© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

6

Mixed Strategies

In the previous chapters we restricted players to using pure strategies and we postponed discussing the option that a player may choose to randomize between several of his pure strategies. You may wonder why anyone would wish to randomize between actions. This turns out to be an important type of behavior to consider, with interesting implications and interpretations. In fact, as we will now see, there are many games for which there will be no equilibrium predictions if we do not consider the players' ability to choose stochastic strategies.

Consider the following classic zero-sum game called Matching Pennies.[1] Players 1 and 2 each put a penny on a table simultaneously. If the two pennies come up the same side (heads or tails) then player 1 gets both; otherwise player 2 does. We can represent this in the following matrix:

		Player 2	
		H	T
Player 1	H	$\underline{1, -1}$	$\overline{-1, 1}$
	T	$\overline{-1, 1}$	$\underline{1, -1}$

The matrix also includes the best-response choices of each player using the method we introduced in Section 5.1.1 to find pure-strategy Nash equilibria. As you can see, this method does not work: Given a belief that player 1 has about player 2's choice, he always wants to match it. In contrast, given a belief that player 2 has about player 1's choice, he would like to choose the opposite orientation for his penny. Does this mean that a Nash equilibrium fails to exist? We will soon see that a Nash equilibrium will indeed exist *if we allow players to choose random strategies,* and there will be an intuitive appeal to the proposed equilibrium.

Matching Pennies is not the only simple game that fails to have a pure-strategy Nash equilibrium. Recall the child's game rock-paper-scissors, in which rock beats

1. A *zero-sum game* is one in which the gains of one player are the losses of another, hence their payoffs always sum to zero. The class of zero-sum games was the main subject of analysis before Nash introduced his solution concept in the 1950s. These games have some very nice mathematical properties and were a central object of analysis in von Neumann and Morgenstern's (1944) seminal book.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

scissors, scissors beats paper, and paper beats rock. If winning gives the player a payoff of 1 and the loser a payoff of −1, and if we assume that a tie is worth 0, then we can describe this game by the following matrix:

		Player 2		
		R	P	S
Player 1	R	0, 0	−1, 1	1, −1
	P	1, −1	0, 0	−1, 1
	S	−1, 1	1, −1	0, 0

It is rather straightforward to write down the best-response correspondence for player 1 when he believes that player 2 will play one of his pure strategies as follows:

$$s_1(s_2) = \begin{cases} P & \text{when } s_2 = R \\ S & \text{when } s_2 = P \\ R & \text{when } s_2 = S, \end{cases}$$

and a similar (symmetric) list would be the best-response correspondence of player 2. Examining the two best-response correspondences immediately implies that there is no pure-strategy equilibrium, just like in the Matching Pennies game. The reason is that, starting with any pair of pure strategies, at least one player is not playing a best response and will want to change his strategy in response.

6.1 Strategies, Beliefs, and Expected Payoffs

We now introduce the possibility that players choose stochastic strategies, such as flipping a coin or rolling a die to determine what they will choose to do. This approach will turn out to offer us several important advances over that followed so far. Aside from giving the players a richer set of actions from which to choose, it will more importantly give them a richer set of possible beliefs that capture an uncertain world. If player i can believe that his opponents are choosing stochastic strategies, then this puts player i in the same kind of situation as a decision maker who faces a decision problem with probabilistic uncertainty. If you are not familiar with such settings, you are encouraged to review Chapter 2, which lays out the simple decision problem with random events.

6.1.1 Finite Strategy Sets

We start with the basic definition of random play when players have finite strategy sets S_i:

Definition 6.1 Let $S_i = \{s_{i1}, s_{i2}, \ldots, s_{im}\}$ be player i's finite set of pure strategies. Define $\triangle S_i$ as the **simplex** of S_i, which is the set of all probability distributions over S_i. A **mixed strategy** for player i is an element $\sigma_i \in \triangle S_i$, so that $\sigma_i = \{\sigma_i(s_{i1}), \sigma_i(s_{i2}), \ldots, \sigma_i(s_{im}))$ is a probability distribution over S_i, where $\sigma_i(s_i)$ is the probability that player i plays s_i.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

That is, a mixed strategy for player i is just a probability distribution over his pure strategies. Recall that any probability distribution $\sigma_i(\cdot)$ over a finite set of elements (a finite state space), in our case S_i, must satisfy two conditions:

1. $\sigma_i(s_i) \geq 0$ for all $s_i \in S_i$, and
2. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

That is, the probability of any event happening must be nonnegative, and the sum of the probabilities of all the possible events must add up to one.[2] Notice that every pure strategy is a mixed strategy with a *degenerate* distribution that picks a single pure strategy with probability one and all other pure strategies with probability zero.

As an example, consider the Matching Pennies game described earlier, with the matrix

		Player 2 H	Player 2 T
Player 1	H	1, −1	−1, 1
	T	−1, 1	1, −1

For each player i, $S_i = \{H, T\}$, and the simplex, which is the set of mixed strategies, can be written as

$$\triangle S_i = \{(\sigma_i(H), \sigma_i(T)) : \sigma_i(H) \geq 0, \sigma_i(T) \geq 0, \sigma_i(H) + \sigma_i(T) = 1\}.$$

We read this as follows: the set of mixed strategies is the set of all pairs $(\sigma_i(H), \sigma_i(T))$ such that both are nonnegative numbers, and they both sum to one.[3] We use the notation $\sigma_i(H)$ to represent the probability that player i plays H and $\sigma_i(T)$ to represent the probability that player i plays T.

Now consider the example of the rock-paper-scissors game, in which $S_i = \{R, P, S\}$ (for rock, paper, and scissors, respectively). We can define the simplex as

$$\triangle S_i = \{(\sigma_i(R), \sigma_i(P), \sigma_i(S)) : \sigma_i(R), \sigma_i(P), \sigma_i(S) \geq 0, \sigma_i(R) + \sigma_i(P) + \sigma_i(S) = 1\},$$

which is now three numbers, each defining the probability that the player plays one of his pure strategies. As mentioned earlier, a pure strategy is just a special case of a mixed strategy. For example, in this game we can represent the pure strategy of playing R with the degenerate mixed strategy: $\sigma(R) = 1$, $\sigma(P) = \sigma(S) = 0$.

From our definition it is clear that when a player uses a mixed strategy, he may choose not to use all of his pure strategies in the mix; that is, he may have some pure strategies that are not selected with positive probability. Given a player's

2. The notation $\sum_{s_i \in S_i} \sigma(s_i)$ means the sum of $\sigma(s_i)$ over all the $s_i \in S_i$. If S_i has m elements, as in the definition, we could write this as $\sum_{k=1}^{m} \sigma_i(s_{ik})$.

3. The simplex of this two-element strategy set can be represented by a single number $p \in [0, 1]$, where p is the probability that player i plays H and $1 - p$ is the probability that player i plays T. This follows from the definition of a probability distribution over a two-element set. In general the simplex of a strategy set with m pure strategies will be in an $(m - 1)$-dimensional space, where each of the $m - 1$ numbers is in $[0, 1]$, and will represent the probability of the first $m - 1$ pure strategies. All sum to a number equal to or less than one so that the remainder is the probability of the mth pure strategy.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

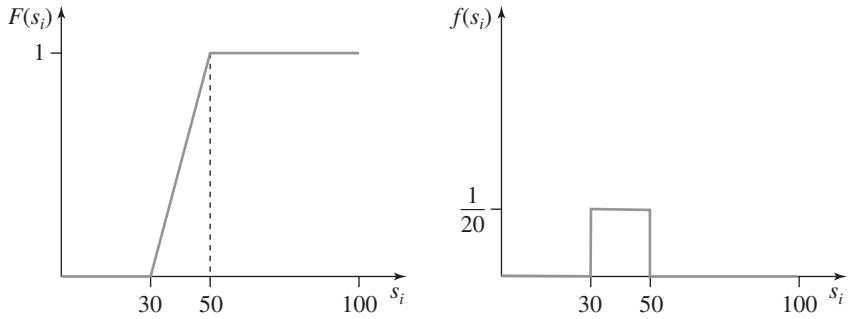

FIGURE 6.1 A continuous mixed strategy in the Cournot game.

mixed strategy $\sigma_i(\cdot)$, it will be useful to distinguish between pure strategies that are chosen with a positive probability and those that are not. We offer the following definition:

Definition 6.2 Given a mixed strategy $\sigma_i(\cdot)$ for player i, we will say that a pure strategy $s_i \in S_i$ is **in the support of** $\sigma_i(\cdot)$ if and only if it occurs with positive probability, that is, $\sigma_i(s_i) > 0$.

For example, in the game of rock-paper-scissors, a player can choose rock or paper, each with equal probability, and not choose scissors. In this case $\sigma_i(R) = \sigma_i(P) = 0.5$ and $\sigma_i(S) = 0$. We will then say that R and P are in the support of $\sigma_i(\cdot)$, but S is not.

6.1.2 Continuous Strategy Sets

As we have seen with the Cournot and Bertrand duopoly examples, or the tragedy of the commons example in Section 5.2.2, the pure-strategy sets that players have need not be finite. In the case in which the pure-strategy sets are well-defined intervals, a mixed strategy will be given by a cumulative distribution function:

Definition 6.3 Let S_i be player i's pure-strategy set and assume that S_i is an interval. A **mixed strategy** for player i is a cumulative distribution function $F_i : S_i \to [0, 1]$, where $F_i(x) = \Pr\{s_i \leq x\}$. If $F_i(\cdot)$ is differentiable with density $f_i(\cdot)$ then we say that $s_i \in S_i$ is in the support of $F_i(\cdot)$ if $f_i(s_i) > 0$.

