4 Mixed Strategy Equilibrium

Prerequisite: Chapter 2

4.1 Introduction

4.1.1 *Stochastic steady states*

A NASH EQUILIBRIUM of a strategic game is an action profile in which every player's action is optimal given every other player's action (Definition 23.1). Such an action profile corresponds to a steady state of the idealized situation in which for each player in the game there is a population of individuals, and whenever the game is played, one player is drawn randomly from each population (see Section 2.6). In a steady state, every player's behavior is the same whenever she plays the game, and no player wishes to change her behavior, knowing (from her experience) the other players' behavior. In a steady state in which each player's "behavior" is simply an action and within each population all players choose the same action, the outcome of every play of the game is the same Nash equilibrium.

More general notions of a steady state allow the players' choices to vary, as long as the pattern of choices remains constant. For example, different members of a given population may choose different actions, each player choosing the same action whenever she plays the game. Or each individual may, on each occasion she plays the game, choose her action probabilistically according to the same, unchanging distribution. These two more general notions of a steady state are equivalent: a steady state of the first type in which the fraction p of the population representing player i chooses the action a corresponds to a steady state of the second type in which each member of the population representing player i chooses a with probability p. In both cases, in each play of the game the probability that the individual in the role of player i chooses a is p. Both these notions of steady state are modeled by a mixed strategy Nash equilibrium, a generalization of the notion of Nash equilibrium. For expository convenience, in most of this chapter I interpret such an equilibrium as a model of the second type of steady state, in which each player chooses her actions probabilistically; such a steady state is called *stochastic* ("involving probability").

4.1.2 Example: Matching Pennies

An analysis of the game *Matching Pennies* (Example 19.1) illustrates the idea of a stochastic steady state. My discussion focuses on the outcomes of this game, given in Figure 100.1, rather than payoffs that represent the players' preferences, as before.

	Head	*Tail*
Head	\$1, −\$1	−\$1, \$1
Tail	−\$1, \$1	\$1, −\$1

Figure 100.1 The outcomes of *Matching Pennies*.

As we saw previously, this game has no Nash equilibrium: no pair of actions is compatible with a steady state in which each player's action is the same whenever the game is played. I claim, however, that the game has a *stochastic* steady state in which each player chooses each of her actions with probability $\frac{1}{2}$. To establish this result, I need to argue that if player 2 chooses each of her actions with probability $\frac{1}{2}$, then player 1 optimally chooses each of her actions with probability $\frac{1}{2}$, and vice versa.

Suppose that player 2 chooses each of her actions with probability $\frac{1}{2}$. If player 1 chooses *Head* with probability p and *Tail* with probability $1 - p$, then each outcome $(Head, Head)$ and $(Head, Tail)$ occurs with probability $\frac{1}{2}p$, and each outcome $(Tail, Head)$ and $(Tail, Tail)$ occurs with probability $\frac{1}{2}(1 - p)$. Thus the probability that the outcome is either $(Head, Head)$ or $(Tail, Tail)$, in which case player 1 gains \$1, is $\frac{1}{2}p + \frac{1}{2}(1 - p)$, which is equal to $\frac{1}{2}$. In the other two outcomes, $(Head, Tail)$ and $(Tail, Head)$, she loses \$1, so the probability of her losing \$1 is also $\frac{1}{2}$. In particular, the probability distribution over outcomes is independent of p! Thus *every* value of p is optimal. In particular, player 1 can do no better than choose *Head* with probability $\frac{1}{2}$ and *Tail* with probability $\frac{1}{2}$. A similar analysis shows that player 2 optimally chooses each action with probability $\frac{1}{2}$ when player 1 does so. We conclude that the game has a stochastic steady state in which each player chooses each action with probability $\frac{1}{2}$.

I further claim that, under a reasonable assumption on the players' preferences, the game has no other steady state. This assumption is that each player wants the probability of her gaining \$1 to be as large as possible. More precisely, if $p > q$, then each player prefers to gain \$1 with probability p and lose \$1 with probability $1 - p$ than to gain \$1 with probability q and lose \$1 with probability $1 - q$.

To show that under this assumption there is no steady state in which the probability of each player's choosing *Head* is different from $\frac{1}{2}$, denote the probability with which player 2 chooses *Head* by q (so that she chooses *Tail* with probability $1 - q$). If player 1 chooses *Head* with probability p, then she gains \$1 with probability $pq + (1 - p)(1 - q)$ (the probability that the outcome is either $(Head, Head)$ or $(Tail, Tail)$) and loses \$1 with probability $(1 - p)q + p(1 - q)$. The first probability is equal to $1 - q + p(2q - 1)$ and the second is equal to $q + p(1 - 2q)$. Thus if $q < \frac{1}{2}$ (player 2 chooses *Head* with probability less than $\frac{1}{2}$), the first probability is decreasing in p and the second is increasing in p, so that the lower is p, the better is the outcome for player 1; the value of p that induces the best probability distribution over outcomes for player 1 is 0. That is, if player 2 chooses *Head* with probability less than $\frac{1}{2}$, then the uniquely best policy for player 1 is to choose *Tail* with certainty. A similar argument shows that if player 2 chooses *Head* with probability greater than $\frac{1}{2}$, the uniquely best policy for player 1 is to choose *Head* with certainty.

Now, if player 1 chooses an action with certainty, a similar analysis leads to the conclusion that the optimal policy of player 2 is to choose an action with certainty (*Head* if player 1 chooses *Tail* and *Tail* if player 1 chooses *Head*).

We conclude that there is no steady state in which the probability that player 2 chooses *Head* differs from $\frac{1}{2}$. A symmetric argument shows that there is no steady state in which the probability that player 1 chooses *Head* differs from $\frac{1}{2}$. Thus in the only stochastic steady state each player chooses each action with probability $\frac{1}{2}$.

As discussed in the opening section (4.1.1), the stable pattern of behavior we have found can be alternatively interpreted as a steady state in which no player randomizes. Instead, half the players in the population of individuals who take the role of player 1 in the game choose *Head* whenever they play the game and half of them choose *Tail* whenever they play the game; similarly half of those who take the role of player 2 choose *Head* and half choose *Tail*. Given that the individuals involved in any given play of the game are chosen randomly from the populations, in each play of the game each individual faces with probability $\frac{1}{2}$ an opponent who chooses *Head*, and with probability $\frac{1}{2}$ an opponent who chooses *Tail*.

? EXERCISE 101.1 (Variant of *Matching Pennies*) Find the steady state(s) of the game that differs from *Matching Pennies* only in that the outcomes of $(Head, Head)$ and of $(Tail, Tail)$ are that player 1 gains \$2 and player 2 loses \$1.

4.1.3 Generalizing the analysis: expected payoffs

The fact that *Matching Pennies* has only two outcomes for each player (gain \$1, lose \$1) makes the analysis of a stochastic steady state particularly simple because it allows us to deduce, under a weak assumption, the players' preferences regarding *lotteries* (probability distributions) over outcomes from their preferences regarding deterministic outcomes (outcomes that occur with certainty). If a player prefers the deterministic outcome a to the deterministic outcome b, it is very plausible that if $p > q$, then she prefers the lottery in which a occurs with probability p (and b occurs with probability $1 - p$) to the lottery in which a occurs with probability q (and b occurs with probability $1 - q$).

In a game with more than two outcomes for some player, we cannot extrapolate in this way from preferences regarding deterministic outcomes to preferences regarding lotteries over outcomes. Suppose, for example, that a game has three possible outcomes, a, b, and c, and that a player prefers a to b to c. Does she prefer the deterministic outcome b to the lottery in which a and c each occur with probability $\frac{1}{2}$, or vice versa? The information about her preferences over deterministic outcomes gives us no clue about the answer to this question. She may prefer b to the lottery in which a and c each occur with probability $\frac{1}{2}$, or she may prefer this lottery to b; both preferences are consistent with her preferring a to b to c. To study her behavior when she is faced with choices between lotteries, we need to add to the model a description of her preferences regarding lotteries over outcomes.

A standard assumption in game theory restricts attention to preferences regarding lotteries over outcomes that may be represented by the expected value of a payoff function over deterministic outcomes. (See Section 17.6.3 if you are unfamiliar with the notion of "expected value".) That is, for every player i there is a payoff function u_i with the property that player i prefers one lottery over outcomes to another if and only if, according to u_i, the expected value of the first lottery exceeds the expected value of the second lottery.

For example, suppose that there are three outcomes, a, b, and c, and lottery P yields a with probability p_a, b with probability p_b, and c with probability p_c, whereas lottery Q yields these three outcomes with probabilities q_a, q_b, and q_c. Then the assumption is that for each player i there are numbers $u_i(a)$, $u_i(b)$, and $u_i(c)$ such that player i prefers lottery P to lottery Q if and only if $p_a u_i(a) + p_b u_i(b) + p_c u_i(c) > q_a u_i(a) + q_b u_i(b) + q_c u_i(c)$. (I discuss the representation of preferences by the expected value of a payoff function in more detail in Section 4.12, an appendix to this chapter.)

The first systematic investigation of preferences regarding lotteries represented by the expected value of a payoff function over deterministic outcomes was undertaken by von Neumann and Morgenstern (1944). Accordingly such preferences are called **vNM preferences**. A payoff function over deterministic outcomes (u_i in the previous paragraph) whose expected value represents such preferences is called a **Bernoulli payoff function** (in honor of Daniel Bernoulli (1700–1782), who appears to have been one of the first persons to use such a function to represent preferences).

The restrictions on preferences regarding deterministic outcomes required for them to be represented by a payoff function are relatively innocuous (see Section 1.2.2). The same is not true of the restrictions on preferences regarding lotteries over outcomes required for them to be represented by the expected value of a payoff function. (I do not discuss these restrictions, but the box at the end of this section gives an example of preferences that violate them.) Nevertheless, we obtain many insights from models that assume that preferences take this form; following standard game theory (and standard economic theory), I maintain the assumption throughout the book.

The assumption that a player's preferences are represented by the expected value of a payoff function does not restrict her attitudes to risk: a person whose preferences are represented by such a function may have an arbitrarily strong liking or dislike for risk.

Suppose, for example, that a, b, and c are three outcomes, and a person prefers a to b to c. If the person is very averse to risky outcomes, then she prefers to obtain b for sure rather than to face the lottery in which a occurs with probability p and c occurs with probability $1 - p$, even if p is relatively large. Such preferences may be represented by the expected value of a payoff function u for which $u(a)$ is close to $u(b)$, which is much larger than $u(c)$. For a case in which a, b, and c are numbers, such payoffs are illustrated in the left panel of Figure 103.1.

If the person is not at all averse to risky outcomes, then she prefers the lottery to the certain outcome b, even if p is relatively small. Such preferences are represented by the expected value of a payoff function u for which $u(a)$ is much larger than $u(b)$, which is close to $u(c)$. For a case in which a, b, and c are numbers, such payoffs are illustrated in the right panel of Figure 103.1. If $u(a) = 10$, $u(b) = 9$, and $u(c) = 0$, for example, then the person prefers the certain outcome b to any lottery between a and c that yields a with probability less than $\frac{9}{10}$. But if $u(a) = 10$, $u(b) = 1$, and $u(c) = 0$, she prefers any lottery between a and c that yields a with probability greater than $\frac{1}{10}$ to the certain outcome b.

Suppose that the outcomes are amounts of money, and a person's preferences are represented by the expected value of a payoff function in which the payoff of each outcome is equal to the amount of money involved. Then we say the person is *risk neutral*. Such a person compares lotteries according to the expected amount of

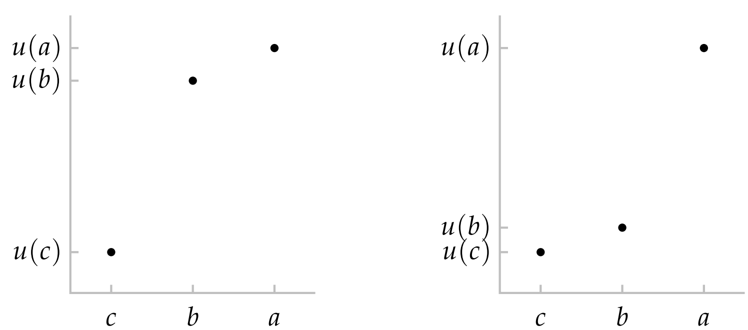

Figure 103.1 The Bernoulli payoffs for the outcomes a, b, and c for persons who prefer a to b to c and are very averse to risk (left panel) and not at all averse to risk (right panel).

money involved. (For example, she is indifferent between receiving \$100 for sure and the lottery that yields \$0 with probability $\frac{9}{10}$ and \$1000 with probability $\frac{1}{10}$.) On the one hand, the fact that people buy insurance suggests that in some circumstances preferences are *risk averse*: people prefer to obtain \$$z$ with certainty than to receive the outcome of a lottery that yields \$$z$ on average. On the other hand, the fact that people buy lottery tickets that pay, on average, much less than their purchase price, suggests that in other circumstances preferences are *risk preferring*. In both cases, preferences over lotteries are not represented by expected *monetary* values, though they still may be represented by the expected value of a *payoff* function (in which the payoffs to outcome are different from the monetary values of the outcomes).

Any given preferences over deterministic outcomes are represented by many different payoff functions (see Section 1.2.2). The same is true of preferences over lotteries; the relation between payoff functions whose expected values represent the same preferences is discussed in Section 4.12.2 in the appendix to this chapter. In particular, we may choose arbitrary payoffs for the outcomes that are best and worst according to the preferences, as long as the payoff to the best outcome exceeds the payoff to the worst outcome. Suppose, for example, that there are three outcomes, a, b, and c, and a person prefers a to b to c, and is indifferent between b and the lottery that yields a with probability $\frac{1}{2}$ and c with probability $\frac{1}{2}$. Then we may choose $u(a) = 3$ and $u(c) = 1$, in which case $u(b) = 2$; or, for example, we may choose $u(a) = 10$ and $u(c) = 0$, in which case $u(b) = 5$, or $u(a) = 1$ and $u(c) = -1$, in which case $u(b) = 0$.

SOME EVIDENCE ON EXPECTED PAYOFF FUNCTIONS

Consider the following two lotteries (the first of which is, in fact, deterministic):

Lottery 1 You receive \$2 million with certainty.

Lottery 2 You receive \$10 million with probability 0.1, \$2 million with probability 0.89, and nothing with probability 0.01.

Which do you prefer? Now consider two more lotteries:

Lottery 3 You receive \$2 million with probability 0.11 and nothing with probability 0.89.

Lottery 4 You receive \$10 million with probability 0.1 and nothing with probability 0.9.

Which do you prefer? A significant fraction of experimental subjects say they prefer lottery 1 to lottery 2, and lottery 4 to lottery 3. (See, for example, Conlisk 1989 and Camerer 1995, 622–623.)

These preferences cannot be represented by an expected payoff function! If they could be, there would exist a payoff function u for which the expected payoff of lottery 1 exceeds that of lottery 2:

$$u(2) > 0.1u(10) + 0.89u(2) + 0.01u(0),$$

where the amounts of money are expressed in millions. Subtracting $0.89u(2)$ and adding $0.89u(0)$ to each side we obtain

$$0.11u(2) + 0.89u(0) > 0.1u(10) + 0.9u(0).$$

But this inequality says that the expected payoff of lottery 3 exceeds that of lottery 4! Thus preferences represented by an expected payoff function that yield a preference for lottery 1 over lottery 2 must also yield a preference for lottery 3 over lottery 4.

Preferences represented by the expected value of a payoff function *are*, however, consistent with a person's being indifferent between lotteries 1 and 2, and between lotteries 3 and 4. Suppose we assume that when a person is almost indifferent between two lotteries, she may make a "mistake". Then a person's expressed preference for lottery 1 over lottery 2 and for lottery 4 over lottery 3 is not directly inconsistent with her preferences' being represented by the expected value of a payoff function in which she is almost indifferent between lotteries 1 and 2 and between lotteries 3 and 4. If, however, we add the assumption that mistakes are distributed symmetrically, then the frequency with which people express a preference for lottery 2 over lottery 1 and for lottery 4 over lottery 3 (also inconsistent with preferences represented by the expected value of a payoff function) should be similar to that with which people express a preference for lottery 1 over lottery 2 and for lottery 3 over lottery 4. In fact, however, the second pattern is significantly more common than the first (Conlisk 1989), so that a more significant modification of the theory is needed to explain the observations.

A limitation of the evidence is that it is based on the preferences expressed by people faced with *hypothetical* choices; understandably (given the amounts of money involved), no experiment has been run in which subjects were paid according to the lotteries they chose! Experiments with stakes consistent with normal research budgets show few choices inconsistent with preferences represented by the expected value of a payoff function (Conlisk 1989). This evidence, however, does not contradict the evidence based on hypothetical choices with large stakes: with larger stakes subjects might make choices in line with the preferences they express when asked about hypothetical choices.

