0.3 = 14 + 24 = 38$ when calculated using the line for CC.

We can now identify the strategy that is best for Navratilova against each possible value of Evert's p. This is just the strategy that yields her the highest expected payoff in each case. For p=0, for example, the diagram shows that Navratilova gets a higher payoff when she plays CC, namely 80, than the payoff she would get by playing DL, namely 10. If Navratilova were to play a 75-25 mixture of DL and CC against Evert's $p = 0$, then Navratilova's expected payoff would be a weighted average of the pure strategy payoffs, or $80 \times 0.75 + 10 \times 0.25 = 62.5$. If her payoff from pure CC is higher than her payoff from pure DL, then her payoff

from pure CC is also higher than that from any such average. Therefore her pure CC is not only better for Navratilova than her pure DL, it is also her best choice among all her strategies, pure and mixed.

This continues to be true for all values of p to the left of the intersection point in the diagram. For all such values of p, pure CC is better than DL for Navratilova when played against Evert's p-mix and is also better than any mixed strategy of her own. Pure CC is Navratilova's best response for this range of values of p. Similarly, to the right of the point of intersection, pure DL is her best response.[4]

Against Evert's specific p-mix at the point of intersection, Navratilova does equally well (expected payoff 38) from her CC and DL. Because the expected payoff associated with a mixture of her CC and DL lies "between" the CC and DL lines, it follows that she also does equally well with any q-mix of her own. The weighted average of 38 and 38 is 38, no matter what the weights used for averaging. Thus, Navratilova's best response to Evert's choice of $p = 0.7$ can be a q-mix for any q in the entire range from 0 to 1.

We can summarize Navratilova's best-response rule as follows:

If $p < 0.7$, choose pure CC ($q = 0$).
If $p = 0.7$, all values of q in the range from 0 to 1 are equal best responses.
If $p > 0.7$, choose pure DL ($q = 1$).

To see the intuition behind this rule, note that p is the probability with which Evert hits her passing shot down the line. Navratilova wants to prepare for Evert's shot. Therefore if p is low (Evert is more likely to go crosscourt), Navratilova does better to defend crosscourt. If p is high (Evert is more likely to go down the line), Navratilova does better to defend down the line. For a critical value of p in between, the two choices are equally good for Navratilova. This break-even point is not $p = 0.5$ because the consequences (measured by the payoff numbers of the various combinations) are not symmetric for the two choices.

Recall that our purpose in drawing Figure 7.2 was to help answer the question of whether Evert's exploitation-proof mixture, $p = 0.7$, is her best mixed strategy, in the sense of giving her the highest expected payoff. The answer is yes. We can see this by looking at the expected payoff that Navratilova can get when she makes her best response to each of Evert's p-mixes. In Figure 7.2, this is shown by the higher of the two lines, which corresponds to Navratilova's coverage of DL or CC. The best-response expected payoffs make up a V-shaped curve. The point of intersection of the two lines in the figure is the lowest point of the V. When Evert chooses $p = 0.7$, Navratilova gets her worst possible expected payoff, and therefore Evert gets her best possible expected payoff.

[4] If, in some numerical problem you are trying to solve, the expected payoff lines for the pure strategies do not intersect, that would indicate that one pure strategy was best for all of the opponent's mixtures. Then this player's best response would always be that pure strategy.

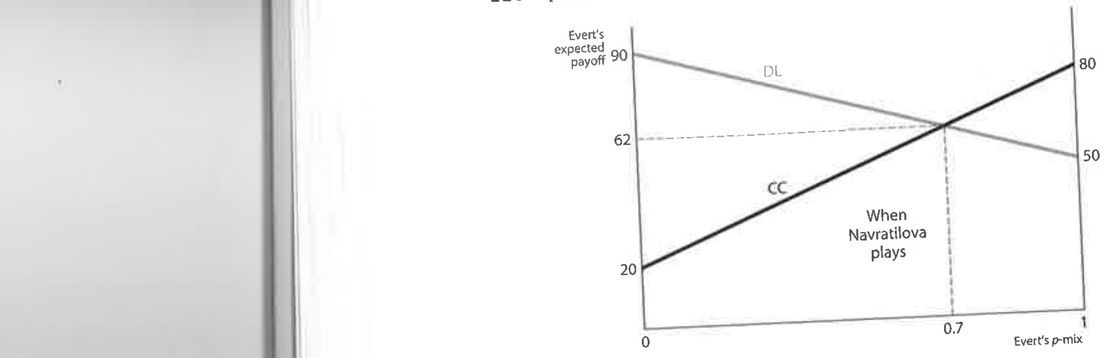

FIGURE 7.3 Navratilova's Best-Response Calculation Using Evert's Payoffs

We can see this even more clearly by graphing Evert's own expected payoffs for each of her p-mixes. We just have to remember that in the constant-sum game, Navratilova does better when Evert gets low payoffs. Figure 7.3 puts Evert's p-mix on the horizontal axis as in Figure 7.2, but Evert's own expected payoffs are on the vertical axis. Since Evert's payoffs are 100 minus Navratilova's payoffs, Figure 7.3 is a vertically inverted version of Figure 7.2. We show two lines again, one corresponding to Navratilova's choice of pure CC and the other for her choice of pure DL. Here, the CC line is rising from left to right. This corresponds to the idea that Evert's expected payoff rises as she uses her DL shot more often (higher p) against Navratilova's covering purely CC. Similarly, the DL line is falling from left to right, as does Evert's expected payoff from using DL more often against an opponent who is always covering for that shot. We see that the point of intersection of the expected payoff lines is again at $p = 0.7$. To the left of this p value, the CC line is lower than the DL line. Therefore, CC is better than DL for Navratilova when p is low; it entails a lower expected payoff for Evert. To the right of $p = 0.7$, the DL line is lower than the CC line, so DL is better than CC for Navratilova. And the expected payoff at the point of intersection is 62 for Evert, which is just 100 minus the 38 that we found for Navratilova in the analysis based on her payoffs. Further, for each of Evert's p-mixes, Navratilova's best response will hold Evert down to the lower of the expected payoffs along the two lines. These best responses make an inverted V in this case, and Evert's highest expected payoff comes at the vertex of the inverted V, that is, at the point of intersection of the two straight lines. Thus, the exploitation-proof mix is indeed Evert's best choice. Our results are the same as we found by using Navratilova's own payoffs; only the basis of the analysis is slightly different.

Let us recapitulate the argument for finding Evert's exploitation-proof mixture. For each p that Evert could choose for her p-mix, we find the lowest payoff to which Navratilova can hold her down by choosing the appropriate response. Then we find the p that gives the highest of these lowest payoffs. This matches exactly the minimax reasoning that we developed for pure strategies in Chapter 4, Section 5, and it has the same justification here. Navratilova is going to choose the strategy that is best for her. The game is zero sum; what is best for Navratilova is worst for Evert. Therefore Evert would be correct to entertain a pessimistic belief about Navratilova's best responses, and that is exactly what minimax reasoning does. We develop this idea in a little more detail in subsection C below.

We can find Navratilova's best mixture by using exactly the same reasoning. First you will need to identify Evert's best choice of p against each of Navratilova's q-mixes. We leave the analysis to you but provide the result. Evert's best-response rule is:

If $q < 0.6$, choose pure DL ($p = 1$).
If $q = 0.6$, all values of p from 0 to 1 are equal best responses.
If $q > 0.6$, choose pure CC ($p = 0$).

To understand the intuition, observe that q is the probability of Navratilova defending down the line. Evert wants to hit her passing shot to the side that Navratilova is not defending. Therefore, when q is low, Evert does better to go down the line and, when q is high, she does better to go crosscourt. Following the same reasoning as we gave above from Evert's perspective, $q = 0.6$ is Navratilova's best mixture, immune to exploitation by Evert.

B. Equilibrium in Mixed Strategies

When each player chooses her best, exploitation-proof, mixture, the outcome will be an equilibrium in mixed strategies. In the next chapter we will verify formally that it is indeed a Nash equilibrium in continuous strategies, regarding the mixture probabilities as the continuous variables in the choices of the two players' strategies. There we will also develop more general theory of mixed strategies and their equilibria when each player has several pure strategies. Here we merely state the procedure you have to follow to calculate such an equilibrium in two-by-two games.

Let the Row player's general mixed strategy be her p-mix, putting probability p on playing her first pure strategy and probability $(1 - p)$ on her second pure strategy. Given this p-mix, calculate the algebraic expressions for the Column player's expected payoffs from each of her pure strategies. As in Section 2.A, these expressions will depend on the parameter p. Then equate the two expressions, and solve the resulting equation for p. This p defines the Row player's optimal, exploitation-proof mixture.

Similarly, for the Column player, let her q-mix put probability q on playing her first pure strategy and probability $(1 - q)$ on her second pure strategy. Find the q that defines the Column player's best q-mix by equating the expressions for the expected payoffs the Row player would get from each of her pure strategies against this q-mix.[5]

We illustrate this method using the tennis point example. We have already seen that against Evert's p-mix, Navratilova's expected payoff is $50p + 10(1 - p)$ if she covers DL, and $20p + 80(1 - p)$ if she covers CC. Equating these two we have

$$50p + 10(1 - p) = 20p + 80(1 - p),$$
$$\text{or} \quad 30p = 70(1 - p), \quad \text{or} \quad 100p = 70, \quad \text{or} \quad p = 0.7.$$

Against Navratilova's q-mix, Evert's expected payoffs are similarly found using Figure 7.1; they are $50q + 80(1 - q)$ if Evert plays DL, and $90q + 20(1 - q)$ if she plays CC. Equating the two, we have

$$50q + 80(1 - q) = 90q + 20(1 - q),$$
$$\text{or} \quad 60(1 - q) = 40q, \quad \text{or} \quad 100q = 60, \quad \text{or} \quad q = 0.6.$$

These values, $p = 0.7$ and $q = 0.6$, describe the equilibrium in mixed strategies in the tennis point game.

Each player's equilibrium (best) mixture probability is found by equating the other player's expected payoffs for her two pure strategies. Thus, the equilibrium mixture of each player satisfies the condition that the other player should be indifferent between her two pure strategies. We call this the **opponent's indifference property** of mixed strategy equilibria.

C. The Minimax Method

In Chapter 4, Section 5, we developed the minimax reasoning for zero-sum games in which players used only pure strategies. That approach can be used for zero-sum games with no equilibrium in pure strategies as well, if we expand the set of strategies to include mixtures for the two players.

In Figure 7.4, we show Evert's payoffs in the tennis point game with pure strategies alone, but augmented to show row minima and column maxima, as we did in Figure 4.8 for the football game. Evert's minimum payoffs (success values) for each of her strategies are shown at the far right of the table: 50 for the DL row and 20 for the CC row. The maximum of these minima, or Evert's maximin, is 50. This is the best that Evert can guarantee for herself by using pure strategies,

[5] If your solution yields values of p or q that are outside the range for probabilities—that is, if one or both of the solutions is either less than 0 or greater than 1—then the game has an equilibrium in pure strategies. Inspect the original payoff matrix to find it.

		NAVRATILOVA		
		DL	CC	
EVERT	DL	50	80	min = 50
	CC	90	20	min = 20
		max = 90	max = 80	

FIGURE 7.4 Minimax Analysis of the Tennis Point with Pure Strategies

knowing that Navratilova responds in her own best interests. To achieve this guaranteed payoff, Evert would play DL.