As an example, consider the Cournot duopoly game with a capacity constraint of 100 units of production, so that $S_i = [0, 100]$ for $i \in \{1, 2\}$. Consider the mixed strategy in which player i chooses a quantity between 30 and 50 using a uniform distribution. That is,

$$F_i(s_i) = \begin{cases} 0 & \text{for } s_i < 30 \\ \frac{s_i - 30}{20} & \text{for } s_i \in [30, 50] \\ 1 & \text{for } s_i > 50 \end{cases} \quad \text{and} \quad f_i(s_i) = \begin{cases} 0 & \text{for } s_i < 30 \\ \frac{1}{20} & \text{for } s_i \in [30, 50] \\ 0 & \text{for } s_i > 50. \end{cases}$$

These two functions are depicted in Figure 6.1.

We will typically focus on games with finite strategy sets to illustrate most of the examples with mixed strategies, but some interesting examples will have infinite strategy sets and will require the use of cumulative distributions and densities to explore behavior in mixed strategies.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

6.1.3 Beliefs and Mixed Strategies

As we discussed earlier, introducing probability distributions not only enriches the set of actions from which a player can choose but also allows us to enrich the beliefs that players can have. Consider, for example, player i, who plays against opponents $-i$. It may be that player i is uncertain about the behavior of his opponents for many reasons. For example, he may believe that his opponents are indeed choosing mixed strategies, which immediately implies that their behavior is not fixed but rather random. An alternative interpretation is the situation in which player i is playing a game against an opponent that he does not know, whose background will determine how he will play. This interpretation will be revisited in Section 12.5, and it is a very appealing justification for beliefs that are random and behavior that is consistent with these beliefs.

To introduce beliefs about mixed strategies formally we define them as follows:

Definition 6.4 A **belief** for player i is given by a *probability distribution* $\pi_i \in \Delta S_{-i}$ over the strategies of his opponents. We denote by $\pi_i(s_{-i})$ the probability player i assigns to his opponents playing $s_{-i} \in S_{-i}$.

Thus a *belief* for player i is a probability distribution over the strategies of his opponents. Notice that the belief of player i lies in the same set that represents the profiles of mixed strategies of player i's opponents. For example, in the rock-paper-scissors game, we can represent the beliefs of player 1 as a triplet, $(\pi_1(R), \pi_1(P), \pi_1(S))$, where by definition $\pi_1(R), \pi_1(P), \pi_1(S) \geq 0$ and $\pi_1(R) + \pi_1(P) + \pi_1(S) = 1$. The interpretation of $\pi_1(s_2)$ is the probability that player 1 assigns to player 2 playing some particular $s_2 \in S_2$. Recall that the strategy of player 2 is a triplet $\sigma_2(R), \sigma_2(P), \sigma_2(S) \geq 0$, with $\sigma_2(R) + \sigma_2(P) + \sigma_2(S) = 1$, so we can clearly see the analogy between π and σ.

6.1.4 Expected Payoffs

Consider the Matching Pennies game described previously, and assume for the moment that player 2 chooses the mixed strategy $\sigma_2(H) = \frac{1}{3}$ and $\sigma_2(T) = \frac{2}{3}$. If player 1 plays H then he will win and get 1 with probability $\frac{1}{3}$ while he will lose and get -1 with probability $\frac{2}{3}$. If, however, he plays T then he will win and get 1 with probability $\frac{2}{3}$ while he will lose and get -1 with probability $\frac{1}{3}$. Thus by choosing different actions player 1 will face different lotteries, as described in Chapter 2.

To evaluate these lotteries we will resort to the notion of expected payoff over lotteries as presented in Section 2.2. Thus we define the expected payoff of a player as follows:

Definition 6.5 The **expected payoff** of player i when he chooses the pure strategy $s_i \in S_i$ and his opponents play the mixed strategy $\sigma_{-i} \in \Delta S_{-i}$ is

$$v_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i}) v_i(s_i, s_{-i}).$$

Similarly the expected payoff of player i when he chooses the mixed strategy $\sigma_i \in \Delta S_i$ and his opponents play the mixed strategy $\sigma_{-i} \in \Delta S_{-i}$ is

$$v_1(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) v_i(s_i, \sigma_{-i}) = \sum_{s_i \in S_i} \left(\sum_{s_{-i} \in S_{-i}} \sigma_i(s_i) \sigma_{-i}(s_{-i}) v_i(s_i, s_{-i}) \right).$$

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

The idea is a straightforward adaptation of definition 2.3 in Section 2.2.1. The randomness that player i faces if he chooses some $s_i \in S_i$ is created by the random selection of $s_{-i} \in S_{-i}$ that is described by the probability distribution $\sigma_{-i}(\cdot)$. Clearly the definition we just presented is well defined only for finite strategy sets S_i. The analog to interval strategy sets is a straightforward adaptation of the second part of definition 2.3.[4]

As an example, recall the rock-paper-scissors game:

		Player 2		
		R	P	S
Player 1	R	0, 0	−1, 1	1, −1
	P	1, −1	0, 0	−1, 1
	S	−1, 1	1, −1	0, 0

and assume that player 2 plays $\sigma_2(R) = \sigma_2(P) = \frac{1}{2}; \sigma_2(S) = 0$. We can now calculate the expected payoff for player 1 from any of his pure strategies,

$$v_1(R, \sigma_2) = \tfrac{1}{2} \times 0 + \tfrac{1}{2} \times (-1) + 0 \times 1 = -\tfrac{1}{2}$$
$$v_1(P, \sigma_2) = \tfrac{1}{2} \times 1 + \tfrac{1}{2} \times 0 + 0 \times (-1) = \tfrac{1}{2}$$
$$v_1(S, \sigma_2) = \tfrac{1}{2} \times (-1) + \tfrac{1}{2} \times 1 + 0 \times 0 = 0.$$

It is easy to see that player 1 has a unique best response to this mixed strategy of player 2. If he plays P, he wins or ties with equal probability, while his other two pure strategies are worse: with R he either loses or ties and with S he either loses or wins. Clearly if his beliefs about the strategy of his opponent are different then player 1 is likely to have a different best response.

It is useful to consider an example in which the players have strategy sets that are intervals. Consider the following game, known as an **all-pay auction,** in which two players can bid for a dollar. Each can submit a bid that is a real number (we are not restricted to penny increments), so that $S_i = [0, \infty)$, $i \in \{1, 2\}$. The person with the higher bid gets the dollar, but the twist is that both bidders have to pay their bids (hence the name of the game). If there is a tie then both pay and the dollar is awarded to each player with an equal probability of 0.5. Thus if player i bids s_i and player $j \neq i$ bids s_j then player i's payoff is

4. Consider a game in which each player has a strategy set given by the interval $S_i = [\underline{s}_i, \bar{s}_i]$. If player 1 is playing s_1 and his opponents, players $j = 2, 3, \ldots, n$, are using the mixed strategies given by the density function $f_j(\cdot)$ then the expected payoff of player 1 is given by

$$\int_{\underline{s}_2}^{\bar{s}_2} \int_{\underline{s}_3}^{\bar{s}_3} \cdots \int_{\underline{s}_n}^{\bar{s}_n} v_i(s_i, s_{-i}) f_2(s_2) f_3(s_3) \cdots f_n(s_n) ds_2 ds_3 \cdots ds_n.$$

For more on this topic see Section 19.4.4.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

$$v_i(s_i, s_{-i}) = \begin{cases} -s_i & \text{if } s_i < s_j \\ \frac{1}{2} - s_i & \text{if } s_i = s_j \\ 1 - s_i & \text{if } s_i > s_j. \end{cases}$$

Now imagine that player 2 is playing a mixed strategy in which he is *uniformly* choosing a bid between 0 and 1. That is, player 2's mixed strategy σ_2 is a uniform distribution over the interval 0 and 1, which is represented by the cumulative distribution function and density

$$F_2(s_2) = \begin{cases} s_2 & \text{for } s_2 \in [0, 1] \\ 1 & \text{for } s_2 > 1 \end{cases} \qquad \text{and} \qquad f_2(s_2) = \begin{cases} 1 & \text{for } s_2 \in [0, 1] \\ 0 & \text{for } s_2 > 1. \end{cases}$$

The expected payoff of player 1 from offering a bid $s_i > 1$ is $1 - s_i < 0$ because he will win for sure, but this would not be wise. The expected payoff from bidding $s_i < 1$ is[5]

$$\begin{aligned} v_1(s_1, \sigma_2) &= \Pr\{s_1 < s_2\}(-s_1) + \Pr\{s_1 = s_2\}\left(\tfrac{1}{2} - s_1\right) + \Pr\{s_1 > s_2\}\left(1 - s_1\right) \\ &= (1 - F_2(s_1))(-s_1) + 0\left(\tfrac{1}{2} - s_1\right) + F_2(s_1)(1 - s_1) \\ &= 0. \end{aligned}$$

Thus when player 2 is using a uniform distribution between 0 and 1 for his bid, then player 1 cannot get any positive *expected* payoff from any bid he offers: any bid less than 1 offers an expected payoff of 0, and any bid above 1 guarantees getting the dollar at an inflated price. This game is one to which we will return later, as it has several interesting features and twists.

