PART II

STATIC GAMES OF COMPLETE INFORMATION

3

Preliminaries

The language and tools of analysis that we have developed so far seem to be ideal to depict and analyze a wide variety of decision problems that a rational individual, or an entity with well-defined objectives, could face. The essence of our framework argues that any decision problem is best understood when we set it up in terms of the three elements of which it is made up: the possible actions, the deterministic or probabilistic relationship between actions and outcomes, and the decision maker's preferences over the possible outcomes. We proceeded to argue that a decision maker will choose those actions that are in his best interest.

This framework offers many attractive features: it is precise, well structured, and generally applicable, and most importantly it lends itself to systematic and consistent analysis. It does, however, suffer from one drawback: the world of a decision problem was described as a world in which the outcomes that determine our well-being are consequences of our own actions and some randomness that is beyond our control.

Let's consider for a moment a decision problem that you may be facing now if you are using this text as part of a university course, which you are taking for a grade. It is, I believe, safe to assume that your objective is some combination of learning the material and obtaining a good grade in the course, with higher grades being preferred over lower ones. This objective determines your preferences over outcomes, which are the set of all possible combinations of how much you learned and what grade you obtained. Your set of possible actions is deciding how hard to study, which includes such elements as deciding how many lectures to attend, how carefully to read the text, how hard to work on your problem sets, and how much time to spend preparing for the exams. Hence you are now facing a well-defined decision problem.

To complete the description of your decision problem, I have yet to explain how the outcome of your success is affected by the amount of work you choose to put into your course work. Clearly as an experienced student you know that the harder you study the more you learn, and you are also more likely to succeed on the exams. There is some uncertainty over how hard a given exam will be; that may depend on many random events, such as how you feel on the day of the exam and what mood the professor was in when the exam was written.

Still something seems to be missing. Indeed you must surely know that grades are often set on a curve, so that your grade relies on your success on the exam as an absolute measure of not only how much you got right but also how much the *other students in the class* got right. In other words, if you're having a bad day on an exam, your only hope is that everyone else in your class is having a worse day!

The purpose of this example is to point out that our framework for a decision problem will be inadequate if your outcomes, and as a consequence your well-being, will depend on the choices made by other decision makers. Perhaps we can just treat the other players in this decision problem as part of the randomness of nature: maybe they'll work hard, maybe not, maybe they'll have a bad day, maybe not, and so on. This, however, would not be part of a rational framework, for it would not be sensible for you to treat your fellow players as mere random "noise." Just as you are trying to optimize your decisions, so are they. Each player is trying to guess what others are doing, and how to act accordingly. In essence, you and your peers are engaged in a *strategic environment* in which you have to think hard about what other players are doing in order to decide what is best for you—knowing that the other players are going through the same difficulties.

We therefore need to modify our decision problem framework to help us describe and analyze strategic situations in which players who interact understand their environment, how their actions affect the outcomes that they and their counterparts will face, and how these outcomes are assessed by the other players. It is useful, therefore, to start with the simplest set of situations possible, and the simplest language that will capture these strategic situations, which we refer to as **games.** We will start with **static games of complete information,** which are the most fundamental games, or environments, in which such strategic considerations can be analyzed.

A static game is similar to the very simple decision problems in which a player makes a once-and-for-all decision, after which outcomes are realized. In a static game, a *set of players* independently choose once-and-for-all actions, which in turn cause the realization of an outcome. Thus a static game can be thought of as having two distinct steps:

Step 1: Each player *simultaneously and independently* chooses an action.

By *simultaneously and independently,* we mean a condition broader and more accommodating than players all choosing their actions at the *exact same* moment. We mean that players must take their actions without observing what actions their counterparts take and without interacting with other players to coordinate their actions. For example, imagine that you have to study for your midterm exam *two days before the midterm* because of an athletic event in which you have to participate on the day before the exam. Assume further that I plan on studying *the day before the midterm,* which will be after your studying effort has ended. If I don't know how much you studied, then by choosing my action after you I have no informational advantage over you; it is *as if* we are making our choices simultaneously and independently of each other. This idea will receive considerable attention as we proceed.

Step 2: Conditional on the players' choices of actions, payoffs are distributed to each player.

That is, once the players have all made their choices, these choices will result in a particular outcome, or probabilistic distribution over outcomes. The players have preferences over the outcomes of the game given by some payoff function over outcomes. For example, if we are playing rock-paper-scissors and I draw paper while you simultaneously draw scissors, then the outcome is that you win and I lose, and the payoffs are what winning and losing mean in our context—something tangible, like $0.10, or just the intrinsic joy of winning versus the suffering of losing.

Steps 1 and 2 settle what we mean by *static*. What do we mean by *complete information*? The loose meaning is that *all players understand the environment they are in*—that is, the game they are playing—in every way. This definition is very much related to our assumptions about rational choice in Section 1.2. Recall that when we had a single-person decision problem we argued that the player must know four things: (1) all his possible actions, A; (2) all the possible outcomes, X; (3) exactly how each action affects which outcome will materialize; and (4) what his preferences are over outcomes. How should this be adjusted to fit a game in which many such players interact?

Games of Complete Information A **game of complete information** requires that the following four components be common knowledge among all the players of the game:

1. all the possible actions of all the players,
2. all the possible outcomes,
3. how each combination of actions of all players affects which outcome will materialize, and
4. the preferences of each and every player over outcomes.

This is by no means an innocuous set of assumptions. In fact, as we will discuss later, they are quite demanding and perhaps almost impossible to justify completely for many real-world "games." However, as with rational choice theory, we use these assumptions because they provide structure and, perhaps surprisingly, describe and predict many phenomena quite well.

You may notice that a new term snuck into the description of games of complete information: *common knowledge*. This is a term that we often use loosely: "it's common knowledge that he gives hard exams" or "it's common knowledge that green vegetables are good for your health." It turns out that *what exactly common knowledge means* is by no means common knowledge. To make it clear,

Definition 3.1 An event E is **common knowledge** if (1) everyone knows E, (2) everyone knows that everyone knows E, and so on *ad infinitum*.

On the face of it, this may seem like an innocuous assumption, and indeed it may be in some cases. For example, if you and I are both walking in the rain together, then it is safe to assume that the event "it is raining" is common knowledge between us. However, if we are both sitting in class and the professor says "tomorrow there is an exam," then the event "there is an exam tomorrow" may not be common knowledge. Despite me knowing that I heard him say it, perhaps you were daydreaming at the time, implying that I *cannot be sure* that you heard the statement as well.

Thus requiring common knowledge is not as innocuous as it may seem, but without this assumption it is quite impossible to analyze games within a structured framework. This difficulty arises because we are seeking to depict a situation in which players can engage in *strategic reasoning*. That is, I want to predict how you will make your choice, given *my belief* that you understand the game. Your understanding incorporates *your belief* about my understanding, and so on. Hence common knowledge will assist us dramatically in our ability to perform this kind of reasoning.

3.1 Normal-Form Games with Pure Strategies

Now that we understand the basic ingredients of a static game of complete information, we develop a formal framework to represent it in a parsimonious and general way, which captures the strategic essence of a game. As with the simple decision problem, the players will have actions from which to choose, and the combination of their choices will result in outcomes over which the players have preferences. For now we will restrict attention to players choosing certain (deterministic) actions that together cause certain (deterministic) outcomes. That is, players will not choose actions stochastically, and there will be no "Nature" player who will randomly select outcomes given a combination of actions that the players will choose.

One of the most common ways of representing a game is described in the following definition of the normal-form game:

A **normal-form game** consists of three features:

1. A set of **players.**
2. A set of **actions** for each player.
3. A **set of payoff functions** for each player that give a payoff value to each combination of the players' chosen actions.

This definition is similar to that of the single-person decision problem that we introduced in Chapter 1, but here we incorporate the fact that many players are interacting. Each has a set of possible actions, the combination (profile) of actions that the players choose will result in an outcome, and each has a payoff from the resulting outcome.

We now introduce the commonly used concept of a **strategy.** A strategy is often defined as *a plan of action intended to accomplish a specific goal.* Imagine a candidate in a local election going to meet a group of potential voters at the home of a neighborhood supporter. Before the meeting, our aspiring politician should have a plan of action to deal with the possible questions he will face. We can think of this plan as a list of the form "if they ask me question q_1 then I will respond with answer a_1; if they ask me question q_2 then I will respond with answer a_2; . . . " and so on. A different candidate may, and often will, have a different strategy of this kind.

The concept of a strategy will escort us throughout this book, and for this reason we now give it both formal notation and a definition:

Definition 3.2 A **pure strategy** for player i is a deterministic plan of action. The set of all **pure strategies** for player i is denoted S_i. A **profile of pure strategies** $s = (s_1, s_2, \ldots, s_n)$, $s_i \in S_i$ for all $i = 1, 2, \ldots, n$, describes a particular combination of pure strategies chosen by all n players in the game.

A brief pause to consider the term "pure" is in order. As mentioned earlier, for the time being and until we reach Chapter 6, we restrict our attention to the case in which players choose deterministic actions. This is what we mean by "pure" strategies: you choose a *certain* plan of action. To illustrate this idea, imagine that you have an exam in three hours, and you must decide how long to study for the exam and how long to just relax, knowing that your classmates are facing the same choice. If, say, you measure time in intervals of 15 minutes, then there are a total of 12 time units in the three-hour window. Your set of pure strategies is then $S_i = \{1, 2, \ldots, 12\}$, where each $s_i \in S_i$ determines how many 15-minute units you will spend studying for the exam. For example, if you choose $s_i = 7$ then you will spend 1 hour and 45 minutes studying and 1 hour and 15 minutes relaxing. An alternative to choosing one of your pure strategies would be for you to choose actions *stochastically.* For example, you can take a die and say "I will roll the die and study for as many 15-minute units as the number on the die indicates." This means that you are stochastically (or randomly) choosing between *any one* of the six pure strategies of studying for 15 minutes, 30 minutes, and so on for up to 1 hour and 30 minutes.

You may wonder why anyone would choose randomly among plans of action. As an example, dwell on the following situation. You meet a friend to go to lunch. Your strategy can be to offer the names of two restaurants that you like and then have your friend decide. But what should you do if he says, "You go ahead and choose"? One option is for you to be prepared with a choice. Another is for you to take out a coin and flip it, so that it is *not you* who is choosing; instead you are randomizing between the two choices.[1] For now, we will restrict attention to pure strategies in which such stochastic play is not possible. That said, stochastic choices play a critical role in game theory. We will introduce *stochastic* or *mixed strategies* in Chapter 6 and continue to use them throughout the rest of the book.

To some extent applying the concept of a strategy or a plan of action to a static game of complete information is overkill, because the players choose actions *once and for all* and *simultaneously.* Thus the only set of relevant plans for player i is the set of his possible actions. This change of focus from actions to strategies may therefore seem redundant. That said, focusing on strategies instead of actions will set the stage for games in which there will be relevance to conditioning one's actions on events that unfold over time, as we will see in Chapter 7. Hence what now seems merely semantic will later be quite useful and important. We now formally define a normal-form game as follows.[2]

Definition 3.3 A **normal-form game** includes three components as follows:

1. A finite **set of players,** $N = \{1, 2, \ldots, n\}$.
2. A **collection of sets of pure strategies,** $\{S_1, S_2, \ldots, S_n\}$.
3. A **set of payoff functions,** $\{v_1, v_2, \ldots, v_n\}$, each assigning a payoff value to each combination of chosen strategies, that is, a set of functions $v_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ for each $i \in N$.

This representation is very general, and it will capture many situations in which each of the players $i \in N$ must simultaneously choose a possible strategy $s_i \in S_i$. Recall again that by *simultaneous* we mean the more general construct in which

1. From my experience, once you offer to take out the coin then your friend is very likely to say, "Oh never mind, let's go to x." By taking out the coin you are effectively telling your friend, "If you have a preference for one of the places, now is your last chance to reveal it." This takes away your friend's option of "being nice" by letting you choose since it is not you who is choosing. I always find this strategy amusing since it works so well.

2. Recall that a finite set of elements will be written as $A = \{a, b, c, d\}$, where A is the set and a, b, c, and d are the elements it includes. Writing $a \in A$ means "a is an element of the set A." If we have two sets, A and B, we define the *Cartesian product* of these sets as $A \times B$. If $a \in A$ and $h \in B$ then we can write $(a, h) \in A \times B$. For more on this subject, refer to Section 19.1 of the mathematical appendix.

each player is choosing a strategy without knowing the choices of the other players. After strategies are selected, each player will realize his payoff, given by $v_i(s_1, s_2, \ldots, s_n) \in \mathbb{R}$, where $(s_1, s_2, \ldots, s_n)$ is the strategy profile that was selected by the agents. Thus from now on the normal-form game will be a triple of sets: $\langle N, \{S_i\}_{i=1}^n, \{v_i(\cdot)\}_{i=1}^n \rangle$, where N is the set of players, $\{S_i\}_{i=1}^n$ is the set of all players' strategy sets, and $\{v_i(\cdot)\}_{i=1}^n$ is the set of all players' payoff functions over the strategy profiles of all the players.[3]

3.1.1 Example: The Prisoner's Dilemma

The Prisoner's Dilemma is perhaps the best-known example in game theory, and it often serves as a parable for many different applications in economics and political science. It is a static game of complete information that represents a situation consisting of two individuals (the players) who are suspects in a serious crime, say, armed robbery. The police have evidence of only petty theft, and to nail the suspects for the armed robbery they need testimony from at least one of the suspects.

The police decide to be clever, separating the two suspects at the police station and questioning each in a different room. Each suspect is offered a deal that reduces the sentence he will get if he confesses, or "finks" (F), on his partner in crime. The alternative is for the suspect to say nothing to the investigators, or remain "mum" (M), so that they do not get the incriminating testimony from him. (As the Mafia would put it, the suspect follows the "omertà"—the code of silence.)

The payoff of each suspect is determined as follows: If both choose mum, then both get 2 years in prison because the evidence can support only the charge of petty theft. If, say, player 1 mums while player 2 finks, then player 1 gets 5 years in prison while player 2 gets only 1 year in prison for being the sole cooperator. The reverse outcome occurs if player 1 finks while player 2 mums. Finally, if both fink then both get only 4 years in prison. (There is some reduction of the 5-year sentence because each would blame the other for being the mastermind behind the robbery.)

Because it is reasonable to assume that more time in prison is worse, we use the payoff representation that equates each year in prison with a value of -1. We can now represent this game in its normal form as follows:

Players: $N = \{1, 2\}$.

Strategy sets: $S_i = \{M, F\}$ for $i \in \{1, 2\}$.

Payoffs: Let $v_i(s_1, s_2)$ be the payoff to player i if player 1 chooses s_1 and player 2 chooses s_2. We can then write payoffs as

$$v_1(M, M) = v_2(M, M) = -2$$
$$v_1(F, F) = v_2(F, F) = -4$$
$$v_1(M, F) = v_2(F, M) = -5$$
$$v_1(F, M) = v_2(M, F) = -1.$$

This completes the normal-form representation of the Prisoner's Dilemma. We will soon analyze how rational players would behave if they were faced with this game.

3. $\{S_i\}_{i=1}^n$ is another way of writing $\{S_1, S_2, \ldots, S_n\}$, and similarly for $\{v_i(\cdot)\}_{i=1}^n$.

3.1.2 Example: Cournot Duopoly

A variant of this example was first introduced by Augustin Cournot (1838). Two identical firms, players 1 and 2, produce some good. Assume that there are no fixed costs of production, and let the variable cost to each firm i of producing quantity $q_i \geq 0$ be given by the cost function, $c_i(q_i) = q_i^2$ for $i \in \{1, 2\}$. Demand is given by the function $q = 100 - p$, where $q = q_1 + q_2$. Cournot starts with the benchmark of firms that operate in a competitive environment in which each firm takes the market price, p, as given, and believes that its behavior cannot influence the market price. Under this assumption, as every economist knows, the solution will be the competitive equilibrium in which each firm produces at a point at which price equals marginal costs, so that the profits on the marginally produced unit are zero. In this particular case, each firm would produce $q_i = 25$, the price would be $p = 50$, and each firm would make 625 in profits.[4]

Cournot then argues that this competitive equilibrium is naive because rational firms should understand that the price is not given, but rather determined by their actions. For example, if firm 1 realizes its effect on the market price, and produces $q_1 = 24$ instead of $q_1 = 25$, then the price will have to increase to $p(49) = 51$ for demand to equal supply because total supply will drop from 50 to 49. The profits of firm 1 will now be $v_1 = 51 \times 24 - 24^2 = 648 > 625$. Of course, if firm 1 realizes that it has such an effect on price, it should not just set $q_1 = 24$ but instead look for the best choice it can make. However, its best choice depends on the quantity that firm 2 will produce—what will that be? Clearly firm 2 should be as sophisticated, and thus we will have to find a solution that considers both *the actions and the counteractions* of these rational and sophisticated firms.