At the foot of each column, we show the maximum of Evert's success percentages attainable in that column. This value represents the worst payoff that Navratilova can get if she plays the pure strategy corresponding to that column. The smallest of these maxima, or the minimax, is 80. This is the best payoff that Navratilova can guarantee for herself by using pure strategies, knowing that Evert responds in her own best interests. To achieve this, Navratilova would play CC.

The maximin and minimax values for the two players are not the same: Evert's maximin success percentage (50) is less than Navratilova's minimax (80). As we explained in Chapter 4, this shows that the game has no equilibrium in the (pure) strategies available to the players in this analysis. In this game, each player can achieve a better outcome—a higher maximin value for Evert and a lower minimax value for Navratilova—by choosing a suitable random mixture of DL and CC.

Let us expand the sets of strategies available to players to include randomization of moves or mixed strategies. We do the analysis here from Evert's perspective; that for Navratilova is similar, and we leave it for you as an exercise. First, in Figure 7.5 we expand the payoff table of Figure 7.4 by adding a row for Evert's p-mix, where she plays DL with probability p and CC with probability $(1 - p)$. We show the expected payoffs of this mix against each of Navratilova's

		NAVRATILOVA		
		DL	CC	
EVERT	DL	50	80	min = 50
	CC	90	20	min = 20
	p-mix	$50p + 90(1 - p)$	$80p + 20(1 - p)$	min = ?

FIGURE 7.5 Minimax Analysis of the Tennis Point with Evert's Mixed Strategies

pure strategies, DL and CC. The row minimum for the p-mix is then just the smaller of these two expressions. But which expression is smaller depends on the value of p, so we leave it as a question mark there.

To find the correct value for p, turn to Figure 7.3, which graphs each of the two payoff expressions from the last row of the table in Figure 7.5 against p. When $p < 0.7$, the minimum payoff occurs along the line corresponding to Navratilova's CC, and when $p > 0.7$, it occurs along the line corresponding to Navratilova's DL. The maximum of these minima for Evert occurs at $p = 0.7$, namely the peak of the inverted V formed by the lower segments of the two lines.[6] Evert's expected payoff at this point is 62. Thus suitable mixing enables Evert to raise her maximin from 50 to 62.

A similar analysis done from Navratilova's perspective shows that suitable mixing (choice of an exploitation-proof q) enables her to lower her minimax. We already calculated this mix; it is $q = 0.6$. The resulting expected payoff to Evert is $50 \times 0.6 + 80 \times 0.4 = 30 + 32 = 62$. Thus Navratilova's best mixing enables her to lower her minimax from 80 to 62.

When the two best mixes are pitted against each other, Evert's maximin equals Navratilova's minimax, and we have a Nash equilibrium in mixed strategies. The equality of maximin and minimax for optimal mixes is a general property of zero-sum games and was proved by John von Neumann and Oskar Morgenstern in 1945.

3 NASH EQUILIBRIUM AS A SYSTEM OF BELIEFS AND RESPONSES

When the moves in a game are simultaneous, neither player can respond to the other's actual choice. Instead, each takes her best action in light of what she thinks the other might be choosing at that instant. In Chapter 4, we called such thinking a player's belief about the other's strategy choice. We then interpreted Nash equilibrium as a configuration where such beliefs are correct, so each chooses her best response to the actual actions of the other. This concept proved useful for understanding the structures and outcomes of many important types of games, most notably zero-sum games and minimax strategies, dominance and the prisoners' dilemma, and focal points and various coordination games, as well as chicken.

[6] We could have replaced the question mark in the table by the whole graph of the inverted V, but that would be complicated to do and difficult to read.

However, in Chapter 4 we considered only pure-strategy Nash equilibria. Therefore a hidden assumption went almost unremarked—namely, that each player was sure or confident in her belief that the other would choose a particular pure strategy. Now that we are considering more general mixed strategies, the concept of belief requires a corresponding reinterpretation.

Players may be unsure about what others might be doing. In the coordination game in Chapter 5, in which Harry wanted to meet Sally, he might be unsure whether she would go to Starbucks or Local Latte, and his belief might be that there was a 50–50 chance that she would go to either one. And in the tennis example, Evert might recognize that Navratilova was trying to keep her (Evert) guessing and would therefore be unsure of which of Navratilova's available actions she would play. In Chapter 2, Section 4, we labeled this as strategic uncertainty, and in Chapter 4 we mentioned that such uncertainty can give rise to mixed-strategy equilibria. Now we develop this idea more fully.

It is important, however, to distinguish between being unsure and having incorrect beliefs. For example, in the tennis example, Navratilova cannot be sure of what Evert is choosing on any one occasion. But she can still have correct beliefs about Evert's mixture—namely, about the probabilities with which Evert chooses between her two pure strategies. Having correct beliefs about mixed actions means knowing or calculating or guessing the correct probabilities with which the other player chooses from among her underlying basic or pure actions. In the equilibrium of our example, it turned out that Evert's equilibrium mixture was 70% DL and 30% CC. If Navratilova believes that Evert will play DL with 70% probability and CC with 30% probability, then her belief, although uncertain, will be correct in equilibrium.

Thus we have an alternative and mathematically equivalent way to define Nash equilibrium in terms of beliefs: each player forms beliefs about the probabilities of the mixture that the other is choosing and chooses her own best response to this. A Nash equilibrium in mixed strategies occurs when the beliefs are correct, in the sense just explained.

In the next section, we consider mixed strategies and their Nash equilibria in non-zero-sum games. In such games, there is no general reason that the other player's pursuit of her own interests should work against your interests. Therefore it is not in general the case that you would want to conceal your intentions from the other player, and there is no general argument in favor of keeping the other player guessing. However, because moves are simultaneous, each player may still be subjectively unsure of what action the other is taking and therefore may have uncertain beliefs that in turn lead her to be unsure about how she should act. This can lead to mixed-strategy equilibria, and their interpretation in terms of subjectively uncertain but correct beliefs proves particularly important.

4 MIXING IN NON-ZERO-SUM GAMES

The same mathematical method used to find mixed-strategy equilibria in zero-sum games—namely, exploitation-proofness or the opponent's indifference property—can be applied to non-zero-sum games as well, and it can reveal mixed strategy equilibria in some of them. However, in such games the players' interests may coincide to some extent. Therefore the fact that the other player will exploit your systematic choice of strategy to her advantage need not work out to your disadvantage, as was the case with zero-sum interactions. In a coordination game of the kind we studied in Chapter 4, for example, the players are better able to coordinate if each can rely on the other's acting systematically; random actions only increase the risk of coordination failure. As a result, mixed-strategy equilibria have a weaker rationale, and sometimes no rationale at all, in non-zero-sum games. Here we examine mixed-strategy equilibria in some prominent non-zero-sum games and discuss their relevance, or lack thereof.

A. Will Harry Meet Sally? Assurance, Pure Coordination, and Battle of the Sexes

We illustrate mixing in non-zero-sum games by using the assurance version of the meeting game. For your convenience, we reproduce its table (Figure 4.12) as Figure 7.6 below. We consider the game from Sally's perspective first. If she is confident that Harry will go to Starbucks, she also should go to Starbucks. If she is confident that Harry will go to Local Latte, so should she. But if she is unsure about Harry's choice, what is her own best choice?

To answer this question, we must give a more precise meaning to the uncertainty in Sally's mind. (The technical term for this uncertainty, in the theory of probability and statistics, is her subjective uncertainty. In the context where the uncertainty is about another player's action in a game, it is also strategic uncertainty; recall the distinctions we discussed in Chapter 2, Section 2.D.) We gain precision by stipulating the probability with which Sally thinks Harry will

		SALLY	
		Starbucks	Local Latte
HARRY	Starbucks	1, 1	0, 0
	Local Latte	0, 0	2, 2

FIGURE 7.6 Assurance Game

choose one café or the other. The probability of his choosing Local Latte can be any real number between 0 and 1 (that is, between 0% and 100%). We can cover all possible cases by using algebra, letting the symbol p denote the probability (in Sally's mind) that Harry chooses Starbucks; the variable p can take on any real value between 0 and 1. Then $(1 - p)$ is the probability (again in Sally's mind) that Harry chooses Local Latte. In other words, we describe Sally's strategic uncertainty as follows: she thinks that Harry is using a mixed strategy, mixing the two pure strategies, Starbucks and Local Latte, in proportions or probabilities p and $(1 - p)$ respectively. We call this mixed strategy Harry's p-mix, even though for the moment it is purely an idea in Sally's mind.

Given her uncertainty, Sally can calculate the expected payoffs from her actions when they are played against her belief about Harry's p-mix. If she chooses Starbucks, it will yield her $1 \times p + 0 \times (1 - p) = p$. If she chooses Local Latte, it will yield her $0 \times p + 2 \times (1 - p) = 2 \times (1 - p)$. When p is high, $p > 2(1 - p)$, so that Sally is fairly sure that Harry is going to Starbucks, then she does better by also going to Starbucks. Similarly, if p is low, $p < 2(1 - p)$ and if Sally is fairly sure that Harry is going to Local Latte, then she does better by going to Local Latte. If $p = 2(1 - p)$, or $3p = 2$, or $p = 2/3$, the two choices give Sally the same expected payoff. Therefore if she believes that $p = 2/3$, she might be unsure about her own choice, so she might dither between the two.

Harry can figure this out, and that makes him unsure about Sally's choice. Thus Harry also faces subjective strategic uncertainty. Suppose in his mind Sally will choose Starbucks with probability q and Local Latte with probability $(1 - q)$. Similar reasoning shows that Harry should choose Starbucks if $q > 2/3$ and Local Latte if $q < 2/3$. If $q = 2/3$, he will be indifferent between the two actions and unsure about his own choice.

Now we have the basis for a mixed-strategy equilibrium with $p = 2/3$ and $q = 2/3$. In such an equilibrium, these p and q values are simultaneously the actual mixture probabilities and the subjective beliefs of each player about the other's mixture probabilities. The correct beliefs sustain each player's own indifference between the two pure strategies and therefore each player's willingness to mix between the two. This matches exactly the concept of a Nash equilibrium as a system of self-fulfilling beliefs and responses described in Section 3.

The key to finding the mixed-strategy equilibrium is that Sally is willing to mix between her two pure strategies only if her subjective uncertainty about Harry's choice is just right—that is, if the value of p in Harry's p-mix is just right. Algebraically, this idea is borne out by solving for the equilibrium value of p by using the equation $p = 2(1 - p)$, which ensures that Sally gets the same expected payoff from her two pure strategies when each is matched against Harry's p-mix. When the equation holds in equilibrium, it is as if Harry's mixture probabilities are doing the job of keeping Sally indifferent. We emphasize the "as if" because in this game, Harry has no reason to keep Sally indifferent; the outcome is merely a property of the equilibrium. Still, the general idea is worth remembering: in a mixed-strategy Nash equilibrium, each person's mixture probabilities keep the other player indifferent between her pure strategies. We called this the opponent's indifference method in the zero-sum discussion above, and now we see that it remains valid even in non-zero-sum games.

However, the mixed-strategy equilibrium has some very undesirable properties in the assurance game. First, it yields both players rather low expected payoffs. The formulas for Sally's expected payoffs from her two actions, p and $2(1 - p)$, both equal 2/3 when $p = 2/3$. Similarly, Harry's expected payoffs against Sally's equilibrium q-mix for $q = 2/3$ are also both 2/3. Thus, each player gets 2/3 in the mixed-strategy equilibrium. In Chapter 4 we found two pure strategy equilibria for this game; even the worse of them (both choosing Starbucks) yields the players 1 each, and the better one (both choosing Local Latte) yields them 2 each.

