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

8

Simultaneous-Move Games with Mixed Strategies II: Some General Theory

In Chapter 7 we developed the basic concept of a mixed strategy and a method for finding mixed-strategy equilibria in two-by-two games based on the *opponent's indifference property.* That analysis will suffice for readers who wish only to acquire basic familiarity with the subject of mixing. But more systematic study of games with mixed-strategy equilibria will give a deeper understanding to those who plan to go on to further study of game theory. We provide that deeper analysis in this chapter.

We should say quite frankly that most of this chapter is more difficult than the rest of the book. But we do not apologize for including this material. First, it has an important place in the history of game theory. Von Neumann and Morgenstern's pioneering classic book and Luce and Raiffa's equally classic textbook[1] devote more than a third of their space to the theory of mixed strategies. Second, this topic has its own constituency. Just as many of our other chapters will have special appeal to students whose interests lie in economics, business, political science, or evolutionary biology, this chapter is for students who come from a more mathematical background. Although only a small amount of college mathematics is needed, readers will require a familiarity with algebraic notation

[1]See John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1944; 2nd ed., 1947; 3rd ed., 1953); R. Duncan Luce and Howard Raiffa, *Games and Decisions* (New York: John Wiley and Sons), 1957.

and logic to appreciate fully the material presented here. Those who lack the interest or the preparation for this material can omit this chapter or cover only parts of it without loss of continuity.

In addition to providing a more general approach to solving two-by-two games with mixed strategies, we consider larger games in this chapter as well—those in which players have three or more pure strategies initially available. A substantive new question arises in such games: When will a player's equilibrium mix include all of his pure strategies, and when will it include only a subset? Finally, the basic result regarding general mixed-strategy equilibria, due to von Neumann and Morgenstern, can claim to be the first true theorem of game theory. We describe that general result in Section 7 and introduce the other ideas through numerical examples in the earlier sections.

1 BEST-RESPONSE ANALYSIS

When we developed the concept of a mixed strategy in Section 7.1, we pointed out that a mixture is a special kind of continuous strategy. The probability of using one of the pure strategies is the continuous variable that characterizes the strategy. In Chapter 7, we then developed the opponent's indifference property as the method for calculating equilibrium-mixture probabilities. But when we studied Nash equilibria with continuous strategies in Chapter 5, we developed a different solution method, namely best-response analysis. Now we reconcile the apparent difference in approaches by using best-response analysis to find Nash equilibria in mixed strategies. Although the opponent's indifference property remains a quick method of calculating mixture probabilities, best-response analysis is the more general method that locates all Nash equilibria (if the game has multiple equilibria), pure and mixed.

A. Best-Response Analysis of the Tennis Point

We develop best-response analysis of games with mixed strategies using the same tennis-point example that we used for developing the concept of mixed strategies and their equilibria in Chapter 7. For your convenience we reproduce as Figure 8.1 the payoff table from Figure 7.1, but we now explicitly include a third, mixed, strategy for each player. For Evert, we refer to the mixture as her p-mix, in which we assign the general probability p to choosing her first strategy, DL; we do likewise for Navratilova and her q-mix. The expressions for the expected payoffs when one player's mixture faces a pure strategy of the other player were derived in Section 7.2. The algebraic expression for the bottom right cell where mix meets mix is more complicated, but we won't need it, so we leave that cell blank.

		NAVRATILOVA		
		DL	CC	q-mix
EVERT	DL	50, 50	80, 20	$50q + 80(1-q)$, $50q + 20(1-q)$
	CC	90, 10	20, 80	$90q + 20(1-q)$, $10q + 80(1-q)$
	p-mix	$50p + 90(1-p)$, $50p + 10(1-p)$	$80p + 20(1-p)$, $20p + 80(1-p)$	

FIGURE 8.1 Expected Payoffs for General Mixtures in the Tennis Point

In Section 7.2 we also derived Navratilova's best response to Evert's p-mix. We restate it here:

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 ($q = 1$).

Just as a reminder, if p is low (Evert is quite likely to choose CC), Navratilova does better by covering CC; if p is high (Evert is quite likely to choose DL), Navratilova does better by covering DL; for a critical value of p in between, namely 0.7, Navratilova gets the same expected payoff from either of her pure strategies and therefore also the same from any mixture of the two.

We show this best response graphically in the left-hand panel of Figure 8.2. Navratilova is choosing q for each given value of Evert's p; thus q is a function of p. Therefore p is on the horizontal axis and q is on the vertical axis. We know that, for $p < 0.7$, Navratilova does better by choosing pure CC ($q = 0$); this segment of her best-response curve is the horizontal solid (green) line along the bottom edge of the graph. For $p > 0.7$, Navratilova does better by choosing pure DL ($q = 1$); this segment of her best-response curve is the horizontal solid line along the top edge of the graph. For $p = 0.7$, Navratilova does equally well with all of her choices, pure and mixed, and so any value of q between 0 and 1 (inclusive) is a best response; the vertical solid line in the graph at $p = 0.7$ shows her best responses for this choice by Evert. Navratilova's best responses are then shown by the three separate line segments joined end to end and shown in green. As we did for general continuous strategies in Chapter 5, we call this construction Navratilova's *best-response curve*. It is conceptually the same as the best-response curves that we drew in Chapter 5; the only difference is that, for mixed strategies, the curve has this special shape. Because it actually consists of three straight-line segments, "curve" is a misnomer, but it is the standard general terminology in this context.

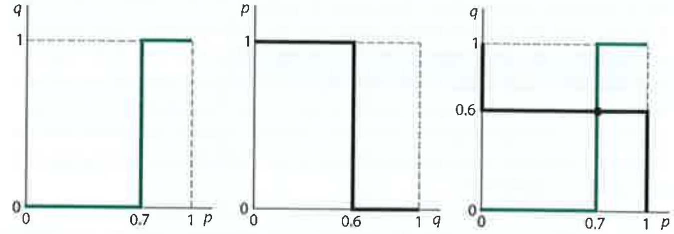

FIGURE 8.2 Best-Response Curves and Nash Equilibrium

Similarly, we recall Evert's best-response rule from Section 7.2:

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 ($p = 0$).

And we depict her best response in the middle panel of Figure 8.2. Here, Evert's best p values are determined in relation to Navratilova's various possible choices for q; so q is on the horizontal axis and p on the vertical axis. For $q < 0.6$, Evert does better playing pure DL ($p = 1$); for $q > 0.6$, she does better with pure CC ($p = 0$); for $q = 0.6$, she does equally well with all choices, pure or mixed. The thick black curve, consisting of three line segments joined end to end, is Evert's best-response curve.

The right-hand panel in Figure 8.2 combines the other two panels by reflecting the middle graph across the diagonal (45° line) so that p is on the horizontal axis and q on the vertical axis and then superimposing this graph on the left-hand graph. Now the green and black curves meet at the point $p = 0.7$ and $q = 0.6$. Here each player's mixture choice is a best response to the other's choice, so we see clearly the derivation of our Nash equilibrium in mixed strategies.

This picture also shows that the best-response curves do not have any other common points. Thus the mixed-strategy equilibrium in this example is the unique Nash equilibrium in the game. What is more, this representation includes pure strategies as special cases corresponding to extreme values of p and q. So we can also see that the best-response curves do not have any points in common at any of the sides of the square where each value of p and q equals either 0 or 1; thus we have another way to verify that the game does not have any pure-strategy equilibria. The mixed-strategy equilibrium in this example is the unique Nash equilibrium in the game.

Observe that each player's best response is a pure strategy for almost all values of her opponent's mixture. Thus Navratilova's best response is pure CC for all of Evert's choices of $p < 0.7$, and it is pure DL for all of Evert's choices of $p > 0.7$. Only for the one crucial value $p = 0.7$ is Navratilova's best response a mixed strategy, as is represented by the vertical part of her three-segment best-response curve in the left panel of Figure 8.2. Similarly, only for the one crucial value $q = 0.6$ of Navratilova's mixture is Evert's best response a mixed strategy—namely, the horizontal segment of her best-response curve in the middle panel of Figure 8.2. But these seemingly exceptional or rare strategies are just the ones that emerge in the equilibrium.

These special values of p and q have an important feature in common. Evert's equilibrium p is where Navratilova's best-response curve has its vertical segment; so Navratilova is indifferent among all her strategies, pure and mixed. Navratilova's equilibrium q is where Evert's best-response curve has its horizontal segment; so Evert is indifferent among all her strategies, pure and mixed. Each player's equilibrium mixture is such that the other player is indifferent among all her mixes. Thus the opponent's indifference property is reconfirmed by our best-response analysis.

The best-response-curve method thus provides a very complete analysis of the game. Like the cell-by-cell inspection method, which examines all the cells of a pure-strategy game, the best-response-curve method is the one to use when we want to locate all of the equilibria, whether in pure or mixed strategies, that a game might have. The best-response-curve diagram can show both types of equilibria in the same place. It could also be used to show equilibria in which one player uses a mixed strategy and the other player uses a pure strategy. Such hybrids occur only in exceptional cases; we give some examples in Section 2.B.

B. Best-Response Analysis in Non-Zero-Sum Games

Best-response analysis can also be used for non-zero-sum games and yields all equilibria, in mixed as well as pure strategies. We illustrate this case using the assurance game, leaving the other two classic games we discussed in Chapter 7 (battle of the sexes and chicken) for you as exercises.

We reproduce the payoff matrix of the assurance version of the meeting game (Figure 4.12 or 7.4) as Figure 8.3, but as in Figure 8.1 for the tennis point, we show an added row and an added column corresponding to Harry's p-mix and Sally's q-mix.[2]

We depict Sally's payoffs from her two pure strategies in relation to Harry's p-mix, and her best-response rule, in the two panels of Figure 8.4. When $p < 2/3$, Sally's best response is pure Local Latte ($q = 0$). The intuition is simple; if Harry is not very likely to go to Starbucks, neither should

[2]Following standard practice, p and q represent the probabilities of choosing the first strategy (Starbucks).

		SALLY		
		Starbucks	Local Latte	q-mix
HARRY	Starbucks	1, 1	0, 0	q, q
	Local Latte	0, 0	2, 2	$2(1 - q)$, $2(1 - q)$
	p-mix	p, p	$2(1 - p)$, $2(1 - p)$	

FIGURE 8.3 Expected Payoffs for General Mixtures in the Assurance Game

Sally. Similarly, when $p > 2/3$, Sally's best response is pure Starbucks ($q = 1$). When $p = 2/3$, all values of q are equally good for Sally, so her best response is any combination of Local Latte and Starbucks. This portion of her best-response curve is shown by the vertical straight line from $q = 0$ to $q = 1$ at $p = 2/3$ in the right panel of Figure 8.4.

We find the mixed-strategy equilibrium by superimposing the two best-response curves, as we did for the tennis-point game in Section 1.A above. Because the payoffs of the two players are symmetric, Harry's best-response curve will look just like Sally's with the two axes interchanged. Figure 8.5 shows both of the best-response curves at the same time. Sally's is the thick green curve; Harry's is the thick black curve.