For now, however, let's focus on the representation of the normal form of the game proposed by Cournot. The actions are choices of quantity, and the payoffs are the profits. Hence the following represents the normal form:

Players: $N = \{1, 2\}$.

Strategy sets: $S_i = [0, \infty]$ for $i \in \{1, 2\}$ and firms choose quantities $s_i \in S_i$.

Payoffs: For $i, j \in \{1, 2\}, \quad i \neq j$,

$$v_i(s_i, s_j) = \begin{cases} (100 - s_i - s_j)s_i - s_i^2 & \text{if } s_i + s_j < 100 \\ -s_i^2 & \text{if } s_i + s_j \geq 100. \end{cases}$$

Notice that the payoff function is a little tricky because it has to be well defined for *any* pair of strategies (quantities) that the players choose. We are implicitly assuming that prices cannot fall below zero, so that if the firms together produce a quantity that is greater than 100, the price will be zero (because $p = 100 - s_1 - s_2$) and each firm's payoffs are its costs.

3.1.3 Example: Voting on a New Agenda

Consider three players on a committee who have to vote on whether to remain at the status quo (whatever it is) or adopt a new policy. For example, they could be three

4. Those who have taken a course in microeconomics know that the marginal cost is the derivative of the cost function and hence is equal to $2q_i$. Equating this to the price gives us each firm's supply function, $2q_i = p$ or $q_i = \frac{p}{2}$, and adding up the two supply functions yields the market supply, $q = p$. Equating this to demand yields $p = 100 - p$, resulting in the competitive price of $p = 50$, and plugging this into the supply function yields $q_i = 25$ for $i = 1, 2$.

housemates who currently have an agreement under which they clean the house once every two weeks (the status quo) and they are considering cleaning it every week (the new policy). They could also be the members of the board of a firm who have to vote on changing the CEO's compensation, or they could be a committee of legislators who must vote on whether to adopt new regulations.

Each can vote "yes" (Y), "no" (N), or "abstain" (A). We can set the payoff from the status quo to be 0 for each player. Players 1 and 2 prefer the new policy, so let their payoff value for it be 1, while player 3 dislikes the new policy, so let his payoff from it be -1. Assume that the choice is made by majority voting as follows: if there is no majority of Y over N then the status quo prevails; otherwise the majority is decisive.

We can represent this game in normal form as follows:

Players: $N = \{1, 2, 3\}$.

Strategy sets: $S_i = \{Y, N, A\}$ for $i \in \{1, 2, 3\}$.

Payoffs: Let P denote the set of strategy profiles for which the new agenda is chosen (at least two "yes" votes), and let Q denote the set of strategy profiles for which the status quo remains (no more than one "yes" vote). Formally,

$$P = \left\{ \begin{array}{ll} (Y, Y, N), & (Y, N, Y), \\ (Y, Y, A), & (Y, A, Y), \\ (Y, A, A), & (A, Y, A), \\ (Y, Y, Y), & (N, Y, Y), \\ (A, Y, Y), & (A, A, Y) \end{array} \right\} \quad \text{and}$$

$$Q = \left\{ \begin{array}{llll} (N, N, N), & (N, N, Y), & (N, Y, N), & (Y, N, N), \\ (A, A, A), & (A, A, N), & (A, N, A), & (N, A, A), \\ (A, Y, N), & (A, N, Y), & (N, A, Y), & (Y, A, N), \\ (N, Y, A), & (Y, N, A), & (N, N, A), & (N, A, N), \\ (A, N, N) & & & \end{array} \right\}.$$

Then payoffs can be written as

$$v_i(s_1, s_2, s_3) = \begin{cases} 1 & \text{if } (s_1, s_2, s_3) \in P \\ 0 & \text{if } (s_1, s_2, s_3) \in Q \end{cases} \quad \text{for } i \in \{1, 2\},$$

$$v_3(s_1, s_2, s_3) = \begin{cases} -1 & \text{if } (s_1, s_2, s_3) \in P \\ 0 & \text{if } (s_1, s_2, s_3) \in Q. \end{cases}$$

This completes the normal-form representation of the voting game.

3.2 Matrix Representation: Two-Player Finite Game

As the voting game demonstrates, games that are easy to describe verbally can sometimes be tedious to describe formally. The value of a formal representation is clarity, because it forces us to specify who the players are, what they can do, and how their actions affect each and every player. We could take some shortcuts to make our life easier, and sometimes we will, but such convenience can come at the cost of misspecifying the game. It turns out that for two-person games in which each player has a finite number of strategies, there is a convenient representation that is easy to read.

In many cases, players may be constrained to choose one of a finite number of actions. This is the case for the Prisoner's Dilemma, rock-paper-scissors, the voting game described previously, and many more strategic situations. In fact, even when players have infinitely many actions to choose from, we may be able to provide a good approximation by restricting attention to a finite number of actions. If we think of the Cournot duopoly example, then for any product that comes in well-defined units (a car, a computer, or a shirt), we can safely assume that we are limited to integer units (an assumption that reduces the strategy set to the natural numbers—after all, fractional shirts will not sell very well). Furthermore, the demand function $p = 100 - q$ suggests that flooding the market with more than 100 units will cause the price of the product to drop to zero. This means that we have effectively restricted the strategy set to a finite number of strategies (101, to be accurate, for the quantities $0, 1, \ldots, 100$).

Being able to distinguish games with finite action sets is useful, so we define a finite game as follows:

Definition 3.4 A **finite game** is a game with a finite number of players, in which the number of strategies in S_i is finite for all players $i \in N$.

As it turns out, any two-player finite game can be represented by a matrix that will capture all the relevant information of the normal-form game. This is done as follows:

Rows Each row represents one of player 1's strategies. If there are k strategies in S_1 then the matrix will have k rows.

Columns Each column represents one of player 2's strategies. If there are m strategies in S_2 then the matrix will have m columns.

Matrix entries Each entry in this matrix contains a two-element vector (v_1, v_2), where v_i is player i's payoff when the actions of both players correspond to the row and column of that entry.

As the following examples show, this is a much simpler way of representing a two-player finite game because all the information will appear in a concise and clear way. Note, however, that neither the Cournot duopoly nor the voting example described earlier can be represented by a matrix. The Cournot duopoly is not a finite game (there are an infinite number of actions for each player), and the voting game has more than two players.[5]

It will be useful to illustrate this with two familiar examples.

3.2.1 Example: The Prisoner's Dilemma

Recall that in the Prisoner's Dilemma each player had two actions, M (mum) and F (fink). Therefore, our matrix will have two rows (for player 1) and two columns (for player 2). Using the payoffs for the prisoner's dilemma given in the example above, the matrix representation of the Prisoner's Dilemma is

5. We can represent the voting game using three 3×3 matrices: the rows of each matrix represent the actions of player 1, the columns those of player 2, and each matrix corresponds to an action of player 3. However, the convenient features of two-player matrix games are harder to use for three-player, multiple-matrix representations—not to mention the rather cumbersome structure of multiple matrices.

		Player 2: M	Player 2: F
Player 1	M	$-2, -2$	$-5, -1$
	F	$-1, -5$	$-4, -4$

Notice that all the relevant information appears in this matrix.

3.2.2 Example: Rock-Paper-Scissors

Consider the famous child's game rock-paper-scissors. Recall that rock (R) beats scissors (S), scissors beats paper (P), and paper beats rock. Let the winner's payoff be 1 and the loser's be -1, and let the payoff for each player from a tie (i.e., they both choose the same action) be 0. This is a game with two players, $N = \{1, 2\}$, and three strategies for each player, $S_i = \{R, P, S\}$. Given the payoffs already described, we can write the matrix representation of this game as follows:

		Player 2: R	Player 2: P	Player 2: S
	R	$0, 0$	$-1, 1$	$1, -1$
Player 1	P	$1, -1$	$0, 0$	$-1, 1$
	S	$-1, 1$	$1, -1$	$0, 0$

Remark Such a matrix is sometimes referred to as a *bi-matrix.* In a traditional matrix, by definition, each entry corresponding to a row-column combination must be a single number, or element, while here each entry has a vector of two elements—the payoffs for each of the two players. Thus we formally have *two matrices,* one for each player. We will nonetheless adopt the common abuse of terminology and call this a matrix.

3.3 Solution Concepts

We have focused our attention on how to describe a game formally and fit it into a well-defined structure. This approach, of course, adds value only if we can use the structure to provide some analysis of what will or should happen in the game. Ideally we would like to be able to either advise players on how to play or try to predict how players will play. To accomplish this, we need some method to *solve* the game, and in this section we outline some criteria that will be helpful in evaluating potential methods to analyze and solve games.

As an example, consider again the Prisoner's Dilemma and imagine that you are player 1's lawyer, and that you wish to advise him about how to behave. The game may be represented as follows:

		Player 2: M	Player 2: F
Player 1	M	$-2, -2$	$-5, -1$
	F	$-1, -5$	$-4, -4$

Being a thoughtful and rational adviser, you make the following observation for player 1: "If player 2 chooses F, then playing F gives you -4, while playing M gives you -5, so F is better." Player 1 will then bark at you, "My buddy will never squeal on me!" You, however, being a loyal adviser, must coolly reply as follows: "If you're right, and player 2 chooses M, then playing F gives you -1, while playing M gives you -2, so F is still better. In fact, it seems like F is always better!"

Indeed if I were player 2's lawyer, then the same analysis would work for him, and this is the "dilemma": each player is better off playing F regardless of his opponent's actions, but this leads the players to receive payoffs of -4 each, while if they could only agree to both choose M, then they would obtain -2 each. Left to their own devices, and to the advocacy of their lawyers, the players should not be able to resist the temptation to choose F. Even if player 1 believes that player 2 will play M, he is better off choosing F (and vice versa).

Perhaps your intuition steers you to a different conclusion. You might want to say that they are friends, having stolen together for some time now, and therefore that they care for one another. In this case one of our assumptions is incorrect: the payoffs in the matrix may not represent their true payoffs, and if taken into consideration, altruism would lead both players to choose M instead of F. For example, to capture the idea of altruism and mutual caring, we can assume that a year in prison for each player is worth -1 to himself and imposes $-\frac{1}{2}$ on the other player's payoff. (You care about your friend, but not as much as you care about yourself.) In this case, if player 1 chooses F and player 2 chooses M then player 1 gets $-3\frac{1}{2}$ $\left(-\frac{1}{2}\right.$ for each of the 5 years player 2 goes to jail, and -1 for player 1's year in jail$\left.\right)$ and player 2 gets $-5\frac{1}{2}$ $\left(-\frac{1}{2}\right.$ for the year player 1 is in jail and -5 for the 5 years he spends in jail$\left.\right)$. The matrix representing the "altruistic" Prisoner's Dilemma is given by the following:

		Player 2: M	Player 2: F
Player 1	M	$-3, -3$	$-5\frac{1}{2}, -3\frac{1}{2}$
	F	$-3\frac{1}{2}, -5\frac{1}{2}$	$-6, -6$

The altruistic game will predict cooperative behavior: regardless of what player 2 does, it is always better for player 1 to play M, and the same holds true for player 2. This shows us that our results will, as they always do, depend crucially on our assumptions.[6] This is another manifestation of the "garbage in, garbage out" caveat—we have to get the game parameters right if we want to learn something from our analysis.

Another classic game is the *Battle of the Sexes,* introduced by R. Duncan Luce and Howard Raiffa (1957) in their seminal book *Games and Decisions.* The story goes as follows. Alex and Chris are a couple, and they need to choose where to meet this evening. The catch is that the choice needs to be made while each is at work, and they have no means of communicating. (There were no cell phones or email in 1957, and even landline phones were not in abundance.) Both players prefer being together over not being together, but Alex prefers opera (O) to football (F), while Chris prefers the opposite. This implies that for each player being together at the venue of choice is

6. Another change in assumptions might be that player 2's brother is a psychopath. If player 1 finks, then player 2's brother will kill him, giving player 1 a utility of, say, $-\infty$ from choosing to fink.

better than being together at the other place, and this in turn is better than being alone. Using the payoffs of 2,1 and 0 to represent this order, the game is summarized in the following matrix:

		Chris	
		O	F
Alex	O	2, 1	0, 0
	F	0, 0	1, 2

What can you recommend to each player now? Unlike the situation in the Prisoner's Dilemma, the best action for Alex depends on what Chris will do and vice versa. If we want to predict or prescribe actions for this game, we need to make assumptions about the *behavior* and the *beliefs* of the players. We therefore need a **solution concept** that will result in predictions or prescriptions.

A solution concept is a method of analyzing games with the objective of restricting the set of *all possible outcomes* to those that are *more reasonable than others*. That is, we will consider some reasonable and consistent assumptions about the behavior and beliefs of players that will divide the space of outcomes into "more likely" and "less likely." Furthermore, we would like our solution concept to apply to a large set of games so that it is widely applicable.

Consider, for example, the solution concept that prescribes that each player choose the action that is always best, regardless of what his opponents will choose. As we saw earlier in the Prisoner's Dilemma, playing F is always better than playing M. Hence this solution concept will predict that in this game both players will choose F. For the Battle of the Sexes, however, there is *no strategy* that is always best: playing F is best if your opponent plays F, and playing O is best if your opponent plays O. Hence for the Battle of the Sexes, this simple solution concept is not useful and offers no guidance.

We will use the term **equilibrium** for any one of the strategy profiles that emerges as one of the solution concept's predictions. We will often think of equilibria as the *actual predictions* of our theory. A more forgiving meaning would be that equilibria are the *likely predictions,* because our theory will often not account for all that is going on. Furthermore, in some cases we will see that more than one equilibrium prediction is possible for the same game. In fact, this will sometimes be a strength, and not a weakness, of the theory.

3.3.1 Assumptions and Setup

To set up the background for equilibrium analysis, it is useful to revisit the assumptions that we will be making throughout:

1. **Players are "rational":** A *rational* player is one who chooses his action, $s_i \in S_i$, to maximize his payoff consistent with his beliefs about what is going on in the game.
2. **Players are "intelligent":** An *intelligent* player knows everything about the game: the actions, the outcomes, and the preferences of all the players.
3. **Common knowledge:** The fact that players are rational and intelligent is common knowledge among the players of the game.

To these three assumptions, which we discussed briefly at the beginning of this chapter, we add a fourth, which constrains the set of outcomes that are reasonable:

4. **Self-enforcement:** Any prediction (or equilibrium) of a solution concept must be *self-enforcing.*

The requirement that any equilibrium must be self-enforcing is at the core of our analysis and at the heart of **noncooperative game theory.** We will assume throughout this book that the players engage in noncooperative behavior in the following sense: each player is in control of his own actions, and he will stick to an action only if he finds it to be in his best interest. That is, if a profile of strategies is to be an equilibrium, we will require each player to be happy with his own choice given how the others make their own choices. As you can probably figure out, the profile (F, F) is self-enforcing in the Prisoner's Dilemma game: each player is happy playing F. Indeed, we will see that this is a very robust outcome in terms of equilibrium analysis.

The requirement of self-enforcing equilibria is a natural one if we take the game to be the complete description of the environment. If there are outside parties that can, through the use of force or sanctions, enforce profiles of strategies, then the game we are using is likely to be an inadequate depiction of the actual environment. In this case we ought to model the third party as a player who has actions (strategies) that describe the enforcement.

3.3.2 Evaluating Solution Concepts

In developing a theory that predicts the behavior of players in games, we must evaluate our theory by how well it does as a *methodological* tool. That is, for our theory to be widely useful, it must describe a method of analysis that applies to a rich set of games, which describe the strategic situations in which we are interested. We will introduce three criteria that will help us evaluate a variety of solution concepts: existence, uniqueness, and invariance.

3.3.2.1 Existence: How Often Does It Apply? A solution concept is valuable insofar as it applies to a wide variety of games, and not just to a small and select family of games. A solution concept should apply generally and should not be developed in an ad hoc way that is specific to a certain situation or game. That is, when we apply our solution concept to different games we require it to result in the *existence* of an equilibrium solution.