In summary, the evidence for an inconsistency with preferences compatible with an expected payoff function is, at a minimum, suggestive. It has spurred the development of alternative theories. Nevertheless, the vast majority of models in game theory (and also in economics) that involve choice under uncertainty currently assume that each decision-maker's preferences are represented by the expected value of a payoff function. I maintain this assumption throughout the book, although many of the ideas I discuss appear not to depend on it.

4.2 Strategic games in which players may randomize

To study stochastic steady states, we extend the notion of a strategic game given in Definition 13.1 by endowing each player with vNM preferences about lotteries over the set of action profiles.

▶ DEFINITION 106.1 (*Strategic game with vNM preferences*) A **strategic game** (with vNM preferences) consists of

- a set of **players**
- for each player, a set of **actions**
- for each player, **preferences** regarding lotteries over action profiles that may be represented by the expected value of a ("Bernoulli") payoff function over action profiles.

A two-player strategic game with vNM preferences in which each player has finitely many actions may be presented in a table like those in Chapter 2. Such a table looks exactly the same as it did before, though the interpretation of the numbers in the boxes is different. In Chapter 2 these numbers are values of payoff functions that represent the players' preferences over deterministic outcomes; here they are the values of (Bernoulli) payoff functions whose expected values represent the players' preferences over lotteries.

Given the change in the interpretation of the payoffs, two tables that represent the same strategic game with ordinal preferences no longer necessarily represent the same strategic game with vNM preferences. For example, the two tables in Figure 107.1 represent the same game with ordinal preferences—namely the *Prisoner's Dilemma* (Section 2.2). In both cases the best outcome for each player is that in which she chooses F and the other player chooses Q, the next best outcome is (Q, Q), then comes (F, F), and the worst outcome is that in which she chooses Q and the other player chooses F. However, the tables represent *different* strategic games with vNM preferences. For example, in the table on the left, player 1's payoff to (Q, Q) is the *same* as her expected payoff to the lottery that yields (F, Q) with probability $\frac{1}{2}$ and (F, F) with probability $\frac{1}{2}$ ($\frac{1}{2}u_1(F, Q) + \frac{1}{2}u_1(F, F) = \frac{1}{2} \cdot 3 + \frac{1}{2} \cdot 1 = 2 = u_1(Q, Q)$), whereas in the table on the right, her payoff to (Q, Q) is *greater than* her expected payoff to this lottery ($3 > \frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 1$). Thus the left-hand table represents a situation in which player 1 is indifferent between the deterministic outcome (Q, Q) and the lottery in which (F, Q) occurs with probability $\frac{1}{2}$ and (F, F) occurs with probability $\frac{1}{2}$. In the right-hand table, however, she prefers the deterministic outcome (Q, Q) to the lottery.

To show, as in this example, that two tables represent different strategic games with vNM preferences, we need only find a pair of lotteries whose expected payoffs are ordered differently by the two tables. To show that they represent the *same* strategic game with vNM preferences is more difficult; see Section 4.12.2.

? EXERCISE 106.2 (Extensions of *BoS* with vNM preferences) Construct a table of payoffs for a strategic game with vNM preferences in which the players' prefer-

	Q	F
Q	$2,2$	$0,3$
F	$3,0$	$1,1$

	Q	F
Q	$3,3$	$0,4$
F	$4,0$	$1,1$

Figure 107.1 Two tables that represent the same strategic game with ordinal preferences but different strategic games with vNM preferences.

ences over deterministic outcomes are the same as they are in *BoS* (Example 18.2), and their preferences over lotteries satisfy the following condition. Each player is indifferent between (*i*) going to her less preferred concert in the company of the other player, and (*ii*) the lottery in which with probability $\frac{1}{2}$ she and the other player go to different concerts and with probability $\frac{1}{2}$ they both go to her more preferred concert. Do the same in the case that each player is indifferent between (*i*) going to her less preferred concert in the company of the other player and (*ii*) the lottery in which with probability $\frac{3}{4}$ she and the other player go to different concerts and with probability $\frac{1}{4}$ they both go to her more preferred concert. (In each case set each player's payoff to the outcome that she least prefers equal to 0 and her payoff to the outcome that she most prefers equal to 2.)

Despite the importance of saying how the numbers in a payoff table should be interpreted, users of game theory sometimes fail to make the interpretation clear. When one interprets discussions of Nash equilibrium in the literature, a reasonably safe assumption is that if the players are not allowed to choose their actions randomly, then the numbers in payoff tables are payoffs that represent the players' ordinal preferences, whereas if the players are allowed to randomize, then the numbers are payoffs whose expected values represent the players' preferences regarding lotteries over outcomes.

4.3 Mixed strategy Nash equilibrium

4.3.1 Mixed strategies

In the generalization of the notion of Nash equilibrium that models a stochastic steady state of a strategic game with vNM preferences, we allow each player to choose a probability distribution over her set of actions rather than restricting her to choose a single deterministic action. We refer to such a probability distribution as a *mixed strategy*.

► DEFINITION 107.1 (*Mixed strategy*) A **mixed strategy** of a player in a strategic game is a probability distribution over the player's actions.

I usually use α to denote a profile of mixed strategies; $\alpha_i(a_i)$ is the probability assigned by player i's mixed strategy α_i to her action a_i. To specify a mixed strategy of player i we need to give the probability it assigns to each of player i's actions. For example, the strategy of player 1 in *Matching Pennies* that assigns probability $\frac{1}{2}$ to each action is the strategy α_1 for which $\alpha_1(\textit{Head}) = \frac{1}{2}$ and $\alpha_1(\textit{Tail}) = \frac{1}{2}$. Because this way of describing a mixed strategy is cumbersome, I often use a shorthand for a game that is presented in a table like those in Figure 107.1: I write a mixed strategy as a list of probabilities, one for each action, *in the order the actions are given in the table*. For example, the mixed strategy $(\frac{1}{3}, \frac{2}{3})$ for player 1 in either of the games in Figure 107.1 assigns probability $\frac{1}{3}$ to Q and probability $\frac{2}{3}$ to F.

A mixed strategy may assign probability 1 to a single action: by *allowing* a player to choose probability distributions, we do not prohibit her from choosing deterministic actions. We refer to such a mixed strategy as a **pure strategy**. Player i's choosing the pure strategy that assigns probability 1 to the action a_i is equivalent to her simply choosing the action a_i, and I denote this strategy simply by a_i.

4.3.2 Equilibrium

The notion of equilibrium that we study is called "mixed strategy Nash equilibrium". The idea behind it is the same as the idea behind the notion of Nash equilibrium for a game with ordinal preferences: a mixed strategy Nash equilibrium is a mixed strategy profile α^* with the property that no player i has a mixed strategy α_i such that she prefers the lottery over outcomes generated by the strategy profile $(\alpha_i, \alpha^*_{-i})$ to the lottery over outcomes generated by the strategy profile α^*. The following definition states this condition using payoff functions whose expected values represent the players' preferences.

► DEFINITION 108.1 (*Mixed strategy Nash equilibrium of strategic game with vNM preferences*) The mixed strategy profile α^* in a strategic game with vNM preferences is a **(mixed strategy) Nash equilibrium** if, for each player i and every mixed strategy α_i of player i, the expected payoff to player i of α^* is at least as large as the expected payoff to player i of $(\alpha_i, \alpha^*_{-i})$ according to a payoff function whose expected value represents player i's preferences over lotteries. Equivalently, for each player i,

$$U_i(\alpha^*) \geq U_i(\alpha_i, \alpha^*_{-i}) \text{ for every mixed strategy } \alpha_i \text{ of player } i, \qquad (108.2)$$

where $U_i(\alpha)$ is player i's expected payoff to the mixed strategy profile α.

The technique of constructing the players' best response functions (Section 2.8), useful in finding Nash equilibria of some strategic games with ordinal preferences, is useful too in finding mixed strategy Nash equilibria of some strategic games with vNM preferences, especially very simple ones. I discuss this technique in the next section. In Section 4.3.4 I discuss a characterization of mixed strategy Nash equilibrium that is an invaluable tool for studying the equilibria of any game.

4.3.3 Best response functions

General definition As before, I denote player i's best response function by B_i. For a strategic game with ordinal preferences, $B_i(a_{-i})$ is the set of player i's best actions when the list of the other players' actions is a_{-i}. For a strategic game with vNM preferences, $B_i(\alpha_{-i})$ is the set of player i's best mixed strategies when the list of the other players' mixed strategies is α_{-i}. From the definition of a mixed strategy equilibrium, a profile α^* of mixed strategies is a mixed strategy Nash equilibrium if and only if every player's mixed strategy is a best response to the other players' mixed strategies (cf. Proposition 36.1):

> *the mixed strategy profile α^* is a mixed strategy Nash equilibrium if and only if α_i^* is in $B_i(\alpha_{-i}^*)$ for every player i.*

Two-player two-action games The analysis of *Matching Pennies* in Section 4.1.2 shows that each player's set of best responses to a mixed strategy of the other player is either a single pure strategy or the set of *all* mixed strategies. (For example, if player 2's mixed strategy assigns probability less than $\frac{1}{2}$ to *Head*, then player 1's unique best response is the pure strategy *Tail*, if player 2's mixed strategy assigns probability greater than $\frac{1}{2}$ to *Head*, then player 1's unique best response is the pure strategy *Head*, and if player 2's mixed strategy assigns probability $\frac{1}{2}$ to *Head*, then all of player 1's mixed strategies are best responses.)

The character of each player's set of best responses in any two-player game in which each player has two actions is similar to the character of each player's set of best responses in *Matching Pennies*: it consists either of a single pure strategy or of all mixed strategies. The reason lies in the form of the payoff functions.

Consider a two-player game in which each player has two actions, T and B for player 1 and L and R for player 2. Denote by u_i, for $i = 1, 2$, a Bernoulli payoff function for player i. (That is, u_i is a payoff function over action pairs whose expected value represents player i's preferences regarding lotteries over action pairs.) Player 1's mixed strategy α_1 assigns probability $\alpha_1(T)$ to her action T and probability $\alpha_1(B)$ to her action B (with $\alpha_1(T) + \alpha_1(B) = 1$). For convenience, let $p = \alpha_1(T)$, so that $\alpha_1(B) = 1 - p$. Similarly, denote the probability $\alpha_2(L)$ that player 2's mixed strategy assigns to L by q, so that $\alpha_2(R) = 1 - q$.

We take the players' choices to be independent, so that when the players use the mixed strategies α_1 and α_2, the probability of any action pair (a_1, a_2) is the product of the probability player 1's mixed strategy assigns to a_1 and the probability player 2's mixed strategy assigns to a_2. (See Section 17.6.2 in the mathematical appendix if you are not familiar with the idea of independence.) Thus the probability distribution generated by the mixed strategy pair (α_1, α_2) over the four possible outcomes of the game has the form given in Figure 109.1: (T, L) occurs with probability pq, (T, R) occurs with probability $p(1 - q)$, (B, L) occurs with probability $(1 - p)q$, and (B, R) occurs with probability $(1 - p)(1 - q)$.

	$L\ (q)$	$R\ (1-q)$
$T\ (p)$	pq	$p(1-q)$
$B\ (1-p)$	$(1-p)q$	$(1-p)(1-q)$

Figure 109.1 The probabilities of the four outcomes in a two-player two-action strategic game when player 1's mixed strategy is $(p, 1 - p)$ and player 2's mixed strategy is $(q, 1 - q)$.

From this probability distribution we see that player 1's expected payoff to the mixed strategy pair (α_1, α_2) is

$$pq \cdot u_1(T,L) + p(1-q) \cdot u_1(T,R) + (1-p)q \cdot u_1(B,L) + (1-p)(1-q) \cdot u_1(B,R),$$

which we can alternatively write as

$$p[q \cdot u_1(T,L) + (1-q) \cdot u_1(T,R)] + (1-p)[q \cdot u_1(B,L) + (1-q) \cdot u_1(B,R)].$$

The first term in square brackets is player 1's expected payoff when she uses a *pure* strategy that assigns probability 1 to T and player 2 uses her mixed strategy α_2; the second term in square brackets is player 1's expected payoff when she uses a *pure* strategy that assigns probability 1 to B and player 2 uses her mixed strategy α_2. Denote these two expected payoffs $E_1(T, \alpha_2)$ and $E_1(B, \alpha_2)$. Then player 1's expected payoff to the mixed strategy pair (α_1, α_2) is

$$pE_1(T, \alpha_2) + (1-p)E_1(B, \alpha_2).$$

That is, player 1's expected payoff to the mixed strategy pair (α_1, α_2) is a weighted average of her expected payoffs to T and B when player 2 uses the mixed strategy α_2, with weights equal to the probabilities assigned to T and B by α_1.

In particular, player 1's expected payoff, given player 2's mixed strategy, is a *linear* function of p—when plotted in a graph, it is a straight line.[1] A case in which $E_1(T, \alpha_2) > E_1(B, \alpha_2)$ is illustrated in Figure 110.1.

? EXERCISE 110.1 (Expected payoffs) Construct diagrams like Figure 110.1 for *BoS* (Figure 19.1) and the game in the right panel of Figure 21.1 (in each case treating the numbers in the tables as Bernoulli payoffs). In each diagram, plot player 1's expected payoff as a function of the probability p that she assigns to her top action in three cases: when the probability q that player 2 assigns to her left action is 0, $\frac{1}{2}$, and 1.

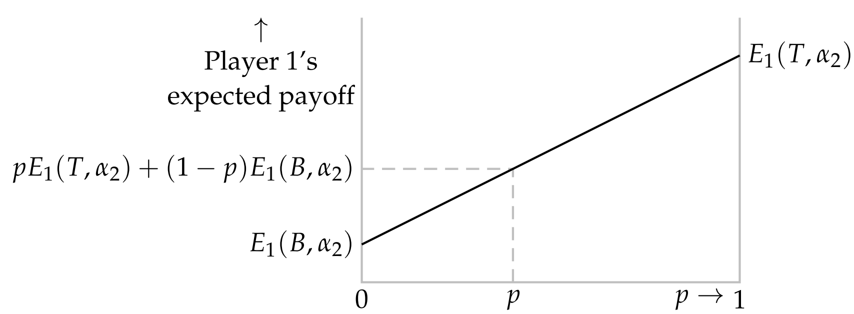

Figure 110.1 Player 1's expected payoff as a function of the probability p she assigns to T in the game in which her actions are T and B, when player 2's mixed strategy is α_2 and $E_1(T, \alpha_2) > E_1(B, \alpha_2)$.

[1]See Section 17.3 (in particular Figure 496.2) for my usage of the term "linear".

A significant implication of the linearity of player 1's expected payoff is that there are three possibilities for her best response to a given mixed strategy of player 2:

- player 1's unique best response is the pure strategy T (if $E_1(T, \alpha_2) > E_1(B, \alpha_2)$, as in Figure 110.1)
- player 1's unique best response is the pure strategy B (if $E_1(B, \alpha_2) > E_1(T, \alpha_2)$, in which case the line representing player 1's expected payoff as a function of p in the analogue of Figure 110.1 slopes down)
- all mixed strategies of player 1 yield the same expected payoff, and hence all are best responses (if $E_1(T, \alpha_2) = E_1(B, \alpha_2)$, in which case the line representing player 1's expected payoff as a function of p in the analogue of Figure 110.1 is horizontal).

In particular, a mixed strategy $(p, 1-p)$ for which $0 < p < 1$ is never the *unique* best response; either it is not a best response or *all* mixed strategies are best responses.

? EXERCISE 111.1 (Best responses) For each game and each value of q in Exercise 110.1, use the graphs you drew in that exercise to find player 1's set of best responses.

Example: Matching Pennies The argument in Section 4.1.2 establishes that *Matching Pennies* has a unique mixed strategy Nash equilibrium, in which each player's mixed strategy assigns probability $\frac{1}{2}$ to *Head* and probability $\frac{1}{2}$ to *Tail*. I now describe an alternative route to this conclusion that uses the method described in Section 2.8.3, which involves explicitly constructing the players' best response functions; this method may be used in other games.

Represent each player's preferences by the expected value of a payoff function that assigns the payoff 1 to a gain of \$1 and the payoff -1 to a loss of \$1. The resulting strategic game with vNM preferences is shown in Figure 111.1.

	Head	*Tail*
Head	$1, -1$	$-1, \ \ 1$
Tail	$-1, \ \ 1$	$1, -1$

Figure 111.1 *Matching Pennies.*

Denote by p the probability that player 1's mixed strategy assigns to *Head*, and by q the probability that player 2's mixed strategy assigns to *Head*. Then, given player 2's mixed strategy, player 1's expected payoff to the pure strategy *Head* is

$$q \cdot 1 + (1-q) \cdot (-1) = 2q - 1$$

and her expected payoff to *Tail* is

$$q \cdot (-1) + (1-q) \cdot 1 = 1 - 2q.$$

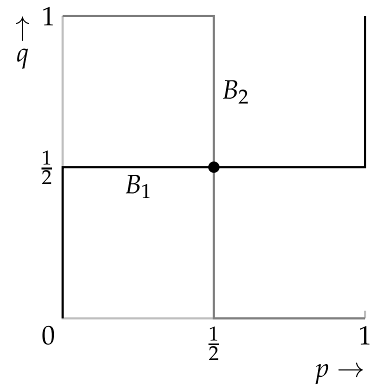

Figure 112.1 The players' best response functions in *Matching Pennies* (Figure 111.1) when randomization is allowed. The probabilities assigned by players 1 and 2 to *Head* are p and q, respectively. The best response function of player 1 is black and that of player 2 is gray. The disk indicates the unique Nash equilibrium.