The two best-response curves meet at three points. One is at the top right where $p = 1$ and $q = 1$. This point corresponds to each player choosing Starbucks

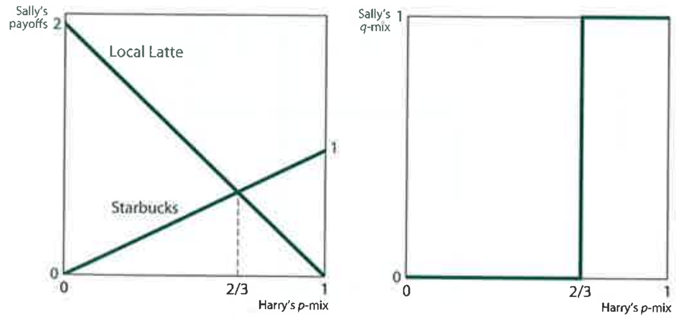

FIGURE 8.4 Sally's Best Responses

for sure, in the correct subjective belief that the other is doing likewise. This is a self-sustaining state of affairs. The second meeting point is at the bottom left, where $p = 0$ and $q = 0$. Here each is not putting any probability on going to Starbucks—that is, is going to Local Latte for sure—in the correct subjective belief that the other is doing likewise. This also is a self-sustaining situation. These are just the two pure-strategy Nash equilibria for this game that we found in Chapter 4.

But there is a third meeting point in Figure 8.5, where $p = 2/3$ and $q = 2/3$. This intersection is just the mixed-strategy equilibrium that we calculated using the opponent's indifference property in Section 7.4.A. As promised, best-response analysis gives us all three of the Nash equilibria for this game, whether in pure or mixed strategies, in one go.

2 MIXING WHEN ONE PLAYER HAS THREE OR MORE PURE STRATEGIES

The discussion of mixed strategies to this point has been confined to games in which each player has only two pure strategies, as well as mixes between them. In many strategic situations, each player has available a larger number of pure strategies, and we should be ready to calculate equilibrium mixes for those cases as well. However, these calculations get complicated quite quickly. For truly complex games, we would turn to a computer to find the mixed-strategy equilibrium. But for some small games, it is possible to calculate equilibria by hand quite easily. The calculation process gives us a better understanding of how the equilibrium works than can be obtained just from looking at a computer-generated solution. Therefore in this section and the next one we solve some larger games.

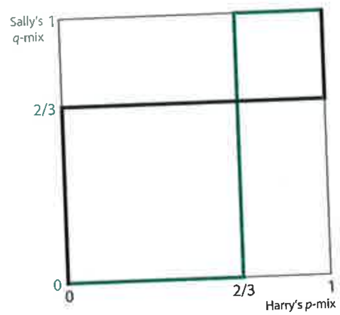

FIGURE 8.5 Mixed-Strategy Equilibrium in the Assurance Game

Here we consider zero-sum games in which one of the players has only two pure strategies, whereas the other has more. In such games, we find that the player who has three (or more) pure strategies typically uses only two of them in equilibrium. The others do not figure in his mix; they get zero probabilities. We must determine which ones are used and which ones are not.[3]

Our example is that of the tennis-point game augmented by giving Evert a third type of return. In addition to going down the line or crosscourt, she now can consider using a lob (a slower but higher and longer return). The equilibrium depends on the payoffs of the lob against each of Navratilova's two defensive stances. We begin with the case that is most likely to arise and then consider a coincidental or exceptional case.

A. A General Case

Evert now has three pure strategies in her repertoire: DL, CC, and Lob. We leave Navratilova with just two pure strategies, Cover DL or Cover CC. The payoff table for this new game can be obtained by adding a Lob row to the table in Figure 8.1. The result is shown in Figure 8.6. Now that you are more familiar with mixing in constant-sum games, we show only Evert's payoffs, and ask you to remember that Navratilova chooses her strategies so as to achieve smaller expected payoffs for Evert.

The payoffs in the first three rows of the table are straightforward. When Evert uses her pure strategies DL and CC, her payoffs against Navratilova's pure strategies or the q-mix are exactly as in Figure 8.1. The third row also is analogous. When Evert uses Lob, we assume that her success percentages against Navratilova's DL and CC are, respectively, 70% and 60%. When Navratilova uses

		NAVRATILOVA		
		DL	CC	q-mix
EVERT	DL	50	80	$50q + 80(1 - q)$
	CC	90	20	$90q + 20(1 - q)$
	Lob	70	60	$70q + 60(1 - q)$
	p-mix	$50p_1 + 90p_2 + 70(1 - p_1 - p_2)$	$80p_1 + 20p_2 + 60(1 - p_1 - p_2)$	

FIGURE 8.6 Payoff Table for Tennis Point with Lob

[3]Even when a player has only two pure strategies, he may not use one of them in equilibrium. The other player then generally finds one of his strategies to be better against the one that the first player does use. In other words, the equilibrium "mixtures" collapse to the special case of pure strategies. But when one or both players have three or more strategies, we can have a genuinely mixed-strategy equilibrium where some of the pure strategies go unused.

her q-mix, using DL a fraction q of the time and CC a fraction $(1 - q)$ of the time, Evert's expected payoff from Lob is $70q + 60(1 - q)$; therefore that is the entry in the cell where Evert's Lob meets Navratilova's q-mix.

The really new feature here is the last row of the table. Evert now has three pure strategies, so she must now consider three different actions in her mix. The mix cannot be described by just one number p. Rather, we suppose that Evert plays DL with probability p_1 and CC with probability p_2, leaving Lob to get the remaining probability, $1 - p_1 - p_2$. Thus we need two numbers, p_1 and p_2, to define Evert's p-mix. Each of them, being a probability, must be between 0 and 1. Moreover, the two must add to something no more than 1; that is, they must satisfy the condition $p_1 + p_2 \leq 1$, because the probability $(1 - p_1 - p_2)$ of using Lob must be nonnegative.

Using this characterization of Evert's p-mix, then, we see that her expected payoff, when Navratilova plays her pure strategy DL, is given by $50p_1 + 90p_2 + 70(1 - p_1 - p_2)$. This is the entry in the first cell of the last row of the table in Figure 8.6. Evert's expected payoff from using her p-mix against Navratilova's CC is similarly $80p_1 + 20p_2 + 60(1 - p_1 - p_2)$. We do not show the expression for the payoff of mix against mix, because it is too long and we do not need it for our calculations.

Technically, before we begin looking for a mixed-strategy equilibrium, we should verify that there is no pure-strategy equilibrium. This is easy to do, however, so we leave it to you and turn to mixed strategies.

The easiest way to solve for a mixed-strategy equilibrium in a constant-sum game where one player has just two pure strategies, and the other has any number, is to use the minimax method from the perspective of the player who has just two pure strategies; here that player is Navratilova.[4] This approach works because Navratilova's mixture can be specified by using just one variable—namely, the single probability (q) used to define her mixed strategy. That probability, of choosing DL in this case, fully specifies her mixed strategy; after q is known, the probability of choosing CC is simply $(1 - q)$.

Figure 8.7 shows Evert's expected payoffs (success percentages) from playing each of her pure strategies DL, CC, and Lob as the q in Navratilova's q-mix varies over its full range from 0 to 1. These graphs are just those of the expressions in the right-hand column of Figure 8.6. Given the usual worst-case assumption that is appropriate in zero-sum games, Navratilova's calculation of her minimax strategy is as follows. For each q, if Navratilova were to choose that q-mix in equilibrium, Evert's best response would be to choose the strategy that gives her (Evert) the highest payoff. Evert's best response, which is also the worst-case

[4]For many amusing uses of this method, see John D. Williams, *The Compleat Strategyst* (New York: McGraw-Hill, 1954; reprint, New York: Dover, 1986).

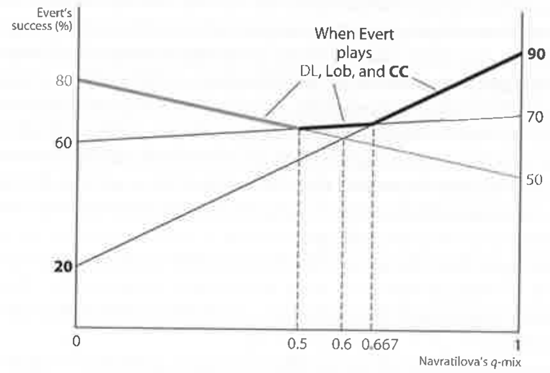

FIGURE 8.7 Diagrammatic Solution for Navratilova's q-Mix

scenario for Navratilova, is then shown on the highest of the three lines at that q. We show this set of worst-case outcomes with the thicker lines in Figure 8.7; these outcomes are formed from the *upper envelope* of the three payoff lines. Navratilova's optimal choice of q will make Evert's payoff as low as possible (thereby making her own as large as possible) from this set of worst-case outcomes.

To be more precise about Navratilova's optimal choice of q, we must calculate the coordinates of the kink points in the line showing her worst-case outcomes. The value of q at the left-most kink in this line makes Evert indifferent between DL and Lob. That q must equate the two payoffs from DL and Lob when used against the q-mix. Setting those two expressions equal gives us $50q + 80(1 - q) = 70q + 60(1 - q)$ and $q = 20/40 = 1/2$. This first kink is thus at $q = 0.5$, or 50%. Evert's expected payoff at this point is $50 \times 0.5 + 80 \times 0.5 = 70 \times 0.5 + 60 \times 0.5 = 65$. At the second (rightmost) kink, Evert is indifferent between CC and Lob. Thus the q value at this kink is the one that equates the CC and Lob payoff expressions. Setting $90q + 20(1 - q) = 70q + 60(1 - q)$, we find $q = 40/60$; the rightmost kink is at $q = 0.667$, or 66.7%. Here, Evert's expected payoff is $90 \times 0.667 + 20 \times 0.333 = 70 \times 0.667 + 60 \times 0.333 = 66.67$.

Now we can explicitly describe Evert's best responses to Navratilova's different choices of q. Evert's best response is DL when $q < 0.5$, CC when $q > 0.667$, and Lob when $0.5 < q < 0.667$. As usual, Evert's best response is pure for most values of q. When $q = 0.5$, Evert is indifferent between DL and Lob and therefore equally indifferent between those two pure strategies and any mixture of them.

When $q = 0.667$, she is indifferent between CC and Lob and therefore equally indifferent between those two pure strategies and any mixture of them.

Figure 8.7 also shows that, of all the worst-case scenarios for Navratilova, the best (or least bad) occurs at the left kink, where $q = 0.5$ and Evert's expected payoff is 65. The thick line shows all the maxima (for each q), and this point represents the minimum among them; this is Navratilova's minimax. At this point, Evert achieves the smallest of the payoffs associated with choosing her best response to each q that Navratilova might pick. Therefore, at $q = 0.5$, Navratilova achieves the largest of her worst-case payoffs, and she should choose this q in equilibrium.

When Navratilova chooses $q = 0.5$, Evert is indifferent between DL and Lob, and either of these choices gives her a better payoff than does CC. Therefore Evert will not use CC at all in equilibrium. CC will be an unused strategy in her equilibrium mix.

Now we can proceed with the equilibrium analysis as if this were a game with just two pure strategies for each player: DL and CC for Navratilova, and DL and Lob for Evert. We are back in familiar territory. Therefore we leave the calculation to you and just tell you the result. Evert's optimal mixture in this game entails her using DL with probability 0.25 and Lob with probability 0.75. Evert's expected payoff from this mixture, taken against Navratilova's DL and CC, respectively, is

$$50 \times 0.25 + 70 \times 0.75 = 80 \times 0.25 + 60 \times 0.75 = 65.$$

This payoff is Evert's maximin value, and it equals Navratilova's minimax, in conformity with the general result on mixed-strategy Nash equilibrium in zero-sum games. Thus, in equilibrium Evert mixes her DL and Lob with probabilities 0.25 and 0.75. Navratilova mixes her DL and CC with probabilities 0.5 each. The maximin (minimax) payoff to each is 65.