For example, consider an ad hoc solution concept that offers the following prediction: "Players always choose the action that they think their opponent will choose." If this is our "theory" of behavior, then it will fail to apply to many—maybe most—strategic situations. In particular when players have different sets of actions (e.g., one chooses a software package and the other a hardware package) then this theory would be unable to predict which outcomes are more likely to emerge as equilibrium outcomes.

Any proposed theory for a solution concept that relies on very specific elements of a game will not be general and will be hard to adapt to a wide variety of strategic situations, making the proposed theory useless beyond the very special situations it was tailored to address. Thus one goal is to have a method that will be general enough to apply to many strategic situations; that is, it will prescribe a solution that will *exist* for most games we can think of.

3.3.2.2 Uniqueness: How Much Does It Restrict Behavior? Just as we require our solution concept to apply broadly, we require that it be meaningful in that it restricts the set of possible outcomes to a smaller set of reasonable outcomes. In fact one might argue that being able to pinpoint a single *unique* outcome as a prediction would be ideal. *Uniqueness* is then an important counterpart to *existence.*

For example, if the proposed solution concept says "anything can happen," then it always exists: regardless of the game we apply this concept to, "anything can happen" will always say that the solution is one of the (sometimes infinite) possible outcomes. Clearly this solution concept is useless. A good solution concept is one that balances existence (so that it works for many games) with uniqueness (so that we can add some intelligent insight into what can possibly happen).

It turns out that the nature of games makes the uniqueness requirement quite hard to meet. The reason, as we will learn to appreciate, lies in the nature of strategic interaction in a noncooperative environment. To foreshadow the reasons behind this observation, notice that a player's best action will often depend on what other players are doing. A consequence is that there will often be several combinations of strategies that will support each other in this way.

3.3.2.3 Invariance: How Sensitive Is It to Small Changes? Aside from existence and uniqueness, a third more subtle criterion is important in qualifying a solution concept as a reasonable one, namely that the solution concept be *invariant* to small changes in the game's structure. However, the term "small changes" needs to be qualified more precisely.

Adding a player to a game, for instance, may not be a small change if that player has actions that can wildly change the outcomes of the game. Thus adding or removing a player cannot innocuously be considered a small change. Similarly if we add or delete strategies from the set of actions that are available to a player, we may hinder his ability to guarantee himself some outcomes, and therefore this too should not be considered a small change to the game. We are left with only one component to fiddle with: the payoff functions of the players. It is reasonable to argue that if the payoffs of a game are modified only slightly, then this is a small change to the game that should not affect the predictions of a "robust" solution concept.

For example, consider the Prisoner's Dilemma. If instead of 5 years in prison, imposing a pain of -5 for the players, it imposed a pain of -5.01 for player 1 and -4.99 for player 2, we should be somewhat discouraged if our solution concept suddenly changed the prediction of what players will or ought to do. Thus *invariance* is a robustness property with which we require a solution concept to comply. In other words, if two games are "close," so that the action sets and players are the same yet the payoffs are only slightly different, then our solution concept should offer predictions that are not wildly different for the two games. Put formally, if for a small enough value $\varepsilon > 0$ we alter the payoffs of every outcome for every player by no more than ε, then the solution concept's prediction should not change.

3.3.3 Evaluating Outcomes

Once we subscribe to any particular solution concept, as social scientists we would like to evaluate the properties of the *solutions,* or predictions, that the solution concept will prescribe. This process will offer insights into what we expect the players of a game to achieve when they are left to their own devices. In turn, these insights can guide us toward possibly changing the environment of the game so as to improve the social outcomes of the players.

We have to be precise about the meaning of "to improve the social outcomes." For example, many people may agree that it would be socially better for the government to take \$10 away from the very rich Bill Gates and give that \$10 to an orphan in Latin America. In fact even Gates himself might have approved of this transfer, especially if the money would have saved the child's life. However, Gates may or may not have *liked* the idea, especially if such government intervention would imply that over time most of his wealth would be dissipated through such transfers.

Economists use a particular criterion for evaluating whether an outcome is *socially undesirable.* An outcome is considered to be socially undesirable if there is a different outcome that would make some people better off *without harming anyone else.* As social scientists we wish to avoid outcomes that are socially undesirable, and we therefore turn to the criterion of *Pareto optimality,* which is in tune with the idea of efficiency or "no waste." That is, we would like all the possible value deriving from a given interaction to be distributed among the players. To put this formally:[7]

Definition 3.5 A strategy profile $s \in S$ **Pareto dominates** strategy profile $s' \in S$ if $v_i(s) \geq v_i(s') \forall\, i \in N$ and $v_i(s) > v_i(s')$ for at least one $i \in N$ (in which case, we will also say that s' is **Pareto dominated** by s). A strategy profile is **Pareto optimal** if it is not Pareto dominated by any other strategy profile.

As social scientists, strategic advisers, or policy makers, we hope that players will act in accordance with the *Pareto criterion* and find ways to coordinate on Pareto-optimal outcomes, or avoid those that are Pareto dominated.[8] However, as we will see time and time again, this result will not be achievable in many games. For example, in the Prisoner's Dilemma we made the case that (F, F) should be considered as a very likely outcome. In fact, as we will argue several times, it is the *only* likely outcome. One can see, however, that it is Pareto dominated by (M, M). (Notice that (M, M) is not the only Pareto-optimal outcome. (M, F) and (F, M) are also Pareto-optimal outcomes because no other profile dominates any of them. Don't confuse Pareto optimality with the best "symmetric" outcome that leaves all players "equally" happy.)

3.4 Summary

- A normal-form game includes a finite set of players, a set of pure strategies for each player, and a payoff function for each player that assigns a payoff value to each combination of chosen strategies.
- Any two-player finite game can be represented by a matrix. Each row represents one of player 1's strategies, each column represents one of player 2's strategies, and each cell in the matrix contains the payoffs for both players.
- A solution concept that proposes predictions of how games will be played should be widely applicable, should restrict the set of possible outcomes to a small set of reasonable outcomes, and should not be too sensitive to small changes in the game.
- Outcomes should be evaluated using the Pareto criterion, yet self-enforcing behavior will dictate the set of reasonable outcomes.

7. The symbol $\forall$ denotes "for all."

8. The criterion is named after the Italian economist Vilfredo Pareto. In general economists and other advocates of rational choice theory view this criterion as noncontroversial. However, this view is not necessarily held by everyone. For example, consider two outcomes: In the first, two players get \$5 each. In the second, player 1 gets \$6 while player 2 gets \$60. The Pareto criterion clearly prefers the second outcome, but some other social criterion with equity considerations may disagree with this ranking.

3.5 Exercises

3.1 **eBay:** Hundreds of millions of people bid on eBay auctions to purchase goods from all over the world. Despite being carried out on line, in spirit these auctions are similar to those that have been conducted for centuries. Is an auction a game? Why or why not?

3.2 **Penalty Kicks:** Imagine a kicker and a goalie who confront each other in a penalty kick that will determine the outcome of a soccer game. The kicker can kick the ball left or right, while the goalie can choose to jump left or right. Because of the speed of the kick, the decisions need to be made simultaneously. If the goalie jumps in the same direction as the kick, then the goalie wins and the kicker loses. If the goalie jumps in the opposite direction of the kick, then the kicker wins and the goalie loses. Model this as a normal-form game and write down the matrix that represents the game you modeled.

3.3 **Meeting Up:** Two old friends plan to meet at a conference in San Francisco, and they agree to meet by "the tower." After arriving in town, each realizes that there are two natural choices: Sutro Tower or Coit Tower. Not having cell phones, each must choose independently which tower to go to. Each player prefers meeting up to not meeting up, and neither cares where this would happen. Model this as a normal-form game and write down the matrix form of the game.

3.4 **Hunting:** Two hunters, players 1 and 2, can each choose to hunt a stag, which provides a rather large and tasty meal, or hunt a hare—also tasty, but much less filling. Hunting stags is challenging and requires mutual cooperation. If either hunts a stag alone, then the stag will get away, while hunting the stag together guarantees that the stag will be caught. Hunting hares is an individualistic enterprise that is not done in pairs, and whoever chooses to hunt a hare will catch one. The payoff from hunting a hare is 1, while the payoff to each from hunting a stag together is 3. The payoff from an unsuccessful stag hunt is 0. Represent this game as a matrix.

3.5 **Matching Pennies:** Players 1 and 2 each put a penny on a table simultaneously. If the two pennies come up the same side (heads or tails) then player 1 gets both pennies; otherwise player 2 gets both pennies. Represent this game as a matrix.

3.6 **Price Competition:** Imagine a market with demand $p(q) = 100 - q$. There are two firms, 1 and 2, and each firm i has to simultaneously choose its price p_i. If $p_i < p_j$, then firm i gets all of the market while no one demands the good of firm j. If the prices are the same then both firms split the market demand equally. Imagine that there are no costs to produce any quantity of the good. (These are two large dairy farms, and the product is manure.) Write down the normal form of this game.

3.7 **Public Good Contribution:** Three players live in a town, and each can choose to contribute to fund a streetlamp. The value of having the streetlamp is 3 for each player, and the value of not having it is 0. The mayor asks each player to contribute either 1 or nothing. If at least two players contribute then the lamp will be erected. If one player or no players contribute then the lamp will not be erected, in which case any person who contributed will not get his money back. Write down the normal form of this game.

4

Rationality and Common Knowledge

In this chapter we study the implications of imposing the assumptions of rationality as well as common knowledge of rationality. We derive and explore some solution concepts that result from these two assumptions and seek to understand the restrictions that each of the two assumptions imposes on the way in which players will play games.

4.1 Dominance in Pure Strategies

It will be useful to begin by introducing some new notation. We denoted the payoff of a player i from a profile of strategies $s = (s_1, s_2, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n)$ as $v_i(s)$. It will soon be very useful to refer specifically to the strategies of a player's opponents in a game. For example, the actions chosen by the players who are *not* player i are denoted by the profile

$$(s_1, s_2, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n) \in S_1 \times S_2 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n.$$

To simplify we will hereafter use a common shorthand notation as follows: We define $S_{-i} \equiv S_1 \times S_2 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n$ as the set of all the strategy sets of *all players who are not* player i. We then define $s_{-i} \in S_{-i}$ as a particular possible profile of strategies for all players who are not i. Hence we can rewrite the payoff of player i from strategy s as $v_i(s_i, s_{-i})$, where $s = (s_i, s_{-i})$.

4.1.1 Dominated Strategies

The Prisoner's Dilemma was easy to analyze: each of the two players has an action that is best *regardless* of what his opponent chooses. Suggesting that each player will choose this action seems natural because it is consistent with the basic concept of rationality. If we assume that the players are rational, then we should expect them to choose whatever they deem to be best for them. If it turns out that a player's best strategy does not depend on the strategies of his opponents then we should be all the more confident that he will choose it.

It is not too often that we will find ourselves in situations in which we have a best action that does not depend on the actions of our opponents. We begin, therefore, with a less demanding concept that follows from rationality. In particular consider the strategy mum in the Prisoner's Dilemma:

		Player 2	
		M	F
Player 1	M	$-2, -2$	$-5, -1$
	F	$-1, -5$	$-4, -4$

As we argued earlier, playing M is worse than playing F for each player *regardless* of what the player's opponent does. What makes it unappealing is that there is another strategy, F, that is better than M regardless of what one's opponent chooses. We say that such a strategy is *dominated*. Formally we have

Definition 4.1 Let $s_i \in S_i$ and $s_i' \in S_i$ be possible strategies for player i. We say that s_i' is **strictly dominated** by s_i if for any possible combination of the other players' strategies, $s_{-i} \in S_{-i}$, player i's payoff from s_i' is strictly less than that from s_i. That is,

$$v_i(s_i, s_{-i}) > v_i(s_i', s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}.$$

We will write $s_i \succ_i s_i'$ to denote that s_i' is strictly dominated by s_i.

Now that we have a precise definition for a dominated strategy, it is straightforward to draw an obvious conclusion:

Claim 4.1 *A rational player will never play a strictly dominated strategy.*

This claim is obvious. If a player plays a dominated strategy then he cannot be playing optimally because, by the definition of a dominated strategy, the player has another strategy that will yield him a higher payoff *regardless of the strategies of his opponents.* Hence knowledge of the game implies that a player should recognize dominated strategies, and rationality implies that these strategies will be avoided.

When we apply the notion of a dominated strategy to the Prisoner's Dilemma we argue that each of the two players has one dominated strategy that he should never use, and hence each player is left with one strategy that is not dominated. Therefore, for the Prisoner's Dilemma, rationality alone is enough to offer a prediction about which outcome will prevail: (F, F) is this outcome.

Many games, however, will not be as special as the Prisoner's Dilemma, and rationality alone will not suggest a clear-cut, unique prediction. As an example, consider the following advertising game. Two competing brands can choose one of three marketing campaigns—low (L), medium (M), and high (H)—with payoffs given by the following matrix:

		Player 2		
		L	M	H
	L	6, 6	2, 8	0, 4
Player 1	M	8, 2	4, 4	1, 3
	H	4, 0	3, 1	2, 2

It is easy to see that each player has one dominated strategy, which is L. However, neither M nor H is dominated. For example, if player 2 plays M then player 1 should also play M, while if player 2 plays H then player 1 should also play H. Hence rationality alone does not offer a unique prediction. It is nonetheless worth spending some time on the extreme cases in which it does.

4.1.2 Dominant Strategy Equilibrium

Because a strictly dominated strategy is one to avoid at all costs,[1] there is a counterpart strategy, represented by F in the Prisoner's Dilemma, that would be desirable. This is a strategy that is *always* the best thing you can do, regardless of what your opponents choose. Formally we have

Definition 4.2 $s_i \in S_i$ is a **strictly dominant strategy** for i if every other strategy of i is strictly dominated by it, that is,

$$v_i(s_i, s_{-i}) > v_i(s_i', s_{-i}) \quad \text{for all } s_i' \in S_i, \quad s_i' \neq s_i, \quad \text{and all } s_{-i} \in S_{-i}.$$

If, as in the Prisoner's Dilemma, every player had such a wonderful dominant strategy, then it would be a very sensible predictor of behavior because it follows from rationality alone. We can introduce this simple idea as our first solution concept:

Definition 4.3 The strategy profile $s^D \in S$ is a **strict dominant strategy equilibrium** if $s_i^D \in S_i$ is a strict dominant strategy for all $i \in N$.

This gives a formal definition for the outcome "both players fink," or (F, F), in the Prisoner's Dilemma: it is a dominant strategy equilibrium. In this equilibrium the payoffs are $(-4, -4)$ for players 1 and 2, respectively.

Caveat Be careful not to make a common error by referring to the pair of payoffs $(-4, -4)$ as the solution. The solution should always be described as the strategies that the players will choose. Strategies are a set of actions by the players, and payoffs are a result of the outcome. When we talk about predictions, or equilibria, we will always refer to *what players do* as the equilibrium, not their payoffs.

Using this solution concept for any game is not that difficult. It basically requires that we identify a strict dominant strategy for each player and then use this profile of strategies to predict or prescribe behavior. If, as in the Prisoner's Dilemma, we are lucky enough to find a dominant strategy equilibrium for other games, then this solution concept has a very appealing property:

Proposition 4.1 *If the game $\Gamma = \langle N, \{S_i\}_{i=1}^n, \{v_i\}_{i=1}^n \rangle$ has a strictly dominant strategy equilibrium s^D, then s^D is the unique dominant strategy equilibrium.*

This proposition is rather easy to prove, and that proof is left as exercise 4.1 at the end of the chapter.

---

1. This is a good point at which to stop and reflect on a very simple yet powerful lesson. In any situation, look first for your dominated strategies and avoid them!

4.1.3 Evaluating Dominant Strategy Equilibrium

Proposition 4.1 is very useful in addressing one of our proposed criteria for evaluating a solution concept: when it exists, the strict-dominance solution concept guarantees uniqueness. However, what do we know about existence? A quick observation will easily convince you that this is a problem.

Consider the Battle of the Sexes game introduced in Section 3.3 and described again in the following matrix:

		Chris	
		O	F
Alex	O	2, 1	0, 0
	F	0, 0	1, 2

Neither player has a dominated strategy, implying that neither has a dominant strategy either. The best strategy for Chris depends on what Alex chooses and vice versa. Thus if we stick to the solution concept of strict dominance we will encounter games, in fact many of them, for which there will be no equilibrium. This unfortunate conclusion implies that the strict-dominance solution concept will often fail to predict the choices that players ought to, or will, choose in many games.