6.2 Mixed-Strategy Nash Equilibrium

Now that we are equipped with a richer space for both strategies and beliefs, we are ready to restate the definition of a Nash equilibrium for this more general setup as follows:

Definition 6.6 The mixed-strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ is a *Nash equilibrium* if for each player σ_i^* is a best response to σ_{-i}^*. That is, for all $i \in N$,

$$v_i(\sigma_i^*, \sigma_{-i}^*) \geq v_i(\sigma_i, \sigma_{-i}^*) \;\; \forall\, \sigma_i \in \Delta S_i.$$

This definition is the natural generalization of definition 5.1. We require that each player be choosing a strategy $\sigma_i^* \in \Delta S_i$ that is (one of) the best choice(s) he can make when his opponents are choosing some profile $\sigma_{-i}^* \in \Delta S_{-i}$.

As we discussed previously, there is another interesting interpretation of the definition of a Nash equilibrium. We can think of σ_{-i}^* as the belief of player i about his opponents, π_i, which captures the idea that player i is uncertain of his opponents' behavior. The profile of mixed strategies σ_{-i}^* thus captures this uncertain belief over all of the pure strategies that player i's opponents can play. Clearly rationality requires

5. If player 2 is using a uniform distribution over [0, 1] then $\Pr\{s_1 = s_2\} = 0$ for any $s_1 \in [0, 1]$.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

that a player play a best response given his beliefs (and this now extends the notion of rationalizability to allow for uncertain beliefs). A Nash equilibrium requires that these beliefs be correct.

Recall that we defined a pure strategy $s_i \in S_i$ to be *in the support of* σ_i if $\sigma_i(s_i) > 0$, that is, if s_i is played with positive probability (see definition 6.2). Now imagine that in the Nash equilibrium profile σ^* the support of i's mixed strategy σ_i^* contains more than one pure strategy—say s_i and s_i' are both in the support of σ_i^*.

What must we conclude about a rational player i if σ_i^* is indeed part of a Nash equilibrium $(\sigma_i^*, \sigma_{-i}^*)$? By definition σ_i^* is a best response against σ_{-i}^*, which means that given σ_{-i}^* player i cannot do better than to randomize between more than one of his pure strategies, in this case, s_i and s_i'. But when would a player be willing to randomize between two alternative pure strategies? The answer is predictable:

Proposition 6.1 *If σ^* is a Nash equilibrium, and both s_i and s_i' are in the support of σ_i^*, then*

$$v_i(s_i, \sigma_{-i}^*) = v_i(s_i', \sigma_{-i}^*) = v_i(\sigma_i^*, \sigma_{-i}^*).$$

The proof is quite straightforward and follows from the observation that if a player is randomizing between two alternatives then he must be indifferent between them. If this were not the case, say $v_i(s_i, \sigma_{-i}^*) > v_i(s_i', \sigma_{-i}^*)$ with both s_i and s_i' in the support of σ_i^*, then by reducing the probability of playing s_i' from $\sigma_i^*(s_i')$ to zero, and increasing the probability of playing s_i from $\sigma_i^*(s_i)$ to $\sigma_i^*(s_i) + \sigma_i^*(s_i')$, player i's expected payoff must go up, implying that σ_i^* could not have been a best response to σ_{-i}^*.

This simple observation will play an important role in computing mixed-strategy Nash equilibria. In particular we know that if a player is playing a mixed strategy then he must be indifferent between the actions he is choosing with positive probability, that is, the actions that are in the support of his mixed strategy. One player's indifference will impose restrictions on the behavior of other players, and these restrictions will help us find the mixed-strategy Nash equilibrium.

For games with many players, or with two players who have many strategies, finding the set of mixed-strategy Nash equilibria is a tedious task. It is often done with the help of computer algorithms, because it generally takes on the form of a linear programming problem. Nevertheless it will be useful to see how one computes mixed-strategy Nash equilibria for simpler games.

6.2.1 Example: Matching Pennies

Consider the Matching Pennies game,

		Player 2	
		H	T
Player 1	H	$1, -1$	$-1, 1$
	T	$-1, 1$	$1, -1$

and recall that we showed that this game does not have a pure-strategy Nash equilibrium. We now ask, does it have a mixed-strategy Nash equilibrium? To answer this, we have to find mixed strategies for both players that are mutual best responses.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

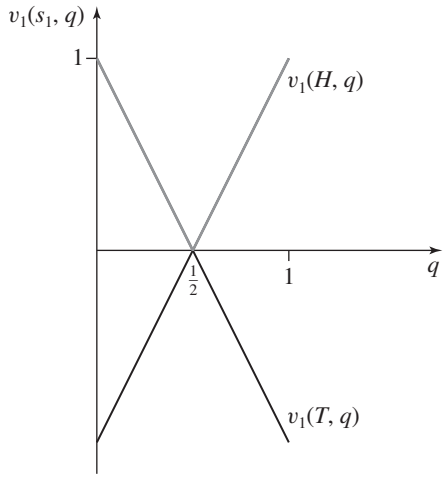

FIGURE 6.2 Expected payoffs for player 1 in the Matching Pennies game.

To simplify the notation, define mixed strategies for players 1 and 2 as follows: Let p be the probability that player 1 plays H and $1 - p$ the probability that he plays T. Similarly let q be the probability that player 2 plays H and $1 - q$ the probability that he plays T.

Using the formulas for expected payoffs in this game, we can write player 1's expected payoff from each of his two pure actions as follows:

$$v_1(H, q) = q \times 1 + (1 - q) \times (-1) = 2q - 1 \tag{6.1}$$

$$v_1(T, q) = q \times (-1) + (1 - q) \times 1 = 1 - 2q. \tag{6.2}$$

With these equations in hand, we can calculate the best response of player 1 for any choice q of player 2. In particular player 1 will prefer to play H over playing T if and only if $v_1(H, q) > v_1(T, q)$. Using (6.1) and (6.2), this will be true if and only if

$$2q - 1 > 1 - 2q,$$

which is equivalent to $q > \frac{1}{2}$. Similarly playing T will be strictly better than playing H for player 1 if and only if $q < \frac{1}{2}$. Finally, when $q = \frac{1}{2}$ player 1 will be indifferent between playing H or T.

It is useful to graph the expected payoff of player 1 from choosing either H or T as a function of player 2's choice of q, as shown in Figure 6.2.

The expected payoff of player 1 from playing H was given by the function $v_1(H, q) = 2q - 1$, as described in (6.1). This is the rising linear function in the figure. Similarly $v_1(T, q) = 1 - 2q$, described in (6.2), is the declining function. Now it is easy to see what determines the best response of player 1. The gray "upper envelope" of the graph will show the highest payoff that player 1 can achieve when player 2 plays any given level of q. When $q < \frac{1}{2}$ this is achieved by playing T; when $q > \frac{1}{2}$ this is achieved by playing H; and when $q = \frac{1}{2}$ both H and T are equally good for player 1, giving him an expected payoff of zero.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

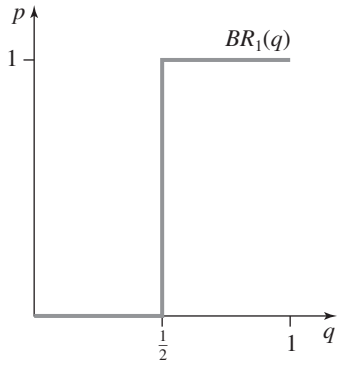

FIGURE 6.3 Player 1's best-response correspondences in the Matching Pennies game.

This simple analysis results in the best-response correspondence of player 1, which is

$$BR_1(q) = \begin{cases} p = 0 & \text{if } q < \frac{1}{2} \\ p \in [0, 1] & \text{if } q = \frac{1}{2} \\ p = 1 & \text{if } q > \frac{1}{2} \end{cases}$$

and is depicted in Figure 6.3. Notice that this is a best-response *correspondence*, and not a function, because at the value of $q = \frac{1}{2}$ any value of $p \in [0, 1]$ is a best response.

In a similar way we can calculate the payoffs of player 2 given a mixed-strategy p of player 1 to be

$$v_2(p, H) = p \times (-1) + (1 - p) \times 1 = 1 - 2p$$
$$v_2(p, T) = p \times 1 + (1 - p) \times (-1) = 2p - 1,$$

and this implies that player 2's best response is

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{2} \\ q \in [0, 1] & \text{if } p = \frac{1}{2} \\ q = 0 & \text{if } p > \frac{1}{2}. \end{cases}$$

To find a Nash equilibrium we are looking for a pair of choices (p, q) for which the two best-response correspondences cross. Were we to superimpose the best response of player 2 onto Figure 6.3 then we would see that the two best-response correspondences cross at $p = q = \frac{1}{2}$. Nevertheless it is worth walking through the logic of this solution.