We could not have started our analysis with this two-by-two game, because we did not know in advance which of her three strategies Evert would not use. But we can be confident that in the general case there will be one such strategy. When the three expected payoff lines take the most general positions, they intersect pair by pair rather than all crossing at a single point. Then the upper envelope has the shape that we see in Figure 8.7. Its lowest point is defined by the intersection of the payoff lines associated with two of the three strategies. The payoff from the third strategy lies below the intersection at this point, so the player choosing among the three strategies does not use that third one.

B. Exceptional Cases

The positions and intersections of the three lines of Figure 8.7 depend on the payoffs specified for the pure strategies. We chose the payoffs for that particular game to show a general configuration of the lines. But if the payoffs stand in very specific relationships to each other, we can get some exceptional configurations with different results. We describe the possibilities here but leave it to you to redraw the diagrams for these cases.

First, if Evert's payoffs from Lob against Navratilova's DL and CC are equal, then the line for Lob is horizontal, and a whole range of q-values achieve Navratilova's minimax. For example, if the two payoffs in the Lob row of the table in Figure 8.6 are 70 each, then it is easy to calculate that the left kink in a revised Figure 8.7 would be at $q = 1/3$ and the right kink at $q = 5/7$. For any q in the range from 1/3 to 5/7, Evert's best response is Lob, and we get an unusual equilibrium in which Evert plays a pure strategy and Navratilova mixes. Further, Navratilova's equilibrium mixture probabilities are indeterminate within the range from $q = 1/3$ to $q = 5/7$.

Second, if Evert's payoffs from Lob against Navaratilova's DL and CC are lower than those of Figure 8.7 by just the right amounts (or those of the other two strategies are higher by just the right amounts), all three lines can meet in one point. For example, if the payoffs of Evert's Lob are 66 and 56 against Navratilova's DL and CC, respectively, instead of 70 and 60, then for $q = 0.6$ Evert's expected payoff from the Lob becomes $66 \times 0.6 + 56 \times 0.4 = 39.6 + 22.6 = 62$, the same as that from DL and CC when $q = 0.6$. Then Evert is indifferent among all three of her strategies when $q = 0.6$ and is equally willing to mix among them.

In this special case, Evert's equilibrium mixture probabilities are not fully determinate. Rather, a whole range of mixtures, including some where all three strategies are used, can do the job of keeping Navratilova indifferent between her DL and CC and therefore willing to mix. However, to achieve her minimax, Navratilova must use the mixture with $q = 0.6$. If she does not, Evert's best response will be to switch to one of her pure strategies, and this will work to Navratilova's detriment. We do not dwell on the determination of the precise range over which Evert's equilibrium mixtures can vary, because this case can only arise for exceptional combinations of the payoff numbers and is therefore relatively unimportant.

C. Case of Domination by a Mixed Strategy

What if Evert's payoffs from using her Lob against Navratilova's DL and CC are even lower than the values that make all three lines intersect in one point? Figure 8.8 illustrates such a payoff matrix.

When we graph the expected payoff lines from Evert's three pure strategies against Navratilova's choice of q, we now find that Lob has shifted down from its position in Figure 8.7. Figure 8.9 shows the new configuration of lines and the determination of Navratilova's minimax. The calculations that give us the positions of the three lines labeled DL, CC, and Lob, and their intersections follow

		NAVRATILOVA	
		DL	CC
EVERT	DL	50	80
	CC	90	20
	Lob	75	30

FIGURE 8.8 Tennis Point When Lob Is Never a Best Response

the same procedures as before; so we omit the details. The line labeled Mix is explained soon.

It is clear that, with these numbers, Lob is not a very good strategy for Evert. In fact, the line showing the payoffs from Lob lies everywhere below either DL or CC—and so below the upper envelope of those lines—as well as below the point of intersection of the DL and CC lines. Thus, for each q in Navratilova's mix, at least one of the pure strategies DL or CC gives Evert a higher payoff than does Lob. Figure 8.9 shows that if $q < 0.667$, DL is a better response for Evert than Lob; if $q > 0.4$, CC is better than Lob; and when $0.4 < q < 0.667$, both DL and CC are better than Lob. In other words, Lob is *never the best response* for Evert.

However, it is also true that Lob is not dominated by either DL or CC. If $q < 0.4$, Lob does better than CC, whereas if $q > 0.667$, Lob does better than

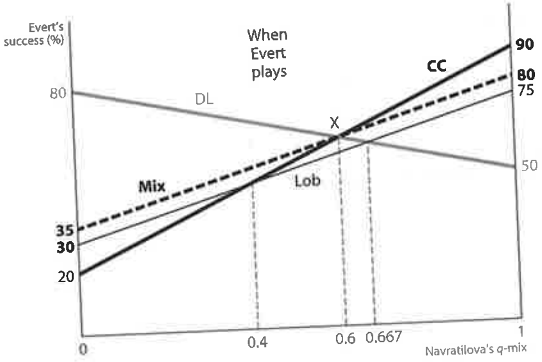

FIGURE 8.9 Domination by a Mixed Strategy

DL. Thus it appears that Lob cannot be eliminated from Evert's consideration by using dominance reasoning. Instead, we use reasoning below that is related to the concept of rationalizability from Chapter 5.

But now that we are allowing mixed strategies, we can consider mixtures of DL and CC as well. Can one such mixture dominate Lob? Yes. Consider a mixture between DL and CC with respective probabilities p and $(1 - p)$. Its payoff against Navratilova's pure DL is $50p + 90(1 - p)$, and its payoff against Navratilova's pure CC is $80p + 20(1 - p)$. Therefore its payoff against Navratilova's q-mix is $[50p + 90(1 - p)]\, q + [80p + 20(1 - p)](1 - q) = [50q + 80(1 - q)]p + [90q + 20(1 - q)](1 - p)$. This expression is an average of the heights of the DL and CC lines, with proportions p and $(1 - p)$.

In other words, Evert's expected payoff from her p-mix between DL and CC against Navratilova's q-mix is a straight line that averages the two separate DL and CC lines. Naturally, the average line must pass through the point of intersection of the two separate DL and CC lines; if, for some q, the two pure strategies DL and CC give equal payoffs, so must any mixture of them.

If $p = 1$ in Evert's mixture, she plays pure DL and the line for her mixture is just the DL line; if $p = 0$, the mixture line is the CC line. As p ranges from 0 to 1, the line for the p-mix rotates from the pure CC to the pure DL line, always passing through the point of intersection at $q = 0.6$. Among all such lines corresponding to different values of p, one will be parallel to the line for Lob. This line is shown in Figure 8.9 as the fourth (dashed) line labeled Mix. Because the line for Lob lies below the point of intersection of DL and CC, it also lies entirely below this parallel line representing a particular mixture of DL and CC. Then, no matter what q may be, Lob yields a lower expected payoff than this particular mixed strategy. In other words, Lob is dominated by this mixed strategy.

It remains to find the probabilities that define the mixture that dominates Lob for Evert. By construction, the Mix line is parallel to the Lob line and passes through the point at which $q = 0.6$ and the expected payoff is 62. Then it is easy to calculate the vertical coordinates of the Mix line. We use the fact that the slope of Lob $= (75 - 30)/(1 - 0) = 45$. Then the vertical coordinate of the Mix line when $q = 0$ is given by v in the equation $45 = (62 - v)/(0.6 - 0)$, or $v = 35$; similarly, when $q = 1$, the vertical coordinate solves $45 = (v - 62)/(1 - 0.6)$, which gives us $v = 80$. For any value of q between 0 and 1, the vertical coordinate of the Mix line is then $80 \times q + 35 \times (1 - q)$. Now compare this result with the expression for the general p-mix line just derived. For the two to be consistent, the part of that p-mix expression that multiplied q—$50p + 90(1 - p)$—must equal 80, and $80p + 20(1 - p)$ must equal 35. Both of them imply that $p = 0.25$.

Thus Evert's mixture of DL with probability 0.25 and CC with probability 0.75 yields the Mix line drawn in Figure 8.9. This mixed strategy gives her a better expected payoff than does her Lob, for each and every value of Navratilova's q. In other words, this particular mixture dominates Lob from Evert's perspective.[5]

Now we have shown that if a strategy is never a best response, then we can find a mixed strategy that dominates it.[6] In the process, we have expanded the scope of the concept of dominance to include domination by mixed strategies. The converse also is true; if a strategy is dominated by another strategy, albeit a mixed one, it can never be a best response to any of the other player's strategies. We can then use all of the other concepts associated with dominance that were developed in Chapter 4, but now allowing for mixed strategies also. We can do successive or iterated elimination of strategies, pure or mixed, that are dominated by other strategies, pure or mixed. If the process leaves just one strategy for each player, the game is dominance solvable and we have found a Nash equilibrium. More typically, the process only narrows down the range of strategies. In Chapter 5, we defined as *rationalizable* the set of strategies that remain after iterated elimination of pure strategies that are never best responses. Now we see that in two-player games we can think of rationalizable strategies as the set that survives after doing all possible iterated elimination of strategies that are dominated by other pure or mixed strategies.[7]

3 MIXING WHEN BOTH PLAYERS HAVE THREE STRATEGIES

As we saw in our two-by-three strategy example in the preceding section, a player mixing among three pure strategies can choose the probabilities of any two of them independently (as long as they are nonnegative and add up to no more than 1); then the probability of the third must equal 1 minus the sum of the probabil-

[5]We constructed the parallel line to guarantee dominance. Other lines through the point of intersection of the DL and CC lines also can dominate Lob as long as their slopes are not too different from that of the Lob line.

[6]In the example of Section 5.4 (Figure 5.6), we saw that Column's strategy C4 is never a best response but it is not dominated by any of the pure strategies C1, C2, or C3. Now we know that we can look for domination by a mixed strategy. In that game, it is easy to see that C4 is (strictly) dominated by an equal-probability mixture of C1 and C3.

[7]This equivalence between "never a best response" and "dominated by a mixed strategy" works fine in two-player games, but an added complication arises in many-player games. Consider a game with three players, A, B, and C. One of A's strategies—say, A1—may never be a best response to any pure strategies or independently mixed strategies of B and C, but A1 may fail to be dominated by any other pure or mixed strategy for A. However, if A1 is never a best response to any pure strategies or arbitrarily correlated mixed strategies of B and C, then it must be dominated by another of A's pure or mixed strategies. A complete treatment of this requires more advanced game theory, so we merely mention it. See Andreu Mas-Colell, Michael Whinston, and Jerry R. Green, *Microeconomic Theory* (New York: Oxford University Press, 1995), pp. 242–245 and 262–263.

ities of the other two. Thus we need two variables to specify a mix.[8] When both players have three strategies, we cannot find a mixed-strategy equilibrium without doing two-variable algebra. In many cases, such algebra is still manageable.