Regarding the invariance criterion, the strict-dominance solution concept does comply. From definition 4.2, s^D is a strictly dominant strategy equilibrium if and only if $v_i(s_i^D, s_{-i}) > v_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and all $s_{-i} \in S_{-i}$. Because the inequality is strict, we can find a small enough value $\varepsilon > 0$ such that if we either add or subtract ε from any payoff $v_i(s_i, s_{-i})$ the inequality will still hold.

We now turn to the Pareto criterion of equilibrium outcomes when a strict-dominance solution exists. The Prisoner's Dilemma has an equilibrium prediction using the strict-dominance solution concept, so we can evaluate the efficiency properties of the unique strictly dominant strategy equilibrium for that game. It is easy to see that the outcome prescribed by this solution is *not Pareto optimal:* both players would be better off if they could each commit to play M, yet left to their own devices they will not do this. Of course, in other games the solution may be Pareto optimal (see, for example, the "altruistic" Prisoner's Dilemma in Section 3.3).

The failure of Pareto optimality is *not* a failure of the solution concept. The assumption that players are rational causes them to fink in the Prisoner's Dilemma if we restrict attention to self-enforcing behavior. The failure of Pareto optimality implies that the players would benefit from modifying the environment in which they find themselves to create other enforcement mechanisms—for example, creating a "mafia" with norms of conduct that enforce implicit agreements so as to punish those who fink.

To explicitly see how this can work, imagine that a mafia member who finks on another member is very seriously reprimanded, which will change the payoff structure of the Prisoner's Dilemma if he is caught. If the pain from mafia punishment is equivalent to z, then we have to subtract z units of payoff for each player who finks. The "mafia-modified" Prisoner's Dilemma is represented by the following matrix:

		Player 2	
		M	F
Player 1	M	$-2, -2$	$-5, -1-z$
	F	$-1-z, -5$	$-4-z, -4-z$

If z is strictly greater than 1 then this punishment will be enough to flip our predicted equilibrium outcome of the game because then M becomes the strict dominant strategy (and (M, M) is Pareto optimal).

This example demonstrates that "institutional design," which changes the game that players play, can be very useful in affecting the well-being of players. By introducing this external enforcer, or institution, we are able to get the players to choose outcomes that make them both better off compared to what they can achieve without the additional institution. Moreover, notice that if the players believe that the mafia will enforce the code of conduct then there is no need to actually enforce it—the players choose not to fink, and the enforcement of punishments need not happen. However, we need to be suspicious of whether such an institution will be *self-enforcing,* that is, whether the mafia will indeed enforce the punishments. And, for it to be self-enforcing, we need to model the behavior of potential punishers and whether they themselves will have the selfish incentives to carry out the enforcement activities. This is something we will explore at length when we consider multistage and repeated games in Chapters Chapters 9 and 10.

Remark A related notion is that of **weak dominance.** We say that s_i' is *weakly dominated* by s_i if, for any possible combination of the other players' strategies, player i's payoff from s_i' is weakly less than that from s_i. That is,

$$v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}.$$

This means that for some $s_{-i} \in S_{-i}$ this weak inequality may hold strictly, while for other $s_{-i}' \in S_{-i}$ it will hold with equality. We define a strategy to be **weakly dominant** in a similar way. This is still useful because if we can find a dominant strategy for a player, be it weak or strict, this seems like the most obvious thing to prescribe. An important difference between weak and strict dominance is that if a weakly dominant equilibrium exists, it need not be unique. To show this is left as exercise 4.2.

4.2 Iterated Elimination of Strictly Dominated Pure Strategies

As we saw in the previous chapter, our requirement that players be rational implied two important conclusions:

1. A rational player will never play a dominated strategy.
2. If a rational player has a strictly dominant strategy then he will play it.

We used this second conclusion to define the solution concept of strict dominance, which is very appealing because, when it exists, it requires only rationality as its driving force. A drawback of the dominant strategy solution concept is, however, that it will often fail to exist. Hence if we wish to develop a predictive theory of behavior in games then we must consider alternative approaches that will apply to a wide variety of games.

4.2.1 Iterated Elimination and Common Knowledge of Rationality

We begin with the premise that players are rational, and we build on the first conclusion in the previous section, which claims that a rational player will never play a dominated strategy. This conclusion is by itself useful in that it rules out what players *will not do*. As a result, we conclude that rationality tells us which strategies will never be played.

Now turn to another important assumption introduced earlier: the structure of the game and the rationality of the players are common knowledge among the players. The introduction of *common knowledge of rationality* allows us to do much more than identify strategies that rational players will avoid. If indeed all the players *know* that each player will never play a strictly dominated strategy, they can effectively ignore those strictly dominated strategies that their opponents will never play, and their opponents can do the same thing. If the original game has some players with some strictly dominated strategies, then all the players know that *effectively* they are facing a "smaller" restricted game with fewer total strategies.

This logic can be taken further. Because it is *common knowledge* that all players are rational, then everyone knows that everyone knows that the game is effectively a smaller game. In this smaller restricted game, everyone knows that players will not play strictly dominated strategies. In fact we may indeed find additional strategies that are dominated in the restricted game that *were not dominated* in the original game. Because it is common knowledge that players will perform this kind of reasoning again, the process can continue until no more strategies can be eliminated in this way.

To see this idea more concretely, consider the following two-player finite game:

		Player 2		
		L	C	R
	U	4, 3	5, 1	6, 2
Player 1	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

(4.1)

A quick observation reveals that there is no strictly dominant strategy, neither for player 1 nor for player 2. Also note that there is no strictly dominated strategy for player 1. There is, however, a strictly dominated strategy for player 2: the strategy C is strictly dominated by R because $2 > 1$ (row U), $6 > 4$ (row M), and $8 > 6$ (row D). Thus, because this is common knowledge, both players know that we can effectively eliminate the strategy C from player 2's strategy set, which results in the following reduced game:

		Player 2	
		L	R
	U	4, 3	6, 2
Player 1	M	2, 1	3, 6
	D	3, 0	2, 8

In this reduced game, both M and D are strictly dominated by U for player 1, allowing us to perform a second round of eliminating strategies, this time for player 1. Eliminating these two strategies yields the following trivial game:

		Player 2	
		L	R
Player 1	U	**4, 3**	6, 2

in which player 2 has a strictly dominated strategy, playing R. Thus for this example the iterated process of eliminating dominated strategies yields a unique prediction: the strategy profile we expect these players to play is (U, L), giving the players the payoffs of (4, 3).

As the example demonstrates, this process of **iterated elimination of strictly dominated strategies** (IESDS) builds on the assumption of common knowledge of rationality. The first step of iterated elimination is a consequence of player 2's rationality; the second stage follows because players *know* that players are rational; the third stage follows because players *know that players know* that they are rational, and this ends in a unique prediction.

More generally we can apply this process to games in the following way. Let S_i^k denote the strategy set of player i that survives k rounds of IESDS. We begin the process by defining $S_i^0 = S_i$ for each i, the original strategy set of player i in the game.

Step 1: Define $S_i^0 = S_i$ for each i, the original strategy set of player i in the game, and set $k = 0$.

Step 2: Are there players for whom there are strategies $s_i \in S_i^k$ that are strictly dominated? If yes, go to step 3. If not, go to step 4.

Step 3: For all the players $i \in N$, remove any strategies $s_i \in S_i^k$ that are strictly dominated. Set $k = k + 1$, and define a new game with strategy sets S_i^k that *do not include* the strictly dominated strategies that have been removed. Go back to step 2.

Step 4: The remaining strategies in S_i^k are reasonable predictions for behavior.

In this chapter we refrain from giving a precise mathematical definition of the process because this requires us to consider richer behavior by the players, in particular, allowing them to choose randomly between their different pure strategies. We will revisit this approach briefly when such stochastic play, or **mixed strategies,** is introduced later.[2]

Using the process of IESDS we can define a new solution concept:

Definition 4.4 We will call any strategy profile $s^{ES} = (s_1^{ES}, \ldots, s_n^{ES})$ that survives the process of IESDS an **iterated-elimination equilibrium.**

Like the concept of a strictly dominant strategy equilibrium, the iterated-elimination equilibrium starts with the premise of rationality. However, in addition to rationality, IESDS requires a lot more: common knowledge of rationality. We will discuss the implications of this requirement later in this chapter.

4.2.2 Example: Cournot Duopoly

Recall the Cournot duopoly example we introduced in Section 3.1.2, but consider instead a simpler example of this problem in which the firms have linear rather than quadratic costs: the cost for each firm for producing quantity q_i is given by $c_i(q_i) = 10q_i$ for $i \in \{1, 2\}$. (Using economics jargon, this is a case of constant marginal cost

2. Just to satisfy your curiosity, think of the Battle of the Sexes, and imagine that Chris can pull out a coin and flip between the decision of opera or football. This by itself introduces a "new" strategy, and we will exploit such strategies to develop a formal definition of IESDS.

equal to 10 and no fixed costs.) Let the demand be given by $p(q) = 100 - q$, where $q = q_1 + q_2$.

Consider first the profit (payoff) function of firm 1:

$$v_1(q_1, q_2) = \overbrace{(100 - q_1 - q_2)q_1}^{\text{Revenue}} - \overbrace{10q_1}^{\text{Costs}}$$
$$= 90q_1 - q_1^2 - q_1q_2.$$

What should firm 1 do? If it knew what quantity firm 2 will choose to produce, say some value of q_2, then the profits of firm 1 would be maximized when the first-order condition is satisfied, that is, when $90 - 2q_1 - q_2 = 0$. Thus, for any given value of q_2, firm 1 maximizes its profits when it sets its own quantity according to the function

$$q_1 = \frac{90 - q_2}{2}. \tag{4.2}$$

Though it is true that the choice of firm 1 depends on what it believes firm 2 is choosing, equation (4.2) implies that firm 1 will *never* choose to produce more than $q_1 = 45$. This follows from the simple observation that q_2 is never negative, in which case equation (4.2) implies that $q_1 \le 45$. In fact, this is equivalent to showing that any quantity $q_1 > 45$ is strictly dominated by $q_1 = 45$. To see this, for any q_2, the profits from setting $q_1 = 45$ are given by

$$v_1(45, q_2) = (100 - 45 - q_2)45 - 450 = 2025 - 45q_2.$$

The profits from choosing any other q_1 are given by

$$v_1(q_1, q_2) = (100 - q_1 - q_2)q_1 - 10q_1 = 90q_1 - q_1q_2 - q_1^2.$$

Thus we can subtract $v_1(q_1, q_2)$ from $v_1(45, q_2)$ and obtain

$$v_1(45, q_2) - v_1(q_1, q_2) = 2025 - 45q_2 - (90q_1 - q_1q_2 - q_1^2)$$
$$= 2025 - q_1(90 - q_1) - q_2(45 - q_1).$$

It is easy to check that for any $q_1 > 45$ this difference is positive regardless of the value of q_2.[3] Hence we conclude that any $q_1 > 45$ is strictly dominated by $q_1 = 45$.

It is easy to see that firm 2 faces exactly the same profit function, which implies that any $q_2 > 45$ is strictly dominated by $q_2 = 45$. This observation leads to our first round of iterated elimination: a rational firm produces no more than 45 units, implying that the effective strategy space that survives one round of elimination is $q_i \in [0, 45]$ for $i \in \{1, 2\}$.

We can now turn to the second round of elimination. Because $q_2 \le 45$, equation (4.2) implies that firm 1 will choose a quantity no less than 22.5, and a symmetric argument applies to firm 2. Hence the second round of elimination implies that the surviving strategy sets are $q_i \in [22.5, 45]$ for $i \in \{1, 2\}$.

The next step of this process will reduce the strategy set to $q_i \in [22.5, 33\frac{3}{4}]$, and the process will continue on and on. Interestingly the set of strategies that survives

3. This follows because if $q_1 > 45$ then $q_1(90 - q_1) < 2025$ and $q_2(45 - q_1) \le 0$ for any $q_2 \ge 0$.

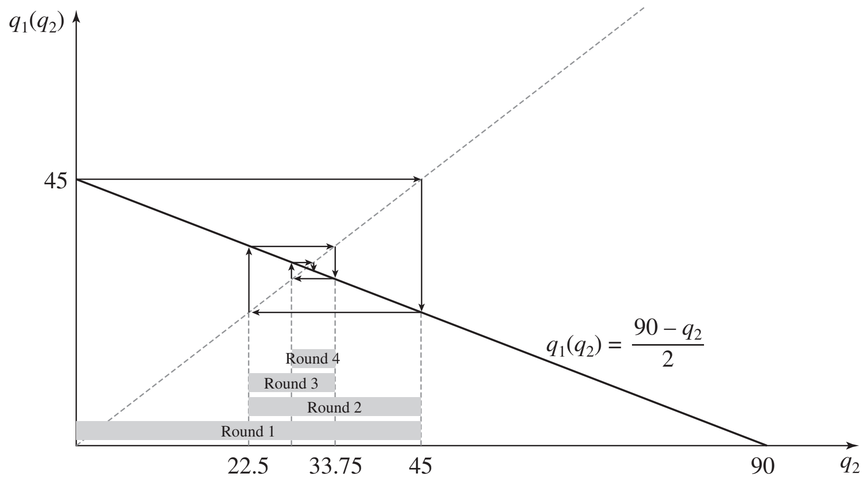

FIGURE 4.1 IESDS convergence in the Cournot game.

this process converges to a single quantity choice of $q_i = 30$. To see this, notice how we moved from one surviving interval to the next. We started by noting that $q_2 \geq 0$, and using equation (4.2) we found that $q_1 \leq 45$, creating the first-round interval of [0, 45]. Then, by symmetry, it follows that $q_2 \leq 45$, and using equation (4.2) again we conclude that $q_1 \geq 22.5$, creating the second-round interval [22.5, 45]. We can see this process graphically in Figure 4.1, where we use the upper (lower) end of the previous interval to determine the lower (upper) end of the next one. If this were to converge to an interval and not to a single point, then by the symmetry between both firms, the resulting interval for each firm would be $[q_{\min}, q_{\max}]$ that simultaneously satisfy two equations with two unknowns: $q_{\min} = \frac{90-q_{\max}}{2}$ and $q_{\max} = \frac{90-q_{\min}}{2}$. However, the only solution to these two equations is $q_{\min} = q_{\max} = 30$. Hence using IESDS for the Cournot game results in a unique predictor of behavior where $q_1 = q_2 = 30$, and each firm earns a profit of $v_1 = v_2 = 900$.

4.2.3 Evaluating IESDS

We turn to evaluate the IESDS solution concept using the criteria we introduced earlier. Start with existence and note that, unlike the concept of strict dominance, we can apply IESDS to *any* game by applying the algorithm just described. It does not require the existence of a strictly dominant strategy, nor does it require the existence of strictly dominated strategies. It is the latter characteristic, however, that gives this concept some bite: when strictly dominated strategies exist, the process of IESDS is able to say something about how common knowledge of rationality restricts behavior.

It is worth noting that this existence result is a consequence of assuming common knowledge of rationality. By doing so we are giving the players the ability to reason through the strategic implications of rationality, and to do so over and over again, while correctly anticipating that other players can perform the same kind of reasoning. Rationality alone does not provide this kind of reasoning.

It is indeed attractive that an IESDS solution always exists. This comes, however, at the cost of uniqueness. In the simple 3×3 matrix game described in (4.1) and the

Cournot duopoly game, IESDS implied the survival of a unique strategy. Consider instead the Battle of the Sexes, given by the following matrix:

		Chris	
		O	F
Alex	O	2, 1	0, 0
	F	0, 0	1, 2

IESDS cannot restrict the set of strategies here for the simple reason that neither O nor F is a strictly dominated strategy for each player. As we can see, this solution concept can be applied (it exists) to any game, but it will often fail to provide a unique solution. For the Battle of the Sexes game, IESDS can only conclude that "anything can happen."

After analyzing the efficiency of the outcomes that can be derived from strict dominance in Section 4.1.3, you may have anticipated the possible efficiency of IESDS equilibria. An easy illustration can be provided by the Prisoner's Dilemma. IESDS leaves (F, F) as the unique survivor, or IESDS equilibrium, after only one round of elimination. As we already demonstrated, the outcome from (F, F) is not Pareto optimal. Similarly, both previous examples (the 3×3 matrix game in (4.1) and the Cournot game) provide further evidence that Pareto optimality need not be achieved by IESDS: In the 3×3 matrix example, both strategy profiles (M, C) and (D, C) yield higher payoffs for both players—(8, 4) and (9, 6), respectively—than the unique IESDS equilibrium, which yields (4, 3). For the Cournot game, producing $q_1 = q_2 = 30$ yields profits of 900 for each firm. If instead they would both produce $q_1 = q_2 = 20$ then each would earn profits of 1000. Thus common knowledge of rationality does not mean that players can guarantee the best outcome for themselves when their own incentives dictate their behavior.