The reason the two players fare so badly in the mixed-strategy equilibrium is that when they choose their actions independently and randomly, they create a significant probability of going to different places; when that happens, they do not meet, and each gets a payoff of 0. Harry and Sally fail to meet if one goes to Starbucks and the other goes to Local Latte, or vice versa. The probability of this happening when both are using their equilibrium mixtures is $2 \times (2/3) \times (1/3) = 4/9$.[7] Similar problems exist in the mixed-strategy equilibria of most non-zero-sum games.

A second undesirable property of the mixed-strategy equilibrium here is that it is very fragile. If either player departs ever so slightly from the exact values $p = 2/3$ or $q = 2/3$, the best choice of the other tips to one pure strategy. Once one player chooses a pure strategy, then the other also does better by choosing the same pure strategy, and play moves to one of the two pure-strategy equilibria. This instability of mixed-strategy equilibria is common to many non-zero-sum games. However, some important non-zero-sum games do have mixed-strategy equilibria that are not so fragile. One example considered later in this chapter and in Chapter 13 is the mixed-strategy equilibrium in Chicken, which has an interesting evolutionary interpretation.

Given the analysis of the mixed-strategy equilibrium in the assurance version of the meeting game, you can now probably guess the mixed-strategy

[7]The probability that each chooses Starbucks in equilibrium is 2/3. The probability that each chooses Local Latte is 1/3. The probability that one chooses Starbucks while the other chooses Local Latte is $2/3 \times 1/3$. But that can happen two different ways (once when Harry chooses Starbucks and Sally chooses Local Latte, and again when the choices are reversed) so the total probability of not meeting is $2 \times 2/3 \times 1/3$. See the Appendix to this chapter for more details on the algebra of probabilities.

equilibria for the related non-zero-sum meeting games. In the pure-coordination version (see Figure 4.11), the payoffs from meeting in the two cafés are the same; so the mixed-strategy equilibrium will have $p = 1/2$ and $q = 1/2$. In the battle-of-the-sexes variant (see Figure 4.13), Sally prefers to meet at Local Latte because her payoff is 2 rather than the 1 that she gets from meeting at Starbucks. Her decision hinges on whether her subjective probability of Harry's going to Starbucks is greater than or less than 2/3. (Sally's payoffs here are similar to those in the assurance version, so the critical p is the same.) Harry prefers to meet at Starbucks, so his decision hinges on whether his subjective probability of Sally's going to Starbucks is greater than or less than 1/3. Therefore the mixed-strategy Nash equilibrium has $p = 2/3$ and $q = 1/3$.

B. Will James Meet Dean? Chicken

The non-zero-sum game of Chicken also has a mixed-strategy equilibrium that can be found using the same method developed above, although its interpretations are slightly different. Recall that this is a game between James and Dean, who are trying to *avoid* a meeting; the game table, originally introduced in Figure 4.14, is reproduced here as Figure 7.7.

If we introduce mixed strategies, James's p-mix will entail a probability p of swerving and a probability $1 - p$ of going straight. Against that p-mix, Dean gets $0 \times p - 1 \times (1 - p) = p - 1$ if he chooses Swerve and $1 \times p - 2 \times (1 - p) = 3p - 2$ if he chooses Straight. Comparing the two, we see that Dean does better by choosing swerve when $p - 1 > 3p - 2$, or when $2p < 1$, or when $p < 1/2$, that is, when p is low and James is more likely to choose Straight. Conversely, when p is high and James is more likely to choose Swerve, then Dean does better by choosing Straight. If James's p-mix has p exactly equal to 1/2, then Dean is indifferent between his two pure actions; he is therefore equally willing to mix between the two. Similar analysis of the game from James's perspective when considering his options against Dean's q-mix yields the same results. Therefore $p = 1/2$ and $q = 1/2$ is a mixed-strategy equilibrium of this game.

		DEAN	
		Swerve (Chicken)	Straight (Tough)
JAMES	Swerve (Chicken)	0, 0	–1, 1
	Straight (Tough)	1, –1	–2, –2

FIGURE 7.7 Chicken

The properties of this equilibrium have some similarities but also some differences when compared with the mixed-strategy equilibria of the meeting game. Here, each player's expected payoff in the mixed-strategy equilibrium is low (−1/2). This is bad, as was the case in the meeting game, but unlike in that game, the mixed-strategy equilibrium payoff is not worse for both players than either of the two pure-strategy equilibria. In fact, because player interests are somewhat opposed here, each player will do strictly better in the mixed-strategy equilibrium than in the pure-strategy equilibrium that entails his choosing Swerve.

This mixed-strategy equilibrium is again unstable, however. If James increases his probability of choosing Straight to just slightly above 1/2, this change tips Dean's choice to pure Swerve. Then (Straight, Swerve) becomes the pure-strategy equilibrium. If James instead lowers his probability of choosing Straight slightly below 1/2, Dean chooses Straight, and the game goes to the other pure-strategy equilibrium.[8]

5 GENERAL DISCUSSION OF MIXED-STRATEGY EQUILIBRIA

Now that we have seen how to find mixed-strategy equilibria in both zero-sum and non-zero-sum games, it is worthwhile to consider some additional features of these equilibria. In particular, we highlight in this section some general properties of mixed-strategy equilibria. We also introduce you to some results that seem counterintuitive at first, until you fully analyze the game in question.

A. Weak Sense of Equilibrium

The opponent's indifference property described in Section 2 implies that in a mixed-strategy equilibrium, each player gets the same expected payoff from each of her two pure strategies, and therefore also gets the same expected payoff from any mixture between them. Thus mixed-strategy equilibria are Nash equilibria only in a weak sense. When one player is choosing her equilibrium mix, the other has no positive reason to deviate from her own equilibrium mix. But she would not do any worse if she chose another mix or even one of her pure strategies. Each player is indifferent between her pure strategies, or indeed between any mixture of them, so long as the other player is playing her correct (equilibrium) mix. This is also a very general property of mixed-strategy Nash equilibria.

[8]In Chapter 13 we consider a different kind of stability, namely evolutionary stability. The question in the evolutionary context is whether a stable mix of Straight and Swerve choosers can arise and persist in a population of Chicken players. The answer is yes, and the proportions of the two types are exactly equal to the probabilities of playing each action in the mixed-strategy equilibrium. Thus, we derive a new and different motivation for that equilibrium in this game.

This property seems to undermine the basis for mixed-strategy Nash equilibria as the solution concept for games. Why should a player choose her appropriate mixture when the other player is choosing her own? Why not just do the simpler thing by choosing one of her pure strategies? After all, the expected payoff is the same. The answer is that to do so would not be a Nash equilibrium; it would not be a stable outcome, because then the other player would not choose to use her mixture. For example, if Harry chooses pure Starbucks in the assurance version of the meeting game, then Sally can get a higher payoff in equilibrium (1 instead of 2/3) by switching from her 50-50 mix to her pure Starbucks as well.

B. Counterintuitive Changes in Mixture Probabilities in Zero-Sum Games

Games with mixed-strategy equilibria may exhibit some features that seem counterintuitive at first glance. The most interesting of them is the change in the equilibrium mixes that follow a change in the structure of a game's payoffs. To illustrate, we return to Evert and Navratilova and their tennis point.

Suppose that Navratilova works on improving her skills covering down the line to the point where Evert's success using her DL strategy against Navratilova's covering DL drops to 30% from 50%. This improvement in Navratilova's skill alters the payoff table, including the mixed strategies for each player, from that illustrated in Figure 7.1. We present the new table in Figure 7.8.

The only change from the table in Figure 7.1 has occurred in the upper-left-hand cell, where our earlier 50 for Evert is now a 30 and the 50 for Navratilova is now a 70. This change in the payoff table does not lead to a game with a pure-strategy equilibrium, because the players still have opposing interests; Navratilova still wants their choices to coincide, and Evert still wants their choices to differ. We still have a game in which mixing will occur.

But how will the equilibrium mixes in this new game differ from those calculated in Section 2? At first glance, many people would argue that Navratilova should cover DL more often now that she has gotten so much better at doing so. Thus, the assumption is that her equilibrium q-mix should be more heavily weighted toward DL and her equilibrium q should be higher than the 0.6 calculated before.

		NAVRATILOVA	
		DL	CC
EVERT	DL	30, 70	80, 20
	CC	90, 10	20, 80

FIGURE 7.8 Changed Payoffs in the Tennis Point

But when we calculate Navratilova's q-mix by using the condition of Evert's indifference between her two pure strategies, we get $30q + 80(1 - q) = 90q + 20(1 - q)$, or $q = 0.5$. The actual equilibrium value for q, 50%, has exactly the opposite relation to the original q of 60% than what many people's intuition predicts.

Although the intuition seems reasonable, it misses an important aspect of the theory of strategy: the interaction between the two players. Evert will also be reassessing her equilibrium mix after the change in payoffs, and Navratilova must take the new payoff structure *and* Evert's behavior into account when determining her new mix. Specifically, because Navratilova is now so much better at covering DL, Evert uses CC more often in her mix. To counter that, Navratilova covers CC more often, too.

We can see this more explicitly by calculating Evert's new mixture. Her equilibrium p must equate Navrtilova's expected payoff from covering DL, $30p + 90(1 - p)$, with her expected payoff from covering CC, $80p + 20(1 - p)$. So we have $30p + 90(1 - p) = 80p + 20(1 - p)$, or $90 - 60p = 20 + 60p$, or $120p = 70$. Thus, Evert's p must be 7/12, which is 0.583, or 58.3%. Comparing this new equilibrium p with the original 70% calculated in Section 2 shows that Evert has significantly decreased the number of times she sends her shot DL in response to Navratilova's improved skills. Evert has taken into account the fact that she is now facing an opponent with better DL coverage, and so she does better to play DL less frequently in her mixture. By virtue of this behavior, Evert makes it better for Navratilova also to decrease the frequency of her DL play. Evert would now exploit any other choice of mix by Navratilova, in particular a mix heavily favoring DL.

So is Navratilova's skill improvement wasted? No, but we must judge it properly—not by how often one strategy or the other gets used but by the resulting payoffs. When Navratilova uses her new equilibrium mix with $q = 0.5$, Evert's success percentage from either of her pure strategies is $(30 \times 0.5) + (80 \times 0.5) = (90 \times 0.5) + (20 \times 0.5) = 55$. This is less than Evert's success percentage of 62 in the original example. Thus, Navratilova's average payoff also rises from 38 to 45, and she does benefit by improving her DL coverage.

Unlike the counterintuitive result that we saw when we considered Navratilova's strategic response to the change in payoffs, we see here that her response is absolutely intuitive when considered in light of her expected payoff. In fact, players' expected payoff responses to changed payoffs can never be counterintuitive, although strategic responses, as we have seen, can be.[9] The most interesting aspect of this counterintuitive outcome in players' strategic responses is the message that it sends to tennis players and to strategic game players more

[9] For a general theory of the effect that changing the payoff in a particular cell has on the equilibrium mixture and the expected payoffs in equilibrium, see Vincent Crawford and Dennis Smallwood, "Comparative Statics of Mixed-Strategy Equilibria in Noncooperative Games," *Theory and Decision*, vol. 16 (May 1984), pp. 225–232.

generally. The result here is equivalent to saying that Navratilova should improve her down-the-line coverage so that she does not have to use it so often.

There are other similar examples of possibly counterintuitive results, but they require slightly more ease with algebra than many readers may have. Therefore we postpone them to the next, optional, chapter (Section 8.6).