Thus if $q < \frac{1}{2}$, then player 1's expected payoff to *Tail* exceeds her expected payoff to *Head*, and hence exceeds also her expected payoff to every mixed strategy that assigns a positive probability to *Head*. (Recall the discussion in the previous section.) Similarly, if $q > \frac{1}{2}$, then her expected payoff to *Head* exceeds her expected payoff to *Tail*, and hence exceeds her expected payoff to every mixed strategy that assigns a positive probability to *Tail*. If $q = \frac{1}{2}$, then both *Head* and *Tail*, and hence all her mixed strategies, yield the same expected payoff. We conclude that player 1's best responses to player 2's strategy are her mixed strategy that assigns probability 0 to *Head* if $q < \frac{1}{2}$, her mixed strategy that assigns probability 1 to *Head* if $q > \frac{1}{2}$, and all her mixed strategies if $q = \frac{1}{2}$. That is, denoting by $B_1(q)$ the set of probabilities player 1 assigns to *Head* in best responses to q, we have

$$B_1(q) = \begin{cases} \{0\} & \text{if } q < \frac{1}{2} \\ \{p\colon 0 \le p \le 1\} & \text{if } q = \frac{1}{2} \\ \{1\} & \text{if } q > \frac{1}{2}. \end{cases}$$

The best response function of player 2 is similar: $B_2(p) = \{1\}$ if $p < \frac{1}{2}$, $B_2(p) = \{q\colon 0 \le q \le 1\}$ if $p = \frac{1}{2}$, and $B_2(p) = \{0\}$ if $p > \frac{1}{2}$. Both players' best response functions are illustrated in Figure 112.1 (the best response function of player 1 is black and that of player 2 is gray).

The set of mixed strategy Nash equilibria of the game corresponds (as before) to the set of intersections of the best response functions in this figure; we see that there is one intersection, corresponding to the equilibrium we found previously, in which each player assigns probability $\frac{1}{2}$ to *Head*.

Matching Pennies has no Nash equilibrium if the players are not allowed to randomize. If a game has a Nash equilibrium when randomization is not allowed, is it possible that it has additional equilibria when randomization is allowed? The following example shows that the answer is positive.

	B	S
B	$2,1$	$0,0$
S	$0,0$	$1,2$

Figure 113.1 A version of the game *BoS* with vNM preferences.

Example: BoS Consider the two-player game with vNM preferences in which the players' preferences over deterministic action profiles are the same as in *BoS* and their preferences over lotteries are represented by the expected value of the payoff functions specified in Figure 113.1. What are the mixed strategy equilibria of this game?

First construct player 1's best response function. Suppose that player 2 assigns probability q to B. Then player 1's expected payoff to B is $2 \cdot q + 0 \cdot (1-q) = 2q$ and her expected payoff to S is $0 \cdot q + 1 \cdot (1-q) = 1-q$. Thus if $2q > 1-q$, or $q > \frac{1}{3}$, then her unique best response is B, while if $q < \frac{1}{3}$, then her unique best response is S. If $q = \frac{1}{3}$, then both B and S, and hence all player 1's mixed strategies, yield the same expected payoffs, so that every mixed strategy is a best response. In summary, player 1's best response function is

$$B_1(q) = \begin{cases} \{0\} & \text{if } q < \frac{1}{3} \\ \{p : 0 \le p \le 1\} & \text{if } q = \frac{1}{3} \\ \{1\} & \text{if } q > \frac{1}{3}. \end{cases}$$

Similarly we can find player 2's best response function. The best response functions of both players are shown in Figure 113.2.

We see that the game has three mixed strategy Nash equilibria, in which $(p,q) = (0,0)$, $(\frac{2}{3}, \frac{1}{3})$, and $(1,1)$. The first and third equilibria correspond to the Nash equilibria of the ordinal version of the game when the players were not allowed to randomize (Section 2.7.2). The second equilibrium is new. In this equilibrium each

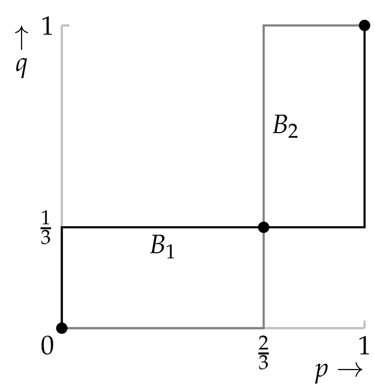

Figure 113.2 The players' best response functions in *BoS* (Figure 113.1) when randomization is allowed. The probabilities assigned by players 1 and 2 to B are p and q, respectively. The best response function of player 1 is black and that of player 2 is gray. The disks indicate the Nash equilibria (two pure, one mixed).

player chooses both B and S with positive probability (so that each of the four outcomes (B, B), (B, S), (S, B), and (S, S) occurs with positive probability).

? EXERCISE 114.1 (Mixed strategy equilibria of *Hawk–Dove*) Consider the two-player game with vNM preferences in which the players' preferences over deterministic action profiles are the same as in *Hawk–Dove* (Exercise 31.2) and their preferences over lotteries satisfy the following two conditions. Each player is indifferent between (*i*) the outcome $(Passive, Passive)$ and (*ii*) the lottery that assigns probability $\frac{1}{2}$ to $(Aggressive, Aggressive)$ and probability $\frac{1}{2}$ to the outcome in which she is aggressive and the other player is passive; each player is indifferent also between (*i*) the outcome in which she is passive and the other player is aggressive and (*ii*) the lottery that assigns probability $\frac{2}{3}$ to the outcome $(Aggressive, Aggressive)$ and probability $\frac{1}{3}$ to the outcome $(Passive, Passive)$. Find payoffs whose expected values represent these preferences (take each player's payoff to $(Aggressive, Aggressive)$ to be 0 and each player's payoff to the outcome in which she is passive and the other player is aggressive to be 1). Find the mixed strategy Nash equilibrium of the resulting strategic game.

Both *Matching Pennies* and *BoS* have finitely many mixed strategy Nash equilibria: the players' best response functions intersect at a finite number of points (one for *Matching Pennies*, three for *BoS*). One of the games in the next exercise has a continuum of mixed strategy Nash equilibria because segments of the players' best response functions coincide.

? EXERCISE 114.2 (Games with mixed strategy equilibria) Find all the mixed strategy Nash equilibria of the strategic games in Figure 114.1.

? EXERCISE 114.3 (A coordination game) Two people can perform a task if, and only if, they both exert effort. They are both better off if they both exert effort and perform the task than if neither exerts effort (and nothing is accomplished); the worst outcome for each person is that she exerts effort and the other person does not (in which case again nothing is accomplished). Specifically, the players' preferences are represented by the expected value of the payoff functions in Figure 115.1, where c is a positive number less than 1 that can be interpreted as the cost of exerting effort. Find all the mixed strategy Nash equilibria of this game. How do the equilibria change as c increases? Explain the reasons for the changes.

?? EXERCISE 114.4 (Swimming with sharks) You and a friend are spending two days at the beach; you both enjoy swimming. Each of you believes that with probabil-

	L	R
T	$6, 0$	$0, 6$
B	$3, 2$	$6, 0$

	L	R
T	$0, 1$	$0, 2$
B	$2, 2$	$0, 1$

Figure 114.1 Two strategic games with vNM preferences.

	No effort	*Effort*
No effort	$0, 0$	$0, -c$
Effort	$-c, 0$	$1-c, 1-c$

Figure 115.1 The coordination game in Exercise 114.3.

ity π the water is infested with sharks. If sharks are present, any swimmer will surely be attacked. Each of you has preferences represented by the expected value of a payoff function that assigns $-c$ to being attacked by a shark, 0 to sitting on the beach, and 1 to a day's worth of undisturbed swimming (where $c > 0$!). If a swimmer is attacked by sharks on the first day, then you both deduce that a swimmer will surely be attacked the next day, and hence do not go swimming the next day. If at least one of you swims on the first day and is not attacked, then you both know that the water is shark-free. If neither of you swims on the first day, each of you retains the belief that the probability of the water's being infested is π, and hence on the second day swims if $-\pi c + 1 - \pi > 0$ and sits on the beach if $-\pi c + 1 - \pi < 0$, thus receiving an expected payoff of $\max\{-\pi c + 1 - \pi, 0\}$. Model this situation as a strategic game in which you and your friend each decide whether to go swimming on your first day at the beach. If, for example, you go swimming on the first day, you (and your friend, if she goes swimming) are attacked with probability π, in which case you stay out of the water on the second day; you (and your friend, if she goes swimming) swim undisturbed with probability $1 - \pi$, in which case you swim on the second day. Thus your expected payoff if you swim on the first day is $\pi(-c + 0) + (1 - \pi)(1 + 1) = -\pi c + 2(1 - \pi)$, independent of your friend's action. Find the mixed strategy Nash equilibria of the game (depending on c and π). Does the existence of a friend make it more or less likely that you decide to go swimming on the first day? (Penguins diving into water where seals may lurk are sometimes said to face the same dilemma; Court (1996) argues that they do not.)

4.3.4 A useful characterization of mixed strategy Nash equilibrium

The method we have used so far to study the set of mixed strategy Nash equilibria of a game involves constructing the players' best response functions. This method is useful in simple games, but is of limited use in more complicated ones. I now present a characterization of mixed strategy Nash equilibrium that is invaluable in the study of general games. The characterization gives us an easy way to check whether a mixed strategy profile is an equilibrium; it is also the basis of a procedure (described in Section 4.10) for finding all equilibria of a game.

The key point is an observation made in Section 4.3.3 for two-player two-action games: a player's expected payoff to a mixed strategy profile is a weighted average of her expected payoffs to her pure strategies, where the weight attached to each pure strategy is the probability assigned to that strategy by the player's mixed strategy. This property holds for any game (with any number of players) in which each player has finitely many actions. We can state it more precisely as follows.

> *A player's expected payoff to the mixed strategy profile α is a weighted average of her expected payoffs to all mixed strategy profiles of the type (a_i, α_{-i}), where the weight attached to (a_i, α_{-i}) is the probability $\alpha_i(a_i)$ assigned to a_i by player i's mixed strategy α_i.* (116.1)

Symbolically we have

$$U_i(\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i) E_i(a_i, \alpha_{-i}),$$

where A_i is player i's set of actions (pure strategies) and $E_i(a_i, \alpha_{-i})$ is (as before) her expected payoff when she uses the pure strategy that assigns probability 1 to a_i and every other player j uses her mixed strategy α_j. (See the end of Section 17.2 in the appendix on mathematics for an explanation of the $\sum$ notation.)

This property leads to a useful characterization of mixed strategy Nash equilibrium. Let α^* be a mixed strategy Nash equilibrium and denote by E_i^* player i's expected payoff in the equilibrium (i.e. $E_i^* = U_i(\alpha^*)$). Because α^* is an equilibrium, player i's expected payoff, given α_{-i}^*, to all her strategies, including all her pure strategies, is at most E_i^*. Now, by (116.1), E_i^* is a weighted average of player i's expected payoffs to the pure strategies to which α_i^* assigns positive probability. Thus player i's expected payoffs to these pure strategies are all equal to E_i^*. (If any were smaller, then the weighted average would be smaller.) We conclude that the expected payoff to each action to which α_i^* assigns positive probability is E_i^*, and the expected payoff to every other action is at most E_i^*. Conversely, if these conditions are satisfied for every player i, then α^* is a mixed strategy Nash equilibrium: the expected payoff to α_i^* is E_i^*, and the expected payoff to any other mixed strategy is at most E_i^*, because by (116.1) it is a weighted average of E_i^* and numbers that are at most E_i^*.

This argument establishes the following result.

■ PROPOSITION 116.2 (Characterization of mixed strategy Nash equilibrium of finite game) *A mixed strategy profile α^* in a strategic game with vNM preferences in which each player has finitely many actions is a mixed strategy Nash equilibrium if and only if, for each player i,*

- *the expected payoff, given α_{-i}^*, to every action to which α_i^* assigns positive probability is the same*
- *the expected payoff, given α_{-i}^*, to every action to which α_i^* assigns zero probability is at most the expected payoff to any action to which α_i^* assigns positive probability.*

Each player's expected payoff in an equilibrium is her expected payoff to any of her actions that she uses with positive probability.

The significance of this result is that it gives conditions for a mixed strategy Nash equilibrium in terms of each player's expected payoffs only to her *pure* strategies. For games in which each player has finitely many actions, it allows us easily

	L (0)	C ($\frac{1}{3}$)	R ($\frac{2}{3}$)
T ($\frac{3}{4}$)	$\cdot, 2$	$3, 3$	$1, 1$
M (0)	$\cdot, \cdot$	$0, \cdot$	$2, \cdot$
B ($\frac{1}{4}$)	$\cdot, 4$	$5, 1$	$0, 7$

Figure 117.1 A partially specified strategic game, illustrating a method of checking whether a mixed strategy profile is a mixed strategy Nash equilibrium. The dots indicate irrelevant payoffs.

to check whether a mixed strategy profile is an equilibrium. For example, in *BoS* (Section 4.3.3) the strategy pair $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$ is a mixed strategy Nash equilibrium because given player 2's strategy $(\frac{1}{3}, \frac{2}{3})$, player 1's expected payoffs to B and S are both equal to $\frac{2}{3}$, and given player 1's strategy $(\frac{2}{3}, \frac{1}{3})$, player 2's expected payoffs to B and S are both equal to $\frac{2}{3}$.

The next example is slightly more complicated.

◆ EXAMPLE 117.1 (Checking whether a mixed strategy profile is a mixed strategy Nash equilibrium) I claim that for the game in Figure 117.1 (in which the dots indicate irrelevant payoffs), the indicated pair of strategies, $(\frac{3}{4}, 0, \frac{1}{4})$ for player 1 and $(0, \frac{1}{3}, \frac{2}{3})$ for player 2, is a mixed strategy Nash equilibrium. To verify this claim, it suffices, by Proposition 116.2, to study each player's expected payoffs to her three pure strategies. For player 1 these payoffs are

$$\begin{aligned} T &: \tfrac{1}{3} \cdot 3 + \tfrac{2}{3} \cdot 1 = \tfrac{5}{3} \\ M &: \tfrac{1}{3} \cdot 0 + \tfrac{2}{3} \cdot 2 = \tfrac{4}{3} \\ B &: \tfrac{1}{3} \cdot 5 + \tfrac{2}{3} \cdot 0 = \tfrac{5}{3}. \end{aligned}$$

Player 1's mixed strategy assigns positive probability to T and B and probability zero to M, so the two conditions in Proposition 116.2 are satisfied for player 1. The expected payoff to each of player 2's pure strategies is $\frac{5}{2}$ $(\frac{3}{4} \cdot 2 + \frac{1}{4} \cdot 4 = \frac{3}{4} \cdot 3 + \frac{1}{4} \cdot 1 = \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 7 = \frac{5}{2})$, so the two conditions in Proposition 116.2 are satisfied also for her.

Note that the expected payoff to player 2's action L, which she uses with probability zero, is the *same* as the expected payoff to her other two actions. This equality is consistent with Proposition 116.2, the second part of which requires only that the expected payoffs to actions used with probability zero be *no greater than* the expected payoffs to actions used with positive probability (not that they necessarily be less). Note also that the fact that player 2's expected payoff to L is the same as her expected payoffs to C and R does *not* imply that the game has a mixed strategy Nash equilibrium in which player 2 uses L with positive probability—it may, or it may not, depending on the unspecified payoffs.

? EXERCISE 117.2 (Choosing numbers) Players 1 and 2 each choose a positive integer up to K. If the players choose the same number, then player 2 pays \$1 to player 1; otherwise no payment is made. Each player's preferences are represented by her expected monetary payoff.

a. Show that the game has a mixed strategy Nash equilibrium in which each player chooses each positive integer up to K with probability $1/K$.

b. (More difficult.) Show that the game has no other mixed strategy Nash equilibria. (Deduce from the fact that player 1 assigns positive probability to some action k that player 2 must do so; then look at the implied restriction on player 1's equilibrium strategy.)

? EXERCISE 118.1 (Silverman's game) Each of two players chooses a positive integer. If player i's integer is greater than player j's integer and less than three times this integer, then player j pays \$1 to player i. If player i's integer is at least three times player j's integer, then player i pays \$1 to player j. If the integers are equal, no payment is made. Each player's preferences are represented by her expected monetary payoff. Show that the game has no Nash equilibrium in pure strategies and that the pair of mixed strategies in which each player chooses 1, 2, and 5 each with probability $\frac{1}{3}$ is a mixed strategy Nash equilibrium. (In fact, this pair of mixed strategies is the unique mixed strategy Nash equilibrium.) (You cannot appeal to Proposition 116.2 because the number of actions of each player is not finite. However, you can use the argument for the "if" part of this result.)