On a final note, it is interesting to observe that there is a simple and quite obvious relationship between the IESDS solution concept and the strict-dominance solution concept:

Proposition 4.2 *If for a game* $\Gamma = \langle N, \{S_i\}_{i=1}^n, \{v_i\}_{i=1}^n \rangle$ s^* *is a* strict dominant strategy equilibrium, *then* s^* uniquely survives IESDS.

Proof If $s^* = (s_1^*, \ldots, s_n^*)$ is a strict dominant strategy equilibrium then, by definition, for every player i all other strategies s_i' are strictly dominated by s_i^*. This implies that after one stage of elimination we will be left with a single profile of strategies, which is exactly s^*, and this concludes the proof. ■

This simple proposition is both intuitive and straightforward. Because rationality is the only requirement needed in order to eliminate all strictly dominated strategies in one round, then if all strategies but one are strictly dominated for each and every player, both IESDS and strict dominance will result in the same outcome. This shows us that whenever strict dominance results in a unique outcome, then IESDS will result in the same unique outcome after one round. However, as we saw earlier, IESDS may offer a fine prediction when strict dominance does not apply. This is exactly what the extra assumption of common knowledge of rationality delivers: a more widely applicable solution concept. However, the assumption of common knowledge of rationality is far from innocuous. It requires the players to be, in some way, extremely intelligent and to possess unusual levels of foresight. For the most part, game theory relies on this strong assumption, and hence it must be applied to the real world with caution. Remember the rule about how assumptions drive conclusions: garbage in, garbage out.

4.3 Beliefs, Best Response, and Rationalizability

Both of the solution concepts we have seen so far, strict dominance and IESDS, are based on eliminating actions that players would never play. An alternative approach is to ask: what possible strategies might players choose to play and under what conditions? When we considered eliminating strategies that no rational player would choose to play, it was by finding some strategy that is *always* better or, as we said, that dominates the eliminated strategies. A strategy that cannot be eliminated, therefore, suggests that *under some conditions* this strategy is the one that the player may like to choose. When we qualify a strategy to be the best one a player can choose *under some conditions,* these conditions must be expressed in terms that are rigorous and are related to the game that is being played.

To set the stage, think about situations in which you were puzzled about the behavior of someone you knew. To consider his choice as irrational, or simply stupid, you would have to consider whether there is a way in which he could defend his action as a good choice. A natural way to determine whether this is the case is to simply ask him, "What were you thinking?" If the response lays out a plausible situation for which his choice was a good one, then you cannot question his rationality. (You may of course question his wisdom, or even his sanity, if his thoughts seem bizarre.)

This is a type of reasoning that we will formalize and discuss in this chapter. If a strategy s_i is not strictly dominated for player i then it must be that there are combinations of strategies of player i's opponents for which the strategy s_i is player i's best choice. This reasoning will allow us to justify or *rationalize* the choice of player i.

4.3.1 The Best Response

As we discussed early in Chapter 3, what makes a game different from a single-player decision problem is that once you understand the actions, outcomes, and preferences of a decision problem, then you can choose your best or optimal action. In a game, however, your optimal decision not only depends on the structure of the game, but it will often depend on what the other players are doing.

Take the Battle of the Sexes as an example:

		Chris	
		O	F
Alex	O	2, 1	0, 0
	F	0, 0	1, 2

As the matrix demonstrates, the best choice of Alex depends on what Chris will do. If Chris goes to the opera then Alex would rather go to the opera instead of going to the football game. If, however, Chris goes to the football game then Alex's optimal action is switched around.

This simple example illustrates an important idea that will escort us throughout this book and that (one hopes) will escort you through your own decision making in strategic situations. In order for a player to be optimizing in a game, he has to choose a best strategy *as a response* to the strategies of his opponents. We therefore introduce the following formal definition:

Definition 4.5 The strategy $s_i \in S_i$ is player i's **best response** to his opponents' strategies $s_{-i} \in S_{-i}$ if

$$v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i}) \quad \forall s_i' \in S_i.$$

I can't emphasize enough how central this definition is to the concept of strategic behavior and rationality. In fact rationality implies that given any belief a player has about his opponents' behavior, he must choose an action that is best for him given his beliefs. That is,

Claim 4.2 *A rational player who believes that his opponents are playing some $s_{-i} \in S_{-i}$ will always choose a best response to s_{-i}.*

For instance, in the Battle of the Sexes, if Chris believes that Alex will go to the opera then Chris's best response is to go to the opera because

$$v_2(O, O) = 1 > 0 = v_2(O, F).$$

Similarly if Chris believes that Alex will go to the football game then Chris's best response is to go to the game as well.

There are some appealing relationships between the concept of playing a best response and the concept of dominated strategies. First, if a strategy s_i is strictly dominated, it means that some other strategy s_i' is *always* better. This leads us to the observation that the strategy s_i could not be a best response to *anything:*

Proposition 4.3 *If s_i is a* strictly dominated strategy *for player i, then it* cannot be a best response *to any $s_{-i} \in S_{-i}$.*

Proof If s_i is strictly dominated, then there exists some $s_i' \succ_i s_i$ such that $v_i(s_i', s_{-i}) > v_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. But this in turn implies that there is no $s_{-i} \in S_{-i}$ for which $v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i})$, and thus that s_i cannot be a best response to any $s_{-i} \in S_{-i}$. ■

A companion to this proposition would explore strictly dominant strategies, which are in some loose way the "opposite" of strictly dominated strategies. You should easily be able to convince yourself that if a strategy s_i^D is a strictly dominant strategy then it must be a best response to *anything* i's opponents can do. This immediately implies the next proposition, which is slightly broader than the simple intuition just provided and requires a bit more work to prove formally:

Proposition 4.4 *If in a finite normal-form game s^* is a strict dominant strategy equilibrium, or if it uniquely survives IESDS, then s_i^* is a best response to $s_{-i}^* \forall i \in N$.*

Proof If s^* is a dominant strategy equilibrium then it uniquely survives IESDS, so it is enough to prove the proposition for strategies that uniquely survive IESDS. Suppose s^* uniquely survives IESDS, and choose some $i \in N$. Suppose in negation to the proposition that s_i^* is not a best response to s_{-i}^*. This implies that there exists an $s'_i \in S_i \backslash \{s^*_i\}$ (this is the set S_i without the strategy s^*_i) such that $v_i(s'_i, s^*_{-i}) > v_i(s^*_i, s^*_{-i})$. Let $S'_i \subset S_i$ be the set of all such s'_i for which $v_i(s'_i, s^*_{-i}) > v_i(s^*_i, s^*_{-i})$. Because s'_i was eliminated while s^*_{-i} was not (recall that s^* uniquely survives IESDS), there must be some s''_i such that $v_i(s''_i, s^*_{-i}) > v_i(s'_i, s^*_{-i}) > v_i(s^*_i, s^*_{-i})$, implying that $s''_i \in S'_i$. Because the game is finite, an induction argument on S'_i then implies that there exists a strategy $s'''_i \in S'_i$ that must survive IESDS. But this is a contradiction to s^* being the unique survivor of IESDS. ■

This is a "proof by contradiction," and the intuition behind the proof may be a bit easier than the formal write-up. The logic goes as follows: If it is true that s^*_i was not a best response to s^*_{-i} then there was some other strategy s'_i that was a best response to s^*_{-i} that was eliminated at some previous round. But then it must be that there was a third strategy that was better than both of i's aforementioned strategies against s^*_{-i} in order to knock s'_i out in some earlier round. But then how could s^*_i knock out this third strategy against s^*_{-i}, which survived? This can't be true, which means that the initial negative premise in the proof, that s^*_i was not a best response to s^*_{-i}, must be false, hence the contradiction.

With the concept of a best response in hand, we need to think more seriously about the following question: to what profile of strategies should a player be playing a best response? Put differently, if my best response depends on what the other players are doing, then how should I choose between all the best responses I can possibly have? This is particularly pertinent because we are discussing static games, in which players choose their actions without knowing what their opponents are choosing. To tackle this important question, we need to give players the ability to form *conjectures* about what others are doing. We have alluded to the next step in the claim made earlier, which stated that "a rational player who believes that his opponents are playing some $s_{-i} \in S_{-i}$ will always choose a best response to s_{-i}." Thus we have to be mindful of what a player *believes* in order to draw conclusions about whether or not the player is choosing a best response.

4.3.2 Beliefs and Best-Response Correspondences

Suppose that s'_i is a best response for player i to his opponents playing s'_{-i}, and assume for the moment that it is not a best response for any other profile of actions that i's opponents can choose. When would a rational player i choose to play s'_i? The answer follows directly from rationality: he will play s'_i *only when his beliefs about other players' behavior* justify the use of s'_i, or in other words when he believes that his opponents will play s'_{-i}.

Introducing the concept of beliefs, and actions that best respond to beliefs, is central to the analysis of strategic behavior. If a player is fortunate enough to be playing in a game in which he has a strictly dominant strategy then his beliefs about the behavior of others play no role. The player's strictly dominant strategy is his best response *independent of* his opponents' play, and hence it is always a best response. But when no strictly dominant strategy exists, a player must ask himself, "What do I think my opponents will do?" The answer to this question should guide his own behavior.

To make this kind of reasoning precise we need to define what we mean by a player's belief. In a well-defined game, the only thing a player should be thinking about is what he believes his opponents are doing. Therefore we offer the following definition:

Definition 4.6 A **belief** of player i is a possible profile of his opponents' strategies, $s_{-i} \in S_{-i}$.

Given that a player has a particular belief about his opponents' strategies, he will be able to formulate a best response to that belief. The best response of a player to a certain strategy of his opponents may be unique, as in many of the games we have seen up to now.

For example, consider the Battle of the Sexes. When Chris believes that Alex is going to the opera, his unique best response is to go to the opera. Similarly, if he believes that Alex will go to the football game, then he should go to the game. For every unique belief there is a best response. Similarly recall that in the Cournot game the best choice of player 1 given any choice of player 2 solved the first-order condition of player 1's maximization problem, resulting in the function

$$q_1(q_2) = \frac{90 - q_2}{2}, \tag{4.3}$$

which assigns a unique value of q_1 to any value of q_2 for $q_2 \in [0, 90]$. Hence the function in (4.3) is the best-response function of firm 1 in the Cournot game.

We can therefore think of a rational player as having a *recipe book* that is a list of instructions as follows: "If I think my opponents are doing s_{-i}, then I should do s_i; if I think they're doing s'_{-i}, then I should do s'_i; . . . " This list should go on until it exhausts all the possible strategies that player i's opponents can choose. If we think of this list of best responses as a plan, then this plan *maps beliefs into a choice of action,* and this choice of action must be a best response to the beliefs. We can think of this as player i's *best-response function*.

There may, however, be games in which for some beliefs a player will have more than one best-response strategy. Consider, for example, the following simple game:

		Player 2		
		L	C	R
Player 1	U	3, 3	5, 1	6, 2
	M	4, 1	8, 4	3, 6
	D	4, 0	9, 6	6, 8

If player 1 believes that player 2 is playing the column R then both U and D are each a best response. Similarly if player 1 believes that player 2 is playing the column L then both M and D are each a best response.

The fact that a player may have more than one best response implies that we can't think of the best-response mapping from opponents' strategies S_{-i} to an action by player i as a *function,* because by definition a function would select only *one action* as a best response (see Section 19.2 of the mathematical appendix). Thus we offer the following definition:

Definition 4.7 The **best-response correspondence** of player i selects for each $s_{-i} \in S_{-i}$ a subset $BR_i(s_{-i}) \subset S_i$ where each strategy $s_i \in BR_i(s_{-i})$ is a best response to s_{-i}.

That is, given a belief player i has about his opponents, s_{-i}, the set of all his possible strategies that are a best response to s_{-i} is denoted by $BR_i(s_{-i})$. If he has a unique best response to s_{-i} then $BR_i(s_{-i})$ will contain only one strategy from S_i.

4.3.3 Rationalizability

Equipped with the idea of beliefs and a player's best responses to his beliefs, the next natural step is to allow the players to reason about which beliefs to have about their opponents. This reasoning must take into account the rationality of all players, common knowledge of rationality, and the fact that all the players are trying to guess the behavior of their opponents.

To a large extent we employed similar reasoning when we introduced the solution concept of IESDS: instead of asking what your opponents might be doing, you asked "What would a rational player *not do*?" Then, assuming that all players follow this process by common knowledge of rationality, we were able to make some prediction about which strategies cannot be eliminated.

In what follows we introduce another way of reasoning that rules out irrational behavior with a similar iterated process that is, in many ways, the mirror image of IESDS. This next solution concept also builds on the assumption of common knowledge of rationality. However, instead of asking "What would a rational player *not do*?" our next concept asks "What *might* a rational player *do*?" A rational player will select only strategies that are a *best response to some profile of his opponents*. Thus we have

Definition 4.8 A strategy $s_i \in S_i$ is *never a best response* if there are no beliefs $s_{-i} \in S_{-i}$ for player i for which $s_i \in BR_i(s_{-i})$.

The next step, as in IESDS, is to use the common knowledge of rationality to build an iterative process that takes this reasoning to the limit. After employing this reasoning one time, we can eliminate all the strategies that are *never a best response,* resulting in a possibly "smaller" reduced game that includes only strategies that can be a best response in the original game. Then we can employ this reasoning again and again, in a similar way that we did for IESDS, in order to eliminate outcomes that should not be played by players who share a common knowledge of rationality.

The solution concept of rationalizability is defined precisely by iterating this thought process. The set of strategy profiles that survive this process is called the set of *rationalizable strategies*. (We postpone offering a definition of rationalizable strategies because the introduction of mixed strategies, in which players can play stochastic strategies, is essential for the complete definition.)

4.3.4 The Cournot Duopoly Revisited

Consider the Cournot duopoly example used to demonstrate IESDS in Section 4.2.2, with demand $p(q) = 100 - q$ and costs $c_i(q_i) = 10q_i$ for both firms. As we showed earlier, firm 1 maximizes its profits $v_1(q_1, q_2) = (100 - q_1 - q_2)q_1 - 10q_1$ by setting the first-order condition $90 - 2q_1 - q_2 = 0$. Now that we have introduced the idea of a best response, it should be clear that this firm's best response is *immediately derived* from the first-order condition. In other words, if firm 1 believes that firm 2 will choose the quantity q_2, then it should choose q_1 according to the best-response function,

$$BR_1(q_2) = \begin{cases} \frac{90-q_2}{2} & \text{if } 0 \le q_2 < 90 \\ 0 & \text{if } q_2 \ge 90. \end{cases}$$

Notice that the best response is indeed a function. For all $0 \le q_2 < 90$ there is a unique positive best response. For $q_2 \ge 90$ the price is guaranteed to be below 10, in which case any quantity firm 1 will choose will yield a negative profit (its costs per unit are 10), and hence the best response is to produce nothing. Similarly we can define the best-response correspondence of firm 2, which is symmetric.

Examining $BR_1(q_2)$ implies that firm 1 will choose to produce *only* quantities between 0 and 45. That is, there will be no beliefs about q_2 for which quantities above 45 are a best response. By symmetry the same is true for firm 2. Thus a first round of rationalizability implies that the only quantities that can be best-response quantities for both firms must lie in the interval [0, 45]. The next round of rationalizability for the game in which $S_i = [0, 45]$ for both firms shows that the best response of firm i is to choose any quantity $q_i \in [22.5, 45]$. Just as with IESDS, this process will continue on and on. The set of rationalizable strategies converges to a single quantity choice of $q_i = 30$ for both firms.