6 HOW TO USE MIXED STRATEGIES IN PRACTICE

There are several important things to remember when finding or using a mixed strategy in a zero-sum game. First, to use a mixed strategy effectively in such a game, a player needs to do more than calculate the equilibrium percentages with which to use each of her actions. Indeed, in our tennis-point game, Evert cannot simply play DL seven-tenths of the time and CC three-tenths of the time by mechanically rotating seven shots down the line and three shots cross court. Why not? Because mixing your strategies is supposed to help you benefit from the element of surprise against your opponent. If you use a recognizable pattern of plays, your opponent is sure to discover it and exploit it to her advantage.

The lack of a pattern means that, after any history of choices, the probability of choosing DL or CC on the next turn is the same as it always was. If a run of several successive DLs happens by chance, there is no sense in which CC is now "due" on the next turn. In practice, many people mistakenly think otherwise, and therefore they alternate their choices too much compared with what a truly random sequence of choices would require; they produce too few runs of identical successive choices. However, detecting a pattern from observed actions is a tricky statistical exercise that the opponents may not be able to perform while playing the game. As we will see in Section 7, analysis of data from grand-slam tennis finals found that servers alternated their serves too much, but receivers were not able to detect this departure from true randomization.

However, to make sure that your opponent cannot exploit your mixed strategy, you need a truly random pattern of actions on each play of the game. For example, you may want to rely on a computer's ability to generate random numbers for you, from which you can determine your appropriate choice of action. If the computer generates numbers between 1 and 100 and you want to mix pure strategies A and B in a 60-40 split, you may decide to play A for any random number between 1 and 60 and to play B for any random number between 61 and 100. Similarly, you could employ a device such as the color-coded spinners provided in many children's games. For the same 60-40 mixture, you would color 60% of the circle on the spinner in blue, for example, and 40% in red; the first 216 degrees of the circle would be blue, the remaining 144 red. Then you would spin the spinner arrow and play A if it landed in blue but B if it landed in red. You can use the second hand on a watch as the same type of device, but it is important that your watch not be so accurate and synchronized that your opponent can use the same watch and figure out what you are going to do.

The importance of avoiding a predictable system of randomization is clearest in ongoing interactions of a zero-sum nature. Because of the diametrically opposed interests of the players in such games, your opponent always benefits from exploiting your choice of action to the greatest degree possible. Thus, if you play the same game against each other on a regular basis, she will always be on the lookout for ways to break the code that you are using to randomize your moves. If she can do so, she has a chance to improve her payoffs in future plays of the game. Mixing is still justified in single-meet (sometimes called one-shot) zero-sum games because the benefit of tactical surprise remains important.

Finally, players must understand and accept the fact that the use of mixed strategies guards you against exploitation, and gives the best possible expected payoff against an opponent who is making her best choices, but that it is only a probabilistic average. On particular occasions, you can get poor outcomes. For example, the long pass on third down with a yard to go, intended to keep the defense honest, may fail on any specific occasion. If you use a mixed strategy in a situation in which you are responsible to a higher authority, therefore, you may need to plan ahead for this possibility. You may need to justify your use of such a strategy ahead of time to your coach or your boss, for example. They need to understand why you have adopted your mixture and why you expect it to yield you the best possible payoff on average, even though it might yield an occasional low payoff as well. Even such advance planning may not work to protect your "reputation," though, and you should prepare yourself for criticism in the face of a bad outcome.

7 EVIDENCE ON MIXING

A. Zero-Sum Games

Early researchers who performed laboratory experiments were generally dismissive of mixed strategies. To quote Douglas Davis and Charles Holt, "Subjects in experiments are rarely (if ever) observed flipping coins, and when told ex post that the equilibrium involves randomization, subjects have expressed surprise and skepticism."[10] When the predicted equilibrium entails mixing two or more pure strategies, experimental results do show some subjects in the group pursuing one of the pure strategies and others pursuing another, but this does not

[10]Douglas D. Davis and Charles A. Holt, *Experimental Economics* (Princeton: Princeton University Press, 1993), p. 99.

constitute true mixing by an individual player. When subjects play zero-sum games repeatedly, individual players often choose different pure strategies over time. But they seem to mistake alternation for randomization, that is, they switch their choices more often than true randomization would require.

Later research has found somewhat better evidence for mixing in zero-sum games. When laboratory subjects are allowed to acquire a lot of experience, they do appear to learn mixing in zero-sum games. However, departures from equilibrium predictions remain significant. To quote Camerer, "The overall picture is that mixed equilibria do not provide bad guesses about how people behave, on average."[11]

An instance of randomization in practice comes from Malaya in the late 1940s.[12] The British army escorted convoys of food trucks to protect the trucks from communist terrorist attacks. The terrorists could either launch a large-scale attack or create a smaller sniping incident intended to frighten the truck drivers and keep them from serving again. The British escort could be either concentrated or dispersed throughout the convoy. For the army, concentration was better to counter a full-scale attack, and dispersal was better against sniping. For the terrorists, a full-scale attack was better if the army escort was dispersed, and sniping was better if the escort was concentrated. This zero-sum game has only a mixed-strategy equilibrium. The escort commander, who had never heard of game theory, made his decision as follows. Each morning, as the convoy was forming, he took a blade of grass and concealed it in one of his hands, holding both hands behind his back. Then he asked one of his troops to guess which hand held the blade, and he chose the form of the convoy according to whether the man guessed correctly. Although the precise payoff numbers are difficult to judge and therefore we cannot say whether 50-50 was the right mixture, the officer had correctly figured out the need for true randomization and the importance of using a fresh randomization procedure every day to avoid falling into a pattern or making too much alternation between the choices

The best evidence in support of mixed strategies in zero-sum games comes from sports, especially from professional sports, in which players accumulate a great deal of experience of such games, and their intrinsic desire to win is buttressed by large financial gains from winning.

Mark Walker and John Wooders examined the serve-and-return play of top-level players at Wimbledon.[13] They model this interaction as a game with two players, the server and the receiver, in which each player has two pure strategies.

[11]For a detailed account and discussion see Chapter 3 of Colin F. Camerer, *Behvioral Game Theory* (Princeton: Princeton University Press, 2003). The quote is from p. 468 of this book.

[12]R. S. Beresford and M. H. Peston, "A Mixed Strategy in Action," *Operations Research*, vol. 6, no. 4 (December 1955), pp. 173–176.

[13]Mark Walker and John Wooders, "Minimax Play at Wimbledon," *American Economic Review*, vol. 91, no. 5 (December 2001), pp. 1521–1538.

The server can serve to the receiver's forehand or backhand, and the receiver can guess to which side the serve will go and move that way. Because serves are so fast at the top levels of men's singles, the receiver cannot react after observing the actual direction of the serve; rather, the receiver must move in anticipation of the serve's direction. Thus, this game has simultaneous moves. Further, because the receiver wants to guess correctly and the server wants to wrong-foot the receiver, this interaction has a mixed-strategy equilibrium.

If the tennis players are using their equilibrium mixtures in the serve-and-return game, the server should win the point with the same probability whether he serves to the receiver's forehand or backhand. An actual tennis match contains a hundred or more points played by the same two players; thus there is enough data to test whether this implication holds. Walker and Wooders tabulated the results of serves in 10 matches. Each match contains four kinds of serve-and-return combinations: A serving to B and vice versa, combined with service from the right or the left side of the court (Deuce or Ad side). Thus they had data on 40 serving situations and found that in 39 of them the server's success rates with forehand and backhand serves were equal to within acceptable limits of statistical error.

The top-level players must have had enough general experience playing the game, as well as particular experience playing against the specific opponents, to have learned the general principle of mixing and the correct proportions to mix against the specific opponents. However, in one respect the servers' choices departed from true mixing. To achieve the necessary unpredictability, there should be no pattern of any kind in a sequence of serves: the choice of side for each serve should be independent of what has gone before. As we said in reference to the practice of mixed strategies, players can alternate too much, not realizing that alternation is a pattern just as much as repeating the same action a few times would be a pattern. And indeed, the data show that the tennis servers alternated too much. But the data also indicate that this departure from true mixing was not enough for the opponents to pick up and exploit.

Penalty kicks in soccer are another good context in which to study mixed strategies. Two such studies find firm support for predictions of the theory.

Kickers usually kick with the inside of the foot. Therefore the natural direction of kicking for a right-footed kicker is to the goalie's right; for a left-footed kicker it is to the goalie's left. For simplicity of writing we will refer to the natural side as "Right." So the choices are Left and Right for each player. When the goalie chooses Right, it means covering the kicker's natural side. Using a large data set from professional soccer leagues in Europe, Ignacio Palacios-Huerta constructed the payoff table of the kicker's average success probabilities shown in Figure 7.9.[14]

[14]See "Professionals Play Minimax," by Ignacio Palacios-Huerta, *Review of Economics Studies*, vol. 70, no. 20 (2003), pp. 395–415.

		GOALIE	
		Left	Right
KICKER	Left	58	95
	Right	93	70

FIGURE 7.9 Soccer Penalty Kick Success Probabilities in European Major Leagues

This game is similar to the tennis-point game and similarly has only a mixed-strategy equilibrium. Using the opponent's indifference property, it is easy to calculate that the kicker *should* choose Left 38.3% of the time and Right 61.7% of the time. This mixture achieves a success rate of 79.6% no matter what the goalie chooses. The goalie *should* choose the probabilities of covering her Left and Right to be 41.7 and 58.3 respectively; this mixture holds the kicker down to a success rate of 79.6%.

What actually happens? Kickers choose Left 40.0% of times, and goalies choose Left 41.3% of the times. These values are startlingly close to the theoretical predictions. The chosen mixtures are almost exploitation proof. The kickers' mix achieves a success rate of 79.0% against the goalie's Left and 80% against the goalie's Right. The goalies' mix holds kickers down to 79.3% if they choose Left and 79.7% if they choose Right.

In an earlier paper, Pierre-Andre Chiappori, Timothy Groseclose, and Steven Levitt used similar data and found similar results.[15] They also analyzed the whole sequence of choices of each kicker and goalie and did not even find too much alternation. Thus these findings suggest that behavior in soccer penalty kicks is even closer to true mixing than behavior in the tennis serve-and-return game.

B. Non-Zero-Sum Games

Laboratory experiments on games with mixed strategies in non-zero-sum games yield even more negative results than experiments involving mixing in zero-sum games. This is not surprising. As we have seen, in such games the property that each player's equilibrium mixture keeps her opponent indifferent among her pure strategies is a logical property of the equilibrium. Unlike in zero-sum games, in general each player in a non-zero-sum game has no positive or

[15]Pierre-André Chiappori, Timothy Groseclose, and Steven Levitt, "Testing Mixed Strategy Equilibria When Players are Heterogeneous: The Case of Penalty Kicks in Soccer," *American Economic Review*, vol. 92, no. 4 (September 2002), pp. 1138–1151.

purposive reason to keep the other players indifferent. Then the reasoning underlying the mixture calculations is more difficult for players to comprehend and learn. This shows up in their behavior.

In a group of experimental subjects playing a non-zero-sum game, we may see some pursuing one pure strategy and others pursuing another. This type of mixing in the population, although it does not fit the theory of mixed-strategy equilibria, does have an interesting evolutionary interpretation, which we examine in Chapter 13.

Other experimental issues concern subjects who play the game many times. When collecting evidence on play in non-zero-sum games, it is important to rotate or randomize each subject's opponents to avoid tacit cooperation in repeated interactions. In such experiments, players change their actions from one play to the next. But each player's mixture probabilities should be generated using the other player's indifference conditon; therefore these probabilities should not change when a player's own payoffs change. (You will find more on this in Chapter 8, Section 6.B.) But in fact they do.[16] Thus the changes of action from one play to the next may not be true mixing, but some other kind of experimentation.