?? EXERCISE 118.2 (Voter participation) Consider the game of voter participation in Exercise 34.2. Assume that $2 \leq k \leq m$ and that each player's preferences are represented by the expectation of her payoffs given in Exercise 34.2. Show that there is a value of p between 0 and 1 such that the game has a mixed strategy Nash equilibrium in which every supporter of candidate A votes with probability p, k supporters of candidate B vote with certainty, and the remaining $m-k$ supporters of candidate B abstain. How do the probability p that a supporter of candidate A votes and the expected number of voters ("turnout") depend upon c? (Note that if every supporter of candidate A votes with probability p, then the probability that exactly $k-1$ of them vote is $kp^{k-1}(1-p)$.)

?? EXERCISE 118.3 (Defending territory) General A is defending territory accessible by two mountain passes against an attack by General B. General A has three divisions at her disposal, and general B has two divisions. Each general allocates her divisions between the two passes. General A wins the battle at a pass if and only if she assigns at least as many divisions to the pass as does General B; she successfully defends her territory if and only if she wins the battle at both passes. Formulate this situation as a strategic game and find all its mixed strategy equilibria. (First argue that in every equilibrium B assigns probability zero to the action of allocating one division to each pass. Then argue that in any equilibrium she assigns probability $\frac{1}{2}$ to each of her other actions. Finally, find A's equilibrium strategies.) In an equilibrium, do the generals concentrate all their forces at one pass, or spread them out?

An implication of Proposition 116.2 is that a nondegenerate mixed strategy equilibrium (a mixed strategy equilibrium that is not also a pure strategy equilibrium) is never a *strict* Nash equilibrium: every player whose mixed strategy assigns positive probability to more than one action is indifferent between her equilibrium mixed strategy and every action to which this mixed strategy assigns positive probability.

Any equilibrium (in mixed strategies or not) that is not strict has less appeal than a strict equilibrium because some (or all) of the players lack a positive incentive to choose their equilibrium strategies, given the other players' behavior. There is no reason for them *not* to choose their equilibrium strategies, but at the same time there is no reason for them not to choose another strategy that is equally good. Many pure strategy equilibria—especially in complex games—are also not strict, but among mixed strategy equilibria the problem is pervasive.

Given that in a mixed strategy equilibrium no player has a positive incentive to choose her equilibrium strategy, what determines how she randomizes in equilibrium? From the examples studied in Section 4.3.3 (e.g. *Matching Pennies* and *BoS*) we see that a player's equilibrium mixed strategy in a two-player game keeps the *other* player indifferent between a set of her actions, so that *she* is willing to randomize. In the mixed strategy equilibrium of *BoS*, for example, player 1 chooses B with probability $\frac{2}{3}$ so that player 2 is indifferent between B and S, and hence is willing to choose each with positive probability. Note, however, that the theory is *not* that the players consciously choose their strategies with this goal in mind! Rather, the conditions for equilibrium are designed to ensure that it is consistent with a steady state. In *BoS*, for example, if player 1 chooses B with probability $\frac{2}{3}$ and player 2 chooses B with probability $\frac{1}{3}$, then neither player has any reason to change her action. We have not yet studied how a steady state might come about, but have rather simply looked for strategy profiles consistent with steady states. In Section 4.9 I briefly discuss some theories of how a steady state might be reached.

4.3.5 Existence of equilibrium in finite games

Every game we have examined has at least one mixed strategy Nash equilibrium. In fact, every game in which each player has *finitely* many actions has at least one such equilibrium.

■ PROPOSITION 119.1 (Existence of mixed strategy Nash equilibrium in finite games) *Every strategic game with vNM preferences in which each player has finitely many actions has a mixed strategy Nash equilibrium.*

This result is of no help in *finding* equilibria. But it is a useful fact: your quest for an equilibrium of a game in which each player has finitely many actions in principle may succeed! Note that the finiteness of the number of actions of each player is only *sufficient* for the existence of an equilibrium, not *necessary*; many games in which the players have infinitely many actions possess mixed strategy Nash equilibria. Note also that a player's strategy in a mixed strategy Nash equilibrium may assign probability 1 to a single action; if every player's strategy does so, then the equilibrium corresponds to a ("pure strategy") equilibrium of the associated game with ordinal preferences. Relatively advanced mathematical tools are needed to prove the result; see, for example, Osborne and Rubinstein (1994, 19–20).

4.4 Dominated actions

In a strategic game with ordinal preferences, one action of a player strictly dominates another action if it is superior, no matter what the other players do (see Definition 45.1). In a game with vNM preferences in which players may randomize, we extend this definition to allow an action to be dominated by a *mixed strategy*.

▶ DEFINITION 120.1 (*Strict domination*) In a strategic game with vNM preferences, player i's mixed strategy α_i **strictly dominates** her action a'_i if

$$U_i(\alpha_i, a_{-i}) > u_i(a'_i, a_{-i}) \text{ for every list } a_{-i} \text{ of the other players' actions,}$$

where u_i is a payoff function whose expected value represents player i's preferences over lotteries and $U_i(\alpha_i, a_{-i})$ is player i's expected payoff under u_i when she uses the mixed strategy α_i and the actions chosen by the other players are given by a_{-i}. We say that the action a'_i is **strictly dominated**.

Figure 120.1 (in which only player 1's payoffs are given) shows that an action not strictly dominated by any pure strategy (i.e. is not strictly dominated in the sense of Definition 45.1) may be strictly dominated by a mixed strategy. The action T of player 1 is not strictly (or weakly) dominated by either M or B, but it is strictly dominated by the mixed strategy that assigns probability $\frac{1}{2}$ to M and probability $\frac{1}{2}$ to B, because if player 2 chooses L, then the mixed strategy yields player 1 the payoff of 2, whereas the action T yields her the payoff of 1, and if player 2 chooses R, then the mixed strategy yields player 1 the payoff of $\frac{3}{2}$, whereas the action T yields her the payoff of 1.

? EXERCISE 120.2 (Strictly dominating mixed strategies) In Figure 120.1, the mixed strategy that assigns probability $\frac{1}{2}$ to M and probability $\frac{1}{2}$ to B is not the only mixed strategy that strictly dominates T. Find all the mixed strategies that do so.

? EXERCISE 120.3 (Strict domination for mixed strategies) Determine whether each of the following statements is true or false. (*a*) A mixed strategy that assigns positive probability to a strictly dominated action is strictly dominated. (*b*) A mixed strategy that assigns positive probability only to actions that are not strictly dominated is not strictly dominated.

In a Nash equilibrium of a strategic game with ordinal preferences, no player uses a strictly dominated action (Section 2.9.1). I now argue that the same is true of

	L	R
T	1	1
M	4	0
B	0	3

Figure 120.1 Player 1's payoffs in a strategic game with vNM preferences. The action T of player 1 is strictly dominated by the mixed strategy that assigns probability $\frac{1}{2}$ to M and probability $\frac{1}{2}$ to B.

a mixed strategy Nash equilibrium of a strategic game with vNM preferences. In fact, I argue that a strictly dominated action is not a best response to any collection of mixed strategies of the other players.

Suppose that player i's action a_i' is strictly dominated by her mixed strategy α_i. Player i's expected payoff $U_i(\alpha_i, \alpha_{-i})$ when she uses the mixed strategy α_i and the other players use the mixed strategies α_{-i} is a weighted average of her payoffs $U_i(\alpha_i, a_{-i})$ as a_{-i} varies over all the collections of actions for the other players, with the weight on each a_{-i} equal to the probability with which it occurs when the other players' mixed strategies are α_{-i}. Player i's expected payoff when she uses the action a_i' and the other players use the mixed strategies α_{-i} is a similar weighted average; the weights are the same, but the terms take the form $u_i(a_i', a_{-i})$ rather than $U_i(\alpha_i, a_{-i})$. The fact that a_i' is strictly dominated by α_i means that $U_i(\alpha_i, a_{-i}) > u_i(a_i', a_{-i})$ for *every* collection a_{-i} of the other players' actions. Hence player i's expected payoff when she uses the mixed strategy α_i exceeds her expected payoff when she uses the action a_i', given α_{-i}. Consequently,

> *a strictly dominated action is not used with positive probability in any mixed strategy Nash equilibrium.*

Thus when looking for mixed strategy equilibria we can eliminate from consideration every strictly dominated action.

As before, we can define the notion of weak domination (see Definition 46.1).

▶ DEFINITION 121.1 (*Weak domination*) In a strategic game with vNM preferences, player i's mixed strategy α_i **weakly dominates** her action a_i' if

$$U_i(\alpha_i, a_{-i}) \geq u_i(a_i', a_{-i}) \text{ for every list } a_{-i} \text{ of the other players' actions}$$

and

$$U_i(\alpha_i, a_{-i}) > u_i(a_i', a_{-i}) \text{ for some list } a_{-i} \text{ of the other players' actions,}$$

where u_i is a payoff function whose expected value represents player i's preferences over lotteries and $U_i(\alpha_i, a_{-i})$ is player i's expected payoff under u_i when she uses the mixed strategy α_i and the actions chosen by the other players are given by a_{-i}. We say that the action a_i' is **weakly dominated**.

We saw that a weakly dominated action may be used in a Nash equilibrium (see Figure 47.2). Thus a weakly dominated action may be used with positive probability in a mixed strategy equilibrium, so that we *cannot* eliminate *weakly* dominated actions from consideration when finding mixed strategy equilibria.

? EXERCISE 121.2 (Eliminating dominated actions when finding equilibria) Find all the mixed strategy Nash equilibria of the game in Figure 122.1 by first eliminating any strictly dominated actions and then constructing the players' best response functions.

	L	M	R
T	$2,2$	$0,3$	$1,2$
B	$3,1$	$1,0$	$0,2$

Figure 122.1 The strategic game with vNM preferences in Exercise 121.2.

The fact that a player's strategy in a mixed strategy Nash equilibrium may be weakly dominated raises the question of whether a game necessarily has a mixed strategy Nash equilibrium in which no player's strategy is weakly dominated. The following result (which is not easy to prove) shows that the answer is affirmative for a finite game.

■ PROPOSITION 122.1 (Existence of mixed strategy Nash equilibrium with no weakly dominated strategies in finite games) *Every strategic game with vNM preferences in which each player has finitely many actions has a mixed strategy Nash equilibrium in which no player's strategy is weakly dominated.*

4.5 Pure equilibria when randomization is allowed

The analysis in Section 4.3.3 shows that the mixed strategy Nash equilibria of *BoS* in which each player's strategy is pure correspond precisely to the Nash equilibria of the version of the game (considered in Section 2.3) in which the players are not allowed to randomize. The same is true for a general game: equilibria when the players are not allowed to randomize remain equilibria when they are allowed to randomize, and any pure equilibria that exist when the players are allowed to randomize are equilibria when they are not allowed to randomize.

To establish this claim, let N be a set of players and let A_i, for each player i, be a set of actions. Consider the following two games.

G: the strategic game with ordinal preferences in which the set of players is N, the set of actions of each player i is A_i, and the preferences of each player i are represented by the payoff function u_i

G': the strategic game with vNM preferences in which the set of players is N, the set of actions of each player i is A_i, and the preferences of each player i are represented by the expected value of u_i.

First I argue that any Nash equilibrium of G corresponds to a mixed strategy Nash equilibrium (in which each player's strategy is pure) of G'. Let a^* be a Nash equilibrium of G, and for each player i let α_i^* be the mixed strategy that assigns probability 1 to a_i^*. Since a^* is a Nash equilibrium of G, we know that in G' no player i has an action that yields her a payoff higher than does a_i^* when all the other players adhere to α_{-i}^*. Thus α^* satisfies the two conditions in Proposition 116.2, so that it is a mixed strategy equilibrium of G', establishing the following result.

■ PROPOSITION 122.2 (Pure strategy equilibria survive when randomization is allowed) *Let a^* be a Nash equilibrium of G and for each player i let α_i^* be the mixed strategy* *of player i that assigns probability one to the action a_i^*. Then α^* is a mixed strategy Nash equilibrium of G'.*

Next I argue that any mixed strategy Nash equilibrium of G' in which each player's strategy is pure corresponds to a Nash equilibrium of G. Let α^* be a mixed strategy Nash equilibrium of G' in which every player's mixed strategy is pure; for each player i, denote by a_i^* the action to which α_i assigns probability one. Then no mixed strategy of player i yields her a payoff higher than does α_i^* when the other players' mixed strategies are given by α_{-i}^*. Hence, in particular, no *pure* strategy of player i yields her a payoff higher than does α_i^*. Thus a^* is a Nash equilibrium of G. In words, if a pure strategy is optimal for a player when she is allowed to randomize, then it remains optimal when she is prohibited from randomizing. (More generally, prohibiting a decision-maker from taking an action that is not optimal does not change the set of actions that are optimal.)

■ PROPOSITION 123.1 (Pure strategy equilibria survive when randomization is prohibited) *Let α^* be a mixed strategy Nash equilibrium of G' in which the mixed strategy of each player i assigns probability one to the single action a_i^*. Then a^* is a Nash equilibrium of G.*

4.6 Illustration: expert diagnosis

I seem to confront the following predicament all too frequently. Something about which I am relatively ill informed (my car, my computer, my body) stops working properly. I consult an expert, who makes a diagnosis and recommends an action. I am not sure that the diagnosis is correct—the expert, after all, has an interest in selling her services. I have to decide whether to follow the expert's advice or to try to fix the problem myself, put up with it, or consult another expert.

4.6.1 Model

A simple model that captures the main features of this situation starts with the assumption that there are two types of problem, *major* and *minor*. Denote the fraction of problems that are major by r, and assume that $0 < r < 1$. An expert knows, on seeing a problem, whether it is *major* or *minor*; a consumer knows only the probability r. (The diagnosis is costly neither to the expert nor to the consumer.) An expert may recommend either a major or a minor repair (regardless of the true nature of the problem), and a consumer may either accept the expert's recommendation or seek another remedy. A major repair fixes both a major problem and a minor one.

Assume that a consumer always accepts an expert's advice to obtain a minor repair—there is no reason for her to doubt such a diagnosis—but may either accept or reject advice to obtain a major repair. Further assume that an expert always recommends a major repair for a major problem—a minor repair does not fix a major problem, so there is no point in an expert's recommending one for a major problem—but may recommend either repair for a minor problem. Suppose that an expert obtains the same profit $\pi > 0$ (per unit of time) from selling a minor repair to a consumer with a minor problem as she does from selling a major repair to a consumer with a major problem, but obtains the profit $\pi' > \pi$ from selling a major repair to a consumer with a minor problem. (The rationale is that in the last case the expert does not in fact perform a major repair, at least not in its entirety.) A consumer pays an expert E for a major repair and $I < E$ for a minor one; the cost she effectively bears if she chooses some other remedy is $E' > E$ if her problem is major and $I' > I$ if it is minor. (Perhaps she consults other experts before proceeding, or works on the problem herself, in either case spending valuable time.) I assume throughout that $E > I'$.

Under these assumptions we can model the situation as a strategic game in which the expert has two actions (recommend a minor repair for a minor problem; recommend a major repair for a minor problem), and the consumer has two actions (accept the recommendation of a major repair; reject the recommendation of a major repair). I name the actions as follows.

Expert *Honest* (recommend a minor repair for a minor problem and a major repair for a major problem) and *Dishonest* (recommend a major repair for both types of problem).

Consumer *Accept* (buy whatever repair the expert recommends) and *Reject* (buy a minor repair but seek some other remedy if a major repair is recommended)

Assume that each player's preferences are represented by the player's expected monetary payoff. Then the players' payoffs to the four action pairs are as follows (the strategic game is given in Figure 125.1).

(H, A): With probability r the consumer's problem is major, so she pays E, and with probability $1 - r$ it is minor, so she pays I. Thus her expected payoff is $-rE - (1 - r)I$. The expert's profit is π.

(D, A): The consumer's payoff is $-E$. The consumer's problem is major with probability r, yielding the expert π, and minor with probability $1 - r$, yielding the expert π', so that the expert's expected payoff is $r\pi + (1 - r)\pi'$.

(H, R): The consumer's cost is E' if her problem is major (in which case she rejects the expert's advice to get a major repair) and I if her problem is minor, so that her expected payoff is $-rE' - (1 - r)I$. The expert obtains a payoff only if the consumer's problem is minor, in which case she gets π; thus her expected payoff is $(1 - r)\pi$.

(D, R): The consumer never accepts the expert's advice, and thus obtains the expected payoff $-rE' - (1 - r)I'$. The expert does not get any business, and thus obtains the payoff of 0.

		Consumer	
		Accept (q)	*Reject* ($1-q$)
Expert	*Honest* (p)	$\pi, -rE-(1-r)I$	$(1-r)\pi, -rE'-(1-r)I$
	Dishonest ($1-p$)	$r\pi+(1-r)\pi', -E$	$0, -rE'-(1-r)I'$

Figure 125.1 A game between an expert and a consumer with a problem.