4.3.5 The "p-Beauty Contest"

Consider a game with n players, so that $N = \{1, \ldots, n\}$. Each player has to choose an integer between 0 and 20, so that $S_i = \{0, 1, 2, \ldots, 19, 20\}$. The winners are the players who choose an integer number that is closest to $\frac{3}{4}$ of the average. For example, if $n = 3$, and if the three players choose $s_1 = 1$, $s_2 = 5$, and $s_3 = 18$, then the average is $(1 + 5 + 18) \div 3 = 8$, and $\frac{3}{4}$ of the average is 6, so the winner is player 2. This game is called the ***p*-beauty contest** because, unlike in a standard beauty contest, you are not trying to guess what everyone else is guessing (beauty is what people choose it to be), but rather you are trying to guess p times the average, in this case $p = \frac{3}{4}$.[4]

Note that it is possible for more than one person to be a winner. If, for example, $s_1 = s_2 = 2$ and $s_3 = 8$, then the average is 3 and $\frac{3}{4}$ of the average is $2\frac{1}{4}$, so that both player 1 and player 2 are winners. Because more than one person can win, define the set of winners $W \subset N$ as those players who are closest to $\frac{3}{4}$ of the average, and the rest are all losers. The set of winners is defined as[5]

$$W = \left\{ i \in N : \arg\min_{i \in N} \left| s_i - \frac{3}{4}\frac{1}{n}\sum_{j=1}^{n} s_j \right| \right\}.$$

4. John Maynard Keynes (1936) described the playing of the stock market as analogous to entering a newspaper beauty-judging contest in which one selects the six prettiest faces out of a hundred photographs, with the prize going to the person whose selections are closest to those of all the other players. Keynes's depiction of a beauty contest is a situation in which you want to guess what others are guessing.

5. The notation $\arg\min_{i \in N} \left| s_i - \frac{3}{4}\frac{1}{n}\sum_{i=1}^{n} s_i \right|$ means that we are selecting all the i's for which the expression $\left| s_i - \frac{3}{4}\frac{1}{n}\sum_{j=1}^{n} s_j \right|$ is minimized. Since this expression is the absolute value of the difference between s_i and $\frac{3}{4}\frac{1}{n}\sum_{j=1}^{n} s_j$, this selection will result in the player (or players) who are closest to $\frac{3}{4}$ of the average.

To introduce payoffs, each player pays 1 to play the game, and winners split the pot equally among themselves. This implies that if there are $m \geq 1$ winners, each gets a payoff of $\frac{N-m}{m}$ (his share of the pot net of his own contribution) while losers get -1 (they lose their contribution).

Turning to rationalizable strategies, we must begin by finding strategies that are never a best response. This is not an easy task, but some simple insights can get us moving in the right direction. In particular the objective is to guess a number closest to $\frac{3}{4}$ of the average, which means that a player would want to guess a number that is generally smaller than the highest numbers that other players may be guessing.

This logic suggests that if there are strategies that are never a best response, they should be the higher numbers, and it is natural to start with 20: can choosing $s_i = 20$ be a best response? If you believe that the average is below 20, then 20 cannot be a best response—a lower number will be the best response. If the average were 20, that means that you and everyone else would be choosing 20, and you would then split the pot with all the other players. If you believe this, and instead of 20 you choose 19, then you will win the whole pot for sure, regardless of the number of players, because for any number of players n if everyone else is choosing 20 then 19 will be a unique winner.[6] This shows that 20 can never be a best response. Interestingly 19 is not the *unique* best response to the belief that all others are playing 20. (This is left as an exercise.) The important point, however, is that 20 *cannot* be a best response to *any* beliefs a player can have.

This analysis shows that only the numbers $S_i^1 = \{0, 1, \ldots, 19\}$ survive the first round of rationalizable behavior. Similarly after each additional round we will “lose” the highest number until we go through 19 rounds and are left with $S_i^{19} = \{0, 1\}$, meaning that after 19 rounds of dropping strategies that cannot be a best response, we are left with two strategies that survive: 1 and 0. If $n > 2$, we cannot reduce this set further: if, for example, player i believes that all the other players are choosing 1 then choosing 1 is a best response for him. (This is left as an exercise.) Similarly regardless of n, if he believes that everyone is choosing 0 then choosing 0 is his best response. Thus we are able to predict using rationalizability that players will not choose a number greater than 1, and if there are only two players then we will predict that both will choose 0.

Will this indeed predict behavior? Only if our *assumptions* about behavior are correct. If you were to play this game and you don’t think that your opponents are doing these steps in their minds, then you may want to choose a number higher than 1. An interesting set of experiments is summarized in Camerer (2003).[7]

Remark By now you must have concluded that IESDS and rationalizability are two sides of the same coin, and you might even think that they are one and the same. This is almost true, and for two-player games it turns out that these two processes indeed result in the same outcomes. We discuss this issue briefly in Section 6.3, after we introduce the concept of mixed strategies. A more complete treatment can be found in Chapter 2 of Fudenberg and Tirole (1991).

6. The average, $\frac{1}{n}[(n-1)20 + 19]$, lies between 19.5 (if $n = 2$) and 20 (for $n \to \infty$), so that $\frac{3}{4}$ of the average lies between $14\frac{5}{8}$ and 15.

7. From my own experience running these games in classes of students, it is rare that the winning number is below 4.

4.3.6 Evaluating Rationalizability

In terms of existence, uniqueness, and implications for Pareto optimality, rationalizability is practically the same as IESDS. It will sometimes have bite, and may even restrict behavior quite dramatically as in the examples given. But if applied to the Battle of the Sexes, rationalizability will say "anything can happen."

4.4 Summary

- Rational players will never play a dominated strategy and will always play a dominant strategy when it exists.
- When players share common knowledge of rationality, the only strategies that are sensible are those that survive IESDS.
- Rational players will always play a best response to their beliefs. Hence any strategy for which there are no beliefs that justify its choice will never be chosen.
- Outcomes that survive IESDS, rationalizability, or strict dominance need not be Pareto optimal, implying that players may not be able to achieve desirable outcomes if they are left to their own devices.

4.5 Exercises

4.1 **Prove Proposition 4.1:** If the game $\Gamma = \langle N, \{S_i\}_{i=1}^n, \{v_i\}_{i=1}^n \rangle$ has a strictly dominant strategy equilibrium s^D, then s^D is the unique dominant strategy equilibrium.

4.2 **Weak Dominance:** We call the strategy profile $s^W \in S$ a **weakly dominant strategy equilibrium** if $s_i^W \in S_i$ is a weakly dominant strategy for all $i \in N$, that is, if $v_i(s_i, s_{-i}) \geq v_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and for all $s_{-i} \in S_{-i}$.

  a. Provide an example of a game in which there is no weakly dominant strategy equilibrium.
  b. Provide an example of a game in which there is more than one weakly dominant strategy equilibrium.

4.3 **Discrete First-Price Auction:** An item is up for auction. Player 1 values the item at 3 while player 2 values the item at 5. Each player can bid either 0, 1, or 2. If player i bids more than player j then i wins the good and pays his bid, while the loser does not pay. If both players bid the same amount then a coin is tossed to determine who the winner is, and the winner gets the good and pays his bid while the loser pays nothing.

  a. Write down the game in matrix form.
  b. Does any player have a strictly dominated strategy?
  c. Which strategies survive IESDS?

4.4 **eBay's Recommendation:** It is hard to imagine that anyone is not familiar with eBay, the most popular auction web site by far. In a typical eBay auction a good is placed for sale, and each bidder places a "proxy bid," which eBay keeps in memory. If you enter a proxy bid that is lower than the current highest bid, then your bid is ignored. If, however, it is higher, then the current bid increases up to one increment (say, \$0.01) above the *second highest* proxy bid. For example, imagine that three people have placed bids on a used laptop of \$55, \$98, and \$112. The current price will be at \$98.01, and if the auction ended the player who bid \$112 would win at a price of \$98.01. If you were to place a bid of \$103.45 then the player who bid \$112 would still win, but at a price of \$103.46, while if your bid was \$123.12 then you would win at a price of \$112.01.

Now consider eBay's historical recommendation that you think hard about the value you impute to the good and that you enter your true value as your bid—no more, no less. Assume that the value of the good for each potential bidder is independent of how much other bidders value it.

a. Argue that bidding more than your valuation is weakly dominated by actually bidding your valuation.
b. Argue that bidding less than your valuation is weakly dominated by actually bidding your valuation.
c. Use your analysis to make sense of eBay's recommendation. Would you follow it?

4.5 **Iterated Elimination:** In the following normal-form game, which strategy profiles survive iterated elimination of strictly dominated strategies?

		Player 2		
		L	C	R
	U	6, 8	2, 6	8, 2
Player 1	M	8, 2	4, 4	9, 5
	D	8, 10	4, 6	6, 7

4.6 **Roommates:** Two roommates each need to choose to clean their apartment, and each can choose an amount of time $t_i \geq 0$ to clean. If their choices are t_i and t_j, then player i's payoff is given by $(10 - t_j)t_i - t_i^2$. (This payoff function implies that the more one roommate cleans, the less valuable is cleaning for the other roommate.)

a. What is the best response correspondence of each player i?
b. Which choices survive one round of IESDS?
c. Which choices survive IESDS?

4.7 **Campaigning:** Two candidates, 1 and 2, are running for office. Each has one of three choices in running his campaign: focus on the positive aspects of one's own platform (call this a positive campaign [P]), focus on the positive aspects of one's own platform while attacking one's opponent's campaign (call this a balanced campaign [B]), or finally focus only on attacking one's opponent (call this a negative campaign [N]). All a candidate cares about is the probability of winning, so assume that if a candidate expects to win with probability $\pi \in [0, 1]$, then his payoff is π. The probability that a candidate wins depends on his choice of campaign and his opponent's choice. The probabilities of winning are given as follows:

If both choose the same campaign then each wins with probability 0.5.
If candidate i uses a positive campaign while $j \neq i$ uses a balanced one then i loses for sure.

If candidate i uses a positive campaign while $j \neq i$ uses a negative one then i wins with probability 0.3.

If candidate i uses a negative campaign while $j \neq i$ uses a balanced one then i wins with probability 0.6.

a. Model this story as a normal-form game. (It suffices to be specific about the payoff function of one player and to explain how the other player's payoff function is different and why.)

b. Write down the game in matrix form.

c. What happens at each stage of elimination of strictly dominated strategies? Will this procedure lead to a clear prediction?

4.8 **Beauty Contest Best Responses:** Consider the p-beauty contest presented in Section 4.3.5.

a. Show that if player i believes that everyone else is choosing 20 then 19 is not the only best response for any number of players n.

b. Show that the set of best-response strategies to everyone else choosing the number 20 depends on the number of players n.

4.9 **Beauty Contest Rationalizability:** Consider the p-beauty contest presented in Section 4.3.5. Show that if the number of players $n > 2$ then the choices $\{0, 1\}$ for each player are both rationalizable, while if $n = 2$ then only the choice of $\{0\}$ by each player is rationalizable.

4.10 **Popsicle Stands:** Five lifeguard towers are lined up along a beach; the leftmost tower is number 1 and the rightmost tower is number 5. Two vendors, players 1 and 2, each have a popsicle stand that can be located next to one of five towers. There are 25 people located next to each tower, and each person will purchase a popsicle from the stand that is closest to him or her. That is, if player 1 locates his stand at tower 2 and player 2 at tower 3, then 50 people (at towers 1 and 2) will purchase from player 1, while 75 (from towers 3, 4, and 5) will purchase from vendor 2. Each purchase yields a profit of \$1.

a. Specify the strategy set for each player. Are there any strictly dominated strategies?

b. Find the set of strategies that survive rationalizability.

5

Pinning Down Beliefs: Nash Equilibrium

We have seen three solution concepts that offer some insights into predicting the behavior of rational players in strategic (normal-form) games. The first, strict dominance, relied only on rationality, and in some cases, like the Prisoner's Dilemma, it predicted a unique outcome, as it would in any game for which a dominant strategy equilibrium exists. However, it often fails to exist. The two sister concepts of IESDS and rationalizability relied on more than rationality by requiring common knowledge of rationality. In return a solution existed for every game, and for some games there was a unique prediction. Moreover, whenever there is a strict dominant equilibrium, it also uniquely survives IESDS and rationalizability. Even for some games for which the strict-dominance solution did not apply, like the Cournot duopoly, we obtained a unique prediction from IESDS and rationalizability.

However, when we consider a game like the Battle of the Sexes, none of these concepts had any bite. Dominant strategy equilibrium did not apply, and both IESDS and rationalizability could not restrict the set of reasonable behavior:

		Chris	
		O	F
Alex	O	2, 1	0, 0
	F	0, 0	1, 2

For example, we cannot rule out the possibility that Alex goes to the opera while Chris goes to the football game, because Alex may behave optimally given his belief that Chris is going to the opera, and Chris may behave optimally given his belief that Alex is going to the football game. Yet there is something troubling about this outcome. If we think of this pair of actions not only as actions, but as a system of actions and beliefs, then there is something of a dissonance: indeed the players are playing best responses to their beliefs, but their beliefs are wrong!

In this chapter we make a rather heroic leap that ties together beliefs and actions and results in the most central and best-known solution concept in game theory. As already mentioned, for dominant strategy equilibrium we required only that players be rational, while for IESDS and rationalizability we required common knowledge of rationality. Now we introduce a much more demanding concept, that of the Nash equilibrium, first put forth by John Nash (1950a), who received the Nobel Prize in Economics for this achievement.[1]

5.1 Nash Equilibrium in Pure Strategies

To cut to the chase, a Nash equilibrium is a system of beliefs and a profile of actions for which each player is playing a best response to his beliefs and, moreover, players have *correct beliefs.* Another common way of defining a Nash equilibrium, which does not refer to beliefs, is as a profile of strategies for which *each* player is choosing a best response to the strategies of *all other* players. Formally we have:

Definition 5.1 The pure-strategy profile $s^* = (s_1^*, s_2^*, \ldots, s_n^*) \in S$ is a **Nash equilibrium** if s_i^* is a best response to s_{-i}^*, for all $i \in N$, that is,

$$v_i(s_i^*, s_{-i}^*) \geq v_i(s_i', s_{-i}^*) \quad \text{for all } s_i' \in S_i \text{ and all } i \in N.$$

Consider as an example the following two-player discrete game, which we used to demonstrate IESDS:

		Player 2		
		L	C	R
	U	$\boxed{4, 3}$	5, 1	6, 2
Player 1	M	2, 1	8, 4	3, 6
	D	3, 0	9, 6	2, 8

In this game, the only pair of pure strategies that survived IESDS is the pair (U, L). As it turns out, this is also the only pair of strategies that constitutes a Nash equilibrium. If player 2 is playing the column L, then player 1's best response is $BR_1(L) = \{U\}$; at the same time, if player 1 is playing the row U, then player 2's best response is $BR_2(U) = \{L\}$.

What about the other games we saw? In the Prisoner's Dilemma, the unique Nash equilibrium is (F, F). This should be easy to see: if each player is playing a dominant strategy then he is by definition playing a best response to *anything* his opponent is choosing, and hence it must be a Nash equilibrium. As we will soon see in Section 5.2.3, and as you may have already anticipated, the unique Nash equilibrium in the Cournot duopoly game that we discussed earlier is $(q_1, q_2) = (30, 30)$.

The relationship between strict-dominance, IESDS, rationalizability, and Nash equilibrium outcomes in many of the examples we have analyzed is no coincidence. There is a simple relationship between the concepts we previously explored and that of Nash equilibrium, as the following proposition clearly states:

Proposition 5.1 *Consider a strategy profile* $s^* = (s_1^*, s_2^*, \ldots, s_n^*)$. *If* s^* *is either*

1. *a strict dominant strategy equilibrium,*
2. *the unique survivor of IESDS, or*

1. Nash's life was the subject of a very successful movie, *A Beautiful Mind,* based on the book by Sylvia Nasar (1998).

3. *the unique rationalizable strategy profile,*

then s^* is the unique Nash equilibrium.

This proposition is not difficult to prove, and the proof is left as exercise 5.1 at the end of the chapter. The intuition is of course quite straightforward: we know that if there is a strict dominant strategy equilibrium then it uniquely survives IESDS and rationalizability, and this in turn must mean that each player is playing a best response to the other players' strategies.

At the risk of being repetitive, let me emphasize the requirements for a Nash equilibrium:

1. Each player is playing a *best response* to his *beliefs.*
2. The *beliefs* of the players about their opponents *are correct.*

The first requirement is a direct consequence of rationality. It is the second requirement that is very demanding and is a tremendous leap beyond the requirements we have considered so far. It is one thing to ask people to behave rationally given their beliefs (play a best response), but it is a totally different thing to ask players to predict the behavior of their opponents correctly.