The overall conclusion is that you should interpret and use mixed-strategy equilibria in non-zero-sum games with, at best, considerable caution.

SUMMARY

Zero-sum games in which one player prefers a coincidence of actions and the other prefers the opposite often have no Nash equilibrium in pure strategies. In these games, each player wants to be unpredictable and thus uses a mixed strategy that specifies a probability distribution over her set of pure strategies. Each player's equilibrium mixture probabilities are calculated using the *opponent's indifference property*, namely that the opponent should get equal *expected payoffs* from all her pure strategies when facing the first player's equilibrium mix. In zero-sum games each player wants to keep the other indifferent in this way, since any clear advantage for the opponent would only work to the first player's disadvantage.

Non-zero-sum games can also have mixed-strategy equilibria that can be calculated from the opponent's indifference property. But here the motivation for keeping the opponent indifferent is weaker or missing; therefore such equilibria have less appeal and are often unstable.

When using mixed strategies, players should remember that their system of randomization should not be predictable in any way. Most important, they should avoid excessive alternation of actions.

Mixed strategies are a special case of continuous strategies but have additional matters that deserve separate study. Mixed-strategy equilibria can be interpreted as outcomes in which each player has correct beliefs about the probabilities with which the other player chooses from among her underlying pure actions. And mixed-strategy equilibria may have some counterintuitive properties when payoffs for players change. Laboratory experiments show only weak support for the use of mixed strategies. But mixed-strategy equilibria give good predictions in many zero-sum situations in sports played by experienced professionals.

[16]Jack Ochs, "Games with Unique Mixed-Strategy Equilibria: An Experimental Study," *Games and Economic Behavior*, vol. 10, no. 1 (July 1995), pp. 202–217.

KEY TERMS

expected payoff (215) **opponent's indifference property (222)**

SOLVED EXERCISES

S1. "When a two-by-two game has a mixed-strategy equilibrium, a player's equilibrium mixture is designed to yield her the same expected payoff when used against each of the other player's pure strategies." True or false? Explain and give an example of a game that illustrates your answer.

S2. Consider the following game:

		COLUMN	
		Safe	Risky
ROW	Safe	4, 4	4, 1
	Risky	1, 4	6, 6

(a) Which game does this most resemble: tennis, assurance, or chicken? Explain.

(b) Find all of this game's Nash equilibria.

S3. The following table illustrates the money payoffs associated with a two-person simultaneous-play game.

		COLUMN	
		Left	Right
ROW	Up	1, 16	4, 6
	Down	2, 20	3, 40

(a) Find the Nash equilibrium in mixed strategies for this game.

(b) What are the players' expected payoffs in this equilibrium?

(c) The two players jointly get the most money when Row plays Down. However, in the equilibrium, Row does not always play Down. Why not? Can you think of ways in which a more cooperative outcome can be sustained?

S4. Recall Exercise S6 from Chapter 4, about an old lady looking for help crossing the street. Only one person is needed to help her; more are okay but no better than one. You and I are the two people in the vicinity who can help; each has to choose simultaneously whether to do so. Each of us will get pleasure worth 3 from her success (no matter who helps her). But each one who goes to help will bear a cost of 1, this being the value of our time taken up in helping. You were asked to set this up as a game and to write the payoff table in Exercise S6 of Chapter 4. If you did that exercise, you also found all of the pure-strategy Nash equilibria of the game. Now find the mixed-strategy equilibrium of this game.

S5. Revisit the tennis game in Section 2.A of this chapter. Recall that the mixed-strategy Nash equilibrium found in that section had Evert playing DL with probability 0.7, while Navratilova played DL with probability 0.6.

Later in the match Evert injures herself, so her DL shots are much slower and easier for Navratilova to defend. The payoffs are now given by the table:

		NAVRATILOVA	
		DL	CC
EVERT	DL	30	60
	CC	90	20

(a) Relative to the game before her injury (see Figure 7.1), Evert's payoffs are reduced when she plays DL and Navratilova plays either DL or CC. Overall, DL seems much less attractive to Evert than before. Would you expect Evert to play DL more or less, or stay the same? Explain.

(b) Find each player's equilibrium mixture for the game above. What is the expected value of the game to Evert?

(c) How do the equilibrium mixtures found in part (b) compare with those of the original game? Explain why each has changed or hasn't changed.

S6. Undeterred by their experiences wih chicken so far (see Section 4.B), James and Dean decide to increase the excitement (and the stakes) by starting their cars farther apart. This way they can keep the crowd in suspense longer, and they'll be able to accelerate to even higher speeds before they may or may not be involved in a much more serious collision. The new game table thus has a higher penalty for collision:

		DEAN	
		Swerve	Straight
JAMES	Swerve	0, 0	−1, 1
	Straight	1, −1	−10, −10

(a) What is the mixed-strategy Nash equilibrium for this more dangerous version of chicken? Do James and Dean play Straight more or less often than the game shown in Figure 7.7?

(b) What is the expected payoff to each player in the mixed-strategy equilibrium found in part (a)?

(c) James and Dean decide to play the chicken game repeatedly (say, in front of different crowds of reckless youths). Moreover, because they don't want to collide, they collude. They alternate between the two pure-strategy equilibria, so that half the time they play (Swerve, Straight) and half the time they play (Straight, Swerve). Assuming they play an even number of games, what is the average payoff to each of them when they alternate between the two pure-strategy equilibria? Is this better or worse than they can expect from playing the mixed-strategy equilibrium? Why?

(d) After several weeks of not playing chicken as in part (c), James and Dean agree to play again. However, it has been so long since their last meeting that each of them has entirely forgotten which pure-strategy Nash equilibrium they played last time. Worse still, they don't realize this until they're revving their engines moments before starting the game, and—sadly—they live decades before the advent of cell phones. Instead of playing the mixed-strategy Nash equilibrium, each of them tosses a separate coin to decide which strategy to play. What is the expected payoff to James and Dean when each mixes 50-50 in this game? How does this compare with their expected payoff when they play their equilibrium mixtures? Explain why these payoffs are the same or different from those found in part (c).

S7. Consider the following game:

		COLUMN	
		Left	Right
ROW	Up	2, 1	−1, 4
	Down	0, 3	3, 2

(a) Show that this game has no pure-strategy Nash equilibrium.
(b) Find the unique mixed-strategy Nash equilibrium to this game.
(c) If the payoffs for the cell (Up, Left) were changed to (4, 1), would Row's equilibrium p-mixture increase, decrease, or remain the same? Explain your answer.
(d) What would happen to the equilibrium if the payoffs for the (Up, Left) cell were instead changed to (−1, 1)?

S8. Exercise S7 in Chapter 6 introduced a simplified version of baseball, and part (e) pointed out that the simultaneous-move game has no Nash equilibrium in pure strategies. This is because pitchers and batters have conflicting goals. Pitchers want to get the ball *past* batters, but batters want to *connect* with pitched balls. The game table is as follows:

		PITCHER	
		Throw fastball	Throw curve
BATTER	Anticipate fastball	0.30	0.20
	Anticipate curve	0.15	0.35

(a) Find the mixed-strategy Nash equilibrium to this simplified baseball game.
(b) What is each player's expected payoff for the game?

S9. Extending Exercise S8, the pitcher wants to improve his expected payoff in the mixed-strategy equilibrium of this game by slowing down his fastball, thereby making it more similar to a curve ball. Assume that slowing down the pitched fastball changes the payoff to the hitter in the "anticipate fastball/throw fastball" cell from 0.30 to 0.25, and the pitcher's payoff adjusts accordingly. Can this modification improve the pitcher's expected payoff as desired? Explain carefully how you determine the answer here and show your work. Also, explain *why* slowing the fastball can or cannot improve the pitcher's expected payoff in the game.

S10. In the last minute of a football game, the home team is down by 5 points. The home team has the ball, and it's third down and goal from the rival team's 20-yard line. The home team thus has two chances (third down and fourth down) to move the ball a total of 20 yards and win the game. If it fails, the rival team will win. For each team, the payoff for winning the football game is 1, and the payoff for losing is 0.

On each down the home team can choose to run either a 10-yard play or a 20-yard play. For its part, the rival team can anticipate (and prepare for) either the 10-yard play or the 20-yard play. The success rate of the home team's play given each team's strategy is as follows:

		RIVAL TEAM	
		Anticipate 10-yard play	Anticipate 20-yard play
HOME TEAM	10-yard play	80%	100%
	20-yard play	100%	50%

The success of a particular play is all or nothing. If it succeeds, it yields exactly the number of yards intended, but if it fails, the home team gains 0 yards.

As in Chapter 6, this game combines simultaneous and sequential moves. There are three possible outcomes on third down: the home team can gain 0 yards if its play fails, it can gain 10 yards if a 10-yard play succeeds, or it can gain 20 yards and win the game if a 20-yard play succeeds. On fourth down (if necessary), the home team will thus either have 10 more yards or 20 more yards to go. To solve the larger two-down game, we use rollback and start with the fourth down.

(a) Suppose the home team's third-down play failed, so that there are still 20 yards to go on fourth down. What strategy must the home team play in this situation? Given that, what's the rival team's best response?
(b) Given your answer in part (a), what is the home team's expected payoff when there are 20 yards to go on fourth down?
(c) Suppose the home team ran a successful 10-yard play on third down, so that there are 10 yards to go on fourth down. Since the end zone is 10 yards deep, note that the home team has both strategies available. What is the mixed-strategy Nash equilibrium on fourth down in this situation?
(d) Given your answer in part (c), what is the home team's expected payoff when there are 10 yards to go on fourth down?

Using the expected payoffs for fourth down with 20 yards to go and fourth down with 10 yards to go as calculated above, we now roll back and look at what the home team might do on third down. We construct a table for the simultaneous-move game at third down and 20.

(e) What are the expected payoffs for each team when the home team runs the 20-yard play on third down while the rival team anticipates the 20-yard play? (Use your answer to part (b) and remember that there is a 50% success rate for the 20-yard play when the rival team anticipates the 20-yard play.)

(f) What are the expected payoffs to each team when the home team runs the 10-yard play on third down and the rival team anticipates the 10-yard play on that down? (Use your answer in part (d) and remember that there is an 80% success rate for the 10-yard play when the rival team anticipates the 10-yard play.)

(g) What are the expected payoffs to each team when the home team runs the 10-yard play on third down while the rival team anticipates the 20-yard play?

(h) Now construct the game table for third down with 20 yards to go. (Use your answers from parts (e), (f), and (g).

(i) What are the equilibrium p-mix and q-mix for each team on third down?

(j) What is the expected payoff to the home team for the overall two-stage game?

S11. The recalcitrant James and Dean are playing their more dangerous variant of chicken again (see Exercise S6). They've noticed that their payoff for being perceived as "tough" varies with the size of the crowd. The larger the crowd on hand, the more glory and praise each receives from driving straight when his opponent swerves. Smaller crowds, of course, have the opposite effect. Let $k > 0$ be the payoff for appearing "tough." The game may now be represented as:

		DEAN	
		Swerve	Straight
JAMES	Swerve	0, 0	−1, k
	Straight	k, −1	−10, −10

(a) Expressed in terms of k, with what probability does each driver play Swerve in the mixed-strategy Nash equilibrium? Do James and Dean play Swerve more or less often as k increases?