A. Full Mixture of All Strategies

Consider a simplified representation of a penalty kick in soccer. Suppose the kicker has just three pure strategies: kick to the left, right, or center. (As in Chapter 7, left and right refer to the goalie's left or right, for a right-footed kicker.) Then he can mix among these strategies, with probabilities denoted by p_L, p_R, and p_C, respectively. Any two of them can be taken to be the independent variables and the third expressed in terms of them. If p_L and p_R are made the two independent choice variables, then $p_C = 1 - p_L - p_R$. The goalie also has three pure strategies—namely, move to the kicker's left (the goalie's own right), move to the kicker's right, or continue to stand in the center—and can mix among them with probabilities q_L, q_R, and q_C, two of which can be chosen independently.

Best-response analysis also works in this context. The goalie would choose his two independent variables, say (q_L, q_R), as his best response to the kicker's two, (p_L, p_R), and vice versa. The Nash equilibrium occurs when each is choosing his best response to the other's mix. However, because each is choosing two magnitudes in response to the other's two, we can't use best-response graphs, because they would have to be in four dimensions.

Instead, we use the principle of the opponent's indifference, which enables us to focus on one player's mixture probabilities. They should be such that the other is indifferent among all the pure strategies that constitute his mixture. This gives us a set of equations that can be solved for the mixture probabilities. In the soccer example, the kicker's (p_L, p_R) would satisfy two equations expressing the requirement that the goalie's expected payoff from using his left should equal that from using his right and that the goalie's expected payoff from using his right should equal that from using his center. (Then the equality of expected payoffs from left and center follows automatically and is not a separate equation.) With more pure strategies, the number of the probabilities to be solved for and the number of equations that they must satisfy also increase.

Figure 8.10 shows the payoff matrix with the kicker as the Row player and the goalie as the Column player. (Unlike in the example in Chapter 7, these are not real data, but similar rounded numbers to simplify calculations.) Because the kicker wants to maximize the percentage probabilities that he successfully scores a goal and the goalie wants to minimize the probability that he lets the

[8]More generally, if a player has N pure strategies, then her mix has $(N - 1)$ independent variables, or "degrees of freedom of choice."

		GOALIE			
		Left	Center	Right	
KICKER	Left	45	90	90	PL
	Center	85	0	85	PC
	Right	95	95	60	PR
		GL	GC	GR	

FIGURE 8.10 Mixing in Soccer Penalty Kick When All Pure Strategies Are Used

goal through, this is a zero-sum game. We show the payoffs from the kicker's (the Row player's) perspective. For example, if the kicker kicks to his left while the goalie moves to the kicker's left (the top-left-corner cell), we suppose that the kicker still succeeds 45% of the time. But if the kicker kicks to his right and the goalie goes to the kicker's left, then the kicker has a 90% chance of scoring; we suppose a 10% probability that he might kick wide or too high. If the kicker kicks to his right (bottom row of the matrix), his probabilities of success are 95% if the goalie guesses wrong and 60% if the goalie guesses correctly. You can experiment with different payoff numbers that you think might be more appropriate.

It is easy to verify that the game has no equilibrium in pure strategies. So suppose the kicker is mixing with probabilities p_L, p_R, and $p_C = 1 - p_L - p_R$. Against each of the goalie's pure strategies, this mixture yields the following payoffs:

$$\begin{aligned}
&\text{Left:} && 45p_L + 85p_C + 95p_R = 45p_L + 85(1 - p_L - p_R) + 95p_R\\
&\text{Center:} && 90p_L + 0p_C + 95p_R = 90p_L + 95p_R\\
&\text{Right:} && 90p_L + 85p_C + 60p_R = 90p_L + 85(1 - p_L - p_R) + 60p_R
\end{aligned}$$

The goalie wants these numbers to be as small as possible. But in a mixed-strategy equilibrium, the kicker's mixture must be such that the goalie is indifferent among his pure strategies. Therefore all three of these expressions must be equal in equilibrium.

Equating the Left and Right expressions and simplifying, we have $45p_L = 35p_R$, or $p_R = (9/7)p_L$. Next, equate the Center and Right expressions and simplify, by using the link between p_L and p_R just obtained. This gives

$$\begin{aligned}
&90p_L + 95(9p_L/7) = 90p_L + 85[1 - p_L - (9p_L/7)] + 60(9p_L/7),\\
&\text{or} \quad [85 + 120(9/7)]\, p_L = 85,\\
&\text{which yields } p_L = 0.355.
\end{aligned}$$

Then we get $p_R = 0.355(9/7) = 0.457$, and finally $p_C = 1 - 0.355 - 0.457 = 0.188$. The kicker's payoff from this mixture against any of the goalie's pure strategies and therefore against any mixture of them can then be calculated by using any of the preceding three payoff lines; the result is 75.4.

The goalie's mixture probabilities can be found by writing down and solving the equations for the kicker's indifference among his three pure strategies against the goalie's mixture. We will do this in detail for a slight variant of the same game in Section 3.B, so we omit the details here and just give you the answer: $q_L = 0.325$, $q_R = 0.561$, and $q_C = 0.113$. The kicker's payoff from any of his pure strategies when played against the goalie's equilibrium mixture is 75.4. Note that this payoff is the same as the number found when calculating the kicker's mix; this is just the maximin = minimax property of zero-sum games.

Now we can interpret the findings. The kicker does better with his pure Right than his pure Left, both when the goalie guesses correctly (60 > 45) and when he guesses incorrectly (95 > 90). (Presumably the kicker is left-footed and can kick harder to his right.) Therefore the kicker chooses his Right with greater probability and, to counter that, the goalie chooses Right with the highest probability, too. However, the kicker should not and does not choose his pure-strategy Right; if he did so, the goalie would then choose his own pure-strategy Right, too, and the kicker's payoff would be only 60, less than the 75.4 that he gets in equilibrium.

B. Equilibrium Mixtures with Some Strategies Unused

In the preceding equilibrium, the probabilities of using Center in the mix are quite low for each player. The (Center, Center) combination would result in a sure save and the kicker would get a really low payoff—namely, 0. Therefore the kicker puts a low probability on this choice. But then the goalie also should put a low probability on it, concentrating on countering the kicker's more likely choices. But if the kicker gets a sufficiently high payoff from choosing Center when the goalie chooses Left or Right, then the kicker will choose Center with some positive probability. If the kicker's payoffs in the Center row were lower, he might then choose Center with zero probability; if so, the goalie would similarly put zero probability on Center. The game would reduce to one with just two basic pure strategies, Left and Right, for each player.

We show such a variant of the soccer game in Figure 8.11. The only difference in payoffs between this variant and the original game of Figure 8.10 is that the kicker's payoffs from (Center, Left) and (Center, Right) have been lowered even further, from 85 to 70. This might be because this kicker has the habit of kicking too high and therefore missing the goal when aiming for the center. Let us try to calculate the equilibrium here by using the same methods as in Section 3.A. This time we do it from the goalie's perspective; we try to find his mixture probabilities q_L, q_R, and q_C by using the condition that the kicker should be indifferent among all three of his pure strategies when played against this mixture.

		GOALIE		
		Left	Center	Right
KICKER	Left	45	90	90
	Center	70	0	70
	Right	95	95	60

FIGURE 8.11 Mixing in Soccer Penalty Kick When Not All Pure Strategies Are Used

The kicker's payoffs from his pure strategies are

Left: $45q_L + 90q_C + 90q_R = 45q_L + 90(1 - q_L - q_R) + 90q_R = 45q_L + 90(1 - q_L)$
Center: $70q_L + 0q_C + 70q_R = 70q_L + 70q_R$
Right: $95q_L + 95q_C + 60q_R = 95q_L + 95(1 - q_L - q_R) + 60\,q_R = 95(1 - q_R) + 60q_R$

Equating the Left and Right expressions and simplifying, we have $90 - 45q_L = 95 - 35q_R$, or $35q_R = 5 + 45q_L$. Next, equate the Left and Center expressions and simplify to get $90 - 45q_L = 70q_L + 70q_R$, or $115q_L + 70q_R = 90$. Substituting for q_R from the first of these equations (after multiplying through by 2 to get $70q_R = 10 + 90q_L$) into the second yields $205q_L = 80$, or $q_L = 0.390$. Then using this value for q_L in either of the equations gives $q_R = 0.644$. Finally, we use both of these values to obtain $q_C = 1 - 0.390 - 0.644 = -0.034$. Because probabilities cannot be negative, something has obviously gone wrong.

To understand what happens in this example, start by noting that Center is now a poorer strategy for the kicker than it was in the original version of the game, where his probability of choosing it was already quite low. But the concept of the opponent's indifference, expressed in the equations that led to the solution, means that the kicker has to be kept willing to use this poor strategy. That can happen only if the goalie is using his best counter to the kicker's Center—namely, the goalie's own Center—sufficiently infrequently. And in this example that logic has to be carried so far that the goalie's probability of Center has to become negative.

As pure algebra, the solution that we derived may be fine, but it violates the requirement of probability theory and real-life randomization that probabilities be nonnegative. The best that can be done in reality is to push the goalie's probability of choosing Center as low as possible—namely, to zero. But that leaves the kicker unwilling to use his own Center. In other words, we get a situation in which each player is not using one of his pure strategies in his mixture, that is, each is using it with zero probability.

Can there then be an equilibrium in which each player is mixing between his two remaining strategies—namely, Left and Right? If we regard this reduced

two-by-two game in its own right, we can easily find its mixed-strategy equilibrium. With all the practice that you have had so far, it is safe to leave the details to you and to state the result:

Kicker's mixture probabilities: $p_L = 0.4375$, $p_R = 0.5625$;
Goalie's mixture probabilities: $q_L = 0.3750$, $q_R = 0.6250$;
Kicker's expected payoff (success percentage): 73.13.

We found this result by simply removing the two players' Center strategies from consideration on intuitive grounds. But we must check that it is a genuine equilibrium of the full three-by-three game. That is, we must check that neither player finds it desirable to bring in his third strategy, given the mixture of two strategies chosen by the other player.

When the goalie is choosing this particular mixture, the kicker's payoff from pure Center is $0.375 \times 70 + 0.625 \times 70 = 70$. This payoff is less than the 73.13 that he gets from either of his pure Left or pure Right or any mixture between the two, so the kicker does not want to bring his Center strategy into play. When the kicker is choosing the two-strategy mixture with the preceding probabilities, his payoff against the goalie's pure Center is $0.4375 \times 90 + 0.5625 \times 95 = 92.8$. This number is higher than the 73.13 that the kicker would get against the goalie's pure Left or pure Right or any mixture of the two and is therefore worse for the goalie. Thus the goalie does not want to bring his Center strategy into play either. The equilibrium that we found for the two-by-two game is indeed an equilibrium of the three-by-three game.

To allow for the possibility that some strategies may go unused in an equilibrium mixture, we must modify or extend the "opponent's indifference" principle. Each player's equilibrium mix should be such that the other player is indifferent among all the strategies *that are actually used in his equilibrium mix*. The other player is not indifferent between these and his unused strategies; he prefers the ones used to the ones unused. In other words, against the opponent's equilibrium mix, all of the strategies used in your own equilibrium mix should give you the same expected payoff, which in turn should be higher than what you would get from any of your unused strategies. This is called the principle of *complementary slackness;* we consider it in greater generality in the next section, where the reason for this strange-sounding name will become clearer.