Then again it may be possible to accept such a strong requirement if we allow for some reasoning that is beyond the physical structure of the game. For example, imagine, in the Battle of the Sexes game, that Alex is an influential person—people just seem to follow Alex, and this is something that Alex knows well. In this case Chris should believe, knowing that Alex is so influential, that Alex would expect Chris to go to the opera. Knowing this, Alex should believe that Chris will indeed believe that Alex is going to the opera, and so Chris will go to the opera too.

It is important to note that this argument *is not* that Chris likes to please Alex—such an argument would change the payoff of the game and increase Chris's payoff from pleasing Alex. Instead this argument is only about beliefs that are "self-fulfilling." That is, if these beliefs have some weight to them, which may be based on past experience or on some kind of deductive reasoning, then they will be self-fulfilling in that they support the behavior that players believe will occur.

Indeed (O, O) is a Nash equilibrium. However, notice that we can make the symmetric argument about Chris being an influential person: (F, F) is also a Nash equilibrium. As the external game theorist, however, we should not say more than "one of these two outcomes is what we predict." (You should be able to convince yourself that no other pair of pure strategies in the Battle of the Sexes game is a Nash equilibrium.)

5.1.1 Pure-Strategy Nash Equilibrium in a Matrix

This short section presents a simple method to find all the pure-strategy Nash equilibria in a matrix game if at least one exists. Consider the following two-person finite game in matrix form:

		Player 2		
		L	C	R
	U	7, 7	4, 2	1, 8
Player 1	M	2, 4	5, 5	2, 3
	D	8, 1	3, 2	0, 0

It is easy to see that no strategy is dominated and thus that strict dominance cannot be applied to this game. For the same reason, IESDS and rationalizability will conclude that anything can happen. However, a pure-strategy Nash equilibrium does exist, and in fact it is unique. To find it we use a simple method that builds on the fact that any Nash equilibrium must call for a pair of strategies in which each of the two players is playing a best response to his opponent's strategy. The procedure is best explained in three steps:

Step 1: For every *column,* which is a strategy of player 2, find the highest payoff entry for player 1. By definition this entry must be in the row that is a best response for the particular column being considered. Underline the pair of payoffs in this row under this column:

		Player 2		
		L	C	R
	U	7, 7	4, 2	1, 8
Player 1	M	2, 4	$\underline{5, 5}$	$\underline{2, 3}$
	D	$\underline{8, 1}$	3, 2	0, 0

Step 1 identifies the best response of player 1 *for each of the pure strategies* (columns) of player 2. For instance, if player 2 is playing L, then player 1's best response is D, and we underline the payoffs associated with this row in column 1. After performing this step we see that there are three pairs of pure strategies at which player 1 is playing a best response: (D, L), (M, C), and (M, R).

Step 2: For every *row,* which is a strategy of player 1, find the highest payoff entry for player 2. By definition this entry must be in the column that is a best response for the particular row being considered. Overline the pair of payoffs in this entry:

		Player 2		
		L	C	R
	U	7, 7	4, 2	$\overline{1, 8}$
Player 1	M	2, 4	$\overline{\underline{5, 5}}$	$\underline{2, 3}$
	D	$\underline{8, 1}$	$\overline{3, 2}$	0, 0

Step 2 similarly identifies the pairs of strategies at which player 2 is playing a best response. For instance, if player 1 is playing D, then player 2's best response is C, and we overline the payoffs associated with this column in row 3. We can continue to conclude that player 2 is playing a best response at three strategy pairs: (D, C), (M, C), and (U, R).

Step 3: If any matrix entry has both an under- and an overline, it is the outcome of a Nash equilibrium in pure strategies.

This follows immediately from the fact that both players are playing a best response at any such pair of strategies. In this example we find that (M, C) is the unique pure-strategy Nash equilibrium—it is the only pair of pure strategies for which both players are playing a best response. If you apply this approach to the Battle of the Sexes, for example, you will find both pure-strategy Nash equilibria, (O, O) and (F, F). For the Prisoner's Dilemma only (F, F) will be identified.

5.1.2 Evaluating the Nash Equilibria Solution

Considering our criteria for evaluating solution concepts, we can see from the Battle of the Sexes example that we may not have a unique Nash equilibrium. However, as alluded to in our earlier discussion of the Battle of the Sexes game, there is no reason to expect that we should. Indeed we may need to entertain other aspects of an environment in which players interact, such as social norms and historical beliefs, to make precise predictions about which of the possible Nash equilibria may result as the more likely outcome.

In Section 5.2.4 we will analyze a price competition game in which a Nash equilibrium may fail to exist. It turns out, however, that for quite general conditions games will have at least one Nash equilibrium. For the interested reader, Section 6.4 discusses some conditions that guarantee the existence of a Nash equilibrium, which was a central part of Nash's Ph.D. dissertation. This fact gives the Nash solution concept its power—like IESDS and rationalizability, the solution concept of Nash is widely applicable. It will, however, usually lead to more refined predictions than those of IESDS and rationalizability, as implied by proposition 5.1.

As with the previous solution concepts, we can easily see that Nash equilibrium does not guarantee Pareto optimality. The theme should be obvious by now: left to their own devices, people in many situations will do what is best for them, at the expense of social efficiency. This point was made quite convincingly and intuitively in Hardin's (1968) "tragedy of the commons" argument, which we explore in Section 5.2.2. This is where our focus on *self-enforcing outcomes* has its bite: our solution concepts took the game as given, and they imposed rationality and common knowledge of rationality to try to see what players would choose to do. If they each seek to maximize their individual well-being then the players may hinder their ability to achieve socially optimal outcomes.

5.2 Nash Equilibrium: Some Classic Applications

The previous section introduced the central pillar of modern noncooperative game theory, the Nash equilibrium solution concept. It has been applied widely in economics, political science, legal studies, and even biology. In what follows we demonstrate some of the best-known applications of the concept.

5.2.1 Two Kinds of Societies

The French philosopher Jean-Jacques Rousseau presented the following situation that describes a trade-off between playing it safe and relying on others to achieve a larger gain. Two hunters, players 1 and 2, can each choose to hunt a stag (S), which provides a rather large and tasty meal, or hunt a hare (H)—also tasty, but much less filling. Hunting stags is challenging and requires mutual cooperation. If either hunts a stag alone, the chance of success is negligible, while hunting hares is an individualistic enterprise that is not done in pairs. Hence hunting stags is most beneficial for society but requires "trust" between the hunters in that each believes that the other is joining forces with him. The game, often referred to as the Stag Hunt game, can be described by the following matrix:

	S	H
S	5, 5	0, 3
H	3, 0	3, 3

It is easy to see that the game has two pure-strategy equilibria: (S, S) and (H, H). However, the payoff from (S, S) Pareto dominates that from (H, H). Why then would (H, H) ever be a reasonable prediction? This is precisely the strength of the Nash equilibrium concept. If each player anticipates that the other will not join forces, then he knows that going out to hunt the stag alone is not likely to be a successful enterprise and that going after the hare will be better. This belief would result in a society of individualists who do not cooperate to achieve a better outcome. In contrast, if the players expect each other to be cooperative in going after the stag, then this anticipation is self-fulfilling and results in what can be considered a cooperative society. In the real world, societies that may look very similar in their endowments, access to technology, and physical environments have very different achievements, all because of self-fulfilling beliefs or, as they are often called, *norms* of behavior.[2]

5.2.2 The Tragedy of the Commons

The *tragedy of the commons* refers to the conflict over scarce resources that results from the tension between individual selfish interests and the common good; the concept was popularized by Hardin (1968). The central idea has proven useful for understanding how we have come to be on the brink of several environmental catastrophes.

Hardin introduces the hypothetical example of a pasture shared by local herders. Each herder wants to maximize his yield, increasing his herd size whenever possible. Each additional animal has a positive effect for its herder, but the cost of that extra animal, namely degradation of the overall quality of the pasture, is shared by all the other herders. As a consequence the individual incentive for each herder is to grow his own herd, and in the end this scenario causes tremendous losses for everyone. To those trained in economics, it is yet another example of the distortion that results from the "free-rider" problem. It should also remind you of the Prisoner's Dilemma, in which individuals driven by selfish incentives cause pain to the group.

In the course of his essay, Hardin develops the theme, drawing on examples of such latter-day commons as the atmosphere, oceans, rivers, fish stocks, national parks, advertising, and even parking meters. A major theme running throughout the essay is the growth of human populations, with the earth's resources being a global commons. (Given that this example concerns the addition of extra "animals" to the population, it is the closest to his original analogy.)

---

2. For an excellent exposition of the role that beliefs play in societies, see Greif (2006). The idea that coordinated changes are needed for developing countries to move out of poverty and into industrial growth dates back to Paul Rosenstein-Rodan's theory of the "big push," which is explored further in Murphy et al. (1989).

Let's put some game theoretic analysis behind this story. Imagine that there are n players, say firms, in the world, each choosing how much to produce. Their production activity in turn consumes some of the clean air that surrounds our planet. There is a total amount of clean air equal to K, and any consumption of clean air comes out of this common resource. Each player i chooses his own consumption of clean air for production, $k_i \geq 0$, and the amount of clean air left is therefore $K - \sum_{i=1}^{n} k_i$. The benefit of consuming an amount $k_i \geq 0$ gives player i a benefit equal to $\ln(k_i)$, and no other player benefits from i's choice. Each player also enjoys consuming the remainder of the clean air, giving each a benefit $\ln(K - \sum_{i=1}^{n} k_i)$. Hence the payoff for player i from the choice $k = (k_1, k_2, \ldots, k_n)$ is equal to

$$v_i(k_i, k_{-i}) = \ln(k_i) + \ln\left(K - \sum_{j=1}^{n} k_j\right). \tag{5.1}$$

To solve for a Nash equilibrium we can compute the best-response correspondences for each player and then find a strategy profile for which all the best-response functions are satisfied together. This is an important point that warrants further emphasis. We know that given k_{-i}, player i will want to choose an element in $BR_i(k_{-i})$. Hence if we find some profile of choices $(k_1^*, k_2^*, \ldots, k_n^*)$ for which $k_i^* = BR_i(k_{-i}^*)$ for all $i \in N$ then this must be a Nash equilibrium.

This means that if we derive all n best-response correspondences, and it turns out that they are functions (unique best responses), then we have a system of n equations, one for each player's best-response function, with n unknowns, the choices of each player. Solving this system will yield a Nash equilibrium. To get player i's best-response function (and we will verify that it is a function), we write down the first-order condition of his payoff function:

$$\frac{\partial v_i(k_i, k_{-i})}{\partial k_i} = \frac{1}{k_i} - \frac{1}{K - \sum_{j=1}^{n} k_j} = 0,$$

and this gives us player i's best response function,[3]

$$BR_i(k_{-i}) = \frac{K - \sum_{j \neq i} k_j}{2}.$$

We therefore have n such equations, one for each player, and if we substitute the choice k_i instead of $BR_i(k_{-i})$ we get the n equations with n unknowns that need to be solved.

We proceed to solve the equilibrium for two players and leave the n-player case as exercise 5.7 at the end of the chapter. Letting $k_i(k_j)$ be the best response of player i, we have two best-response equations:

$$k_1(k_2) = \frac{K - k_2}{2} \qquad \text{and} \qquad k_2(k_1) = \frac{K - k_1}{2}.$$

These two equations are plotted in Figure 5.1. As the figure illustrates, the more player j consumes, the less player i wants to consume. In particular if player 2 consumes nothing (effectively not existing), then player 1 will consume $k_1 = \frac{K}{2}$, and as player

3. Of course, we are implicitly assuming that $\sum_{j=1}^{n} k_j \leq K$.

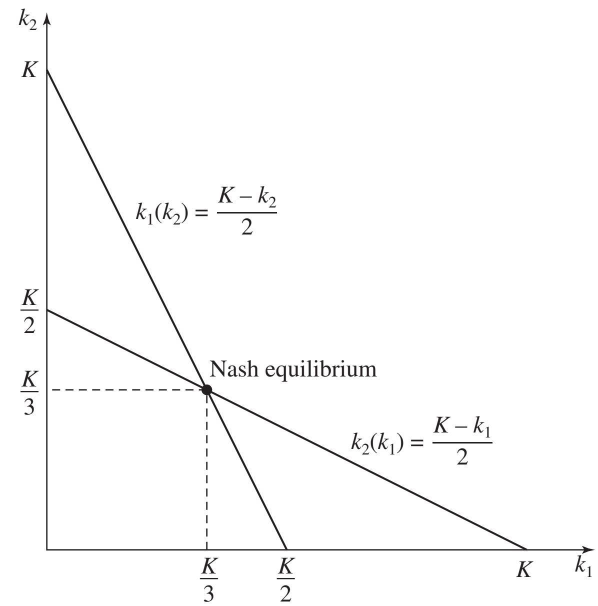

FIGURE 5.1 Best-response functions: two-player tragedy of the commons.

2's consumption increases toward K, player 1's consumption decreases toward zero. If we solve the two best-response functions simultaneously, we find the unique Nash equilibrium, which has both players playing $k_1 = k_2 = \frac{K}{3}$, as shown in Figure 5.1.

Now we can ask whether this two-player society could do better. Is consuming $\frac{K}{3}$ for each player too much or too little? The right way to answer these questions is using the Pareto criterion: can we find another consumption profile that will make everyone better off? If we can, we can compare that with the Nash equilibrium to answer this question. To find such a profile we'll use a little trick: we will maximize the sum of all the payoff functions, which we can think of as the "world's payoff function," $w(k_1, k_2)$. I won't go into the moral justification for using this approach, but it will turn out to be a useful tool.[4] The function we are maximizing is, therefore,

$$\max_{k_1,k_2} \; w(k_1, k_2) = \sum_{i=1}^{2} v_i(k_1, k_2) = \sum_{i=1}^{2} \ln(k_i) + 2 \ln \left(K - \sum_{i=1}^{2} k_i \right).$$

The first-order conditions for this problem are

$$\frac{\partial w(k_1, k_2)}{\partial k_1} = \frac{1}{k_1} - \frac{2}{K - k_1 - k_2} = 0$$

and

$$\frac{\partial w(k_1, k_2)}{\partial k_2} = \frac{1}{k_2} - \frac{2}{K - k_1 - k_2} = 0.$$

4. In general, maximizing the sum of utility functions, or maximizing total welfare, will result in a Pareto-optimal outcome, but it need not be the only one. In this example, this maximization gives us the symmetric Pareto-optimal consumption profile because the payoff function of each player is concave in his own consumption with $\frac{\partial v_i}{\partial k_i} > 0$, $\frac{\partial^2 v_i}{\partial k_i^2} < 0$, and $\lim_{k_i \to 0} \frac{\partial v_i}{\partial k_i} = \infty$.

Solving these two equations simultaneously will result in Pareto-optimal choices for k_1 and k_2. The unique solution to these two equations yields $k_1 = k_2 = \frac{K}{4}$, which means that from a social point of view the Nash equilibrium has the two players each consuming too much clean air. Indeed they would both be better off if each consumed $k_i = \frac{K}{4}$ instead of $k_i = \frac{K}{3}$.[5] In exercise 5.7 at the end of the chapter you are asked to show the consequences of having more than two players.

Thus, as Hardin puts it, giving people the freedom to make choices may make them all worse off than if those choices were somehow regulated. Of course the counterargument is whether we can trust a regulator to keep things under control; if not, the question remains which is the better of the two evils—an answer that game theory cannot offer!

5.2.3 Cournot Duopoly

Let's revisit the Cournot game with demand $P = a - bq$ and cost functions $c_i(q_i) = c_i q_i$ for firms $i \in \{1, 2\}$. The maximization problem that firm i faces when it believes that its opponent chooses quantity q_j is

$$\max_{q_i} \; v_i(q_i, q_j) = (a - bq_i - bq_j)q_i - c_i q_i.$$

Recall that the best-response function for each firm is given by the first-order condition, so that

$$BR_i(q_j) = \frac{a - bq_j - c_i}{2b}.$$

This means that each firm chooses quantities as follows:

$$q_1 = \frac{a - bq_2 - c_1}{2b} \quad \text{and} \quad q_2 = \frac{a - bq_1 - c_2}{2b}. \tag{5.2}$$

A pair of quantities (q_1, q_2) that are *mutual best responses* will be a Cournot-Nash equilibrium, which occurs when we solve both best-response functions (5.2) simultaneously. The best-response functions shown in Figure 5.2 depict the special case we solved earlier, in which $a = 100$, $b = 1$, and $c_1 = c_2 = 10$, in which case the unique Nash equilibrium is $q_1 = q_2 = 30$.