We know from proposition 6.1 that when player 1 is mixing between H and T, both with positive probability, then it must be the case that his payoffs from H and from T are identical. This, it turns out, imposes a restriction on the behavior of player 2, given by the choice of q. Player 1 is willing to mix between H and T if and only if $v_1(H, q) = v_1(T, q)$, which will hold if and only if $q = \frac{1}{2}$. This is the way in which *the indifference of player 1 imposes a restriction on player 2:* only when player 2 is playing $q = \frac{1}{2}$ will player 1 be willing to mix between his actions H and T. Similarly player 2 is willing to mix between H and T only when $v_2(p, H) = v_2(p, T)$, which

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

is true only when $p = \frac{1}{2}$. We have come to the conclusion of our quest for a Nash equilibrium in this game. We can see that there is indeed a pair of mixed strategies that form a Nash equilibrium, and these are precisely when $(p, q) = \left(\frac{1}{2}, \frac{1}{2}\right)$.

There is a simple logic, which we can derive from the Matching Pennies example, that is behind the general method for finding mixed-strategy equilibria in games. The logic relies on a fact that we have already discussed: if a player is mixing several strategies then he *must be indifferent* between them. What a particular player i is willing to do depends on the strategies of his opponents. Therefore, to find out when player i is willing to mix some of his pure strategies, we must find strategies of his opponents, $-i$, that make him indifferent between some of his pure actions.

For the Matching Pennies game this can be easily illustrated as follows. First, we ask *which strategy of player 2* will make *player 1 indifferent* between playing H and T. The answer to this question (assuming it is unique) *must be player 2's strategy in equilibrium*. The reason is simple: if player 1 is to mix in equilibrium, then player 2 must be playing a strategy for which player 1's best response is mixing, and player 2's strategy must therefore make player 1 indifferent between playing H and T. Similarly we ask which strategy of player 1 will make player 2 indifferent between playing H and T, and this must be player 1's equilibrium strategy.

Remark The game of Matching Pennies is representative of situations in which one player wants to match the actions of the other, while the other wants to avoid that matching. One common example is penalty goals in soccer. The goalie wishes to jump in the direction that the kicker will kick the ball, while the kicker wishes to kick the ball in the opposite direction from the one in which the goalie chooses to jump. When they go in the same direction then the goalie wins and the kicker loses, while if they go in different directions then the opposite happens. As you can see, this is exactly the structure of the Matching Pennies game. Other common examples of such games are bosses monitoring their employees and the employees' decisions about how hard to work, or police monitoring crimes and the criminals who wish to commit them.

6.2.2 Example: Rock-Paper-Scissors

When we have games with more than two strategies for each player, then coming up with quick ways to solve mixed-strategy equilibria is a bit more involved than in 2×2 games, and it will usually involve more tedious algebra that solves several equations with several unknowns. If we consider the game of rock-paper-scissors, for example, there are many mixing combinations for each player, and we can't simply draw graphs the way we did for the Matching Pennies game.

		Player 2		
		R	P	S
Player 1	R	0, 0	−1, 1	1, −1
	P	1, −1	0, 0	−1, 1
	S	−1, 1	1, −1	0, 0

To find the Nash equilibrium of the rock-paper-scissors game we proceed in three steps. First we show that there is no Nash equilibrium in which at least one player

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

plays a pure strategy. Then we show that there is no Nash equilibrium in which at least one player mixes only between two pure strategies. These steps will imply that in any Nash equilibrium, both players must be mixing with all three pure strategies, and this will lead to the solution.

Claim 6.1 *There can be no Nash equilibrium in which one player plays a pure strategy and the other mixes.*

To see this, suppose that player i plays a pure strategy. It's easy to see from looking at the payoff matrix that player j always receives different payoffs from each of his pure strategies whenever i plays a pure strategy. Therefore player j cannot be indifferent between any of his pure strategies, so j cannot be playing a mixed strategy if i plays a pure strategy. But we know that there are no pure-strategy equilibria, and hence we conclude that there are no Nash equilibria where either player plays a pure strategy.

Claim 6.2 *There can be no Nash equilibrium in which at least one player mixes only between two pure strategies.*

To see this, suppose that i mixes between R and P. Then j always gets a strictly higher payoff from playing P than from playing R, so no strategy requiring j to play R with positive probability can be a best response for j, and j can't play R in any Nash equilibrium. But if j doesn't play R then i gets a strictly higher payoff from S than from P, so no strategy requiring i to play P with positive probability can be a best response to j not playing R. But we assumed that i was mixing between R and P, so we've reached a contradiction. We conclude that in equilibrium i cannot mix between R and P. We can apply similar reasoning to i's other pairs of pure strategies. We conclude that in any Nash equilibrium of this game, no player can play a mixed strategy in which he only plays two pure strategies with positive probability.

If by now you've guessed that the mixed strategies $\sigma_1^* = \sigma_2^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ form a Nash equilibrium then you are right. If player i plays σ_i^* then j will receive an expected payoff of 0 from every pure strategy, so j will be indifferent between all of his pure strategies. Therefore $BR_j(\sigma_i^*)$ includes all of j's mixed strategies and in particular $\sigma_j^* \in BR_j(\sigma_i^*)$. Similarly $\sigma_i^* \in BR_i(\sigma_j^*)$. We conclude that σ_1^* and σ_2^* form a Nash equilibrium. We will prove that (σ_1^*, σ_2^*) is *the unique Nash equilibrium.*

Suppose player i plays R with probability $\sigma_i(R) \in (0, 1)$, P with probability $\sigma_i(P) \in (0, 1)$, and S with probability $1 - \sigma_i(R) - \sigma_i(P)$. Because we proved that both players have to mix with all three pure strategies, it follows that $\sigma_i(R) + \sigma_i(P) < 1$ so that $1 - \sigma_i(R) - \sigma_i(P) \in (0, 1)$. It follows that player j receives the following payoffs from his three pure strategies:

$$v_j(R, \sigma_i) = -\sigma_i(P) + 1 - \sigma_i(R) - \sigma_i(P) = 1 - \sigma_i(R) - 2\sigma_i(P)$$

$$v_j(P, \sigma_i) = \sigma_i(R) - (1 - \sigma_i(R) - \sigma_i(P)) = 2\sigma_i(R) + \sigma_i(P) - 1$$

$$v_j(S, \sigma_i) = -\sigma_i(R) + \sigma_i(P).$$

In any Nash equilibrium in which j plays all three of his pure strategies with positive probability, he must receive the same expected payoff from all strategies. Therefore, in any equilibrium, we must have $v_j(R, \sigma_i) = v_j(P, \sigma_i) = v_j(S, \sigma_i)$. If we set these

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

payoffs equal to each other and solve for $\sigma_i(R)$ and $\sigma_i(P)$, we get $\sigma_i(R) = \sigma_i(P) = 1 - \sigma_i(R) - \sigma_i(P) = \frac{1}{3}$. We conclude that j is willing to include all three of his pure strategies in his mixed strategy if and only if i plays $\sigma_i^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Similarly i will be willing to play all his pure strategies with positive probability if and only if j plays $\sigma_j^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Therefore there is no other Nash equilibrium in which both players play all their pure strategies with positive probability.

6.2.3 Multiple Equilibria: Pure and Mixed

In the Matching Pennies and rock-paper-scissors games, the unique Nash equilibrium was a mixed-strategy Nash equilibrium. It turns out that mixed-strategy equilibria need not be unique when they exist. In fact when a game has multiple pure-strategy Nash equilibria, it will almost always have other Nash equilibria in mixed strategies. Consider the following game:

		Player 2	
		C	R
Player 1	M	0, 0	3, 5
	D	4, 4	0, 3

It is easy to check that (M, R) and (D, C) are both pure-strategy Nash equilibria. It turns out that in 2×2 matrix games like this one, when there are two distinct pure-strategy Nash equilibria then there will almost always be a third one in mixed strategies.[6]

For this game, let player 1's mixed strategy be given by $\sigma_1 = (\sigma_1(M), \sigma_1(D))$, with $\sigma_1(M) = p$ and $\sigma_1(D) = 1 - p$, and let player 2's mixed strategy be given by $\sigma_2 = (\sigma_2(C), \sigma_2(R))$, with $\sigma_2(C) = q$ and $\sigma_2(R) = 1 - q$. Player 1 will mix when $v_1(M, q) = v_1(D, q)$, or when

$$q \times 0 + (1 - q) \times 3 = q \times 4 + (1 - q) \times 0$$
$$\Rightarrow q = \tfrac{3}{7},$$

and player 2 will mix when $v_2(p, C) = v_2(p, R)$, or when

$$p \times 0 + (1 - p) \times 4 = p \times 5 + (1 - p) \times 3$$
$$\Rightarrow p = \tfrac{1}{6}.$$

This yields our third Nash equilibrium: $(\sigma_1^*, \sigma_2^*) = \left(\left(\frac{1}{6}, \frac{5}{6}\right), \left(\frac{3}{7}, \frac{4}{7}\right)\right)$.

6. The statement "almost always" is not defined here, but it effectively means that if we draw numbers at random from some set of distributions to fill a game matrix, and it will result in more than one pure-strategy Nash equilibrium, then with probability 1 it will also have at least one mixed-strategy equilibrium. In fact a game will typically have an odd number of equilibria. This result is known as an *index theorem* and is far beyond the scope of this text.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

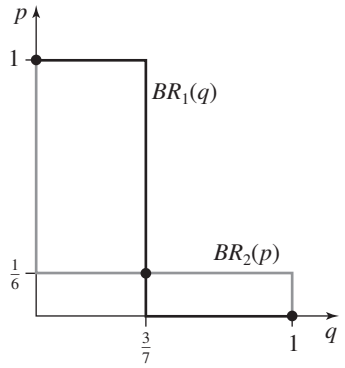

FIGURE 6.4 Best-response correspondences and Nash equilibria.