4 MORE COUNTERINTUITIVE PROPERTIES OF MIXED STRATEGIES

In Chapter 7, Section 5, we pointed out, discussed, and explained some features of mixed-strategy equilibria that appear counterintuitive at first sight. We postponed the discussion of some other counterintuitive outcomes

		OPPONENT EXPECTS	
		P	R
YOU PLAY	P	c	b
	R	a	d

FIGURE 8.12 Table of Success Probabilities of Risky and Percentage Plays

because they required somewhat more mathematics. We take these additional examples up now.

A. Risky and Safe Choices in Zero-Sum Games

In sports, there are always some strategies that are relatively safe; they do not fail disastrously even if anticipated by the opponent but do not do very much better even if unanticipated. Other strategies are risky; they do brilliantly if the other side is not ready for them but fail miserably if the other side is ready. Thus in football, on third down with a yard to go, a run up the middle is safe and a long pass is risky. The following example incorporates this idea of safe–versus–risky strategies. In addition, although most of our examples use illustrative numbers for payoffs, here we change that practice to emphasize the generality of the problem. Therefore, we let the payoffs be general algebraic symbols, subject only to some conditions concerning the basic strategy being considered.

Consider any zero-sum game in which you have two pure strategies. Let us call the relatively safe strategy (the percentage play) P and the riskier strategy R. The opponent has two pure strategies that we also call P and R; his P is his best response to your P, as is his R to your R. Figure 8.12 shows the table of probabilities that your play succeeds; these are not your payoffs. The sense of "safe" and "risky" is captured by requiring $a > b > c > d$. The risky play does really well if the opponent is not prepared for it (your success probability is a) but really badly if he is (your success probability is d); the percentage play does moderately well in either case (you succeed with probability b or c) but a little worse if the opponent expects it ($c < b$).

Let your payoff or utility be W if your play succeeds and L if it fails. A "really big occasion" is when W is much bigger than L. Note that W and L are not necessarily money amounts, so they can be utilities that capture any aversion to risk, as explained in the Appendix to Chapter 7. Now we can write down the table of expected payoffs from the various strategy combinations as in Figure 8.13. Note how this table is constructed. For example, if you play P and your

		OPPONENT EXPECTS	
		P	R
YOU PLAY	P	$cW + (1 - c)L$	$bW + (1 - b)L$
	R	$aW + (1 - a)L$	$dW + (1 - d)L$

FIGURE 8.13 Payoff Table with Risky and Percentage Plays

opponent expects R, then you get utility W with probability b and utility L with probability $(1 - b)$; your expected payoff is $bW + (1 - b)L$. This game is zero-sum; so in each cell your opponent's payoffs are just the negative of yours.

In the mixed-strategy equilibrium, your probability p of choosing P is defined by the opponent's indifference property; therefore

$$p[cW + (1 - c)L] + (1 - p)[aW + (1 - a)L] = \\ p[bW + (1 - b)L] + (1 - p)[dW + (1 - d)L]$$

This equation simplifies to $p = (a - d)/[(a - d) + (b - c)]$. Because $(b - c)$ is small in relation to $(a - d)$, we see that p is close to 1. That is exactly why the strategy P is called the percentage play; it is the normal play in these situations, and the risky strategy R is played only occasionally to keep the opponent guessing or, in football commentators' terminology, "to keep the defense honest."

The interesting part of this result is that the expression for p is completely independent of W and L. That is, the theory says that you should mix the percentage play and the risky play in exactly the same proportions on a big occasion as you would on a minor occasion. This runs against the intuition of many people. They think that the risky play should be engaged in less often when the occasion is more important. Throwing a long pass on third down with a yard to go may be fine on an ordinary Sunday afternoon in October, but doing so in the Super Bowl is too risky.

So which is right: theory or intuition? We suspect that readers will be divided on this issue. Some will think that the sports commentators are wrong and will be glad to have found a theoretical argument to refute their claims. Others will side with the commentators and argue that bigger occasions call for safer play. Still others may think that bigger risks should be taken when the prizes are bigger, but even they will find no support in the theory, which says that the size of the prize or the loss should make no difference to the mixture probabilities.

On many previous occasions when discrepancies between theory and intuition arose, we argued that the discrepancies were only apparent, that they were the result of failing to make the theory sufficiently general or rich enough to capture all the features of the situation that created the intuition, and that improving the theory removed the discrepancy. This one is different; the problem is fundamental to the calculation of payoffs from mixed strategies as probability-weighted averages or expected payoffs. And almost all of existing game theory has this starting point.[9]

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

In Chapter 7, Section 5.C, we described a counterintuitive property of mixed strategies in zero-sum games. If a player improves the payoffs from one of his pure strategies, the probability of using that strategy in the equilibrium mixture can go down. Here we demonstrate an even more general and more surprising result in general non-zero-sum games. One player's equilibrium mixture probabilities depend only on the other player's payoffs, not on his own. Consider a general two-by-two non-zero-sum game with the payoff table shown in Figure 8.14. In actual games, the payoffs would be actual numbers and the strategies would have particular names. In this example, we again use general algebraic symbols for payoffs so that we can examine how the probabilities of the equilibrium mixtures depend on them. Similarly, we use arbitrary generic labels for the strategies.

Suppose the game has a mixed-strategy equilibrium in which Row plays Up with probability p and Down with probability $(1 - p)$. To guarantee that Column also mixes in equilibrium, Row's p-mix must keep Column indifferent between

		COLUMN	
		Left	Right
ROW	Up	a, A	b, B
	Down	c, C	d, D

FIGURE 8.14 General Algebraic Payoff Matrix for Two-by-Two Non-Zero-Sum Game

[9]Vincent P. Crawford, "Equilibrium Without Independence," *Journal of Economic Theory*, vol. 50, no. 1 (February 1990), pp. 127–154; and James Dow and Sergio Werlang, "Nash Equilibrium Under Knightian Uncertainty," *Journal of Economic Theory*, vol. 64, no. 2 (December 1994), pp. 305–324, are among the few research papers that suggest alternative foundations for game theory. And our exposition of this problem in the first edition of this book inspired an article that uses such new methods on it: Simon Grant, Atsushi Kaji, and Ben Polak, "Third Down and a Yard to Go: Recursive Expected Utility and the Dixit-Skeath Conundrum," *Economic Letters*, vol. 73, no. 3 (December 2001), pp. 275–286. Unfortunately, it uses more advanced concepts than those available at the introductory level of this book.

his two pure strategies, Left and Right. Equating Column's expected payoffs from these two strategies when played against Row's mixture, we have[10]

$$pA + (1 - p)C = pB + (1 - p)D, \quad \text{or}$$
$$p = (D - C)/[(A - B) + (D - C)]$$

The surprising thing about the expression for p is not what it contains, but what it does *not* contain. None of Row's own payoffs, a, b, c, or d, appear on the right-hand side. Row's mixture probabilities are totally independent of his own payoffs!

Similarly, the equilibrium probability q of Column playing Left is given by

$$q = (d - b)/[(a - c) + (d - b)].$$

Column's equilibrium mixture also is determined independently of his own payoffs.

The surprising or counterintuitive aspect of these results is resolved if you remember the general principle of the opponent's indifference. Because each player's mixture probabilities are solved by requiring the opponent to be indifferent between his pure strategies, it is natural that these probabilities should depend on the opponent's payoffs, not on one's own. But remember also that it is only in zero-sum games that a player has a genuine reason to keep the opponent indifferent. There, any clear preference of the opponent for one of his pure strategies would work to one's own disadvantage. In non-zero-sum games, the opponent's indifference does not have any such purposive explanation; it is merely a logical property of equilibrium in mixed strategies.

5 MIXING AMONG ANY NUMBER OF STRATEGIES: GENERAL THEORY

We conclude this chapter with some general theory of mixed-strategy equilibria, to unify all of the ideas introduced in the various examples in Chapters 7 and 8 so far. Such general theory unavoidably requires some algebra and some abstract thinking. Readers unprepared for such mathematics or averse to it can omit this section without loss of continuity.

Suppose the Row player has available the pure strategies $R_1, R_2, \ldots, R_m$, and the Column player has strategies $C_1, C_2, \ldots C_n$. Write the Row player's payoff from the strategy combination (i, j) as A_{ij}, and Column's as B_{ij}, where the index i ranges from 1 to m, and the index j ranges from 1 to n. We allow each player

[10]For there to be a mixed-strategy equilibrium, the probability p must be between 0 and 1. This requires that $(A - B)$ and $(D - C)$ have the same sign; if A is bigger than B, then D must be bigger than C, and, if A is smaller than B, then D must be smaller than C. (Otherwise, one of the pure strategies, Left or Right, would dominate the other.)

to mix the available pure strategies. Suppose the Row player's p-mix has probabilities P_i and the Column player's q-mix has Q_j. All these probabilities must be nonnegative, and each set must add to 1, so

$$P_1 + P_2 + \cdots + P_m = 1 = Q_1 + Q_2 + \cdots + Q_n.$$

We write V_i for Row's expected payoff from using his pure strategy i against Column's q-mix. Using the reasoning that we have already seen in several examples, we have

$$V_i = A_{i1}Q_1 + A_{i2}Q_2 + \cdots + A_{in}Q_n = \sum_{j=1}^{n} A_{ij}\,Q_j,$$

where the last expression on the right uses the mathematical notation for summation of a collection of terms. When Row plays his p-mix and it is matched against Column's q-mix, Row's expected payoff is

$$P_1V_1 + \cdots + P_mV_m = \sum_{i=1}^{m} P_iV_i = \sum_{i=1}^{m}\sum_{j=1}^{n} P_iA_{ij}Q_j.$$

The Row player chooses his p-mix to maximize this expression.

Similarly, writing W_j for Column's expected payoff when his pure strategy j is pitted against Row's p-mix, we have

$$W_j = P_1B_{1j} + P_2B_{2j} + \cdots + P_mB_{mj} = \sum_{i=1}^{m} P_iB_{ij}.$$

Pitting mix against mix, we have Column's expected payoff:

$$Q_1W_1 + \cdots + Q_nW_n = \sum_{j=1}^{n} Q_jW_j = \sum_{i=1}^{m}\sum_{j=1}^{n} P_iB_{ij}Q_j,$$

and he chooses his q-mix to maximize this expression.

We have a Nash equilibrium when each player simultaneously chooses his best mix, given that of the other. That is, Row's equilibrium p-mix should be his best response to Column's equilibrium q-mix, and vice versa. Let us begin by finding Row's best-response rule. That is, let us temporarily fix Column's q-mix and consider Row's choice of his p-mix.

Suppose that, against Column's given q-mix, Row has $V_1 > V_2$. Then Row can increase his expected payoff by shifting some probability from strategy R_2 to R_1; that is, Row reduces his probability P_2 of playing R_2 and increases the probability P_1 of playing R_1 by the same amount. Because the expressions for V_1 and V_2 do not include any of the probabilities P_i at all, this is true no matter what the original values of P_1 and P_2 were. Therefore Row should reduce the probability P_2 of playing R_2 as much as possible—that is, all the way to zero.

The idea generalizes immediately. Row should rank the V_i in descending order. At the top there may be just one strategy, in which case it should be the only one used; that is, Row should then use a pure strategy. Or there may be a tie among two or more strategies at the top, in which case Row should mix solely among these strategies and not use any of the others.[11] When there is such a tie, all mixtures between these strategies give Row the same expected payoff. Therefore this consideration alone does not fix Row's equilibrium p-mix. We show later how, in a way that may seem somewhat strange at first sight, Column's indifference condition does that job.