4.6.2 Nash equilibrium

To find the Nash equilibria of the game we can construct the best response functions, as before. Denote by p the probability the expert assigns to H and by q the probability the consumer assigns to A.

Expert's best response function If $q = 0$ (i.e. the consumer chooses R with probability one), then the expert's best response is $p = 1$ (since $(1-r)\pi > 0$). If $q = 1$ (i.e. the consumer chooses A with probability one), then the expert's best response is $p = 0$ (since $\pi' > \pi$, so that $r\pi + (1-r)\pi' > \pi$). For what value of q is the expert indifferent between H and D? Given q, the expert's expected payoff to H is $q\pi + (1-q)(1-r)\pi$ and her expected payoff to D is $q[r\pi + (1-r)\pi']$, so she is indifferent between the two actions if

$$q\pi + (1-q)(1-r)\pi = q[r\pi + (1-r)\pi'].$$

Upon simplification, this yields $q = \pi/\pi'$. We conclude that the expert's best response function takes the form shown in both panels of Figure 126.1.

Consumer's best response function If $p = 0$ (i.e. the expert chooses D with probability one), then the consumer's best response depends on the relative sizes of E and $rE' + (1-r)I'$. If $E < rE' + (1-r)I'$, then the consumer's best response is $q = 1$, whereas if $E > rE' + (1-r)I'$, then her best response is $q = 0$; if $E = rE' + (1-r)I'$, then she is indifferent between R and A.

If $p = 1$ (i.e. the expert chooses H with probability one), then the consumer's best response is $q = 1$ (given $E < E'$).

We conclude that if $E < rE' + (1-r)I'$, then the consumer's best response to every value of p is $q = 1$, as shown in the left panel of Figure 126.1. If $E > rE' + (1-r)I'$, then the consumer is indifferent between A and R if

$$p[rE + (1-r)I] + (1-p)E = p[rE' + (1-r)I] + (1-p)[rE' + (1-r)I'],$$

which reduces to

$$p = \frac{E - [rE' + (1-r)I']}{(1-r)(E-I')}.$$

In this case the consumer's best response function takes the form shown in the right panel of Figure 126.1.

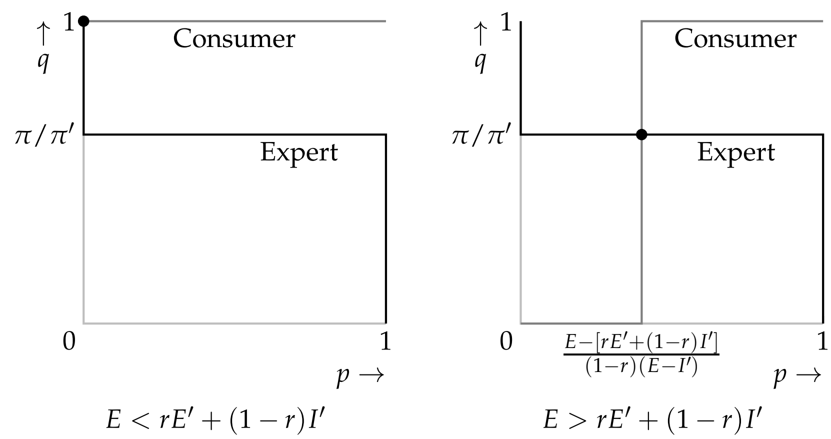

Figure 126.1 The players' best response functions in the game of expert diagnosis. The probability assigned by the expert to H is p and the probability assigned by the consumer to A is q.

Equilibrium Given the best response functions, if $E < rE' + (1-r)I'$, then the pair of pure strategies (D, A) is the unique Nash equilibrium. The condition $E < rE' + (1-r)I'$ says that the cost of a major repair by an expert is less than the *expected* cost of an alternative remedy; the only equilibrium yields the dismal outcome for the consumer in which the expert is always dishonest and the consumer always accepts her advice.

If $E > rE' + (1-r)I'$, then the unique equilibrium of the game is in mixed strategies, with $(p, q) = (p^*, q^*)$, where

$$p^* = \frac{E - [rE' + (1-r)I']}{(1-r)(E-I')} \quad \text{and} \quad q^* = \frac{\pi}{\pi'}.$$

In this equilibrium the expert is sometimes honest, sometimes dishonest, and the consumer sometimes accepts her advice to obtain a major repair and sometimes ignores such advice.

As discussed in the introduction to the chapter, a mixed strategy equilibrium can be given more than one interpretation as a steady state. In the game we are studying, and the games studied earlier in the chapter, I have focused on the interpretation in which each player chooses her action randomly, with probabilities given by her equilibrium mixed strategy, every time she plays the game. In the game of expert diagnosis a different interpretation fits well: among the population of individuals who may play the role of each given player, every individual chooses the same action whenever she plays the game, but different individuals choose different actions; the fraction of individuals who choose each action is equal to the equilibrium probability that that action is used in a mixed strategy equilibrium. Specifically, if $E > rE' + (1-r)I'$, then the fraction p^* of experts is honest (recommending minor repairs for minor problems) and the fraction $1 - p^*$ is dishonest (recommending major repairs for minor problems), while the fraction q^* of consumers is credulous (accepting any recommendation) and the fraction $1 - q^*$ is wary (accepting only a recommendation of a minor repair). Honest and dishonest experts obtain the same expected payoff, as do credulous and wary consumers.

? EXERCISE 127.1 (Equilibrium in the expert diagnosis game) Find the set of mixed strategy Nash equilibria of the game when $E = rE' + (1-r)I'$.

4.6.3 Properties of the mixed strategy Nash equilibrium

Studying how the equilibrium is affected by changes in the parameters of the model helps us understand the nature of the strategic interaction between the players. I consider the effects of three changes.

Suppose that major problems become less common (cars become more reliable, more resources are devoted to preventive health care). If we rearrange the expression for p^* to

$$p^* = 1 - \frac{r(E' - E)}{(1-r)(E - I')},$$

we see that p^* increases as r decreases (the numerator of the fraction decreases and the denominator increases). Thus in a mixed strategy equilibrium, the experts are more honest when major problems are less common. Intuitively, if a major problem is less likely, a consumer has less to lose from ignoring an expert's advice, so the probability of an expert's being honest has to rise for her advice to be heeded. The value of q^* is not affected by the change in r: the probability of a consumer's accepting an expert's advice remains the same when major problems become less common. *Given* the expert's behavior, a decrease in r increases the consumer's payoff to rejecting the expert's advice more than it increases her payoff to accepting this advice, so that she prefers to reject the advice. But this partial analysis is misleading: in the equilibrium that exists after r decreases, the consumer is exactly as likely to accept the expert's advice as she was before the change.

Now suppose that major repairs become less expensive relative to minor ones (technological advances reduce the cost of complex equipment). We see that p^* decreases as E decreases (with E' and I' constant): when major repairs are less costly, experts are less honest. As major repairs become less costly, a consumer has more potentially to lose from ignoring an expert's advice, so that she heeds the advice even if experts are less likely to be honest.

Finally, suppose that the profit π' from an expert's fixing a minor problem with an alleged major repair falls (the government requires experts to return replaced parts to the consumer, making it more difficult for an expert to fraudulently claim to have performed a major repair). Then q^* increases—consumers become less wary. Experts have less to gain from acting dishonestly, so that consumers can be more confident of their advice.

? EXERCISE 127.2 (Incompetent experts) Consider a (realistic?) variant of the model, in which the experts are not entirely competent. Assume that each expert always correctly recognizes a major problem but correctly recognizes a minor problem with probability $s < 1$: with probability $1 - s$ she mistakenly thinks that a minor problem is major, and, if the consumer accepts her advice, performs a major repair and obtains the profit π. Maintain the assumption that each consumer believes (correctly) that the probability her problem is major is r, and assume that a consumer who does not give the job of fixing her problem to an expert bears the cost E' if it is major and I' if it is minor.

Suppose, for example, that an expert is honest and a consumer rejects advice to obtain a major repair. With probability r the consumer's problem is major, so that the expert recommends a major repair, which the consumer rejects; the consumer bears the cost E'. With probability $1 - r$ the consumer's problem is minor. In this case with probability s the expert correctly diagnoses it as minor, and the consumer accepts her advice and pays I; with probability $1 - s$ the expert diagnoses it as major, and the consumer rejects her advice and bears the cost I'. Thus the consumer's expected payoff in this case is $-rE' - (1 - r)[sI + (1 - s)I']$.

Construct the payoffs for every pair of actions and find the mixed strategy equilibrium (equilibria?) when $E > rE' + (1 - r)I'$. Does incompetence breed dishonesty? More wary consumers?

? EXERCISE 128.1 (Choosing a seller) Each of two sellers has available one indivisible unit of a good. Seller 1 posts the price p_1 and seller 2 posts the price p_2. Each of two buyers would like to obtain one unit of the good; they simultaneously decide which seller to approach. If both buyers approach the same seller, each trades with probability $\frac{1}{2}$; the disappointed buyer does not subsequently have the option to trade with the other seller. (This assumption models the risk faced by a buyer that a good is sold out when she patronizes a seller with a low price.) Each buyer's preferences are represented by the expected value of a payoff function that assigns the payoff 0 to not trading and the payoff $1 - p$ to purchasing one unit of the good at the price p. (Neither buyer values more than one unit.) For any pair (p_1, p_2) of prices with $0 \leq p_i \leq 1$ for $i = 1, 2$, find the Nash equilibria (in pure and in mixed strategies) of the strategic game that models this situation. (There are three main cases: $p_2 < 2p_1 - 1$, $2p_1 - 1 < p_2 < \frac{1}{2}(1 + p_1)$, and $p_2 > \frac{1}{2}(1 + p_1)$.)

4.7 Equilibrium in a single population

In Section 2.10 I discussed deterministic steady states in situations in which the members of a single population interact. I now discuss stochastic steady states in such situations.

First extend the definitions of a symmetric strategic game and a symmetric Nash equilibrium (Definitions 51.1 and 52.1) to a game with vNM preferences. Recall that a two-player strategic game with ordinal preferences is symmetric if each player has the same set of actions and each player's evaluation of an outcome depends only on her action and that of her opponent, not on whether she is player 1 or player 2. A symmetric game with vNM preferences satisfies the same conditions; its definition differs from Definition 51.1 only because a player's evaluation of an outcome is given by her expected payoff rather than her ordinal preferences.

	Left	*Right*
Left	1, 1	0, 0
Right	0, 0	1, 1

Figure 129.1 Approaching pedestrians.

▶ DEFINITION 129.1 (*Symmetric two-player strategic game with vNM preferences*) A two-player strategic game with vNM preferences is **symmetric** if the players' sets of actions are the same and the players' preferences are represented by the expected values of payoff functions u_1 and u_2 for which $u_1(a_1, a_2) = u_2(a_2, a_1)$ for every action pair (a_1, a_2).

A Nash equilibrium of a strategic game with ordinal preferences in which every player's set of actions is the same is symmetric if all players take the same action. This notion of equilibrium extends naturally to strategic games with vNM preferences. (As before, it does not depend on the game's having only two players, so I define it for a game with any number of players.)

▶ DEFINITION 129.2 (*Symmetric mixed strategy Nash equilibrium*) A profile α^* of mixed strategies in a strategic game with vNM preferences in which each player has the same set of actions is a **symmetric mixed strategy Nash equilibrium** if it is a mixed strategy Nash equilibrium and α_i^* is the same for every player i.

Now consider again the game of approaching pedestrians (Figure 52.1, reproduced in Figure 129.1), interpreting the payoff numbers as Bernoulli payoffs whose expected values represent the players' preferences over lotteries. We found that this game has two deterministic steady states, corresponding to the two symmetric Nash equilibria in pure strategies, $(Left, Left)$ and $(Right, Right)$. The game also has a symmetric mixed strategy Nash equilibrium, in which each player assigns probability $\frac{1}{2}$ to *Left* and probability $\frac{1}{2}$ to *Right*. This equilibrium corresponds to a steady state in which half of all encounters result in collisions! (Player 1 chooses *Left* and player 2 chooses *Right* with probability $\frac{1}{4}$, and player 1 chooses *Right* and player 2 chooses *Left* with probability $\frac{1}{4}$.)

In this example not only is the game symmetric, but the players' interests coincide. The game in Figure 129.2 is symmetric, but the players prefer to take different actions rather than the same actions. This game has no pure symmetric equilibrium, but has a symmetric mixed strategy equilibrium, in which each player chooses each action with probability $\frac{1}{2}$.

These examples show that a symmetric game may have no symmetric *pure* strategy equilibrium. But both games have a symmetric mixed strategy Nash

	X	*Y*
X	0, 0	1, 1
Y	1, 1	0, 0

Figure 129.2 A symmetric game.

	Stop	*Continue*
Stop	$1, 1$	$1-\epsilon, 2$
Continue	$2, 1-\epsilon$	$0, 0$

Figure 130.1 The game in Exercise 130.2.

equilibrium, as does any symmetric game in which each player has finitely many actions, by the following result (the proof of which requires relatively advanced mathematical tools).

■ PROPOSITION 130.1 (Existence of symmetric mixed strategy Nash equilibrium in symmetric finite games) *Every symmetric strategic game with vNM preferences in which each player's set of actions is finite has a symmetric mixed strategy Nash equilibrium.*

? EXERCISE 130.2 (Approaching cars) Members of a single population of car drivers are randomly matched in pairs when they simultaneously approach intersections from different directions. In each interaction, each driver can either stop or continue. The drivers' preferences are represented by the expected value of the payoff functions given in Figure 130.1; the parameter ϵ, with $0 < \epsilon < 1$, reflects the fact that each driver dislikes being the only one to stop. Find the symmetric Nash equilibrium (equilibria?) of the game (find both the equilibrium strategies and the equilibrium payoffs).

Now suppose that drivers are (re)educated to feel guilty about choosing *Continue*. Assume that their payoffs when choosing *Continue* fall by $\delta > 0$, so that the entry $(2, 1-\epsilon)$ in Figure 130.1 is replaced by $(2-\delta, 1-\epsilon)$, the entry $(1-\epsilon, 2)$ is replaced by $(1-\epsilon, 2-\delta)$, and the entry $(0, 0)$ is replaced by $(-\delta, -\delta)$. Show that all drivers are *better off* in the symmetric equilibrium of this game than they are in the symmetric equilibrium of the original game. Why is the society better off if everyone feels guilty about being aggressive? (The equilibrium of this game, like that of the game of expert diagnosis in Section 4.6, may attractively be interpreted as representing a steady state in which some members of the population always choose one action and other members always choose the other action.)

? EXERCISE 130.3 (Bargaining) Pairs of players from a single population bargain over the division of a pie of size 10. The members of a pair simultaneously make demands; the possible demands are the nonnegative *even* integers up to 10. If the demands sum to 10, then each player receives her demand; if the demands sum to less than 10, then each player receives her demand plus half of the pie that remains after both demands have been satisfied; if the demands sum to more than 10, then neither player receives any payoff. Find all the symmetric mixed strategy Nash equilibria in which each player assigns positive probability to at most two demands. (Many situations in which each player assigns positive probability to two actions, say a' and a'', can be ruled out as equilibria because when one player uses such a strategy, some action yields the other player a payoff higher than does one or both of the actions a' and a''.)

4.8 Illustration: reporting a crime

A crime is observed by a group of n people. Each person would like the police to be informed but prefers that someone else make the phone call. Specifically, suppose that each person attaches the value v to the police being informed and bears the cost c if she makes the phone call, where $v > c > 0$. Then the situation is modeled by the following strategic game with vNM preferences.

Players The n people.

Actions Each player's set of actions is $\{Call, Don't\ call\}$.

Preferences Each player's preferences are represented by the expected value of a payoff function that assigns 0 to the profile in which no one calls, $v - c$ to any profile in which she calls, and v to any profile in which at least one person calls, but she does not.

This game is a variant of the one in Exercise 33.1, with $k = 1$. It has n pure Nash equilibria, in each of which exactly one person calls. (If that person switches to not calling, her payoff falls from $v - c$ to 0; if any other person switches to calling, her payoff falls from v to $v - c$.) If the members of the group differ in some respect, then these asymmetric equilibria may be compelling as steady states. For example, the social norm in which the oldest person in the group makes the phone call is stable.

If the members of the group either do not differ significantly or are not aware of any differences among themselves—if they are drawn from a single homogeneous population—then there is no way for them to coordinate, and a symmetric equilibrium, in which every player uses the same strategy, is more compelling.

The game has no symmetric pure Nash equilibrium. (If everyone calls, then any person is better off switching to not calling. If no one calls, then any person is better off switching to calling.)