It is interesting to see that all three equilibria would show up in a careful drawing of the best-response functions. Using the payoff functions $v_1(M, q)$ and $v_1(D, q)$ we have

$$BR_1(q) = \begin{cases} p = 1 & \text{if } q < \frac{3}{7} \\ p \in [0, 1] & \text{if } q = \frac{3}{7} \\ p = 0 & \text{if } q > \frac{3}{7}. \end{cases}$$

Similarly using the payoff functions $v_2(p, C)$ and $v_2(p, R)$ we have

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{6} \\ q \in [0, 1] & \text{if } p = \frac{1}{6} \\ q = 0 & \text{if } p > \frac{1}{6}. \end{cases}$$

We can draw the two best-response correspondences as they appear in Figure 6.4. Notice that all three Nash equilibria are revealed in Figure 6.4: $(p, q) \in \{(1, 0), \left(\frac{1}{6}, \frac{3}{7}\right), (0, 1)\}$ are all Nash equilibria, where $(p, q) = (1, 0)$ corresponds to the pure strategy (M, R), and $(p, q) = (0, 1)$ corresponds to the pure strategy (D, C).

6.3 IESDS and Rationalizability Revisited

By introducing mixed strategies we offered two advancements: players can have richer beliefs, and players can choose a richer set of actions. This can be useful when we reconsider the concepts of IESDS and rationalizability, and in fact present them in their precise form using mixed strategies. In particular we can now state the following two definitions:

Definition 6.7 Let $\sigma_i \in \Delta S_i$ and $s_i' \in S_i$ be possible strategies for player i. We say that s_i' is **strictly dominated** by σ_i if

$$v_i(\sigma_i, s_{-i}) > v_i(s_i', s_{-i}) \ \ \forall s_{-i} \in S_{-i}.$$

Definition 6.8 A strategy $\sigma_i \in \Delta S_i$ is **never a best response** if there are no beliefs $\sigma_{-i} \in \Delta S_{-i}$ for player i for which $\sigma_i \in BR_i(\sigma_{-i})$.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

That is, to consider a strategy as strictly dominated, we no longer require that some other *pure strategy* dominate it, but allow for mixed strategies to dominate it as well. The same is true for strategies that are never a best response. It turns out that this approach allows both concepts to have more bite. For example, consider the following game:

		Player 2		
		L	C	R
Player 1	U	5, 1	1, 4	1, 0
	M	3, 2	0, 0	3, 5
	D	4, 3	4, 4	0, 3

and denote mixed strategies for players 1 and 2 as triplets, $(\sigma_1(U), \sigma_1(M), \sigma_1(D))$ and $(\sigma_2(L), \sigma_2(C), \sigma_2(R))$, respectively.

Starting with IESDS, it is easy to see that no pure strategy is strictly dominated by another pure strategy for any player. Hence if we restrict attention to pure strategies then IESDS has no bite and suggests that anything can happen in this game. However, if we allow for mixed strategies, we can find that the strategy L for player 2 is strictly dominated by a strategy that mixes between the pure strategies C and R. That is, $(\sigma_2(L), \sigma_2(C), \sigma_2(R)) = \left(0, \frac{1}{2}, \frac{1}{2}\right)$ strictly dominates choosing L for sure because this mixed strategy gives player 2 an expected payoff of 2 if player 1 chooses U, of 2.5 if player 1 chooses M, and of 3.5 if player 1 chooses D.

Effectively it is as if we are increasing the number of columns from which player 2 can choose to infinity, and one of these columns is the strategy in which player 2 mixes between C and R with equal probability, as the following diagram suggests:

	L	C	R
U	5, 1	1, 4	1, 0
M	3, 2	0, 0	3, 5
D	4, 3	4, 4	0, 3

Player 2's expected payoff from mixing C and R $\Rightarrow$

$\left(0, \frac{1}{2}, \frac{1}{2}\right)$
2
2.5
3.5

Hence we can perform the first step of IESDS with mixed strategies relying on the fact that $\left(0, \frac{1}{2}, \frac{1}{2}\right) \succ_2 L$, and now the game reduces to the following:

		Player 2	
		C	R
Player 1	U	1, 4	1, 0
	M	0, 0	3, 5
	D	4, 4	0, 3

$\left(0, \frac{1}{2}, \frac{1}{2}\right)$	2	1.5

In this reduced game there still are no strictly dominated pure strategies, but careful observation will reveal that the strategy U for player 1 is strictly dominated by a strategy that mixes between the pure strategies M and D. That is, $(\sigma_1(U), \sigma_1(M), \sigma_1(D))$

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

$= \left(0, \frac{1}{2}, \frac{1}{2}\right)$ strictly dominates choosing U for sure because this mixed strategy gives player 1 an expected payoff of 2 if player 2 chooses C and 1.5 if player 2 chooses R. We can then perform the second step of IESDS with mixed strategies relying on the fact that $\left(0, \frac{1}{2}, \frac{1}{2}\right) \succ_1 U$ in the reduced game, and now the game reduces further to the following:

	C	R
M	0, 0	3, 5
D	4, 4	0, 3

This last 2×2 game cannot be further reduced.

A question you must be asking is, how did we find these dominated strategies? Well, a good eye for numbers is what it takes—short of a computer program or brute force. Notice also that there are other mixed strategies that would work, because strict dominance implies that if we add a small $\varepsilon > 0$ to one of the probabilities, and subtract it from another, then the resulting expected payoff from the new mixed strategies can be made arbitrarily close to that of the original one; thus it too would dominate the dominated strategy.

Turning to rationalizability, in Section 4.3.3 we introduced the concept that after eliminating all the strategies that are *never a best response,* and employing this reasoning again and again in a way similar to what we did for IESDS, the strategies that remain are called the set of rationalizable strategies. If we use this concept to analyze the game we just solved with IESDS, the result will be the same. Starting with player 2, there is no belief that he can have for which playing L will be a best response. This is easy to see because either C or R will be a best response to one of player 1's pure strategies, and hence, even if player 1 mixes then the best response of player 2 will either be to play C, to play R, or to mix with both. Then after reducing the game a similar argument will work to eliminate U from player 1's strategy set.

As we mentioned briefly in Section 4.3.3, the concepts of IESDS and rationalizability are closely related. To see one obvious relation, the following fact is easy to prove:

Fact If a strategy σ_i is strictly dominated then it is never a best response.

The reason this is obvious is because if σ_i is strictly dominated then there is some other strategy σ_i' for which $v_i(\sigma_i', \sigma_{-i}) > v_i(\sigma_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Delta S_{-i}$. As a consequence, there is no belief about σ_{-i} that player i can have for which σ_i yields a payoff as good as or better than σ_i'.

This fact is useful, and it implies that the set of a player's rationalizable strategies is *no larger* than the set of a player's strategies that survive IESDS. This is true because if a strategy was eliminated using IESDS then it must have been eliminated through the process of rationalizability. Is the reverse true as well?

Proposition 6.2 *For any two-player game a strategy σ_i is strictly dominated if and only if it is never a best response.*

Hence for two-player games the set of strategies that survive IESDS is the same as the set of strategies that are rationalizable. Proving this is not that simple and is beyond the scope of this text. The eager and interested reader is encouraged to read Chapter 2 of Fudenberg and Tirole (1991), and the daring reader can refer to the original research

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

papers by Bernheim (1984) and Pearce (1984), which simultaneously introduced the concept of rationalizability.[7]

6.4 Nash's Existence Theorem

Section 5.1.2 argued that the Nash equilibrium solution concept is powerful because on the one hand, like IESDS and rationalizability, a Nash equilibrium will exist for most games of interest and hence will be widely applicable. On the other hand, the Nash solution concept will usually lead to more refined predictions than those of IESDS and rationalizability, yet the reverse is never true (see proposition 5.1). In his seminal Ph.D. dissertation, which laid the foundations for game theory as it is used and taught today and earned him a Nobel Prize, Nash defined the solution concept that now bears his name and showed some very general conditions under which the solution concept will exist. We first state Nash's theorem:

Theorem (Nash's Existence Theorem) *Any n-player normal-form game with finite strategy sets S_i for all players has a (Nash)* equilibrium in mixed strategies.[8]

Despite its being a bit technical, we will actually prove a restricted version of this theorem. The ideas that Nash used to prove the existence of his equilibrium concept have been widely used by game theorists, who have developed related solution concepts that refine the set of Nash equilibria, or generalize it to games that were not initially considered by Nash himself. It is illuminating to provide some basic intuition first. The central idea of Nash's proof builds on what is known in mathematics as a *fixed-point theorem*. The most basic of these theorems is known as Brouwer's fixed-point theorem:

Theorem (Brouwer's Fixed-Point Theorem) *If $f(x)$ is a continuous function from the domain $[0, 1]$ to itself then there exists at least one value $x^* \in [0, 1]$ for which $f(x^*) = x^*$.*

That is, if $f(x)$ takes values from the interval $[0, 1]$ and generates results from this same interval (or $f : [0, 1] \rightarrow [0, 1]$) then there has to be some value x^* in the interval $[0, 1]$ for which the operation of $f(\cdot)$ on x^* will give back the same value, $f(x^*) = x^*$. The intuition behind the proof of this theorem is actually quite simple.