The same argument applies to Column. He should use only that pure strategy which gives him the highest W_j or should mix only among those of his pure strategies C_j whose W_j are tied at the top. If there is such a tie, then all mixtures are equally good from Column's perspective; the probabilities of the mix are not fixed by this consideration alone.

In general, for most values of $(Q_1, Q_2, \ldots, Q_n)$ that we hold fixed in Column's q-mix, Row's $V_1, V_2, \ldots, V_m$ will not have any ties at the top, and therefore Row's best response will be a pure strategy. Conversely, Column's best response will be one of his pure strategies for most values of $(P_1, P_2, \ldots, P_m)$ that we hold fixed in Row's p-mix. We saw this several times in the examples of Chapters 7 and 8; for example, in Figure 8.2, for most values of p in Row's p-mix, the best q for Column was either 0 or 1, and vice versa. For only one critical value of Row's p was it optimal for Column to mix (choose any q between 0 and 1), and vice versa.

All of these conditions—ties at the top, and worse outcomes from the other strategies—constitute the complicated set of equations and inequalities that, when simultaneously satisfied, defines the mixed-strategy Nash equilibrium of the game. To understand it better, suppose for the moment that we have done all the work and found which strategies are used in the equilibrium mix. We can always relabel the strategies so that Row uses, say, the first g pure strategies, $R_1, R_2, \ldots, R_g$, and does not use the remaining $(m - g)$ pure strategies, $R_{g+1}, R_{g+2}, \ldots, R_m$, while Column uses his first h pure strategies, $C_1, C_2, \ldots, C_h$, and does not use the remaining $(n - h)$ pure strategies, $C_{h+1}, C_{h+2}, \ldots, C_n$. Write V for the tied value of Row's top expected payoffs V_i and, similarly, W for the tied value of Column's top expected payoffs W_j. Then the equations and inequalities can be written as follows. First, for each player, we set the probabilities of the unused strategies equal to zero and require those of the used strategies to sum to 1:

$$P_1 + P_2 + \cdots + P_g = 1, \quad P_{g+1} = P_{g+2} = \cdots P_m = 0 \tag{8.1}$$

and

$$Q_1 + Q_2 + \cdots + Q_h = 1, \quad Q_{h+1} = Q_{h+2} = \cdots Q_n = 0 \tag{8.2}$$

[11] In technical mathematical terms, the expression $\Sigma_i P_iV_i$ is *linear* in the P_i; therefore its maximum must be at an extreme point of the set of permissible P_i.

Next we set Row's expected payoffs for the pure strategies that he uses equal to the top tied value:

$$V = A_{i1} Q_1 + A_{i2} Q_2 + \cdots + A_{ih} Q_h \quad \text{for } i = 1, 2, \ldots, g, \tag{8.3}$$

and note that his expected payoffs from his unused strategies must be smaller (that is why they are unused):

$$V > A_{i1} Q_1 + A_{i2} Q_2 + \cdots + A_{ih} Q_h \quad \text{for } i = g+1, g+2, \ldots, n. \tag{8.4}$$

Next, we do the same for Column, writing W for his top tied payoff value:

$$W = P_1 B_{1j} + P_2 B_{2j} + \cdots + P_g B_{gj} \quad \text{for } j = 1, 2, \ldots, h, \tag{8.5}$$

and

$$W > P_1 B_{1j} + P_2 B_{2j} + \cdots + P_g B_{gj} \quad \text{for } j = h+1, h+2, \ldots, n. \tag{8.6}$$

To find the equilibrium, we must take this whole system; regard the choice of g and h as well as the probabilities $P_1, P_2, \ldots, P_g$ and $Q_1, Q_2, \ldots, Q_h$ as unknowns; and attempt to solve for them.

There is always the exhaustive search method. Try a particular selection of g and h; that is, choose a particular set of pure strategies as candidates for use in equilibrium. Then take Eqs. (8.1) and (8.5) as a set of $(h + 1)$ simultaneous linear equations regarding $P_1, P_2, \ldots, P_g$ and W as $(g + 1)$ unknowns; solve for them; and check if the solution satisfies all the inequalities in Eq. (8.6). Similarly, take Eqs. (8.2) and (8.3) as a set of $(g + 1)$ simultaneous linear equations in the $(h + 1)$ unknowns $Q_1, Q_2, \ldots, Q_h$ and V; solve for them; and check if the solution satisfies all the inequalities in Eq. (8.4). If all these things check out, we have found an equilibrium. If not, take another selection of pure strategies as candidates and try again. There is only a finite number of selections. Row can use $(2^m - 1)$ possible pure strategies in his mix and Column can use $(2^n - 1)$ possible pure strategies in his mix. Therefore the process must end successfully after a finite number of attempts.

When m and n are reasonably small, exhaustive search is manageable. Even then, shortcuts suggest themselves in the course of the calculation for each specific problem. Thus, in the second variant of our soccer penalty kick, the way in which the attempted solution with all strategies used failed told us which strategy to discard in the next attempt.

Even for moderately large problems, however, solutions based on exhaustive search or ad hoc methods become too complex. That is when one must resort to more systematic computer searches or algorithms. What these computer algorithms do is to search simultaneously for solutions to two linear maximization (or linear programming, in the terminology of decision theory) problems: given a q-mix, and therefore all the V_i values, choose a p-mix to maximize Row's expected payoff $\Sigma_i P_i V_i$; and given a p-mix and therefore all the W_j values, choose a q-mix to maximize Column's expected payoff $\Sigma_j Q_j W_j$. However, for a typical q-mix, all the V_i values will be unequal. If Column were to play this q-mix in an actual game, Row would not mix but would instead play just the one pure strategy that gave him the highest V_i. But in our numerical solution method we should not adjust Row's strategy in this drastic fashion. If we did, then at the next step of our algorithm, Column's best q-mix also would change drastically, and Row's chosen pure strategy would no longer look so good. Instead, the algorithm should take a more gradual step, adjusting the p-mix a little bit to improve Row's expected payoff. Then, with the use of this new p-mix for Row, the algorithm should adjust Column's q-mix a little bit to improve his expected payoff. Then back again to another adjustment in the p-mix. The method proceeds in this way until no improvements can be found; that is the equilibrium.

We do not need the details of such procedures, but the general ideas that we have developed above already tell us a lot about equilibrium. Here are some important lessons of this kind.

1. We solve for Row's equilibrium mix probabilities $P_1, P_2, \ldots, P_g$ from Eqs. (8.1) and (8.5). The former is merely the adding-up requirement for probabilities. The more substantive equation is (8.5), which gives the conditions under which Column gets the same payoff from all the pure strategies that he uses against the p-mix. It might seem puzzling that Row adjusts his mix so as to keep Column indifferent when Row is concerned about his own payoffs, not Column's. Actually the puzzle is only apparent. We derived those conditions [Eq. (8.5)] by thinking about Column's choice of his q-mix, motivated by concerns about his own payoffs. We argued that Column would use only those strategies that gave him the best (tied) payoffs against Row's p-mix. This is the requirement embodied in Eq. (8.5). Even though it appears *as if* Row is deliberately choosing his p-mix so as to keep Column indifferent, the actual force that produces this outcome is Column's own purposive strategic choice.

In Chapter 7, we gave the name "the opponent's indifference property" to the idea that each player's indifference conditions constitute the equations that can be solved for the other player's equilibrium mix. We now have a proof of this principle for general games, zero-sum and non-zero-sum.

2. However, in the zero-sum case, the idea that each player chooses his mixture to keep the other indifferent is not just an *as if* matter; there is some genuine reason that a player should behave in this way. When the game is zero-sum, we have a natural link between the two players' payoffs: $B_{ij} = -A_{ij}$ for all i and j, and then similar relations hold among all the combinations and expected payoffs, too. Therefore we can multiply Eqs. (8.5) and (8.6) by -1 to write them in

terms of Row's payoffs rather than Column's. We write these "zero-sum versions" of the conditions as Eqs. (8.5z) and (8.6z):

$$V = P_1 A_{1j} + P_2 A_{2j} + \cdots + P_g A_{gj} \quad \text{for } j = 1, 2, \ldots, h \tag{8.5z}$$

and

$$V < P_1 A_{1j} + P_2 A_{2j} + \cdots + P_g A_{gj} \quad \text{for } j = h+1, h+2, \ldots, n. \tag{8.6z}$$

(Note that multiplying by -1 to go from Eq. (8.6) to Eq. (8.6z) reverses the direction of the inequality.)

Of these, Eqs. (8.5) and (8.5z) tell us that as long as Row is using his equilibrium mix, Column (and therefore Row, too, in this zero-sum game) gets the same payoff from any of the pure strategies that he actually uses in equilibrium. Column cannot do any better for himself—and therefore in this zero-sum game cannot cause any harm to Row—by choosing one of those strategies rather than another. What is more, Eq. (8.6z) tells us that were Column to use any of the other strategies, Row would do even better. In other words, these conditions tell us that Column cannot exploit Row's equilibrium mix. Thus we see in a more general setting the purposive role of mixing in zero-sum games that we saw in the examples of Chapter 7; we also see more explicitly why it works only for zero-sum games.

3. Now we return to the general, non-zero-sum, case. Note that the system comprising Eqs. (8.1) and (8.5) has $(h + 1)$ linear equations and $(g + 1)$ unknowns. In general, such a system has no solution if $h > g$, has exactly one solution if $h = g$, and has many solutions if $h < g$. Conversely, the system comprising Eqs. (8.2) and (8.3) has $(g + 1)$ linear equations and $(h + 1)$ unknowns. In general, such a system has no solution if $g > h$, has exactly one solution if $g = h$, and has multiple solutions if $g < h$. Because in equilibrium we want both systems to be satisfied, in general we need $g = h$. Thus in a mixed-strategy equilibrium, the two players use equal numbers of pure strategies.

We keep on saying "in general" because exceptions can arise for fortuitous combinations of coefficients and right-hand sides of equations. In particular, it is possible for too many equations in too few unknowns to have solutions. This is just what happens in the "exceptional cases" mentioned in Section 2.B of this chapter.

4. We observe a very particular relation between the use of strategies and their payoffs. Row uses strategies P_1 to P_g with positive probabilities, and Eq. (8.3) shows that he gets exactly the payoff V when any one of these pure strategies is played against Column's equilibrium mix. For the remaining pure strategies in Row's armory, Eq. (8.4) shows that they yield a lower payoff than V when played against Column's equilibrium mix, and then they are not used; that is, P_{h+1} to P_m are all zero. In other words, for any i, it is *impossible* to have

$$\text{both } V > A_{i1}Q_1 + A_{i2}Q_2 + \cdots + A_{ih}Q_h \quad \text{and} \quad P_i > 0.$$

At least one of these inequalities must collapse into equality. This is known as the principle of **complementary slackness,** and it is of great importance in the

general theory of games and equilibria, as well as in mathematical optimization (programming).