(b) In terms of k, what is the expected value of the game to each player when both are playing the mixed-strategy Nash equilibrium found in part (a)?

(c) At what value of k do both James and Dean mix 50-50 in the mixed-strategy equilibrium?

(d) How large must k be for the average payoff to be positive under the alternating scheme discussed in part (c) of Exercise S6?

S12. Consider the following zero-sum game:

		COLUMN	
		Left	Right
ROW	Up	0	A
	Down	B	C

The entries are the Row player's payoffs, and the numbers A, B, and C are all positive. What other relations among these numbers (for example, $A < B < C$) must be valid for each of the following cases to arise?

(a) At least one of the players has a dominant strategy.

(b) Neither player has a dominant strategy, but there is a Nash equilibrium in pure strategies.

(c) There is no Nash equilibrium in pure strategies, but there is one in mixed strategies.

(d) Given that case (c) holds, write a formula for Row's probability of choosing Up. Call this probability p, and write it as a function of A, B, and C.

S13. **(Optional)** Return to the game between Evert and Navratilova as shown in Figure 7.1. Suppose that Evert is risk-averse, as discussed in Appendix 2 to this chapter, so that she dislikes uncertainty in outcomes. In particular, suppose that Evert has a square-root utility function, so that her utility equals the square root of the payoff listed in the table, and suppose that Navratilova remains risk neutral, so that her utility equals her payoff. Suppose that the players know each other's utility functions, and each player wishes to maximize her expected utility.

(a) Find the mixed-strategy Nash equilibrium of this game.

(b) How did the two players' mixing proportions change, relative to the original case where both players were risk neutral? Explain why this might have happened.

UNSOLVED EXERCISES

U1. Find the Nash equilibria in mixed strategies for the following games.

(a)

		COLUMN	
		Left	Right
ROW	Up	4	−1
	Down	1	2

(b)

		COLUMN	
		Left	Right
ROW	Up	3	2
	Down	1	4

U2. Exercise S11 in Chapter 4 introduced the game Evens or Odds, which has no Nash equilibrium in pure strategies. There is, however, an equilibrium in mixed strategies.

(a) If Anne plays 1 (that is, she puts in one finger) with probability p, what is the expected payoff to Bruce from playing 1, in terms of p? What is his expected payoff from playing 2?

(b) What level of p will make Bruce indifferent between playing 1 and playing 2?

(c) If Bruce plays 1 with probability q, what level of q will make Anne indifferent between playing 1 and playing 2?

(d) Write the mixed-strategy equilibrium of this game. What is the expected payoff of the game to each player?

U3. In football the offense can either run the ball or pass the ball, whereas the defense can either anticipate (and prepare for) a run or anticipate (and prepare for) a pass. The defense wants to guess correctly to reduce the yards gained by the offense, whereas the offense wants its opponents to guess incorrectly so that it can gain more yards. Assume that the expected payoffs (in yards) for the two teams on any given down are as follows:

		DEFENSE	
		Anticipate Run	Anticipate Pass
OFFENSE	Run	1	5
	Pass	9	−3

(a) Show that this game has no pure-strategy Nash equilibrium.

(b) Find the unique mixed-strategy Nash equilibrium to this game.

(c) Explain why the mixture used by the offense is different from the mixture used by the defense.

(d) How many yards is the offense expected to gain per down in equilibrium?

U4. On the eve of a problem-set due date, a professor receives an e-mail from one of her students, who claims to be stuck on one of the problems after working on it for more than an hour. The professor would rather help the student if he has sincerely been working at the problem, but she would rather not render aid if the student is just fishing for hints. Given the timing of the request, she could simply pretend not to have read the e-mail until later. Obviously, the student would rather receive help whether or not he has been working on the problem. But if help isn't coming, he would rather be working instead of slacking, since the problem set *is* due the next day. Assume the payoffs are as follows:

		STUDENT	
		Work and ask for help	Slack and fish for hints
PROFESSOR	Help student	3, 3	−1, 4
	Ignore e-mail	−2, 1	0, 0

(a) What is the mixed-strategy Nash equilibrium to this game?

(b) What is the expected payoff to each of the players?

U5. Return again to the tennis rivals Evert and Navratilova, discussed in Section 2.A. Months later, they meet again in a new tournament. Evert has healed from her injury (see Exercise S5), but during that same time Navratilova has worked very hard on improving her defense against DL serves. The payoffs of the game are now given by the table:

		NAVRATILOVA	
		DL	CC
EVERT	DL	25	80
	CC	90	20

(a) Find each player's equilibrium mixture for the game above.
(b) What happened to Evert's p-mixture compared to the game presented in Section 2.A? Why?
(c) What is the expected value of the game to Evert? Why is it different from the expected value of the original game in Section 2.A?

U6. Section 4.A of this chapter discussed mixing in the battle of the sexes game between Harry and Sally.
(a) What do you expect to happen to the equilibrium values of p and q found in the chapter if Sally decides she really likes Local Latte a lot more than Starbucks, so that the payoffs in the Local Latte, Local Latte cell are now 1, 3? Explain your reasoning.
(b) Now find the new mixed-strategy equilibrium. What are the new equilibrium values of p and q? How do they compare with those of the original game?
(c) What is the expected payoff to each player in the new mixed-strategy equilibrium?
(d) Do you think Harry and Sally might play the mixed-strategy equilibrium in this new version of the game? Explain why or why not.

U7. Consider the following variant of chicken, in which James's payoff from being "tough" when Dean is "chicken" is 2, rather than 1.

		DEAN	
		Swerve	Straight
JAMES	Swerve	0, 0	−1, 1
	Straight	2, −1	−2, −2

(a) Find the mixed-strategy equilibrium in this game, including the expected payoffs for the players.
(b) Compare the results with those of the original game in Section 4.B of this chapter. Is Dean's probability of playing Straight (being tough) higher now than before? What about James's probability of playing Straight?
(c) What has happened to the two players' expected payoffs? Are these differences in the equilibrium outcomes paradoxical in light of the new payoff structure? Explain how your findings can be understood in light of the opponent's-indifference principle.

U8. Lucy offers to play the following game with Charlie: "Let us show pennies to each other, each choosing either heads or tails. If we both show heads, I pay you \$3. If we both show tails, I pay you \$1. If the two don't match, you pay me \$2." Charlie reasons as follows. "The probability of both heads is 1-4, in which case I get \$3. The probability of both tails is 1-4, in which case I get \$1. The probability of no match is 1-2, and in that case I pay \$2. So it is a fair game." Is he right? If not, (a) why not, and (b) what is Lucy's expected profit from the game?

U9. Consider the following game:

		COLUMN	
		Yes	No
ROW	Yes	x, x	0, 1
	No	1, 0	1, 1

(a) For what values of x does this game have a unique Nash equilibrium? What is that equilibrium?
(b) For what values of x does this game have a mixed-strategy Nash equilibrium? With what probability, expressed in terms of x, does each player play Yes in this mixed-strategy equilibrium?
(c) For the values of x found in part (b), is the game an example of an assurance game, a game of chicken, or a game similar to tennis? Explain.

U10. Consider the following game:

		COLUMN	
		L	R
ROW	U	3, 1	2, 2
	D	0, 2	2, 3

(a) Find and describe all pure-strategy Nash equilibria.
(b) If Row plays U with probability p and Column plays L with probability q, demonstrate that $p = 0.75$, $q = 0$ is a mixed-strategy Nash equilibrium for this game.
(c) Is $p = 0.4$, $q = 0$ a mixed-strategy equilibrium for this game? Explain why or why not.
(d) Is $p = 1$, $q = 0.5$ a mixed-strategy equilibrium for this game? Explain why or why not.
(e) How many mixed-strategy equilibria does this game have? Explain.

U11. Consider another simplified version of baseball. The pitcher can throw either a fastball or a curveball; the batter can either swing at the pitch or take (not swing). These choices are simultaneous for each pitch. On the first pitch, if the batter swings at a curveball or takes a fastball, he strikes out and gets a 0. If the batter swings at a fastball, he has a probability of 0.75 of hitting a home run and getting a 1, and a probability of hitting a fly ball and getting a 0. If the batter takes a curveball, there is a second pitch.

On the second pitch, the first three combinations (swing at a curveball, take a fastball, and swing at a fastball) work as before; if the batter takes a curveball on the second pitch, he walks and earns 0.25.

This is a zero-sum game; the batter tries to maximize his expected score (his probability-weighted average payoff), and the pitcher tries to minimize the batter's expected score. Note that the two pitches constitute a sequential-move game, whereas each individual pitch is a simultaneous-move game.

(a) Use the techniques of Chapter 6 to draw a game tree to represent this two-stage game.

(b) Solve this game using rollback; construct a table of payoffs for the second pitch, and use them to determine the table of payoffs for the first pitch. Show that on the first pitch, the batter should take with a probability of 0.8.

(c) What is the pitcher's strategy in the subgame-perfect equilibrium?

(d) What is the batter's expected score in this equilibrium?

(e) Explain intuitively why the batter's probability of swinging is so small.

U12. **(Optional)** Exercises S5 and U5 demonstrate that in zero-sum games such as the Evert-Navratilova tennis rivalry, changes in a player's payoffs can sometimes lead to unexpected or unintuitive changes to her equilibrium mixture. But what happens to the expected value of the game? Consider the following general form of a two-player zero-sum game:

		COLUMN	
		L	R
ROW	U	*a*	*b*
	D	*c*	*d*

Assume that there is no Nash equilibrium in pure strategies, and assume that *a*, *b*, *c*, and *d* are all greater than or equal to zero. Can an *increase* in any one of *a*, *b*, *c*, or *d* lead to a *lower* expected value of the game for Row? If not, prove why not. If so, provide an example.

U13. **(Optional)** Return to the game in Exercise S3. Suppose that both players are risk averse, as discussed in Appendix 2 of this chapter. In particular, suppose that each player has a square-root utility function, and that each player knows the other's utility function. Each player wants to maximize his expected utility.

(a) Find the mixed-strategy Nash equilibrium of this game.

(b) How did the two players' mixing proportions change relative to Exercise S3, where both players were risk neutral? Explain why this might have happened.

Appendix: Probability and Expected Utility

To calculate the expected payoffs and mixed-strategy equilibria of games in this chapter, we had to do some simple manipulation of probabilities. Some simple rules govern calculations involving probabilities. Many of you may be familiar with them, but we give a brief statement and explanation of the basics here by way of reminder or remediation, as appropriate. We also state how to calculate expected values of random numerical values.

We also consider the expected utility approach to calculating expected payoffs. When the outcomes of your action in a particular game are not certain, either because your opponent is mixing strategies or because of some uncertainty in nature, you may not want to maximize your expected monetary payoff as we have generally assumed in our analysis to this point; rather, you may want to give some attention to the *riskiness* of the payoffs. As mentioned in Chapter 2, such situations can be handled by using the expected values (which are probability-weighted averages) of an appropriate nonlinear rescaling of the monetary payoffs. We offer here a brief discussion of how this can be done.

You should certainly read this material, but to get real knowledge and mastery of it, the best thing to do is to use it. The chapters to come, especially Chapters 8, 9, 13, and 14, will give you plenty of opportunity for practice.