First, because $f : [0, 1] \rightarrow [0, 1]$ maps the interval $[0, 1]$ onto itself, then $0 \leq f(x) \leq 1$ for any $x \in [0, 1]$. Second, note that if $f(0) = 0$ then $x^* = 0$, while if $f(1) = 1$ then $x^* = 1$ (as shown by the function $f_1(x)$ in Figure 6.5). We need to show, therefore, that if $f(0) > 0$ and $f(1) < 1$ then when $f(x)$ is continuous there must be some value x^* for which $f(x^*) = x^*$. To see this consider the two functions, $f_2(x)$ and $f_3(x)$, depicted in Figure 6.5, both of which map the interval $[0, 1]$ onto itself, and for which $f(0) > 0$ and $f(1) < 1$. That is, these functions start above the 45° line and end below it. The function $f_2(x)$ is continuous, and hence if it starts above

7. When there are more than two players, the set of rationalizable strategies is sometimes smaller and more refined than the set of strategies that survive IESDS. There are some conditions on the way players randomize that restore the equivalence result to many-player games, but that subject is also way beyond the scope of this text.

8. Recall that a pure strategy is a degenerate mixed strategy; hence there may be a Nash equilibrium in pure strategies.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

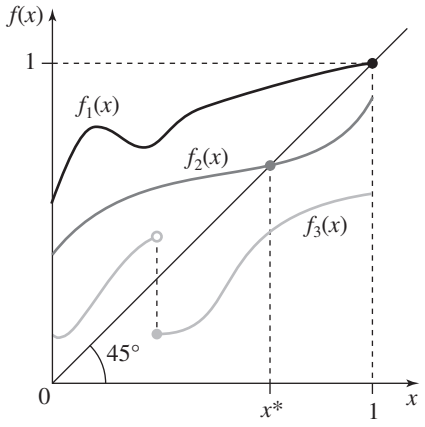

FIGURE 6.5 Brouwer's fixed-point theorem.

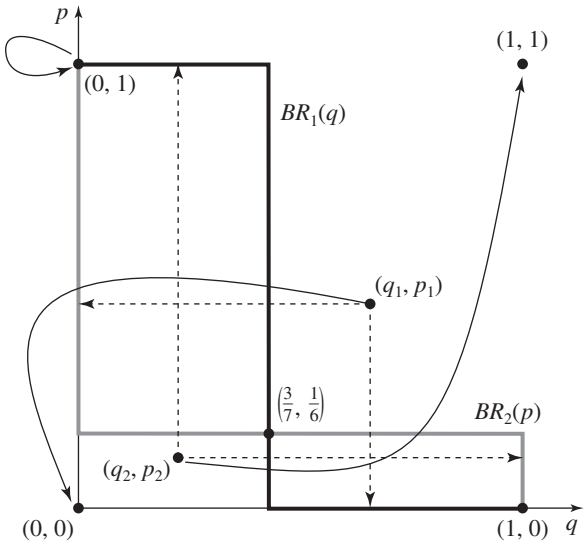

FIGURE 6.6 Mapping mixed strategies using the best-response correspondence.

the 45° line and ends below it, it must cross it at least once. In the figure, this happens at the value of x^*. To see why the continuity assumption is important, consider the function $f_3(x)$ depicted in Figure 6.5. Notice that it "jumps" down from above the 45° line to right below it, and hence this function does not cross the 45° line, in which case there is no value x for which $f(x) = x$.

You might wonder how this relates to the existence of a Nash equilibrium. What Nash showed is that something like continuity is satisfied for a mapping that uses the *best-response correspondences* of all the players at the same time to show that there must be at least one mixed-strategy profile for which each player's strategy is itself a best response to this profile of strategies. This conclusion needs some more explanation, though, because it requires a more powerful fixed-point theorem and a bit more notation and definition.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

Consider the 2×2 game used in Section 6.2.3, described in the following matrix:

		Player 2	
		C	R
Player 1	M	0, 0	3, 5
	D	4, 4	0, 3

A mixed strategy for player 1 is to choose M with probability $p \in [0, 1]$ and for player 2 to choose C with probability $q \in [0, 1]$. The analysis in Section 6.2.3 showed that the best-response correspondences for each player are

$$BR_1(q) = \begin{cases} p = 1 & \text{if } q < \frac{3}{7} \\ p \in [0, 1] & \text{if } q = \frac{3}{7} \\ p = 0 & \text{if } q > \frac{3}{7} \end{cases} \tag{6.3}$$

and

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{6} \\ q \in [0, 1] & \text{if } p = \frac{1}{6} \\ q = 0 & \text{if } p > \frac{1}{6}, \end{cases} \tag{6.4}$$

which are both depicted in Figure 6.6. We now define the **collection of best-response correspondences** as the correspondence that simultaneously represents all of the best-response correspondences of the players. This correspondence maps profiles of mixed strategies into subsets of the possible set of mixed strategies for all the players. Formally we have

Definition 6.9 The collection of best-response correspondences, $BR \equiv BR_1 \times BR_2 \times \cdots \times BR_n$, maps $\Delta S = \Delta S_1 \times \cdots \times \Delta S_n$, the set of profiles of mixed strategies, onto itself. That is, $BR : \Delta S \rightrightarrows \Delta S$ takes every element $\sigma \in \Delta S$ and converts it into a subset $BR(\sigma') \subset \Delta S$.

For a 2×2 matrix game like the one considered here, the BR correspondence can be written as[9] $BR : [0, 1]^2 \rightrightarrows [0, 1]^2$ because it takes pairs of mixed strategies of the form $(q, p) \in [0, 1]^2$ and maps them, using the best-response correspondences of the players, back to these mixed-strategy spaces, so that $BR(q, p) = (BR_2(p), BR_1(q))$. For example, consider the pair of mixed strategies (q_1, p_1) in Figure 6.6. Looking at player 1's best response, $BR_1(q) = 0$, and looking at player 2's best response, $BR_2(p) = 0$ as well. Hence $BR(q_1, p_1) = (0, 0)$, as shown by the curve that takes (q_1, p_1) and maps it onto $(0, 0)$. Similarly (q_2, p_2) is mapped onto $(1, 1)$.

Note that the point $(q, p) = (0, 1)$ is special in that $BR(0, 1) = (0, 1)$. This should be no surprise because, as we have shown in Section 6.2.3, $(q, p) = (0, 1)$ is one of the game's three Nash equilibria, so it must belong to the BR correspondence of itself. The same is true for the point $(q, p) = (1, 0)$. The third interesting point is

9. The space $[0, 1]^2$ is the two-dimensional square $[01] \times [01]$. It is the area in which all the action in Figure 6.6 is happening.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

$\left(\frac{3}{7}, \frac{1}{6}\right)$, because $BR\left(\frac{3}{7}, \frac{1}{6}\right) = ([0, 1], [0, 1])$, which means that the BR correspondence of this point is a pair of sets. This results from the fact that when player 2 mixes with probability $q = \frac{3}{7}$ then player 1 is indifferent between his two actions, causing any $p \in [0, 1]$ to be a best response, and similarly for player 2 when player 1 mixes with probability $p = \frac{1}{6}$. As a consequence, $\left(\frac{3}{7}, \frac{1}{6}\right) \in BR\left(\frac{3}{7}, \frac{1}{6}\right)$, which is the reason it is the third Nash equilibrium of the game. Indeed by now you may have anticipated the following fact, which is a direct consequence of the definition of a Nash equilibrium:

Fact A mixed-strategy profile $\sigma^* \in \Delta S$ is a Nash equilibrium if and only if it is a fixed point of the collection of best-response correspondences, $\sigma^* \in BR(\sigma^*)$.

Now the connection to fixed-point theorems should be more apparent. What Nash figured out is that when the collection of best responses BR is considered, then once it is possible to prove that it has a fixed point, it immediately implies that a Nash equilibrium exists. Nash continued on to show that for games with finite strategy sets for each player it is possible to apply the following theorem:

Theorem 6.1 (Kakutani's Fixed-Point Theorem) *A correspondence $C : X \rightrightarrows X$ has a fixed point $x \in C(x)$ if four conditions are satisfied: (1) X is a non-empty, compact, and convex subset of $\mathbb{R}^n$; (2) $C(x)$ is non-empty for all x; (3) $C(x)$ is convex for all x; and (4) C has a closed graph.*

This may surely seem like a mouthful because we have not defined any of the four qualifiers required by the theorem. For the sake of completeness, we will go over them and conclude with an intuition of why the theorem is true. First, recall that a correspondence can assign more than one value to an input, whereas a function can assign only one value to any input. Now let's introduce the definitions:

- A set $X \subseteq \mathbb{R}^n$ is **convex** if for any two points $x, y \in X$ and any $\alpha \in [0, 1]$, $\alpha x + (1 - \alpha)y \in X$. That is, any point in between x and y that lies on the straight line connecting these two points lies inside the set X.
- A set $X \subseteq \mathbb{R}^n$ is **closed** if for any converging sequence $\{x_n\}_{n=1}^{\infty}$ such that $x_n \in X$ for all n and $\lim_{n\to\infty} x_n \to x^*$ then $x^* \in X$. That is, if an infinite sequence of points are all in X and this sequence converges to a point x^* then x^* must be in X. For example, the set $(0, 1]$ that does not include 0 is not closed because we can construct a sequence of points $\left\{\frac{1}{n}\right\}_{n=1}^{\infty} = \left\{1, \frac{1}{2}, \frac{1}{3}, \ldots\right\}$ that are all in the set $[0, 1)$ and that converge to the point 0, but 0 is not in $(0, 1]$.
- A set $X \subseteq \mathbb{R}^n$ is **compact** if it is both closed and bounded. That is, there is a "largest" and a "smallest" point in the set that do not involve infinity. For example, the set $[0, 1]$ is closed and bounded; the set $(0, 1]$ is bounded but not closed; and the set $[0, \infty)$ is closed but not bounded.
- The **graph** of a correspondence $C : X \rightrightarrows X$ is the set $\{(x, y) \mid x \in X, \ y \in C(x)\}$. The correspondence $C : X \rightrightarrows X$ has a **closed graph** if the graph of C is a closed set: for any sequence $\{(x_n, y_n)\}_{n=1}^{\infty}$ such that $x_n \in X$ and $y_n \in C(x_n)$ for all n, and $\lim_{n\to\infty}(x_n, y_n) = (x^*, y^*)$, then $x^* \in X$ and $y^* \in C(x^*)$. For example, if $C(x) = x^2$ then the graph is the set $\{(x, y) \mid x \in \mathbb{R}, \ y = x^2\}$, which is exactly the plot of the function. The plot of any continuous function is therefore a closed graph. (This is true whenever $C(x)$ is a real continuous

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

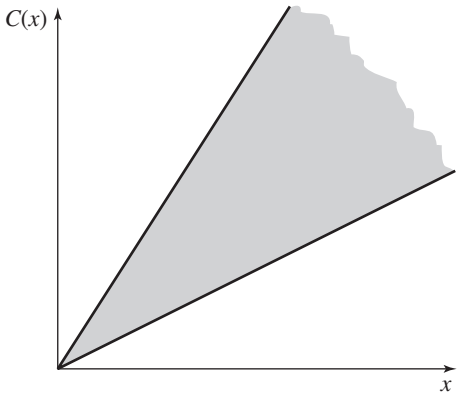

FIGURE 6.7 A correspondence with a closed graph.

function.) Another example is the correspondence $C(x) = \left[\frac{x}{2}, \frac{3x}{2}\right]$ that is depicted in Figure 6.7. In contrast the correspondence $C(x) = \left(\frac{x}{2}, \frac{3x}{2}\right)$ does not have a closed graph (it does not include the "boundaries" that are included in Figure 6.7).

The intuition for Kakutani's fixed-point theorem is somewhat similar to that for Brouwer's theorem. Brouwer's theorem was stated using two qualifiers: first, the function $f(x)$ was continuous, and second, it operated from the domain $[0, 1]$ to itself. This implied that if we draw any such function in $[0, 1]$, we will have to cross the 45° line at at least one point, which is the essence of the fixed-point theorem.

Now let's consider Kakutani's four conditions. His first condition, that X is a non-empty, compact, and convex subset of $\mathbb{R}^n$, is just the more general version of the $[0, 1]$ qualifier in Brouwer's theorem. In fact Brouwer's theorem works for $[0, 1]$ precisely because it is a non-empty, compact, and convex subset of $\mathbb{R}$.[10] His other three conditions basically guarantee that a form of "continuity" is satisfied for the correspondence $C(x)$. If we consider any continuous real function from $[0, 1]$ to itself, it satisfies all three conditions of being non-empty (it has to be well defined), convex (it is always just one point), and closed (again, just one point). Hence the four conditions identified by Kakutani guarantee that a correspondence will cross the relevant 45° line and generate at least one fixed point.

We can now show that for the 2×2 game described earlier, and in fact for any 2×2 game, the four conditions of Kakutani's theorem are satisfied:

1. $BR : [0, 1]^2 \rightrightarrows [0, 1]^2$ operates on the square $[0, 1]^2$, which is a non-empty, convex, and compact subset of $\mathbb{R}$.

10. If instead we consider $(0, 1)$, which is not closed and hence not compact, then the function $f(x) = \sqrt{x}$ does not have a fixed point because within the domain $(0, 1)$ it is everywhere above the 45° line. If we consider the domain $\left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right]$, which is not convex because it is has a gap equal to $\left(\frac{1}{3}, \frac{2}{3}\right)$, then the function $f(x) = \frac{3}{4}$ for all $x \in \left[0, \frac{1}{3}\right]$ and $f(x) = \frac{1}{4}$ for all $x \in \left[\frac{2}{3}, 1\right]$ (which is continuous) will not have a fixed point precisely because of this gap.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

2. $BR(\sigma)$ is non-empty for any $\sigma \in [0, 1]^2$. This is obvious for the example given earlier because both $BR_1(q)$ and $BR_2(p)$ are non-empty, as shown in (6.3) and (6.4). More generally for 2×2 games: each player's strategy is in the compact set $[0, 1]$; each player's expected payoff is a weighted average of the four possible payoffs he can achieve (weighted by the mixed strategies); and his expected payoff is therefore continuous in his strategy. As a consequence each player has at least one best response for any choice of his opponent. (This is because a continuous function operating on a compact set will achieve both a maximum and a minimum over that set. This is known as the **extreme value theorem.**)
3. $BR(\sigma)$ is convex for any $\sigma \in [0, 1]^2$. This is obvious for the example given earlier because both $BR_1(q)$ and $BR_2(p)$ are convex, as shown in (6.3) and (6.4). This follows more generally directly from proposition 6.1, which states that if a player is mixing between two pure strategies then he must be indifferent between them. This in turn implies that he is willing to mix between them in any way, and as a consequence, if two mixed strategies σ_1' and σ_1'' are in $BR_1(q)$ then any mixed strategy is in BR_1, so that BR_1 is convex. Because this argument works for both players, if any two mixed-strategy profiles σ and σ' are in BR then any convex combination of them is also in BR.
4. $BR(\sigma)$ has a closed graph. Again we can see this from (6.3) and (6.4). For each of the two players, $BR_i(\sigma_j)$ is equal to either 0, 1, or the whole interval $[0, 1]$. More generally, for any 2×2 game consider a sequence of mixed-strategy profiles $\{(q_n, p_n)\}_{n=1}^{\infty}$ and a sequence of best responses $\{(q_n', p_n')\}_{n=1}^{\infty}$, where $(q_n', p_n') \in BR((q_n, p_n))$ for all n. Let $\lim_{n\to\infty}(q_n, p_n) = (q^*, p^*)$ and $\lim_{n\to\infty}(q_n', p_n') = (q', p')$. To conclude that $BR(\sigma)$ has a closed graph we need to show that $(q', p') \in BR(q^*, p^*)$. For player 2 it must be that $v_2(q_n', p_n) \geq v_2(q, p_n)$ for any $q \in [0, 1]$ because q_n' is a best response to p_n. Because the (expected) payoff function is linear in q and p, it is continuous in both arguments and, as a consequence, we can take limits on both sides of the inequality while preserving the inequality, so that $\lim_{n\to\infty} v_2(q_n', p_n) \geq \lim_{n\to\infty} v_2(q, p_n)$ for all $q \in [0, 1]$, implying that $v_2(q', p^*) \geq v_2(q, p^*)$ for all $q \in [0, 1]$. But this implies that $q' \in BR_2(p^*)$, and a symmetric argument for player 1 implies that $p' \in BR_1(q^*)$, which together prove that $(q', p') \in BR(q^*, p^*)$.

We conclude therefore that all four conditions of Kakutani's fixed-point theorem are satisfied for the (mixed-) strategy sets and the best-response correspondence of any 2×2 game. Hence the best-response correspondence $BR(\sigma)$ has a fixed point, which in turn implies that any 2×2 game has at least one Nash equilibrium.

Recall that Nash's theorem referred to any finite n-player game and not just 2×2 games. As Nash showed, the basic application of Kakutani's fixed-point theorem to finite games holds for any finite number of pure strategies for each player. If, say, player i has a strategy set consisting of m pure strategies $\{s_{i1}, s_{i2}, \ldots, s_{im}\}$ then his set of mixed strategies is in the simplex $\Sigma_i = \{(\sigma_i(s_{i1}), \sigma_i(s_{i2}), \ldots, \sigma_i(s_{im})) \mid \sigma_i(s_{ik}) \in [0, 1]$ for all $k = 1, 2, \ldots, m$, and $\sum_{k=1}^{m} \sigma_i(s_{ik}) = 1\}$. It is easy to show that the set Σ_i is a non-empty, compact, and convex subset of $\mathbb{R}^m$, meaning that the first condition of Kakutani's theorem is satisfied. Using the same ideas as in points 1–4, it is not too

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

difficult to show that the three other conditions of Kakutani's theorem hold, and as a result that the best-response correspondence BR has a fixed point and that any such game has a Nash equilibrium.

6.5 Summary

- Allowing for mixed strategies enriches both what players can choose and what they can believe about the choices of other players.
- In games for which players have opposing interests, like the Matching Pennies game, there will be no pure-strategy equilibrium but a mixed-strategy equilibrium will exist.
- Allowing for mixed strategies enhances the power of IESDS and of rationalizability.
- Nash proved that for finite games there will always be at least one Nash equilibrium.