5. Back to the zero-sum case. When both players choose their equilibrium mix, Row's expected payoff is

$$V = \sum_{i=1}^{m} \sum_{j=1}^{n} P_i A_{ij} Q_j,$$

and Column's is just the negative of this payoff. Moreover, the equilibrium comes about when Row, for the given q-mix, chooses his p-mix to maximize this expression, and simultaneously Column, for the given p-mix, chooses his q-mix to maximize the negative of the same expression, or to minimize the same expression. If we regard the expression as a function of all the P_i and the Q_j, therefore, and graph it in a sufficiently high-dimensional space, it will look like a saddle. The front-to-back cross section of a saddle looks like a valley or a U, with its minimum at the middle, while the side-to-side cross section looks like a peak or an inverted U, with its maximum at the middle. If each player has just two pure strategies, the p-mix and q-mix can each be described by a single number—say, the probability of choosing the first strategy. (For each player, the probability of choosing his second pure strategy is then just one minus that of choosing his first pure strategy.) We can then draw a graph in three dimensions, where the x and y axes are in a horizontal plane and the z axis points vertically upward. The p-mix is shown along the x-axis, the q-mix along the y-axis, and the value V along the z-axis. The cross section of this

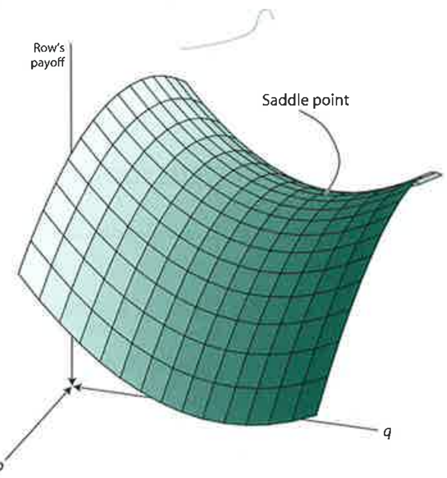

FIGURE 8.15 Saddle Point

saddle-shaped surface along the x direction will show the maximization of V with respect to p, therefore a peak. And the cross section along the y direction will show the minimization of V with respect to q, therefore a valley. Thus the graph will look like a saddle, as illustrated in Figure 8.15. Such an equilibrium is called a **saddle point.**

The value of V in equilibrium—that is, the simultaneous maximum with respect to the P_i and the minimum with respect to the Q_j—is called the *minimax* value of the zero-sum game. The idea of such an equilibrium, as well as the formulation of the conditions such as Eqs. (8.5z) and (8.6z) that define it, was the first important achievement of game theory and appeared in the work of von Neumann and Morgenstern in the 1940s. It is called their minimax theorem.

SUMMARY

Best-response analysis can be used to solve for mixed-strategy equilibria. The best-response-curve diagram can be used to show all mixed-strategy as well as pure-strategy equilibria of a game.

If one player has three strategies and the other has only two, the player with three available pure strategies will generally use only two in her equilibrium mix. In some exceptional cases, equilibrium mixtures may also be indeterminate.

When one or both players have three (or more) strategies, equilibrium mixed strategies may put positive probability on all pure strategies or may include only a subset of the pure strategies. All strategies that are actively used in the mixture yield equal expected payoff against the opponent's equilibrium mix; all the ones unused yield lower expected payoff. This is the principle of complementary slackness.

KEY TERMS

complementary slackness (290) **saddle point** (292)

SOLVED EXERCISES

S1. "When a zero-sum 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.

S2. Sections 1.A and 1.B illustrate how to use a best-response graph to find all of the Nash equilibria of the tennis-point game and the assurance game respectively. For the battle of the sexes game from Figure 4.13, graph the best responses of Harry and Sally on a p-q coordinate plane. Label all of the Nash equilibria.

S3. Revisit Exercise U9 from Chapter 7. For parts (a) and (b) below, graph the best-response curves of Row and Column on a p-q coordinate plane. On the graphs, identify and label the pure-strategy Nash equilibria and any mixed-strategy Nash equilibria.

(a) Let $x = 3$.

(b) Let $x = 1$.

S4. Recall the exceptional cases of the three-by-two tennis-point game described in Section 2.B.

(a) Draw the diagram of Evert's success rate relative to Navratilova's q-mix (similar to Figure 8.7) for the case where Evert's payoff from playing Lob is 70 when Navratilova plays either DL or CC. Use this diagram to explain why it is an equilibrium for Evert to play pure Lob and for Navratilova to mix 50-50.

(b) Describe all possible Nash equilibria for the game given in part (a).

(c) Draw the diagram of Evert's success rate as a function of Navratilova's q-mix (similar to Figure 8.7) for the case where the the payoffs of Evert's Lob are 66 and 56 against Navratilova's DL and CC, respectively. Use this diagram to illustrate that Navratilova's equilibrium q must be 0.6.

(d) As the text notes, there are a number of different Nash equilibria for the game in part (c) involving different mixtures for Evert. Write down the indifference equation for Navratilova in terms of p_1 and p_2. (Hint: Look at Figure 8.6.) Note that this equation can be satisfied by many values of p_1 and p_2. Find the largest and smallest possible values for p_1 and p_2.

S5. Consider the following game:

		OLLISTAN	
		Laurel	Hardy
KARL	Groucho	9	2
	Harpo	7	5
	Chico	5	6
	Zeppo	4	9
	Gummo	1	8

(a) On a single graph, plot the expected payoffs from each of Karl's strategies as a function of Ollistan's q-mix.

(b) Which strategies appear on the upper envelope of the graph in part (a)?
(c) Which strategies are never a best response for Karl? Why?
(d) What is the mixed-strategy Nash equilibrium of this game?

S6. Consider the following game:

		PROFESSOR PLUM		
		Revolver	Knife	Wrench
MRS. PEACOCK	Conservatory	1, 3	2, −2	0, 6
	Ballroom	3, 1	1, 4	5, 0

(a) Graph the expected payoffs from each of Professor Plum's strategies as a function of Mrs. Peacock's p-mix.
(b) Over what range of p does Revolver yield a higher expected payoff for Professor Plum than Knife?
(c) Over what range of p does Revolver yield a higher expected payoff than Wrench?
(d) Which pure strategies will Professor Plum use in his equilibrium mixture? Why?
(e) What is the mixed-strategy Nash equilibrium of this game?

S7. Find all Nash equilibria of the following game. (Hint: Look for dominated strategies.)

		GREEK			
		α	β	γ	δ
ROMAN	A	6	−1	5	4
	B	5	−2	2	1
	C	3	7	4	8

S8. Find all Nash equilibria of the Roman-Greek game in Exercise S7 when the payoff to (A, γ) changes from 5 to 3.

S9. Many of you will be familiar with the children's game rock–paper–scissors. In rock–paper–scissors, two people simultaneously choose either "rock," "paper," or "scissors," usually by putting their hands into the shape of one of the three choices. The game is scored as follows. A person choosing Scissors beats a person choosing Paper (because scissors cut paper). A person choosing Paper beats a person choosing Rock (because paper covers rock). A person choosing Rock beats a person choosing Scissors (because rock breaks scissors). If two players choose the same object, they tie. Suppose that each individual play of the game is worth 10 points. The following matrix shows the possible outcomes in the game:

		PLAYER 2		
		Rock	Scissors	Paper
PLAYER 1	Rock	0	10	−10
	Scissors	−10	0	10
	Paper	10	−10	0

(a) Find the mixed-strategy equilibrium of this rock–paper–scissors game.
(b) Suppose that Player 2 announced that she would use a mixture in which her probability of choosing Rock would be 40%, her probability of choosing Scissors would be 30%, and her probability of Paper, 30%. What is Player 1's best response to this strategy choice by Player 2? Explain why your answer makes sense, given your knowledge of mixed strategies.

S10. Harry and Sally have planned to go out to eat tonight. Harry strongly prefers pasta, which Sally finds disagreeable. Sally really likes a particular sandwich place, but Harry's not a fan. Another option is a buffet, which would allow each of them to get a lower-quality version of the food they'd most like to eat. Harry and Sally enjoy one another's company and would rather eat together than apart, but they are abysmally poor communicators. Once again they have neglected to decide beforehand where to go, and they have both absentmindedly allowed their cell-phone batteries to die. Their payoffs are as follows:

		SALLY		
		Pasta	Sandwich	Buffet
HARRY	Pasta	5, 1	0, 0	0, 0
	Sandwich	0, 0	1, 5	0, 0
	Buffet	0, 0	0, 0	2, 2

The pure-strategy Nash equilibria are easy to see, but this game also has multiple mixed-strategy Nash equilibria.
(a) Show that it is a mixed-strategy equilibrium when Harry plays Pasta 5/6 of the time, Sandwich 1/6 of the time, and Buffet never, whereas Sally plays Pasta 1/6 of the time, Sandwich 5/6 of the time, and Buffet never. What is the expected payoff to each player in this equilibrium?

(b) What mixed-strategy equilibrium results when both Harry and Sally mix only over Pasta and Buffet? What is the expected payoff to each player in this equilibrium? Explain why Sally would never play Sandwich if she knew that Harry was playing his mixed strategy for this equilibrium.

(c) What mixed-strategy equilibrium results when both Harry and Sally mix only over Sandwich and Buffet? What is the expected payoff to each player in this equilibrium? Explain why Harry would never play Pasta if he knew that Sally was playing her mixed strategy for this equilibrium.

(d) There is also a mixed-strategy equilibrium where both Harry and Sally play all three of their strategies with positive probability. What is it? What is the expected payoff to each player?

(e) Which of the seven equilibria of this game is focal? Explain your reasoning.

S11. Recall the game between ice-cream vendors on a beach from Exercise U6 in Chapter 6. In that game, we found two asymmetric pure-strategy equilibria. There is also a symmetric mixed-strategy equilibrium to the game.

(a) Write down the five-by-five table for the game.

(b) Eliminate dominated strategies, and explain why they should not be used in the equilibrium.

(c) Use your answer to part (b) to help you find the mixed-strategy equilibrium to the game.

S12. Suppose that the soccer penalty-kick game of Section 5.A in this chapter is expanded to include a total of six distinct strategies for the kicker: to shoot high and to the left (HL), low and to the left (LL), high and in the center (HC), low and in the center (LC), high right (HR), and low right (LR). The goalkeeper continues to have three strategies: to move to the kicker's left (L) or right (R) or to stay in the center (C). The kicker's success percentages are shown in the following table.

		GOALIE		
		L	C	R
KICKER	HL	0.50	0.85	0.85
	LL	0.40	0.95	0.95
	HC	0.85	0	0.85
	LC	0.70	0	0.70
	HR	0.85	0.85	0.50
	LR	0.95	0.95	0.40

These payoffs incorporate the following information. Shooting high runs some risk of missing the goal even if the goalie goes the wrong way (hence $0.85 < 0.95$). If the goalie guesses correctly, she has a better chance of collecting or deflecting a low shot than a high one (hence $0.40 < 0.50$). And if the shot is to the center while the goalie goes to one side, she has a better chance of using her feet to deflect a low shot than a high one (hence $0.70 < 0.85$).

In this problem, you will verify the following mixed-strategy equilibrium of this game. The goalie uses L and R each 42.2% of the time, and C 15.6% of the time. The kicker uses LL and LR each 37.8% of the time, and HC 24.4% of the time.

(a) Given the goalie's proposed mixed strategy, compute the expected payoff to the kicker for each her six pure strategies. (Use only three significant digits, in order to keep things simple.)