However, it has a symmetric mixed strategy equilibrium in which each person calls with positive probability less than one. In any such equilibrium, each person's expected payoff to calling is equal to her expected payoff to not calling. Each person's payoff to calling is $v - c$, and her payoff to not calling is 0 if no one else calls and v if at least one other person calls, so the equilibrium condition is

$$v - c = 0 \cdot \Pr\{\text{no one else calls}\} + v \cdot \Pr\{\text{at least one other person calls}\},$$

or

$$v - c = v \cdot (1 - \Pr\{\text{no one else calls}\}),$$

or

$$c/v = \Pr\{\text{no one else calls}\}. \tag{131.1}$$

Denote by p the probability with which each person calls. The probability that no one else calls is the probability that every one of the other $n - 1$ people does not call, namely $(1 - p)^{n-1}$. Thus the equilibrium condition is $c/v = (1 - p)^{n-1}$, or

$$p = 1 - (c/v)^{1/(n-1)}.$$

This number p is between 0 and 1, so we conclude that the game has a unique symmetric mixed strategy equilibrium, in which each person calls with probability $1 - (c/v)^{1/(n-1)}$. That is, there is a steady state in which whenever a person is in a group of n people facing the situation modeled by the game, she calls with probability $1 - (c/v)^{1/(n-1)}$.

How does this equilibrium change as the size of the group increases? We see that as n increases, the probability p that any given person calls decreases. (As n increases, $1/(n-1)$ decreases, so that $(c/v)^{1/(n-1)}$ increases.) What about the probability that *at least* one person calls? Fix any player i. Then the event "no one calls" is the same as the event "i does not call and no one *other than* i calls". Thus

$$\Pr\{\text{no one calls}\} = \Pr\{i \text{ does not call}\} \cdot \Pr\{\text{no one else calls}\}. \qquad (132.1)$$

Now, the probability that any given person calls decreases as n increases, or equivalently the probability that she does not call increases as n increases. Further, from the equilibrium condition (131.1), $\Pr\{\text{no one else calls}\}$ is equal to c/v, *independent of* n. We conclude that the probability that no one calls *increases* as n increases. That is, the larger the group, the *less* likely the police are informed of the crime!

The condition defining a mixed strategy equilibrium is responsible for this result. For any given person to be indifferent between calling and not calling, this condition requires that the probability that no one else calls be independent of the size of the group. Thus each person's probability of not calling is larger in a larger group, and hence, by the laws of probability reflected in (132.1), the probability that no one calls is larger in a larger group.

The result that the larger the group, the less likely any given person calls is not surprising. The result that the larger the group, the less likely at least one person calls is a more subtle implication of the notion of equilibrium. In a larger group no individual is any less concerned that the police should be called, but in a steady state the behavior of the group drives down the chance that the police are notified of the crime.

? EXERCISE 132.2 (Reporting a crime when the witnesses are heterogeneous) Consider a variant of the model studied in this section in which n_1 witnesses incur the cost c_1 to report the crime, and n_2 witnesses incur the cost c_2, where $0 < c_1 < v$, $0 < c_2 < v$, and $n_1 + n_2 = n$. Show that if c_1 and c_2 are sufficiently close, then the game has a mixed strategy Nash equilibrium in which every witness's strategy assigns positive probabilities to both reporting and not reporting.

? EXERCISE 132.3 (Contributing to a public good) Consider an extension of the analysis in this section to the game in Exercise 33.1 for $k \geq 2$. (In this case a player may contribute even though the good is not provided; the player's payoff in this case is $-c$.) Denote by $Q_{n-1,m}(p)$ the probability that exactly m of a group of $n-1$ players contribute when each player contributes with probability p. What condition must be satisfied by $Q_{n-1,k-1}(p)$ in a symmetric mixed strategy equilibrium (in which each player contributes with the same probability)? (When does a player's contribution make a difference to the outcome?) For the case $v = 1$, $n = 4$, $k = 2$, and $c = \frac{3}{8}$, find the equilibria explicitly. (You need to use the fact that $Q_{3,1}(p) = 3p(1-p)^2$, and do a bit of algebra.)

Reporting a crime: social psychology and game theory

Thirty-eight people witnessed the brutal murder of Catherine ("Kitty") Genovese over a period of half an hour in New York City in March 1964. During this period, no one significantly responded to her screams for help; no one even called the police. Journalists, psychiatrists, sociologists, and others subsequently struggled to understand the witnesses' inaction. Some ascribed it to apathy engendered by life in a large city: "Indifference to one's neighbor and his troubles is a conditioned reflex of life in New York as it is in other big cities" (Rosenthal 1964, 81–82).

The event particularly interested social psychologists. It led them to try to understand the circumstances under which a bystander would help someone in trouble. Experiments quickly suggested that, contrary to the popular theory, people—even those living in large cities—are not in general apathetic to others' plights. An experimental subject who is the lone witness of a person in distress is very likely to try to help. But as the size of the group of witnesses increases, there is a decline not only in the probability that any given one of them offers assistance, but also in the probability that at least one of them offers assistance. Social psychologists hypothesize that three factors explain these experimental findings. First, "diffusion of responsibility": the larger the group, the lower the psychological cost of not helping. Second, "audience inhibition": the larger the group, the greater the embarrassment suffered by a helper in case the event turns out to be one in which help is inappropriate (because, for example, it is not in fact an emergency). Third, "social influence": a person infers the appropriateness of helping from others' behavior, so that in a large group everyone else's lack of intervention leads any given person to think intervention is less likely to be appropriate.

In terms of the model in Section 4.8, these three factors raise the expected cost and/or reduce the expected benefit of a person's intervening. They all seem plausible. However, they are not needed to explain the phenomenon: our game-theoretic analysis shows that even if the cost and benefit are *independent* of group size, a decrease in the probability that at least one person intervenes is an implication of equilibrium. This game-theoretic analysis has an advantage over the socio-psychological one: it derives the conclusion from the same principles that underlie all the other models studied so far (oligopoly, auctions, voting, and elections, for example), rather than positing special features of the specific environment in which a group of bystanders may come to the aid of a person in distress.

The critical element missing from the socio-psychological analysis is the notion of an *equilibrium*. Whether any given person intervenes depends on the probability she assigns to some other person's intervening. In an equilibrium each person must be indifferent between intervening and not intervening, and as we have seen this condition leads inexorably to the conclusion that an increase in group size reduces the probability that at least one person intervenes.

4.9 The formation of players' beliefs

In a Nash equilibrium, each player chooses a strategy that maximizes her expected payoff, *knowing* the other players' strategies. So far we have not considered how players may acquire such information. Informally, the idea underlying the previous analysis is that the players have learned each other's strategies from their experience playing the game. In the idealized situation to which the analysis corresponds, for each player in the game there is a large population of individuals who may take the role of that player; in any play of the game, one participant is drawn randomly from each population. In this situation, a new individual who joins a population that is in a steady state (i.e. is using a Nash equilibrium strategy profile) can learn the other players' strategies by observing their actions over many plays of the game. As long as the turnover in players is small enough, existing players' encounters with neophytes (who may use nonequilibrium strategies) will be sufficiently rare that their beliefs about the steady state will not be disturbed, so that a new player's problem is simply to learn the other players' actions.

This analysis leaves open the question of what might happen if new players simultaneously join more than one population in sufficient numbers that they have a significant chance of facing opponents who are themselves new. In particular, can we expect a steady state to be reached when no one has experience playing the game?

4.9.1 Eliminating dominated actions

In some games the players may reasonably be expected to choose their Nash equilibrium actions from an introspective analysis of the game. At an extreme, each player's best action may be independent of the other players' actions, as in the *Prisoner's Dilemma* (Example 14.1). In such a game no player needs to worry about the other players' actions. In a less extreme case, some player's best action may depend on the other players' actions, but the actions the other players will choose may be clear because each of these players has an action that strictly dominates all others. For example, in the game in Figure 135.1, player 2's action R strictly dominates L, so that no matter what player 2 thinks player 1 will do, she should choose R. Consequently, player 1, who can deduce by this argument that player 2 will choose R, may reason that she should choose B. That is, even inexperienced players may be led to the unique Nash equilibrium (B, R) in this game.

This line of argument may be extended. For example, in the game in Figure 135.2, player 1's action T is strictly dominated, so player 1 may reason that

	L	R
T	$1,2$	$0,3$
B	$0,0$	$1,1$

Figure 135.1 A game in which player 2 has a strictly dominant action and player 1 does not.

	L	R
T	$0,2$	$0,0$
M	$2,1$	$1,2$
B	$1,1$	$2,2$

Figure 135.2 A game in which player 1 may reason that she should choose B because player 2 will reason that player 1 will not choose T, so that player 2 will choose R.

player 2 will deduce that player 1 will not choose T. Consequently player 1 may deduce that player 2 will choose R, making B a better action for her than M.

The set of action profiles that remain at the end of such a reasoning process contains all Nash equilibria; for many games (unlike these examples) the set contains many other action profiles as well. In fact, in many games no action profiles are eliminated, because no player has a strictly dominated action. Nevertheless, in some classes of games the process is powerful; its logical consequences are explored in Chapter 12.

4.9.2 Learning

Another approach to the question of how a steady state might be reached assumes that each player starts with an unexplained "prior" belief about the other players' actions, and changes these beliefs—"learns"—in response to information she receives. She may learn, for example, from observing the fortunes of other players like herself, from discussing the game with such players, or from her own experience playing the game. Here I briefly discuss two theories in which the same set of participants repeatedly play a game, each participant changing her beliefs about the others' strategies in response to her observations of their actions.

Best response dynamics A particularly simple theory assumes that in each period after the first, each player believes that the other players will choose the actions they chose in the previous period. In the first period, each player chooses a best response to an arbitrary deterministic belief about the other players' actions. In every subsequent period, each player chooses a best response to the other players' actions in the *previous* period. This process is known as *best response dynamics*. An action profile that remains the same from period to period is a pure Nash equilibrium of the game. Further, a pure Nash equilibrium in which each player's action is her only best response to the other players' actions is an action profile that remains the same from period to period.

In some games the sequence of action profiles generated best response dynamics converges to a pure Nash equilibrium, regardless of the players' initial beliefs. The example of Cournot's duopoly game studied in Section 3.1.3 is such a game. Looking at the best response functions in Figure 59.1, you can convince yourself that from arbitrary initial actions, the players' actions approach the Nash equilibrium (q_1^*, q_2^*).

? EXERCISE 136.1 (Best response dynamics in Cournot's duopoly game) Find the sequence of pairs of outputs chosen by the firms in Cournot's duopoly game under the assumptions of Section 3.1.3 if both firms initially choose 0. (If you know how to solve a first-order difference equation, find a formula for the outputs in each period; if not, find the outputs in the first few periods.)

? EXERCISE 136.2 (Best response dynamics in Bertrand's duopoly game) Consider Bertrand's duopoly game in which the set of possible prices is discrete, under the assumptions of Exercise 67.2. Does the sequences of prices under best response dynamics converge to a Nash equilibrium when both prices initially exceed $c + 1$? What happens when both prices are initially equal to c?

For other games there are initial beliefs for which the sequence of action profiles generated by the process does not converge. In *BoS* (Example 18.2), for example, if player 1 initially believes that player 2 will choose *Stravinsky* and player 2 initially believes that player 1 will choose *Bach*, then the players' choices will subsequently alternate indefinitely between the action pairs (*Bach*, *Stravinsky*) and (*Stravinsky*, *Bach*). This example highlights the limited extent to which a player is assumed to reason in the model, which does not consider the possibility that she cottons on to the fact that her opponent's action is always a best response to her own previous action.

Fictitious play Under best response dynamics, the players' beliefs are continually revealed to be incorrect unless the starting point is a Nash equilibrium: the players' actions change from period to period. Further, each player believes that every other player is using a pure strategy: a player's belief does not admit the possibility that her opponents' actions are realizations of mixed strategies.

Another theory, known as *fictitious play*, assumes that players consider actions in all the previous periods when forming a belief about their opponents' strategies. They treat these actions as realizations of mixed strategies. Consider a two-player game. Each player begins with an arbitrary probabilistic belief about the other player's action. In the first play of the game she chooses a best response to this belief and observes the other player's action, say A. She then changes her belief to one that assigns probability 1 to A; in the second period, she chooses a best response to this belief and observes the other player's action, say B. She then changes her belief to one that assigns probability $\frac{1}{2}$ to both A and B, and chooses a best response to this belief. She continues to change her belief each period; in any period she adopts the belief that her opponent is using a mixed strategy in

	Head	*Tail*
Head	$1, -1$	$-1, \quad 1$
Tail	$-1, \quad 1$	$1, -1$

Figure 137.1 *Matching Pennies.*

which the probability of each action is proportional to the frequency with which her opponent chose that action in the previous periods. (If, for example, in the first six periods player 2 chooses A twice, B three times, and C once, player 1's belief in period 7 assigns probability $\frac{1}{3}$ to A, probability $\frac{1}{2}$ to B, and probability $\frac{1}{6}$ to C.)

In the game *Matching Pennies* (Example 19.1), reproduced in Figure 137.1, this process works as follows. Suppose that player 1 begins with the belief that player 2's action will be *Tail*, and player 2 begins with the belief that player 1's action will be *Head*. Then in period 1 both players choose *Tail*. Thus in period 2 both players believe that their opponent will choose *Tail*, so that player 1 chooses *Tail* and player 2 chooses *Head*. Consequently in period 3, player 1's belief is that player 2 will choose *Head* with probability $\frac{1}{2}$ and *Tail* with probability $\frac{1}{2}$, and player 2's belief is that player 1 will definitely choose *Tail*. Thus in period 3, both *Head* and *Tail* are best responses of player 1 to her belief, so that she may take either action; the unique best response of player 2 is *Head*. The process continues similarly in subsequent periods.

In two-player games like *Matching Pennies*, in which the players' interests are directly opposed, and in any two-player game in which each player has two actions, this process converges to a mixed strategy Nash equilibrium from any initial beliefs. That is, after a sufficiently large number of periods, the frequencies with which each player chooses her actions are close to the frequencies induced by her mixed strategy in the Nash equilibrium. For other games there are initial beliefs for which the process does not converge. (The simplest example is too complicated to present compactly.)

The people involved in an interaction that we model as a game may form beliefs about their opponents' strategies from an analysis of the structure of the players' payoffs, from their observations of their opponents' actions, and from information they obtain from other people involved in similar interactions. The models I have outlined in this section explore the logical implications of two ways in which players may draw inferences from their opponents' actions. Models that assume the players to be more sophisticated may give more insights into the circumstances in which a Nash equilibrium is likely to be attained; this topic is an active area of current research.

4.10 Extension: finding all mixed strategy Nash equilibria

We can find all the mixed strategy Nash equilibria of a two-player game in which each player has two actions by constructing the players' best response functions, as we have seen. In more complicated games, this method is usually not practical.

The following systematic method of finding all mixed strategy Nash equilibria of a game is suggested by the characterization of an equilibrium in Proposition 116.2.

- For each player i, choose a subset S_i of her set A_i of actions.
- Check whether there exists a mixed strategy profile α such that (i) the set of actions to which each strategy α_i assigns positive probability is S_i and (ii) α satisfies the conditions in Proposition 116.2.
- Repeat the analysis for every collection of subsets of the players' sets of actions.

The following example illustrates this method for a two-player game in which each player has two actions.

◆ EXAMPLE 138.1 (Finding all mixed strategy equilibria of a two-player game in which each player has two actions) Consider a two-player game in which each player has two actions. Denote the actions and payoffs as in Figure 139.1. Each player's set of actions has three nonempty subsets: two each consisting of a single action, and one consisting of both actions. Thus there are nine (3×3) pairs of subsets of the players' action sets. For each pair (S_1, S_2), we check if there is a pair (α_1, α_2) of mixed strategies such that each strategy α_i assigns positive probability only to actions in S_i and the conditions in Proposition 116.2 are satisfied.

- Checking the four pairs of subsets in which each player's subset consists of a single action amounts to checking whether any of the four pairs of actions is a pure strategy equilibrium. (For each player, the first condition in Proposition 116.2 is automatically satisfied because there is only one action in each subset.)
- Consider the pair of subsets $\{T, B\}$ for player 1 and $\{L\}$ for player 2. The second condition in Proposition 116.2 is automatically satisfied for player 1, who has no actions to which she assigns probability 0, and the first condition is automatically satisfied for player 2, because she assigns positive probability to only one action. Thus for there to be a mixed strategy equilibrium in which player 1's probability of using T is p we need $u_{11} = u_{21}$ (player 1's payoffs to her two actions must be equal) and

  $$pv_{11} + (1-p)v_{21} \geq pv_{12} + (1-p)v_{22}$$

  (L must be at least as good as R, given player 1's mixed strategy). If $u_{11} \neq u_{21}$, or if there is no probability p satisfying the inequality, then there is no equilibrium of this type. A similar argument applies to the three other pairs of subsets in which one player's subset consists of both her actions and the other player's subset consists of a single action.
- To check whether there is a mixed strategy equilibrium in which the subsets are $\{T, B\}$ for player 1 and $\{L, R\}$ for player 2, we need to find a pair of

	L	R
T	u_{11}, v_{11}	u_{12}, v_{12}
B	u_{21}, v_{21}	u_{22}, v_{22}

Figure 139.1 A two-player strategic game.

mixed strategies that satisfies the first condition in Proposition 116.2 (the second condition is automatically satisfied because both players assign positive probability to both their actions). That is, we need to find probabilities p and q (if any such exist) for which

$$\begin{aligned} qu_{11} + (1-q)u_{12} &= qu_{21} + (1-q)u_{22} \\ pv_{11} + (1-p)v_{21} &= pv_{12} + (1-p)v_{22}. \end{aligned}$$

For example, in *BoS* we find the two pure equilibria when we check pairs of subsets in which each subset consists of a single action, we find no equilibria when we check pairs in which one subset consists of a single action and the other consists of both actions, and we find the mixed strategy equilibrium when we check the pair $(\{B,S\},\{B,S\})$.