1 THE BASIC ALGEBRA OF PROBABILITIES

The basic intuition about the probability of an event comes from thinking about the frequency with which this event occurs by chance among a larger set of possibilities. Usually any one element of this larger set is just as likely to occur by chance as any other, so finding the probability of the event in which we are interested is simply a matter of counting the elements corresponding to that event and dividing by the total number of elements in the whole large set.[1]

In any standard deck of 52 playing cards, for instance, there are four suits (clubs, diamonds, hearts, and spades) and 13 cards in each suit (ace through 10 and the face cards—jack, queen, king). We can ask a variety of questions about the likelihood that a card of a particular suit or value—or suit *and* value—might be drawn from this deck of cards: How likely are we to draw a spade? How likely are we to draw a black card? How likely are we to draw a 10? How likely are we to draw the queen of spades? and so on. We would need to know something about the calculation and manipulation of probabilities to answer such questions. If we had two decks of cards, one with blue backs and one with green backs, we could ask even more complex questions ("How likely are we to draw one card from each deck and have them both be the jack of diamonds?"), but we would still use the algebra of probabilities to answer them.

In general, a **probability** measures the likelihood of a particular event or set of events occurring. The likelihood that you draw a spade from a deck of cards is just the probability of the event "drawing a spade." Here the large set has 52 elements—the total number of equally likely possibilities—and the event "drawing a spade" corresponds to a subset of 13 particular elements. Thus you have 13 chances out of the 52 to get a spade, which makes the probability of getting a spade in a single draw equal to 13/52 = 1/4 = 25%. To see this another way, consider the fact that there are four suits of 13 cards each, so your chance of drawing a card from any particular suit is one out of four, or 25%. If you made a large number of such draws (each time from a complete deck), then out of 52 times you will not always draw exactly 13 spades; by chance you may draw a few more or a few less. But the chance averages out over different such occasions—over different sets of 52 draws. Then the probability of 25% is the average of the frequencies of spades drawn in a large number of observations.[2]

[1]When we say "by chance," we simply mean that a systematic order cannot be detected in the outcome or that it cannot be determined by using available scientific methods of prediction and calculation. Actually, the motions of coins and dice are fully determined by laws of physics, and highly skilled people can manipulate decks of cards but, for all practical purposes, coin tosses, rolls of dice, or card shuffles are devices of chance that can be used to generate random outcomes. However, randomness can be harder to achieve than you think. For example, a perfect shuffle, where a deck of cards is divided exactly in half and then interleaved by dropping cards one at a time alternately from each, may seem a good way to destroy the initial order of the deck. But Cornell mathematician Persi Diaconis has shown that, after eight of the shuffles, the original order is fully restored. For slightly imperfect shuffles that people carry out in reality, he finds that some order persists through six, but randomness suddenly appears on the seventh! See "How to Win at Poker, and Other Science Lessons," *The Economist*, October 12, 1996. For an interesting discussion of such topics, see Deborah J. Bennett, *Randomness* (Cambridge: Harvard University Press, 1998), chaps. 6–9.

[2]Bennett, *Randomness*, chaps. 4 and 5, offers several examples of such calculations of probabilities.

The algebra of probabilities simply develops such ideas in general terms and obtains formulas that you can then apply mechanically instead of having to do the thinking from scratch every time. We will organize our discussion of these probability formulas around the types of questions that one might ask when drawing cards from a standard deck (or two: blue backed and green backed).[3] This method will allow us to provide both specific and general formulas for you to use later. You can use the card-drawing analogy to help you reason out other questions about probabilities that you encounter in other contexts. One other point to note: In ordinary language, it is customary to write probabilities as percentages, but the algebra requires that they be written as fractions or decimals; thus instead of 25% the mathematics works with 13/52 or 0.25. We will use one or the other, depending on the occasion; be aware that they mean the same thing.

A. The Addition Rule

The first questions that we ask are: If we were to draw one card from the blue deck, how likely are we to draw a spade? And how likely are we to draw a card that is not a spade? We already know that the probability of drawing a spade is 25% because we determined that earlier. But what is the probability of drawing a card that is not a spade? It is the same likelihood of drawing a club or a diamond or a heart instead of a spade. It should be clear that the probability in question should be larger than any of the individual probabilities of which it is formed; in fact, the probability is 13/52 (clubs) + 13/52 (diamonds) + 13/52 (hearts) = 0.75. The *or* in our verbal interpretation of the question is the clue that the probabilities should be added together, because we want to know the chances of drawing a card from any of those three suits.

We could more easily have found our answer to the second question by noting that not getting a spade is what happens the other 75% of the time. Thus the probability of drawing "not a spade" is 75% (100% − 25%) or, more formally, 1 − 0.25 = 0.75. As is often the case with probability calculations, the same result can be obtained here by two different routes, entailing different ways of thinking about the event for which we are trying to find the probability. We will see other examples of this later in this Appendix, where it will become clear that the different methods of calculation can sometimes require vastly different amounts of effort. As you develop experience, you will discover and remember the easy ways or shortcuts. In the meantime, be comforted that each of the different routes, when correctly followed, leads to the same final answer.

[3]If you want a more detailed exposition of the following addition and multiplication rules, as well as more exercises to practice these rules, we recommend David Freeman, Robert Pisani, and Robert Purves, *Statistics*, 3rd ed. (New York: Norton, 1998), chaps. 13 and 14.

To generalize our preceding calculation, we note that, if you divide the set of events, X, in which you are interested into some number of subsets, $Y, Z, \ldots$, none of which overlap (in mathematical terminology, such subsets are said to be **disjoint**), then the probabilities of each subset occurring must sum to the probability of the full set of events; if that full set of events includes all possible outcomes, then its probability is 1. In other words, if the occurrence of X requires the occurrence of *any one* of several disjoint $Y, Z, \ldots$, then the probability of X is the sum of the separate probabilities of $Y, Z, \ldots$ Using Prob(X) to denote the probability that X occurs and remembering the caveats on X (that it requires any one of $Y, Z, \ldots$) and on $Y, Z, \ldots$ (that they must be disjoint), we can write the **addition rule** in mathematical notation as $\text{Prob}(X) = \text{Prob}(Y) + \text{Prob}(Z) + \cdots$.

EXERCISE Use the addition rule to find the probability of drawing two cards, one from each deck, such that the two cards have identical faces.

B. The Modified Addition Rule

Our analysis in Section A of this appendix covered only situations in which a set of events could be broken down into disjoint, nonoverlapping subsets. But suppose we ask, What is the likelihood, if we draw one card from the blue deck, that the card is either a spade *or* an ace? The *or* in the question suggests, as before, that we should be adding probabilities, but in this case the two categories "spade" and "ace" are not mutually exclusive, because one card, the ace of spades, is in both subsets. Thus "spade" and "ace" are not disjoint subsets of the full deck. So if we were to sum only the probabilities of drawing a spade (13/52) and of drawing an ace (4/52), we would get 17/52. This would suggest that we had 17 different ways of finding either an ace or a spade when in fact we have only 16—there are 13 spades (including the ace) and three additional aces from the other suits. The incorrect answer, 17/52, comes from counting the ace of spades twice. To get the correct probability in the nondisjoint case, then, we must subtract the probability associated with the overlap of the two subsets. The probability of drawing an ace or a spade is the probability of drawing an ace plus the probability of drawing a spade minus the probability of drawing the overlap, the ace of spades; that is, $13/52 + 4/52 - 1/52 = 16/52 = 0.31$.

To make this more general, if you divide the set of events, X, in which you are interested into some number of subsets $Y, Z, \ldots$, which may overlap, then the sum of the probabilities of each subset occurring minus the probability of the overlap yields the probability of the full set of events. More formally, the **modified addition rule** states that, if the occurrence of X requires the occurrence of any one of the nondisjoint Y and Z, then the probability of X is the sum of the separate probabilities of Y and Z minus the probability that *both* Y and Z occur: $\text{Prob}(X) = \text{Prob}(Y) + \text{Prob}(Z) - \text{Prob}(Y \text{ and } Z)$.

EXERCISE Use the modified addition rule to find the probability of drawing two cards, one from each deck, and getting at least one face card.

C. The Multiplication Rule

Now we ask, What is the likelihood that when we draw two cards, one from each deck, both of them will be spades? This event occurs if we draw a spade from the blue deck *and* a spade from the green deck. The switch from *or* to *and* in our interpretation of what we are looking for indicates a switch in mathematical operations from addition to multiplication. Thus the probability of two spades, one from each deck, is the product of the probabilities of drawing a spade from each deck, or $(13/52) \times (13/52) = 1/16 = 0.0625$, or 6.25%. Not surprisingly, we are much less likely to get two spades than we were in Section A to get one spade. (Always check to make sure that your calculations accord in this way with your intuition regarding the outcome.)

In much the same way as the addition rule requires events to be disjoint, the multiplication rules requires them to be independent; if we break down a set of events, X, into some number of subsets $Y, Z, \ldots$, those subsets are independent if the occurrence of one does not affect the probability of the other. Our events—a spade from the blue deck and a spade from the green deck—satisfy this condition of independence; that is, drawing a spade from the blue deck does nothing to alter the probability of getting a spade from the green deck. If we were drawing both cards from the same deck, however, then after we had drawn a spade (with a probability of 13/52), the probability of drawing another spade would no longer be 13/52 (in fact, it would be 12/51); drawing one spade and then a second spade from the *same* deck are not **independent events**.

The formal statement of the **multiplication rule** tells us that, if the occurrence of X requires the simultaneous occurrence of *all* the several independent $Y, Z, \ldots$, then the probability of X is the *product* of the separate probabilities of $Y, Z, \ldots$: $\text{Prob}(X) = \text{Prob}(Y) \times \text{Prob}(Z) \times \ldots$.

EXERCISE Use the multiplication rule to find the probability of drawing two cards, one from each deck, and getting a red card from the blue deck and a face card from the green deck.

D. The Modified Multiplication Rule

What if we are asking about the probability of an event that depends on two nonindependent occurrences? For instance, suppose that we ask, What is the likelihood that with one draw we get a card that is both a spade *and* an ace? If we think about this for a moment, we realize that the probability of this event is just the probability of drawing a spade *and* the probability that our card is an ace *given that* it is a spade. The probability of drawing a spade is $13/52 = 1/4$, and the probability of drawing an ace, given that we have a spade, is 1/13. The *and* in our question tells us to take the product of these two probabilities: (13/52)(1/13) = 1/52.

We could have gotten the same answer by realizing that our question was the same as asking, What is the likelihood of drawing the ace of spades? The calculation of that probability is straightforward; only 1 of the 52 cards is the ace of spades, so the probability of drawing it must be 1/52. As you see, how you word the question affects how you go about looking for an answer.

In the technical language of probabilities, the probability of a particular event occurring (such as getting an ace), given that another event has already occurred (such as getting a spade) is called the **conditional probability** of drawing an ace, for example, conditioned on having drawn a spade. Then the formal statement of the **modified multiplication rule** is that, if the occurrence of X requires the occurrence of both Y and Z, then the probability of X equals the product of two things: (1) the probability that Y alone occurs, and (2) the probability that Z occurs given that Y has already occurred, or the *conditional probability* of Z, conditioned on Y having already occurred: Prob(X) = Prob(Y alone) × Prob (Z given Y).

A third way would be to say that the probability of drawing an ace is 4/52, and the conditional probability of the suit being a spade, given that the card is an ace, is 1/4; so the overall probability of getting an ace of spades is (4/52) × 1/4. More generally, using the terminology just introduced, we have Prob(X) = Prob(Z) Prob(Y given Z).

EXERCISE Use the modified multiplication rule to find the probability that, when you draw two cards from a deck, the second card is the jack of hearts.