(b) Use your answer to part (a) to explain why the kicker's proposed mixed strategy is a best response to the goalie's proposed mixed strategy.

(c) Given the kicker's proposed mixed strategy, compute the expected payoff to the goalie for each her three pure strategies. (Again, use only three significant digits, in order to keep things simple.)

(d) Use your answer to part (a) to explain why the goalie's proposed mixed strategy is a best response to the kicker's proposed mixed strategy.

(e) Using your previous answers, explain why the proposed strategies are indeed a Nash equilibrium.

(f) Compute the equilibrium payoff to the kicker.

S13. (Optional) Recall the three-player game among Marta's sons in Exercise S9 of Chapter 4. In that game, we found three asymmetric Nash equilibria in pure strategies. In this exercise, you will find a symmetric equilibrium in mixed strategies. Note that with three players, we need three different variables (p, q, and r) to stand for the mixing probabilities. We next need to understand how to compute expected payoffs for one player when both of the other players are mixing.

(a) Suppose that Bernardo plays Yes with probability q and No with probability $1 - q$. Further suppose that Carlos independently plays Yes with probability r and No with probability $1 - r$. Then from Arturo's perspective, what is the probability that Bernardo plays Yes *and* Carlos plays No? (Hint: See Section 1.C of the Appendix to Chapter 7.)

(b) What is Arturo's expected payoff from playing Yes, in terms of Bernardo's q and Carlos's r?

(c) Write down an indifference equation for Arturo in terms of q and r.

(d) Write down the indifference equations for Bernardo and Carlos.

(e) Solve the system of three (nonlinear) equations in three unknowns to find the mixed-strategy equilibrium.

S14. **(Optional)** Recall Exercise S12 of Chapter 4, which was based on the bar scene from the film *A Beautiful Mind*. Here we consider the mixed-strategy equilibria of that game when played by $n > 2$ young men.

(a) Begin by considering the symmetric case in which each of the n young men independently goes after the solitary blonde with some probability P. This probability is determined by the condition that each young man should be indifferent between the pure strategies Blonde and Brunette, given that everyone else is mixing. What is the condition that guarantees the indifference of each player? What is the equilibrium value of P in this game?

(b) There are also some asymmetric mixed-strategy equilibria in this game. In these equilibria, $m < n$ young men each go for the blonde with probability Q, and the remaining $n - m$ young men go after the brunettes. What is the condition that guarantees that each of the m young men is indifferent, given what everyone else is doing? What condition must hold so that the remaining $n - m$ players don't want to switch from the pure strategy of choosing a brunette? What is the equilibrium value of Q in the asymmetric equilibrium?

UNSOLVED EXERCISES

U1. For the chicken game from Figure 4.14, graph the best responses of James and Dean on a p-q coordinate plane. Label all of the Nash equilibria.

U2. Revisit Exercise U10 from Chapter 7.

(a) Graph the best-response curves of Row and Column on a p-q coordinate plane.

(b) Identify and label the pure-strategy Nash equilibria.

(c) Identify and label the set of mixed-strategy Nash equilibria.

U3. (a) Find all pure-strategy Nash equilibria of the following non-zero-sum game.

		COLUMN			
		A	B	C	D
ROW	1	1, 1	2, 2	3, 4	9, 3
	2	2, 5	3, 3	1, 2	7, 1

(b) Now find a mixed-strategy equilibrium of the game. What are the players' expected payoffs in the equilibrium?

U4. Consider the following game:

		PROFESSOR PLUM		
		Revolver	Knife	Wrench
MRS. PEACOCK	Conservatory	1, 3	2, −2	0, 6
	Ballroom	3, 2	1, 4	5, 0

(a) Graph the expected payoffs from each of Professor Plum's strategies as a function of Mrs. Peacock's p-mix.

(b) Which strategies will Professor Plum use in his equilibrium mixture? Why?

(c) What is the mixed-strategy Nash equilibrium of this game?

(d) Note that this game is only slightly different from the game in Exercise S6. How are the two games different? Explain why you intuitively think the equilibrium outcome has changed from Exercise S6.

U5. Find all Nash equilibria of the Roman-Greek game in Exercise S7 when the payoff to (A, γ) changes from 5 to 4.

U6. Find all Nash equilibria of the Roman-Greek game in Exercise S7 when the payoff to (A, γ) changes from 5 to 2.

U7. Consider the following game:

		MAXWELL		
		Air	Sea	Land
JAMES	Air	0, 3	2, 0	1, 7
	Sea	2, 4	0, 6	2, 0
	Land	1, 3	2, 4	0, 3

(a) Does this game have a pure-strategy Nash equilibrium? If so, what is it?

(b) Find a mixed-strategy equilibrium to this game.

(c) Actually, this game has two mixed-strategy equilibria. Find the one you didn't find in part (b). (Hint: In one of these equilibria, one of the players plays a mixed strategy, whereas the other plays a pure strategy.)

U8. Consider a slightly different version of rock–paper–scissors in which Player 1 has an advantage. If Player 1 picks Rock and Player 2 picks Scissors, Player 1 wins 20 points from Player 2 (rather than 10). The new payoff matrix is:

		PLAYER 2		
		Rock	Scissors	Paper
PLAYER 1	Rock	0	20	−10
	Scissors	−10	0	10
	Paper	10	−10	0

(a) What is the mixed-strategy equilibrium in this version of the game?
(b) Compare your answer here with your answer for the mixed-strategy equilibrium in Exercise S9. How can you explain the differences in the equilibrium strategy choices?

U9. Section 2.C of Chapter 1 mentioned the story of two chemistry students at Duke who had opted to party hard instead of studying for their final. In the hopes of obtaining a makeup final at a later date, they lied about getting a flat tire on their return trip. Their professor agreed to the makeup, but the students were unpleasantly surprised by the second (and last) question, worth 90 points: "Which tire?" The students hadn't previously decided this part of their story, and they can't communicate during the exam.

(a) Write the game table for the tire-guessing game. (Note that each student has four pure strategies.) Giving the same answer yields a payoff of 90 for each student, while all other outcomes are worth a payoff of 0 to each.
(b) How many pure-strategy Nash equilibria are there? What are they?

As in the restaurant-choice game between Harry and Sally in Exercise S10, the number of zeroes in the table makes it relatively straightforward to find all of the game's mixed-strategy equilibria. There are eleven of them.

(c) There is one mixed-strategy Nash equilibrium where the two students play each of their four strategies with positive probability. What is this equilibrium? What is the expected value to each student from playing this equilibrium?
(d) How many mixed-strategy equilibria are there where each student plays three of the four strategies with positive probability (and never plays the fourth one)? What are they? What is the expected value to each student from playing each of these equilibria?

(e) How many mixed-strategy equilibria are there where each student plays two of the four strategies with positive probability (and never plays the other two)? What are they? What is the expected value to each student from playing each of these equilibria?
(f) If the students can't coordinate on one of the pure-strategy equilibria, can they at least improve their expected payoffs by mixing over a coordinated subset of their strategies? Explain.

U10. Barry and Neill sit down to play a relatively simple card game. They each hold one card of each suit. Each privately selects one of the suits and pushes the card of that suit toward the middle of the table. If the cards are both diamonds, then Barry wins. If the cards match suit but are not diamonds, then Neill wins. If the cards do not match suit and neither card is a diamond, then Barry wins. If the cards do not match suit and one is a diamond, then Neill wins. That is, the payoff table is as follows:

		NEILL			
		Club	Heart	Spade	Diamond
BARRY	Club	0, 1	1, 0	1, 0	0, 1
	Heart	1, 0	0, 1	1, 0	0, 1
	Spade	1, 0	1, 0	0, 1	0, 1
	Diamond	0, 1	0, 1	0, 1	1, 0

(a) What is the mixed-strategy Nash equilibrium of this game?
(b) What is the expected value of this game to each player?

U11. Recall the duel game between Renard and Chagrin in Exercise U12 in Chapter 6. Remember that the duelists start 10 steps apart and walk toward one another at the same pace, 1 step at a time, and either may fire his gun after each step. When one duelist shoots, the probability of scoring a hit depends on the distance; after k steps, it is $k/5$. Each gets a payoff of −1 if he himself is killed and 1 if the other duelist is killed. If neither or both are killed, each gets zero. Now, however, suppose that the duelists have guns with silencers. If one duelist fires and misses, the other does not know that this has happened and cannot follow the strategy of then holding his fire until the final step to get a sure shot. Each must formulate a strategy at the outset that is not conditional on the other's intermediate actions. Thus we have a simultaneous-move game, with strategies of the form "Shoot after n steps if still alive." Each player has five such strategies corresponding to the five steps that can be taken toward his adversary in the duel.

(a) The five-by-five payoff table for this game is shown below. Demonstrate how to calculate the payoffs for Row 2 of the table.

		CHAGRIN				
		1	2	3	4	5
RENARD	1	0	−0.12	−0.28	−0.44	−0.6
	2	0.12	0	0.04	−0.08	−0.2
	3	0.28	−0.04	0	0.28	0.2
	4	0.44	0.08	−0.28	0	0.6
	5	0.6	0.2	−0.2	−0.6	0

The mixed-strategy Nash equilibrium involves each player playing strategies 2, 3, and 5, with proportions 5/11, 5/11, and 1/11, respectively. Strategies 1 and 4 go unused. You will now verify that this is a Nash equilibrium.

(b) Compute the expected payoff to each of the five strategies for Renard, given Chagrin's proposed equilibrium mixture.

(c) Explain how your answer to part (b) demonstrates that the proposed mixed-strategy equilibrium really is a Nash equilibrium.

(d) What is the expected payoff for each player in equilibrium?

U12. **(Optional)** Recall the game from Exercise S10 in Chapter 4, where Larry, Moe, and Curly can choose to buy tickets toward a prize worth $30. We found six pure-strategy Nash equilibria in that game. In this problem you will find a symmetric equilibrium in mixed strategies.

(a) Eliminate the weakly dominated strategy for each player. Explain why a player would never use this weakly dominated strategy in his equilibrium mixture.

(b) Find the equilibrium in mixed strategies.

U13. **(Optional)** Revisit the three-player version of evens or odds played by Roxanne, Sara, and Ted in Exercise U11 of Chapter 4. In addition to the two pure-strategy Nash equilibria found in that problem, the game also has a mixed-strategy equilibrium.

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

(b) Once again, is this a fair game? Explain.

U14. **(Optional)** Find all Nash equilibria of the Roman-Greek game in Exercise S7 when the payoff to (A, γ) changes from 5 to 1.

U15. **(Optional)** Revist the soccer-shootout problem from Exercise S12 to see what happens to the equilibrium mixtures when the payoff in a particular cell is changed slightly from the original game table. You may wish to use software such as Gambit (freely available at http://gambit.sourceforge.net/) to compute the answers to these questions. You can also try to compute the answers by hand, but this is harder because of the number of possibilities to consider.

What is the new mixed-strategy equilibrium, given each of the following changes to the original payoff matrix? What is the expected payoff to the kicker? Compare your answers in each case with the answers for the original game.

(a) Change the payoff for (HC, C) to 0.10, making it slightly more likely for the shooter to score on a high-center shot.

(b) Change the payoff for (HL, L) to 0.70, making it more likely for the shooter to score on a high-left shot.

(c) Change the payoff for (LC, C) to 0.50, making it more likely for the shooter to score on a low-center shot.