? EXERCISE 139.1 (Finding all mixed strategy equilibria of two-player games) Use the method described in Example 138.1 to find all of the mixed strategy equilibria of the games in Figure 114.1.

In a game in which each player has two actions, for any subset of any player's set of actions at most one of the two conditions in Proposition 116.2 is relevant (the first if the subset contains both actions and the second if it contains only one action). When a player has three or more actions and we consider a subset of her set of actions that contains two actions, both conditions are relevant, as the next example illustrates.

◆ EXAMPLE 139.2 (Finding all mixed strategy equilibria of a variant of *BoS*) Consider the variant of *BoS* given in Figure 139.2. First, by inspection we see that the game has two pure strategy Nash equilibria, namely (B, B) and (S, S).

Now consider the possibility of an equilibrium in which player 1's strategy is pure whereas player 2's strategy assigns positive probability to two or more actions. If player 1's strategy is B, then player 2's payoffs to her three actions (2, 0, and 1) are all different, so the first condition in Proposition 116.2 is not satisfied. Thus there is no equilibrium of this type. Similar reasoning rules out an equilibrium in which player 1's strategy is S and player 2's strategy assigns positive

	B	S	X
B	$4,2$	$0,0$	$0,1$
S	$0,0$	$2,4$	$1,3$

Figure 139.2 A variant of the game *BoS*.

probability to more than one action, and also an equilibrium in which player 2's strategy is pure and player 1's strategy assigns positive probability to both of her actions.

Next consider the possibility of an equilibrium in which player 1's strategy assigns positive probability to both her actions and player 2's strategy assigns positive probability to two of her three actions. Denote by p the probability player 1's strategy assigns to B. There are three possibilities for the pair of player 2's actions that have positive probability.

B and S: For the conditions in Proposition 116.2 to be satisfied we need player 2's expected payoff to B to be equal to her expected payoff to S and at least her expected payoff to X. That is, we need

$$2p = 4(1-p) \geq p + 3(1-p).$$

The equation implies that $p = \frac{2}{3}$, which does not satisfy the inequality. (That is, if p is such that B and S yield the same expected payoff, then X yields a higher expected payoff.) Thus there is no equilibrium of this type.

B and X: For the conditions in Proposition 116.2 to be satisfied we need player 2's expected payoff to B to be equal to her expected payoff to X and at least her expected payoff to S. That is, we need

$$2p = p + 3(1-p) \geq 4(1-p).$$

The equation implies that $p = \frac{3}{4}$, which satisfies the inequality. For the first condition in Proposition 116.2 to be satisfied for player 1 we need player 1's expected payoffs to B and S to be equal: $4q = 1 - q$, where q is the probability player 2 assigns to B, or $q = \frac{1}{5}$. Thus the pair of mixed strategies $((\frac{3}{4}, \frac{1}{4}), (\frac{1}{5}, 0, \frac{4}{5}))$ is a mixed strategy equilibrium.

S and X: For every strategy of player 2 that assigns positive probability only to S and X, player 1's expected payoff to S exceeds her expected payoff to B. Thus there is no equilibrium of this sort.

The final possibility is that there is an equilibrium in which player 1's strategy assigns positive probability to both her actions and player 2's strategy assigns positive probability to all three of her actions. Let p be the probability player 1's strategy assigns to B. Then for player 2's expected payoffs to her three actions to be equal we need

$$2p = 4(1-p) = p + 3(1-p).$$

For the first equality we need $p = \frac{2}{3}$, violating the second equality. That is, there is no value of p for which player 2's expected payoffs to her three actions are equal, and thus no equilibrium in which she chooses each action with positive probability.

We conclude that the game has three mixed strategy equilibria: $((1,0),(1,0,0))$ (i.e. the pure strategy equilibrium (B,B)), $((0,1),(0,1,0))$ (i.e. the pure strategy equilibrium (S,S)), and $((\frac{3}{4}, \frac{1}{4}), (\frac{1}{5}, 0, \frac{4}{5}))$.

	L	M	R
T	$2,2$	$0,3$	$1,3$
B	$3,2$	$1,1$	$0,2$

Figure 141.1 The strategic game with vNM preferences in Exercise 141.1.

? EXERCISE 141.1 (Finding all mixed strategy equilibria of a two-player game) Use the method described in Example 139.2 to find all the mixed strategy Nash equilibria of the strategic game in Figure 141.1.

As you can see from the examples, this method has the disadvantage that for games in which each player has several strategies, or in which there are several players, the number of possibilities to examine is huge. Even in a two-player game in which each player has three actions, each player's set of actions has seven nonempty subsets (three each consisting of a single action, three consisting of two actions, and the entire set of actions), so that there are 49 (7×7) possible collections of subsets to check. In a symmetric game, like the one in the next exercise, many cases involve the same argument, reducing the number of distinct cases to be checked.

? EXERCISE 141.2 (*Rock, Paper, Scissors*) Each of two players simultaneously announces either *Rock*, or *Paper*, or *Scissors*. *Paper* beats (wraps) *Rock*, *Rock* beats (blunts) *Scissors*, and *Scissors* beats (cuts) *Paper*. The player who names the winning object receives $1 from her opponent; if both players make the same choice, then no payment is made. Each player's preferences are represented by the expected amount of money she receives. (An example of the variant of Hotelling's model of electoral competition considered in Exercise 75.3 has the same payoff structure. Suppose there are three possible positions, A, B, and C, and three citizens, one of whom prefers A to B to C, one of whom prefers B to C to A, and one of whom prefers C to A to B. Two candidates simultaneously choose positions. If the candidates choose different positions, each citizen votes for the candidate whose position she prefers; if both candidates choose the same position, they tie for first place.)

a. Formulate this situation as a strategic game and find all its mixed strategy equilibria (give both the equilibrium strategies and the equilibrium payoffs).

b. Find all the mixed strategy equilibria of the modified game in which player 1 is prohibited from announcing *Scissors*.

?? EXERCISE 141.3 (Election campaigns) A new political party, A, is challenging an established party, B. The race involves three localities of different sizes. Party A can wage a strong campaign in only one locality; B must commit resources to defend its position in one of the localities, without knowing which locality A has targeted. If A targets district i and B devotes its resources to some other district, then A gains a_i votes at the expense of B; let $a_1 > a_2 > a_3 > 0$. If B devotes

	A	B
A	$1,1,1$	$0,0,0$
B	$0,0,0$	$0,0,0$

A

	A	B
A	$0,0,0$	$0,0,0$
B	$0,0,0$	$4,4,4$

B

Figure 142.1 The three-player game in Exercise 142.1.

resources to the district that A targets, then A gains no votes. Each party's preferences are represented by the expected number of votes it gains. (Perhaps seats in a legislature are allocated proportionally to vote shares.) Formulate this situation as a strategic game and find its mixed strategy equilibria.

Although games with many players cannot in general be conveniently represented in tables like those we use for two-player games, three-player games can be accommodated. We construct one table for each of player 3's actions; player 1 chooses a row, player 2 chooses a column, and player 3 chooses a *table*. The next exercise is an example of such a game.

? EXERCISE 142.1 (A three-player game) Find the mixed strategy Nash equilibria of the three-player game in Figure 142.1, in which each player has two actions.

4.11 Extension: games in which each player has a continuum of actions

In all the games studied so far in this chapter each player has finitely many actions. In Chapter 3 we saw that many situations may conveniently be modeled as games in which each player has a continuum of actions. (For example, in Cournot's model the set of possible outputs for a firm is the set of nonnegative numbers, and in Hotelling's model the set of possible positions for a candidate is the set of nonnegative numbers.) The principles involved in finding mixed strategy equilibria of such games are the same as those involved in finding mixed strategy equilibria of games in which each player has finitely many actions, though the techniques are different.

Proposition 116.2 says that a strategy profile in a game in which each player has finitely many actions is a mixed strategy Nash equilibrium if and only if, for each player, (*a*) every action to which her strategy assigns positive probability yields the same expected payoff, and (*b*) no action yields a higher expected payoff. Now, a mixed strategy of a player who has a continuum of actions is determined by the probabilities it assigns to sets of actions, not by the probabilities it assigns to single actions (all of which may be zero, for example). Thus (*a*) does not fit such a game. However, the following restatement of the result, equivalent to Proposition 116.2 for a game in which each player has finitely many actions, does fit.

■ PROPOSITION 142.2 (Characterization of mixed strategy Nash equilibrium) *A mixed strategy profile α^* in a strategic game with vNM preferences is a mixed strategy Nash equilibrium if and only if, for each player i,*

- *α_i^* assigns probability zero to the set of actions a_i for which the action profile (a_i, α_{-i}^*) yields player i an expected payoff less than her expected payoff to α^**
- *for no action a_i does the action profile (a_i, α_{-i}^*) yield player i an expected payoff greater than her expected payoff to α^*.*

A significant class of games in which each player has a continuum of actions consists of games in which each player's set of actions is a one-dimensional interval of numbers. Consider such a game with two players; let player i's set of actions be the interval from $\underline{a}_i$ to $\overline{a}_i$, for $i = 1, 2$. Identify each player's mixed strategy with a *cumulative* probability distribution on this interval. (See Section 17.6.4 in the appendix on mathematics if you are not familiar with this notion.) That is, the mixed strategy of each player i is a nondecreasing function F_i for which $0 \leq F_i(a_i) \leq 1$ for every action a_i; the number $F_i(a_i)$ is the probability that player i's action is at most a_i.

The form of a mixed strategy Nash equilibrium in such a game may be very complex. Some such games, however, have equilibria of a particularly simple form, in which each player's equilibrium mixed strategy assigns probability zero except in an interval. Specifically, consider a pair (F_1, F_2) of mixed strategies that satisfies the following conditions for $i = 1, 2$.

- There are numbers x_i and y_i such that player i's mixed strategy F_i assigns probability zero except in the interval from x_i to y_i: $F_i(z) = 0$ for $z < x_i$, and $F_i(z) = 1$ for $z \geq y_i$.
- Player i's expected payoff when her action is a_i and the other player uses her mixed strategy F_j takes the form

  $$\begin{cases} = c_i & \text{for } x_i \leq a_i \leq y_i \\ \leq c_i & \text{for } a_i < x_i \text{ and } a_i > y_i, \end{cases}$$

  where c_i is a constant.

(The second condition is illustrated in Figure 144.1.) By Proposition 142.2, such a pair of mixed strategies, if it exists, is a mixed strategy Nash equilibrium of the game, in which player i's expected payoff is c_i, for $i = 1, 2$.

The next example illustrates how a mixed strategy equilibrium of such a game may be found. The example is designed to be very simple; be warned that in most such games an analysis of the equilibria is, at a minimum, somewhat more complex. Further, my analysis is not complete: I merely find an equilibrium, rather than studying all equilibria. (In fact, the game has no other equilibria.)

◆ EXAMPLE 143.1 (All-pay auction) Two people submit sealed bids for an object worth $\$K$ to each of them. Each person's bid may be any nonnegative number up to $\$K$. The winner is the person whose bid is higher; in the event of a tie each person receives half of the object, which she values at $\$K/2$. (Note that this tie-breaking rule differs from the one considered in Section 3.5.) Each person pays her

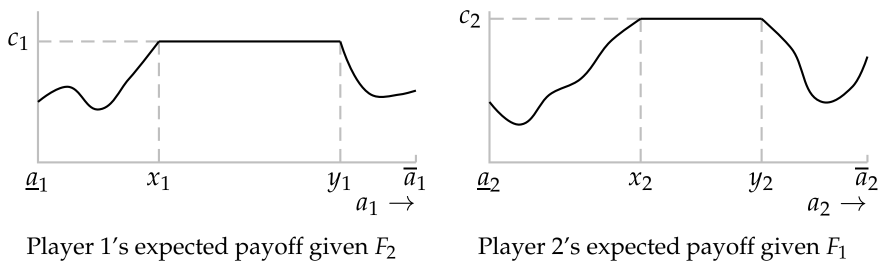

Figure 144.1 If (*i*) F_1 assigns positive probability only to actions in the interval from x_1 to y_1, (*ii*) F_2 assigns positive probability only to actions in the interval from x_2 to y_2, (*iii*) given player 2's mixed strategy F_2, player 1's expected payoff takes the form shown in the left panel, and (*iv*) given player 1's mixed strategy F_1, player 2's expected payoff takes the form shown in the right panel, then (F_1, F_2) is a mixed strategy equilibrium.

bid, *regardless of whether she wins*, and has preferences represented by the expected amount of money she receives.

This situation may be modeled by the following strategic game, known as an **all-pay auction**. (Variants of this game are considered in Exercise 89.1.)

Players The two bidders.

Actions Each player's set of actions is the set of possible bids (nonnegative numbers up to K).

Payoff functions Each player i's preferences are represented by the expected value of the payoff function given by

$$u_i(a_1, a_2) = \begin{cases} -a_i & \text{if } a_i < a_j \\ K/2 - a_i & \text{if } a_i = a_j \\ K - a_i & \text{if } a_i > a_j, \end{cases}$$

where j is the other player.

One situation that may be modeled as such an auction is a lobbying process in which each of two interest groups spends resources to persuade a government to carry out the policy it prefers, and the group that spends the most wins. Another situation that may be modeled as such an auction is the competition between two firms to develop a new product by some deadline, where the firm that spends the most develops a better product, which captures the entire market.

An all-pay auction has no pure strategy Nash equilibrium, by the following argument.

- No pair of actions (x, x) with $x < K$ is a Nash equilibrium because either player can increase her payoff by slightly increasing her bid.
- (K, K) is not a Nash equilibrium because either player can increase her payoff from $-K/2$ to 0 by reducing her bid to 0.
- No pair of actions (a_1, a_2) with $a_1 \neq a_2$ is a Nash equilibrium because the player whose bid is higher can increase her payoff by reducing her bid (and

the player whose bid is lower can, if her bid is positive, increase her payoff by reducing her bid to 0).

Consider the possibility that the game has a mixed strategy Nash equilibrium. Denote by F_i the mixed strategy (i.e. cumulative probability distribution over the interval of possible bids) of player i. I look for an equilibrium in which neither mixed strategy assigns positive probability to any *single* bid. (Remember that there are infinitely many possible bids.) In this case $F_i(a_i)$ is both the probability that player i bids at most a_i and the probability that she bids less than a_i. I further restrict attention to strategy pairs (F_1, F_2) for which, for $i = 1, 2$, there are numbers x_i and y_i such that F_i assigns positive probability only to the interval from x_i to y_i.

To investigate the possibility of such an equilibrium, consider player 1's expected payoff when she uses the action a_1, given player 2's mixed strategy F_2.

- If $a_1 < x_2$, then a_1 is less than player 2's bid with probability one, so that player 1's payoff is $-a_1$.
- If $a_1 > y_2$, then a_1 exceeds player 2's bid with probability one, so that player 1's payoff is $K - a_1$.
- If $x_2 \le a_1 \le y_2$, then player 1's expected payoff is calculated as follows. With probability $F_2(a_1)$ player 2's bid is less than a_1, in which case player 1's payoff is $K - a_1$; with probability $1 - F_2(a_1)$ player 2's bid exceeds a_1, in which case player 1's payoff is $-a_1$; and, by assumption, the probability that player 2's bid is exactly equal to a_1 is zero. Thus player 1's expected payoff is

$$(K - a_1)F_2(a_1) + (-a_1)(1 - F_2(a_1)) = KF_2(a_1) - a_1.$$

We need to find values of x_1 and y_1 and a strategy F_2 such that player 1's expected payoff satisfies the condition illustrated in the left panel of Figure 144.1: it is constant on the interval from x_1 to y_1, and less than this constant for $a_1 < x_1$ and $a_1 > y_1$. The constancy of the payoff on the interval from x_1 to y_1 requires that $KF_2(a_1) - a_1 = c_1$ for $x_1 \le a_1 \le y_1$, for some constant c_1. We also need $F_2(x_2) = 0$ and $F_2(y_2) = 1$ (because I am restricting attention to equilibria in which neither player's strategy assigns positive probability to any single action), and F_2 must be nondecreasing (so that it is a cumulative probability distribution). Analogous conditions must be satisfied by x_2, y_2, and F_1.

We see that if $x_1 = x_2 = 0$, $y_1 = y_2 = K$, and $F_1(z) = F_2(z) = z/K$ for all z with $0 \le z \le K$, then all these conditions are satisfied. Each player's expected payoff is constant, equal to 0 for all her actions.

Thus the game has a mixed strategy Nash equilibrium in which each player randomizes "uniformly" over all her actions. In this equilibrium each player's expected payoff is 0: on average, the amount a player spends is exactly equal to the value of the object. (A more involved argument shows that this equilibrium is the *only* mixed strategy Nash equilibrium of the game.)