6.6 Exercises

6.1 **Best Responses in the Battle of the Sexes:** Use the best-response correspondences in the Battle of the Sexes game to find all the Nash equilibria. (Follow the approach used for the example in Section 6.2.3.)

6.2 **Mixed Dominance 1:** Let σ_i be a mixed strategy of player i that puts positive weight on one strictly dominated pure strategy. Show that there exists a mixed strategy σ_i' that puts no weight on any dominated pure strategy and that dominates σ_i.

6.3 **Mixed Dominance 2:** Consider the game used in Section 6.3:

		Player 2		
		L	C	R
Player 1	U	5, 1	1, 4	1, 0
	M	3, 2	0, 0	3, 5
	D	4, 3	4, 4	0, 3

a. Find a strategy different from $(\sigma_2(L), \sigma_2(C), \sigma_2(R)) = \left(0, \frac{1}{2}, \frac{1}{2}\right)$ that strictly dominates the pure strategy L for player 2. Argue that you can find an infinite number of such strategies.

b. Find a strategy different from $(\sigma_1(U), \sigma_1(M), \sigma_1(D)) = \left(0, \frac{1}{2}, \frac{1}{2}\right)$ that strictly dominates the pure strategy U for player 1 in the game remaining after one stage of elimination. Argue that you can find an infinite number of such strategies.

6.4 **Monitoring:** An employee (player 1) who works for a boss (player 2) can either work (W) or shirk (S), while his boss can either monitor the employee (M) or ignore him (I). As in many employee-boss relationships, if the employee is working then the boss prefers not to monitor, but if the boss is not

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

monitoring then the employee prefers to shirk. The game is represented by the following matrix:

		Player 2	
		M	*I*
Player 1	*W*	1, 1	1, 2
	S	0, 2	2, 1

a. Draw the best-response function of each player.
b. Find the Nash equilibrium of this game. What kind of game does this game remind you of?

6.5 **Cops and Robbers:** Player 1 is a police officer who must decide whether to patrol the streets or to hang out at the coffee shop. His payoff from hanging out at the coffee shop is 10, while his payoff from patrolling the streets depends on whether he catches a robber, who is player 2. If the robber prowls the streets then the police officer will catch him and obtain a payoff of 20. If the robber stays in his hideaway then the officer's payoff is 0. The robber must choose between staying hidden or prowling the streets. If he stays hidden then his payoff is 0, while if he prowls the streets his payoff is -10 if the officer is patrolling the streets and 10 if the officer is at the coffee shop.

a. Write down the matrix form of this game.
b. Draw the best-response function of each player.
c. Find the Nash equilibrium of this game. What kind of game does this game remind you of?

6.6 **Declining Industry:** Consider two competing firms in a declining industry that cannot support both firms profitably. Each firm has three possible choices, as it must decide whether or not to exit the industry immediately, at the end of this quarter, or at the end of the next quarter. If a firm chooses to exit then its payoff is 0 from that point onward. Each quarter that both firms operate yields each a loss equal to -1, and each quarter that a firm operates alone yields it a payoff of 2. For example, if firm 1 plans to exit at the end of this quarter while firm 2 plans to exit at the end of the next quarter then the payoffs are $(-1, 1)$ because both firms lose -1 in the first quarter and firm 2 gains 2 in the second. The payoff for each firm is the sum of its quarterly payoffs.

a. Write down this game in matrix form.
b. Are there any strictly dominated strategies? Are there any weakly dominated strategies?
c. Find the pure-strategy Nash equilibria.
d. Find the unique mixed-strategy Nash equilibrium. (Hint: you can use your answer to (b) to make things easier.)

6.7 **Grad School Competition:** Two students sign up to prepare an honors thesis with a professor. Each can invest time in his own project: either no time, one week, or two weeks (these are the only three options). The cost of time is 0 for no time, and each week costs 1 unit of payoff. The more time a student puts in the better his work will be, so that if one student puts in more time than the other there will be a clear "leader." If they put in the same amount of time then their thesis projects will have the same quality. The professor, however, will give out only one grade of A. If there is a clear leader then he will get

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

the A, while if they are equally good then the professor will toss a fair coin to decide who gets the A. The other student will get a B. Since both wish to continue on to graduate school, a grade of A is worth 3 while a grade of B is worth 0.

a. Write down this game in matrix form.
b. Are there any strictly dominated strategies? Are there any weakly dominated strategies?
c. Find the unique mixed-strategy Nash equilibrium.

6.8 **Market Entry:** Three firms are considering entering a new market. The payoff for each firm that enters is $\frac{150}{n}$, where n is the number of firms that enter. The cost of entering is 62.

a. Find all the pure-strategy Nash equilibria.
b. Find the symmetric mixed-strategy equilibrium in which all three players enter with the same probability.

6.9 **Discrete All-Pay Auction:** In Section 6.1.4 we introduced a version of an all-pay auction that worked as follows: Each bidder submits a bid. The highest bidder gets the good, but *all bidders pay their bids.* Consider an auction in which player 1 values the item at 3 while player 2 values the item at 5. Each player can bid either 0, 1, or 2. If player i bids more than player j then i wins the good and both pay. If both players bid the same amount then a coin is tossed to determine who gets the good, but again both pay.

a. Write down the game in matrix form. Which strategies survive IESDS?
b. Find the Nash equilibria for this game.

6.10 **Continuous All-Pay Auction:** Consider an all-pay auction for a good worth 1 to each of the two bidders. Each bidder can choose to offer a bid from the unit interval so that $S_i = [0, 1]$. Players care only about the expected value they will end up with at the end of the game (i.e., if a player bids 0.4 and expects to win with probability 0.7 then his payoff is $0.7 \times 1 - 0.4$).

a. Model this auction as a normal-form game.
b. Show that this game has no pure-strategy Nash equilibrium.
c. Show that this game cannot have a Nash equilibrium in which each player is randomizing over a finite number of bids.
d. Consider mixed strategies of the following form: Each player i chooses an interval $[\underline{x}_i, \overline{x}_i]$ with $0 \leq \underline{x}_i < \overline{x}_i \leq 1$ together with a cumulative distribution $F_i(x)$ over the interval $[\underline{x}_i, \overline{x}_i]$. (Alternatively you can think of each player choosing $F_i(x)$ over the interval $[0, 1]$ with two values $\underline{x}_i$ and $\overline{x}_i$ such that $F_i(\underline{x}_i) = 0$ and $F_i(\overline{x}_i) = 1$.)
    i. Show that if two such strategies are a mixed-strategy Nash equilibrium then it must be that $\underline{x}_1 = \underline{x}_2$ and $\overline{x}_1 = \overline{x}_2$.
    ii. Show that if two such strategies are a mixed-strategy Nash equilibrium then it must be that $\underline{x}_1 = \underline{x}_2 = 0$.
    iii. Using your answers to (i) and (ii), argue that if two such strategies are a mixed-strategy Nash equilibrium then both players must be getting an expected payoff of zero.
    iv. Show that if two such strategies are a mixed-strategy Nash equilibrium then it must be that $\overline{x}_1 = \overline{x}_2 = 1$.
    v. Show that $F_i(x)$ being uniform over $[0, 1]$ is a symmetric Nash equilibrium of this game.

For general queries, contact webmaster@press.princeton.edu

© Copyright, Princeton University Press. No part of this book may be distributed, posted, or reproduced in any form by digital or mechanical means without prior written permission of the publisher.

6.11 **Bribes:** Two players find themselves in a legal battle over a patent. The patent is worth 20 to each player, so the winner would receive 20 and the loser 0. Given the norms of the country, it is common to bribe the judge hearing a case. Each player can offer a bribe secretly, and the one whose bribe is the highest will be awarded the patent. If both choose not to bribe, or if the bribes are the same amount, then each has an equal chance of being awarded the patent. If a player does bribe, then the bribe can be valued at either 9 or 20. Any other number is considered very unlucky, and the judge would surely rule against a party who offered a different number.

a. Find the unique pure-strategy Nash equilibrium for this game.
b. If the norm were different, so that a bribe of 15 were also acceptable, is there a pure-strategy Nash equilibrium?
c. Find the symmetric mixed-strategy Nash equilibrium for the game with possible bribes of 9, 15, and 20.

6.12 **The Tax Man:** A citizen (player 1) must choose whether to file taxes honestly or to cheat. The tax man (player 2) decides how much effort to invest in auditing and can choose $a \in [0, 1]$; the cost to the tax man of investing at a level a is $c(a) = 100a^2$. If the citizen is honest then he receives the benchmark payoff of 0, and the tax man pays the auditing costs without any benefit from the audit, yielding him a payoff of $-100a^2$. If the citizen cheats then his payoff depends on whether he is caught. If he is caught then his payoff is -100 and the tax man's payoff is $100 - 100a^2$. If he is not caught then his payoff is 50 while the tax man's payoff is $-100a^2$. If the citizen cheats and the tax man chooses to audit at level a then the citizen is caught with probability a and is not caught with probability $(1 - a)$.

a. If the tax man believes that the citizen is cheating for sure, what is his best-response level of a?
b. If the tax man believes that the citizen is honest for sure, what is his best-response level of a?
c. If the tax man believes that the citizen is honest with probability p, what is his best-response level of a as a function of p?
d. Is there a pure-strategy Nash equilibrium for this game? Why or why not?
e. Is there a mixed-strategy Nash equilibrium for this game? Why or why not?

For general queries, contact webmaster@press.princeton.edu