E. The Combination Rule

We could also ask questions of an even more complex nature than we have tried so far, in which it becomes necessary to use both the addition (or modified addition) and the multiplication (or modified multiplication) rules simultaneously. We could ask, What is the likelihood, if we draw one card from each deck, that we draw *at least one* spade? As usual, we could approach the calculation of the necessary probability from several angles, but suppose that we come at it first by considering all of the different ways in which we could draw at least one spade when drawing one card from each deck. There are three possibilities: either we could get one spade from the blue deck and none from the green deck ("spade *and* none") *or* we could get no spade from the blue deck and a spade from the green deck ("none *and* spade") *or* we could get a spade from each deck ("spade *and* spade"); our event requires that one of these three possibilities occurs, each of which entails the occurrence of both of two independent events.

It should be obvious now, by using the *or*s and *and*s as guides, how to calculate the necessary probability. We find the probability of each of the three possible ways of getting at least one spade (which entails three products of two probabilities each) and sum these probabilities together: (1/4 × 3/4) + (3/4 × 1/4) + (1/4 × 1/4) = 7/16 = 43.75%.

The second approach entails recognizing that "at least one spade" and "not any spades" are disjoint events; together they constitute a sure thing. Therefore the probability of "at least one spade" is just 1 minus the probability of "not any spades." And the event "not any spades" occurs only if the blue card is not a spade 3/4 *and* the green card is not a spade 3/4; so its probability is 3/4 × 3/4 = 9/16. The probability of "at least one spade" is then 1 − 9/16 = 7/16, as we found in the preceding paragraph.

Finally, we can formally state the **combination rule** for probabilities: if the occurrence of X requires the occurrence of *exactly one* of a number of *disjoint* $Y, Z, \ldots$, the occurrence of Y requires that of *all* of a number of *independent* $Y_1, Y_2, \ldots$ the occurrence of Z requires that of *all* of a number of *independent* $Z_1, Z_2, \ldots$, and so on, then the probability of X is the sum of the probabilities of $Y, Z, \ldots$, which are the products of the probabilities $Y_1, Y_2, \ldots, Z_1, Z_2, \ldots$: or

$$\begin{aligned}\text{Prob}(X) &= \text{Prob}(Y) + \text{Prob}(Z) + \cdots \\ &= \text{Prob}(Y_1) \times \text{Prob}(Y_2) \times \cdots + \text{Prob}(Z_1) \times \text{Prob}(Z_2) \times \cdots + \cdots\end{aligned}$$

EXERCISE Suppose we now have a third (orange) deck of cards. Find the probability of drawing at least one spade when you draw one card from each of the three decks.

F. Expected Values

If a numerical magnitude (such as money winnings or rainfall) is subject to chance and can take on any one of n possible values $X_1, X_2, \ldots, X_n$ with respective probabilities $p_1, p_2, \ldots, p_n$, then the **expected value** is defined as the weighted average of all its possible values using the probabilities as weights; that is, as $p_1X_1 + p_2X_2 + \cdots + p_nX_n$. For example, suppose you bet on the toss of two fair coins. You win $5 if both coins come up heads, $1 if one shows heads and the other tails, and nothing if both come up tails. Using the rules for manipulating probabilities discussed earlier in this section, you can see that the probabilities of these events are, respectively, 0.25, 0.50, and 0.25. Therefore your expected winnings are (0.25 × $5) + (0.50 × $1) + (0.25 × $0) = $1.75.

In game theory, the numerical magnitudes that we need to average in this way are payoffs, measured in numerical ratings, or money, or, as we will see later in this appendix, utilities. We will refer to the expected values in each context appropriately, for example, as *expected payoffs* or **expected utilities**.

2 ATTITUDES TOWARD RISK AND EXPECTED UTILITY

In Chapter 2, we pointed out a difficulty about using probabilities to calculate the average or expected payoff for players in a game. Consider a game where players gain or lose money, and suppose we measure payoffs simply in money amounts. If a player has a 75% chance of getting nothing and a 25% chance of getting $100, then the expected payoff is calculated as a *probability-weighted average*; the expected payoff is the average of the different payoffs with the probabilities of each as weights. In this case, we have $0 with a probability of 75%, which yields $0.75 \times 0 = 0$ on average, added to $100 with a probability of 25%, which yields $0.25 \times 100 = 25$ on average. That is the same payoff as the player would get from a simple nonrandom outcome that guaranteed him $25 every time he played. People who are indifferent between two alternatives with the same average monetary value but different amounts of risk are said to be **risk-neutral.** In our example, one prospect is riskless ($25 for sure), while the other is risky, yielding either $0 with a probability of 0.75 or $100 with a probability of 0.25, for the same average of $25. In contrast are **risk-averse** people—those who, given a pair of alternatives each with the same average monetary value, would prefer the less risky option. In our example, they would rather get $25 for sure than face the risky $100-or-nothing prospect and, given the choice, would pick the safe prospect. Such risk-averse behavior is quite common; we should therefore have a theory of decision making under uncertainty that takes it into account.

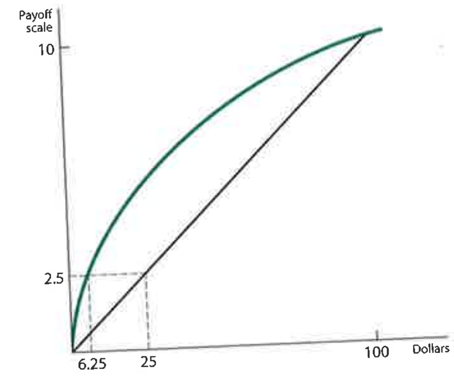

FIGURE 7A.1 Concave Scale: Risk Aversion

We also said in Chapter 2 that a very simple modification of our payoff calculation can get us around this difficulty. We said that we could measure payoffs not in money sums but by using a nonlinear rescaling of the dollar amounts. Here we show explicitly how that rescaling can be done and why it solves our problem for us.

Suppose that, when a person gets D dollars, we define the payoff to be something other than just D, perhaps $\sqrt{D}$. Then the payoff number associated with $0 is 0, and that for $100 is 10. This transformation does not change the way in which the person rates the two payoffs of $0 and $100; it simply rescales the payoff numbers in a particular way.

Now consider the risky prospect of getting $100 with probability 0.25 and nothing otherwise. After our rescaling, the expected payoff (which is the average of the two payoffs with the probabilities as weights) is $(0.75 \times 0) + (0.25 \times 10) = 2.5$. This expected payoff is equivalent to the person's getting the dollar amount whose square root is 2.5; because $2.5 = \sqrt{6.25}$, a person getting $6.25 for sure would also receive a payoff of 2.5. In other words, the person with our square-root payoff scale would be just as happy getting $6.25 for sure as he would getting a 25% chance at $100. This indifference between a guaranteed $6.25 and a 1 in 4 chance of $100 indicates quite a strong aversion to risk; this person is willing to give up the difference between $25 and $6.25 to avoid facing the risk. Figure 7A.1 shows this nonlinear scale (the square root), the expected payoff, and the person's indifference between the sure prospect and the gamble.

What if the nonlinear scale that we use to rescale dollar payoffs is the cube root instead of the square root? Then the payoff from $100 is 4.64, and the

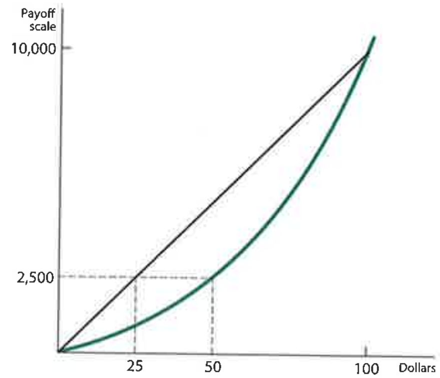

FIGURE 7A.2 Convex Scale: Risk Loving

expected payoff from the gamble is $(0.75 \times 0) + (0.25 \times 4.64) = 1.16$, which is the cube root of 1.56. Therefore a person with this payoff scale would accept only \$1.56 for sure instead of a gamble that has a money value of \$25 on average; such a person is extremely risk-averse indeed. (Compare a graph of the cube root of x with a graph of the square root of x to see why this should be so.)

And what if the rescaling of payoffs from x dollars is done by using the function x^2? Then the expected payoff from the gamble is $(0.75 \times 0) + (0.25 \times 10{,}000) = 2{,}500$, which is the square of 50. Therefore a person with this payoff scale would be indifferent between getting \$50 for sure and the gamble with an expected money value of only \$25. This person must be a risk lover because he is not willing to give up any money to get a reduction in risk; on the contrary, he must be given an extra \$25 in compensation for the loss of risk. Figure 7A.2 shows the nonlinear scale associated with a function such as x^2.

So by using different nonlinear scales instead of pure money payoffs, we can capture different degrees of risk-averse or risk-loving behavior. A concave scale like that of Figure 7A.1 corresponds to risk aversion, and a convex scale like that of Figure 7A.2 to risk-loving behavior. You can experiment with different simple nonlinear scales—for example, logarithms, exponentials, and other roots and powers—to see what they imply about attitudes toward risk.[4]

This method of evaluating risky prospects has a long tradition in decision theory; it is called the expected utility approach. The nonlinear scale that gives payoffs as functions of money values is called the **utility function;** the square root, cube root, and square functions referred to earlier are simple examples. Then the mathematical expectation, or probability-weighted average, of the utility values of the different money sums in a random prospect is called the expected utility of that prospect. And different random prospects are compared with one another in terms of their expected utilities; prospects with higher expected utility are judged to be better than those with lower expected utility.

Almost all of game theory is based on the expected utility approach, and it is indeed very useful, although is not without flaws. We will adopt it in this book, leaving more detailed discussions to advanced treatises.[5] However, we will indicate the difficulties that it leaves unresolved by means of a simple example in Chapter 8.

[4]Additional information on the use of expected utility and risk attitudes of players can be found in many intermediate microeconomic texts; for example, Hal Varian, *Intermediate Microeconomics*, 7th ed. (New York: Norton, 2006), ch. 12; Walter Nicholson and Christopher Snyder, *Microeconomic Theory*, 10th ed. (New York: Dryden Press, 2008), ch. 7.

[5]See R. Duncan Luce and Howard Raiffa, *Games and Decisions* (New York: Wiley, 1957), chap. 2 and app. 1, for an exposition; and Mark Machina, "Choice Under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, vol. 1, no. 1 (Summer 1987), pp. 121–154, for a critique and alternatives. Although decision theory based on these alternatives has made considerable progress, it has not yet influenced game theory to any significant extent.

SUMMARY

The *probability* of an event is the likelihood of its occurrence by chance from among a larger set of possibilities. Probabilities can be combined by using some rules. The *addition rule* says that the probability of any one of a number of *disjoint* events occurring is the sum of the probabilities of these events; the *modified addition rule* generalizes the addition rule to overlapping events. According to the *multiplication rule*, the probability that all of a number of *independent events* will occur is the product of the probabilities of these events; the *modified multiplication rule* generalizes the multiplication rule to allow for lack of independence, by using *conditional probabilities*.

Judging consequences by taking expected monetary payoffs assumes *risk-neutral* behavior. *Risk aversion* can be allowed, by using the *expected-utility* approach, which requires the use of a *utility function*, which is a concave rescaling of monetary payoffs, and taking its probability-weighted average as the measure of expected payoff.

KEY TERMS

addition rule (254)
combination rule (257)
conditional probability (256)
disjoint (254)
expected utility (257)
expected value (257)
independent events (255)
modified addition rule (254)
modified multiplication rule (256)
multiplication rule (255)
probability (252)
risk-averse (258)
risk-neutral (258)
utility function (260)