? EXERCISE 145.1 (All-pay auction with many bidders) Consider the generalization of the game considered in the previous example in which there are $n \ge 2$ bidders.

Find a mixed strategy Nash equilibrium in which each player uses the same mixed strategy. (If you know how, find each player's mean bid in the equilibrium.)

? EXERCISE 146.1 (Bertrand's duopoly game) Consider Bertrand's oligopoly game (Section 3.2) when there are two firms. Assume that each firm's preferences are represented by its expected profit. Show that if the function $(p-c)D(p)$ is increasing in p, and increases without bound as p increases without bound, then for every $\underline{p} > c$, the game has a mixed strategy Nash equilibrium in which each firm uses the same mixed strategy F, with $F(\underline{p}) = 0$ and $F(p) > 0$ for $p > \underline{p}$.

In the games in the example and exercises, each player's payoff depends only on her action and whether this action is greater than, equal to, or less than the other players' actions. The limited dependence of each player's payoff on the other players' actions makes the calculation of a player's expected payoff straightforward. In many games, each player's payoff is affected more substantially by the other players' actions, making the calculation of expected payoff more complex; more sophisticated mathematical tools are required to analyze such games.

4.12 Appendix: representing preferences by expected payoffs

4.12.1 Expected payoffs

Suppose that a decision-maker has preferences over a set of deterministic outcomes, and that each of her actions results in a *lottery* (probability distribution) over these outcomes. To determine the action she chooses, we need to know her preferences over these lotteries. As argued in Section 4.1.3, we cannot *derive* these preferences from her preferences over deterministic outcomes; rather, we must specify them as part of the model.

So assume that we are given the decision-maker's preferences over lotteries. As in the case of preferences over deterministic outcomes, under some fairly weak assumptions we can represent these preferences by a payoff function. (Refer to Section 1.2.2.) That is, when there are K deterministic outcomes we can find a function, say U, over lotteries such that

$$U(p_1, \ldots, p_K) > U(p'_1, \ldots, p'_K)$$

if and only if the decision-maker prefers the lottery $(p_1, \ldots, p_K)$ to the lottery $(p'_1, \ldots, p'_K)$ (where $(p_1, \ldots, p_K)$ is the lottery in which outcome 1 occurs with probability p_1, outcome 2 occurs with probability p_2, and so on).

For many purposes, however, we need more structure: we cannot get very far without restricting ourselves to preferences for which there is a more specific representation. The standard approach, developed by von Neumann and Morgenstern (1944), is to impose an additional assumption—the "independence axiom"—that allows us to conclude that the decision-maker's preferences can be represented by an *expected payoff function*. More precisely, the independence axiom (which I do not describe) allows us to conclude that there is a payoff function

u over *deterministic* outcomes such that the decision-maker's preference relation over lotteries is represented by the function $U(p_1, \ldots, p_K) = \sum_{k=1}^{K} p_k u(a_k)$, where a_k is the kth outcome of the lottery:

$$\sum_{k=1}^{K} p_k u(a_k) > \sum_{k=1}^{K} p'_k u(a_k) \tag{147.1}$$

if and only if the decision-maker prefers the lottery $(p_1, \ldots, p_K)$ to the lottery $(p'_1, \ldots, p'_K)$. That is, the decision-maker evaluates a lottery by its *expected payoff* according to the function u, which is known as the decision-maker's *Bernoulli payoff function*.

Suppose, for example, that there are three possible deterministic outcomes: the decision-maker may receive \$0, \$1, or \$5, and naturally she prefers \$5 to \$1 to \$0. Suppose that she prefers the lottery $(\frac{1}{2}, 0, \frac{1}{2})$ to the lottery $(0, \frac{3}{4}, \frac{1}{4})$ (where the first number in each list is the probability of \$0, the second number is the probability of \$1, and the third number is the probability of \$5). This preference is consistent with preferences represented by the expected value of a payoff function u for which $u(0) = 0$, $u(1) = 1$, and $u(5) = 4$ because

$$\tfrac{1}{2} \cdot 0 + \tfrac{1}{2} \cdot 4 > \tfrac{3}{4} \cdot 1 + \tfrac{1}{4} \cdot 4.$$

(Many other payoff functions are consistent with a preference for $(\frac{1}{2}, 0, \frac{1}{2})$ over $(0, \frac{3}{4}, \frac{1}{4})$. Among those in which $u(0) = 0$ and $u(5) = 4$, for example, any function for which $u(1) < \frac{4}{3}$ does the job.) Suppose, on the other hand, that the decision-maker prefers the lottery $(0, \frac{3}{4}, \frac{1}{4})$ to the lottery $(\frac{1}{2}, 0, \frac{1}{2})$. This preference is consistent with preferences represented by the expected value of a payoff function u for which $u(0) = 0$, $u(1) = 3$, and $u(5) = 4$ because

$$\tfrac{1}{2} \cdot 0 + \tfrac{1}{2} \cdot 4 < \tfrac{3}{4} \cdot 3 + \tfrac{1}{4} \cdot 4.$$

? EXERCISE 147.2 (Preferences over lotteries) There are three possible outcomes; in the outcome a_i a decision-maker gains $\$a_i$, where $a_1 < a_2 < a_3$. The decision-maker prefers a_3 to a_2 to a_1 and she prefers the lottery $(0.3, 0, 0.7)$ to $(0.1, 0.4, 0.5)$ to $(0.3, 0.2, 0.5)$ to $(0.45, 0, 0.55)$. Is this information consistent with the decision-maker's preferences' being represented by the expected value of a payoff function? If so, find a payoff function consistent with the information. If not, show why not. Answer the same questions when, alternatively, the decision-maker prefers the lottery $(0.4, 0, 0.6)$ to $(0, 0.5, 0.5)$ to $(0.3, 0.2, 0.5)$ to $(0.45, 0, 0.55)$.

Preferences represented by the expected value of a (Bernoulli) payoff function have the great advantage that they are completely specified by that payoff function. Once we know $u(a_k)$ for each possible outcome a_k we know the decision-maker's preferences among all lotteries. This significant advantage does, however, carry with it a small price: it is very easy to confuse a Bernoulli payoff function with a payoff function that represents the decision-maker's preferences over deterministic outcomes.

To describe the relation between the two, suppose that a decision-maker's preferences over lotteries are represented by the expected value of the Bernoulli payoff function u. Then certainly u is a payoff function that represents the decision-maker's preferences over deterministic outcomes (which are special cases of lotteries, in which a single outcome is assigned probability 1). However, the converse is *not* true: if the decision-maker's preferences over deterministic outcomes are represented by the payoff function u (i.e. the decision-maker prefers a to a' if and only if $u(a) > u(a')$), then u is *not* necessarily a Bernoulli payoff function whose expected value represents the decision-maker's preferences over lotteries. For instance, suppose that the decision-maker prefers \$5 to \$1 to \$0 and prefers the lottery $(\frac{1}{2}, 0, \frac{1}{2})$ to the lottery $(0, \frac{3}{4}, \frac{1}{4})$. Then her preferences over deterministic outcomes are consistent with the payoff function u for which $u(0) = 0$, $u(1) = 3$, and $u(5) = 4$. However, her preferences over lotteries are *not* consistent with the expected value of this function (since $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 4 < \frac{3}{4} \cdot 3 + \frac{1}{4} \cdot 4$). The moral is that you should be careful to determine the type of payoff function with which you are dealing.

4.12.2 Equivalent Bernoulli payoff functions

If a decision-maker's preferences in a deterministic environment are represented by the payoff function u, then they are represented also by any payoff function that is an increasing function of u (see Section 1.2.2). The analogous property is not satisfied by Bernoulli payoff functions. Consider the example discussed above. A Bernoulli payoff function u for which $u(0) = 0$, $u(1) = 1$, and $u(5) = 4$ is consistent with a preference for the lottery $(\frac{1}{2}, 0, \frac{1}{2})$ over $(0, \frac{3}{4}, \frac{1}{4})$, but the function v defined by $v(x) = \sqrt{u(x)}$ for all x, for which $v(0) = 0$, $v(1) = 1$, and $v(5) = 2$, is not consistent with such a preference $(\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2 < \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 2)$, though the square root function is increasing (larger numbers have larger square roots).

Under what circumstances do the expected values of two Bernoulli payoff functions represent the same preferences? The next result shows that they do so if and only if one payoff function is an increasing *linear* function of the other.

■ LEMMA 148.1 (Equivalence of Bernoulli payoff functions) *Suppose there are at least three possible outcomes. The expected values of the Bernoulli payoff functions u and v represent the same preferences over lotteries if and only if there exist numbers η and θ with $\theta > 0$ such that $u(x) = \eta + \theta v(x)$ for all x.*

If the expected value of u represents a decision-maker's preferences over lotteries, then so, for example, do the expected values of $2u$, $1 + u$, and $-1 + 4u$; but the expected values of u^2 and of $\sqrt{u}$ do not.

Part of the lemma is easy to establish. Let u be a Bernoulli payoff function whose expected value represents a decision-maker's preferences, and let $v(x) = \eta + \theta u(x)$ for all x, where η and θ are constants with $\theta > 0$. I argue that the expected values of u and of v represent the same preferences. Suppose that the decision-maker prefers the lottery $(p_1, \ldots, p_K)$ to the lottery $(p'_1, \ldots, p'_K)$. Then her expected payoff to $(p_1, \ldots, p_K)$ exceeds her expected payoff to $(p'_1, \ldots, p'_K)$, or, as given in (147.1),

$$\sum_{k=1}^{K} p_k u(a_k) > \sum_{k=1}^{K} p'_k u(a_k). \tag{149.1}$$

Now,

$$\sum_{k=1}^{K} p_k v(a_k) = \sum_{k=1}^{K} p_k \eta + \sum_{k=1}^{K} p_k \theta u(a_k) = \eta + \theta \sum_{k=1}^{K} p_k u(a_k),$$

using the fact that the sum of the probabilities p_k is 1. Similarly,

$$\sum_{k=1}^{K} p'_k v(a_k) = \eta + \theta \sum_{k=1}^{K} p'_k u(a_k).$$

Substituting for u in (149.1) we obtain

$$\frac{1}{\theta}\left(\sum_{k=1}^{K} p_k v(a_k) - \eta\right) > \frac{1}{\theta}\left(\sum_{k=1}^{K} p'_k v(a_k) - \eta\right),$$

which, given $\theta > 0$, is equivalent to

$$\sum_{k=1}^{K} p_k v(a_k) > \sum_{k=1}^{K} p'_k v(a_k):$$

according to v, the expected payoff of $(p_1, \ldots, p_K)$ exceeds the expected payoff of $(p'_1, \ldots, p'_K)$. We conclude that if u represents the decision-maker's preferences, then so does the function v defined by $v(x) = \eta + \theta u(x)$.

I omit the more difficult argument that if the expected values of the Bernoulli payoff functions u and v represent the same preferences over lotteries, then $v(x) = \eta + \theta u(x)$ for some constants η and $\theta > 0$.

? EXERCISE 149.2 (Normalized Bernoulli payoff functions) Suppose that a decision-maker's preferences can be represented by the expected value of the Bernoulli payoff function u. Find a Bernoulli payoff function whose expected value represents the decision-maker's preferences and assigns a payoff of 1 to the best outcome and a payoff of 0 to the worst outcome.

4.12.3 Equivalent strategic games with vNM preferences

Turning to games, consider the three payoff tables in Figure 150.1. All three tables represent the same strategic game with deterministic preferences: in each case, player 1 prefers (B, B) to (S, S) to (B, S), which she regards as indifferent to (S, B), and player 2 prefers (S, S) to (B, B) to (B, S), which she regards as indifferent to (S, B). However, only the left and middle tables represent the same strategic game with vNM preferences. The reason is that the payoff functions in the middle table are linear functions of the payoff functions in the left table, whereas the payoff functions in the right table are not. Specifically, denote the Bernoulli payoff functions of player i in the three games by u_i, v_i, and w_i. Then

$$v_1(a) = 2u_1(a) \text{ and } v_2(a) = -3 + 3u_2(a),$$

	B	S
B	$2,1$	$0,0$
S	$0,0$	$1,2$

	B	S
B	$4,\ \ 0$	$0,-3$
S	$0,-3$	$2,\ \ 3$

	B	S
B	$3,2$	$0,1$
S	$0,1$	$1,4$

Figure 150.1 All three tables represent the same strategic game with ordinal preferences, but only the left and middle games, not the right one, represent the same strategic game with vNM preferences.

so that the left and middle tables represent the same strategic game with vNM preferences. However, w_1 is not a linear function of u_1. If it were, there would exist constants η and $\theta > 0$ such that $w_1(a) = \eta + \theta u_1(a)$ for each action pair a, or equivalently

$$\begin{aligned} 0 &= \eta + \theta \cdot 0 \\ 1 &= \eta + \theta \cdot 1 \\ 3 &= \eta + \theta \cdot 2. \end{aligned}$$

However, these three equations have no solution. Thus the left and right tables represent different strategic games with vNM preferences. (As you can check, w_2 is not a linear function of u_2 either; but for the games not to be equivalent it is sufficient that *one* player's preferences be different.) Another way to see that player 1's vNM preferences in the left and right games are different is to note that in the left table player 1 is indifferent between the certain outcome (S,S) and the lottery in which (B,B) occurs with probability $\frac{1}{2}$ and (S,B) occurs with probability $\frac{1}{2}$ (each yields an expected payoff of 1), whereas in the right table she prefers the latter (since it yields an expected payoff of 1.5).

? EXERCISE 150.1 (Games equivalent to the *Prisoner's Dilemma*) Which of the tables in Figure 150.2 represents the same strategic game with vNM preferences as the *Prisoner's Dilemma* as specified in the left panel of Figure 107.1, when the numbers are interpreted as Bernoulli payoffs?

	C	D
C	$3,3$	$0,4$
D	$4,0$	$2,2$

	C	D
C	$6,\ \ 0$	$0,\ \ 2$
D	$9,-4$	$3,-2$

Figure 150.2 The payoff tables for Exercise 150.1.

Notes

The ideas behind mixed strategies and preferences represented by expected payoffs date back in Western thought at least to the eighteenth century (see Guilbaud (1961) and Kuhn (1968), and Bernoulli (1738), respectively). The modern formulation of a mixed strategy is due to Borel (1921; 1924, 204–221; 1927); the model of the representation of preferences by an expected payoff function is due to von Neumann and Morgenstern (1944, 15–31; 1947, 617–632). The model of a mixed strategy Nash equilibrium and Proposition 119.1 on the existence of a mixed strategy Nash equilibrium in a finite game are due to Nash (1950a, 1951). Proposition 122.1 is an implication of the existence of a "trembling hand perfect equilibrium", due to Selten (1975, Theorem 5).

The example in the box on page 104 is taken from Allais (1953). Conlisk (1989) discusses some of the evidence on the theory of expected payoffs; Machina (1987) and Hey (1997) survey the subject. (The purchasing power of the largest prize in Allais' example was roughly U.S.$6.6 million in 1989 (the date of Conlisk's paper, in which the prize is U.S.$5 million) and roughly U.S.$8 million in 1999.) The model in Section 4.6 is due to Pitchik and Schotter (1987). The model in Section 4.8 is a special case of the one in Palfrey and Rosenthal (1984); the interpretation and analysis that I describe is taken from an unpublished 1984 paper of William F. Samuelson. The box on page 133 draws upon Rosenthal (1964), Latané and Nida (1981), Brown (1986), and Aronson (1995). Best response dynamics were first studied by Cournot (1838, Chapter VII), in the context of his duopoly game. Fictitious play was suggested by Brown (1951). Robinson (1951) shows that the process converges to a mixed strategy Nash equilibrium in any two-player game in which the players' interests are opposed; Shapley (1964, Section 5) exhibits a game outside this class in which the process does not converge. Recent work on learning in games is surveyed by Fudenberg and Levine (1998).

The game in Exercise 118.1 is due to David L. Silverman (see Silverman 1981–82 and Heuer 1995). Exercise 118.2 is based on Palfrey and Rosenthal (1983). Exercise 118.3 is taken from Shubik (1982, 226) (who finds only one of the continuum of equilibria of the game).

The model in Exercise 128.1 is taken from Peters (1984). Exercise 130.2 is a variant of an exercise of Moulin (1986b, 167, 185). Exercise 132.3 is based on Palfrey and Rosenthal (1984). The game *Rock, Paper, Scissors* (Exercise 141.2) was first studied by Borel (1924) and von Neumann (1928). Exercise 141.3 is based on Karlin (1959a, 92–94), who attributes the game to an unpublished paper by Dresher.

Exercise 145.1 is based on a result in Baye, Kovenock, and de Vries (1996). The mixed strategy Nash equilibria of Bertrand's model of duopoly (Exercise 146.1) are studied in detail by Baye and Morgan (1999).

The method of finding all mixed strategy equilibrium described in Section 4.10 is computationally very intense in all but the simplest games. Some computationally more efficient methods are implemented in the freely available computer program `GAMBIT`.