Notice that the Nash equilibrium coincides with the unique strategies that survive IESDS and rationalizability, which is the conclusion of proposition 5.1. An exercise that is left for you (exercise 5.8) is to explore the Cournot model with many firms.

5. To see this we can calculate Δv_i, the difference between a player's payoff when we maximize total surplus $\left(\text{which we solved as } k_i = \frac{K}{4}\right)$ and his Nash equilibrium payoff:

$$\begin{aligned}\Delta v_i &= \ln\left(\frac{K}{4}\right) + \ln\left(\frac{K}{2}\right) - \ln\left(\frac{K}{3}\right) - \ln\left(\frac{K}{3}\right)\\ &= \ln(K) - \ln(4) + \ln(K) - \ln(2) - \ln(K) + \ln(3) - \ln(K) + \ln(3)\\ &= 2\ln(3) - \ln(2) - \ln(4)\\ &= 0.051 > 0.\end{aligned}$$

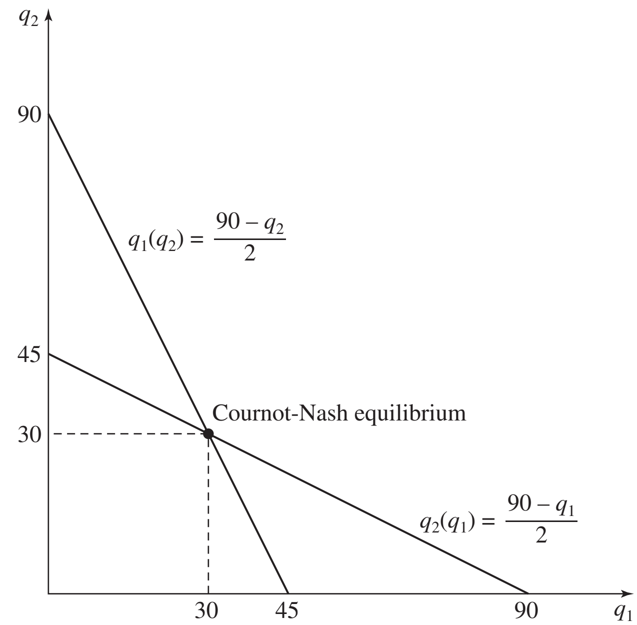

FIGURE 5.2 Cournot duopoly game: best-response functions and Nash equilibrium.

5.2.4 Bertrand Duopoly

The Cournot model assumed that the firms choose quantities and the market price adjusts to clear the demand. However, one can argue that firms often set prices and let consumers choose from where they will purchase, rather than setting quantities and waiting for the market price to equilibrate demand. We now consider the game in which each firm posts a price for otherwise identical goods. This was the situation modeled and analyzed by Joseph Bertrand (1883).

As before, assume that demand is given by $p = 100 - q$ and cost functions $c_i(q_i) = 10q_i$ for firms $i \in \{1, 2\}$. Clearly we would expect all buyers to buy from the firm whose price is the lowest. What happens if there is a tie? Let's assume that the market splits equally between the two firms. This gives us the following normal-form game:

Players: $N = \{1, 2\}$.

Strategy sets: $S_i = [0, \infty]$ for $i \in \{1, 2\}$, and firms choose prices $p_i \in S_i$.

Payoffs: To calculate payoffs we need to know what the quantities will be for each firm. Given our assumption on ties, the quantities are given by

$$q_i(p_i, p_j) = \begin{cases} 100 - p_i & \text{if } p_i < p_j \\ 0 & \text{if } p_i > p_j \\ \frac{100 - p_i}{2} & \text{if } p_i = p_j, \end{cases}$$

which in turn means that the payoff function is given by

$$v_i(p_i, p_j) = \begin{cases} (100 - p_i)(p_i - 10) & \text{if } p_i < p_j \\ 0 & \text{if } p_i > p_j \\ \frac{100 - p_i}{2}(p_i - 10) & \text{if } p_i = p_j. \end{cases}$$

Now that the description of the game is complete, we can try to calculate the best-response functions of both firms. To do this we will start with a slight modification that is motivated by reality: assume that prices cannot be *any* real number $p \geq 0$, but instead are limited to increments of some small fixed value, say $\varepsilon > 0$, which implies that the strategy (price) sets are $S_i = \{0, \varepsilon, 2\varepsilon, 3\varepsilon, \ldots\}$. For example, if we are considering cents as the price increment, so that $\varepsilon = 0.01$, then the strategy set will be $S_i = \{0, 0.01, 0.02, 0.03, \ldots\}$. We will very soon see what happens when this small denomination ε becomes very small and approaches zero.

We derive the best response of a firm by exhausting the relevant situations that it can face. It is useful to start with the situation in which only one monopolistic firm is in the market. We can calculate the monopoly price, which is the price that would maximize a single firm's profits if there were no competitors. This would be obtained by maximizing $v_i(p) = pq - 10q = (100 - p)(p - 10)$ and the first-order condition is $110 - 2p = 0$, resulting in an optimal price of $p = 55$, in a quantity of $q = 45$, and in profits equal to \$2025.

Let us now turn back to the duopoly with two firms and consider the case in which $p_j > 55$. It is easy to see that firm i can act as if there was no competition: just set the monopoly price of 55 and get the whole market. Hence we conclude that if $p_j > 55$ then the best response of firm i is to set $p_i = 55$.

It is also easy to see that in the case in which $p_j < 10$ then the best response of firm i is to set a price that is higher than that set by firm j. If it charges a price $p_i \leq p_j$ then it will sell a positive quantity at a price that is lower than its costs, causing firm i to lose money. If it charges a price $p_i > p_j$ then it sells nothing and loses nothing.

Now consider the case $55 \geq p_j \geq 10.02$. Firm i can choose one of three options: either set $p_i > p_j$ and get nothing, set $p_i = p_j$ and split the market, or set $p_i < p_j$ and get the whole market. It is not too hard to establish that firm i wants to just undercut firm j and capture the whole market, a goal that can be accomplished by setting a price of $p_i = p_j - 0.01$. To see this, observe that if $p_j > 10.01$ then by setting $p_i = p_j$ firm i gets $v_i = \frac{100-p_j}{2}(p_j - 10)$, while if it sets $p_i = p_j - 0.01$ it will get $v_i' = (100 - (p_j - 0.01))((p_j - 0.01) - 10)$. We can calculate the difference between the two as follows:

$$\begin{aligned}\Delta v_i = v_i' - v_i &= (100 - (p_j - 0.01))((p_j - 0.01) - 10) - \frac{100 - p_j}{2}(p_j - 10) \\ &= 55.02p_j - 0.5p_j^2 - 501.1.\end{aligned}$$

It is easy to check that Δv_i is positive at $p_j = 10.02$ (it is equal to 0.0002). To see that it is positive for all values of $p \in [10.02, 55]$, we show that Δv_i has a positive derivative for any $p_j \in [10.02, 55]$, which implies that this difference grows even more positive as p_j increases in this domain. That is,

$$\frac{d\Delta v_i}{dp_j} = 55.02 - p_j > 0 \quad \text{for all } p < 55.02.$$

Thus we conclude that when $p_j \in [10.02, 55]$ the best response of firm i is to charge \$0.01 less, that is, $p_i = p_j - 0.01$.

To complete the analysis, we have to explore two final cases: $p_j = 10.01$ and $p_j = 10$. The three options to consider are setting $p_i = p_j$, $p_i > p_j$, or $p_i < p_j$. When $p_j = 10.01$ then undercutting j's price means setting $p_i = 10$, which gives i zero profits and is the same as setting $p_i > p_j$. Thus the best response is setting $p_i = p_j = 10.01$ and splitting the market with very low profits. Finally, if $p_j = 10$

then any choice of price $p_i \geq p_j$ will yield firm i zero profits, whereas setting $p_i < p_j$ causes losses. Therefore any price $p_i \geq p_j$ is a best response when $p_j = 10$.

In summary we calculated:

$$BR_i(p_j) = \begin{cases} 55 & \text{if } p_j > 55 \\ p_j - 0.01 & \text{if } 55 \geq p_j \geq 10.02 \\ 10.01 & \text{if } p_j = 10.01 \\ p_i \in \{10, 10.01, 10.02, 10.03, \ldots\} & \text{if } p_j = 10. \end{cases}$$

Now given that firm j's best response is exactly symmetric, it should not be hard to see that there are two Nash equilibria that follow immediately from the form of the best-response functions: *The best response* to $p_j = 10.01$ is $BR_i(10.01) = 10.01$, and *a best response* to $p_j = 10$ is $p_i = 10$ or $10 \in BR_i(10)$. Thus the two Nash equilibria are

$$(p_1, p_2) \in \{(10, 10), (10.01, 10.01)\}.$$

It is worth pausing here for a moment to address a common point of confusion, which often arises when a player has more than one best response to a certain action of his opponents. In this example, when $p_2 = 10$, player 1 is indifferent regarding *any price at or above 10 that he chooses:* if he splits the market with $p_1 = 10$ he gets half the market with no profits, and if he sets $p_1 > 10$ he gets no customers and has no profits. One may be tempted to jump to the following conclusion: if player 2 is choosing $p_2 = 10$ then any choice of $p_1 \geq 10$ together with $p_2 = 10$ will be a Nash equilibrium. This is incorrect! It is true that player 1 is playing a best response with any one of his choices, but if $p_1 > 10$ then $p_2 = 10$ is *not* a best response of player 2 to p_1, as we can observe from the foregoing analysis.

Comparing the outcome of the Bertrand game to that of the Cournot game is an interesting exercise. Notice that when firms choose quantities (Cournot), the unique Nash equilibrium is $q_1 = q_2 = 30$. A quick calculation shows that for the aggregate quantity of $q = q_1 + q_2 = 60$ we get a demand price of $p = 40$ and each firm makes a profit of \$900. When instead these firms compete on prices, the two possible equilibria have either zero profits when both choose $p_1 = p_2 = \$10$ or negligible profits (about \$0.45) when they each choose $p_1 = p_2 = \$10.01$. Interestingly, for both the Cournot and Bertrand games, if we had only one player then he would maximize profits by choosing the monopoly price (or quantity) of \$55 (or 45 units) and earn a profit of \$2025.

The message of this analysis is quite striking: one firm may have monopoly power, but when we let one more firm compete, and they compete on prices, then the market will behave competitively—if both choose a price of \$10 then price will equal marginal costs! Notice that if we add a third and a fourth firm this will not change the outcome; prices will have to be \$10 (or practically the same at \$10.01) for all firms in the Nash (Bertrand) equilibrium. This is not the case for Cournot competition, in which firms manage to obtain some market power as long as the number of firms is not too large.

A quick observation should lead you to realize that if we let ε be smaller than \$0.01, the conclusions we reached earlier will be sustained, and we will have two Nash equilibria, one with $p_1 = p_2 = 10$ and one with $p_1 = p_2 = 10 + \varepsilon$. Clearly these two equilibria become "closer" in profits as ε becomes smaller, and they converge to each other as ε approaches zero.

It turns out that if we assume that prices can be chosen as *any* real number, we get a very "clean" result: the unique Bertrand-Nash equilibrium will have prices equal to marginal costs, implying a competitive outcome. We prove this for the more general symmetric case in which $c_i(q_i) = cq_i$ for both firms and demand is equal to $p = a - bq$ with $a > c$.[6]

Proposition 5.2 *For $\varepsilon = 0$ (prices can be any real number) there is a unique Nash equilibrium: $p_1 = p_2 = c$.*

Proof First note that in any equilibrium $p_i \geq c$ for both firms—otherwise at least one firm offering a price lower than c will lose money (pay the consumers to take its goods!). We therefore need to show that $p_i > c$ cannot be part of any equilibrium. We can see this in two steps:

1. If $p_1 = p_2 = \hat{p} > c$ then each would benefit from lowering its price to some price $\hat{p} - \varepsilon$ (ε very small) and get the whole market for almost the same price.
2. If $p_1 > p_2 \geq c$ then player 2 would want to deviate to $p_1 - \varepsilon$ (ε very small) and earn higher profits.

It is easy to see that $p_1 = p_2 = c$ is an equilibrium: Firm i's best response to $p_j = c$ is $p_i(c) \geq c$. That is, any price at or above marginal costs c is a best response to the other player charging c. Hence $p_1 = p_2 = c$ is the unique equilibrium because each is playing a best response to the other's choice and neither wants to deviate to a different price. ■

We will now see an interesting variation of the Bertrand game. Assume that $c_i(q_i) = c_i q_i$ represents the cost of each firm as before. Now, however, let $c_1 = 1$ and $c_2 = 2$ so that the two firms are not identical: firm 1 has a cost advantage. Let the demand still be $p = 100 - q$.

Now consider the case with discrete price jumps with $\varepsilon = 0.01$. The firms are not symmetric in that firm 1 has a lower marginal cost than firm 2, and unlike the example in which both had the same costs we cannot have a Nash equilibrium in which both charge the same price. To see this, imagine that $p_1 = p_2 = 2.00$. Firm 2 has no incentive to deviate, but this is no longer true for firm 1, which will be happy to cut its price by 0.01. We know that firm 2 will not be willing to sell at a price below $p = 2$, so one possible Nash equilibrium (you are asked to find more in exercise 5.12 at the end of this chapter) is

$$(p_1^*, p_2^*) = (1.99, 2.00).$$

Now we can ask ourselves what happens if $\varepsilon = 0$. If we would think of using a "limit" approach to answer this question then we may expect a result similar to the one we saw before: if we focus on the equilibrium pair $(p_1^*, p_2^*) = (2 - \varepsilon, 2)$ then as $\varepsilon \to 0$ we must get the Nash equilibrium $p_1 = p_2 = 2$. But is this really an equilibrium? Interestingly, the answer is no! To see this, consider the best response of firm 1. Its payoff function is not continuous when firm 2 offers a price of 2 (or any other positive price). The profit function of firm 1, as a function of p_1 when $p_2 = 2$, is depicted in Figure 5.3. The figure first draws the profits of firm 1 as if it were a monopolist with

6. The condition $a > c$ is necessary for firms to be able to produce positive quantities and not lose money.

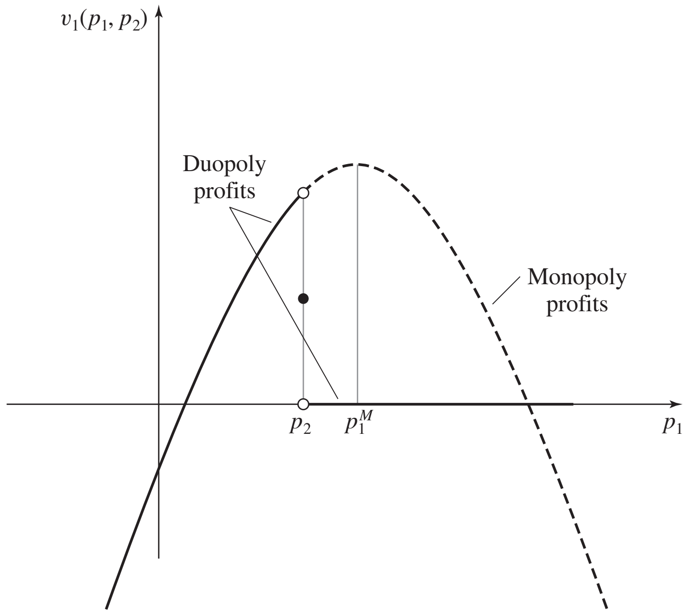

FIGURE 5.3 The profit function in the Bertrand duopoly game.

no competition (the hump-shaped curve), and if this were the case it would charge its monopoly price $p_1^M = 50.5$.[7] If firm 2 charges more than the monopoly price, this will have no impact on the choice of firm 1—it will still charge the monopoly price. If, however, firm 2 charges a price p_2 that is less than the monopoly price then there is a **discontinuity in the payoff function** of firm 1: as its price p_1 approaches p_2 from below, its profits rise. However, when it hits p_2 exactly then it will split the market and see its profits drop by half. Algebraically the profit function of firm 1 when $p_2 = 2$ is given by

$$v_i(p_1, 2) = \begin{cases} (100 - p_1)(p_1 - 1) & \text{if } p_1 < 2 \\ \frac{(100-p_1)(p_1-1)}{2} & \text{if } p_1 = 2 \\ 0 & \text{if } p_1 > 2. \end{cases}$$

This discontinuity causes firm 1 to *not* have a well-defined best response correspondence when $p_2 < 50.5$. Firm 1 wants to set a price *as close to* p_2 *as it can,* but it does not want to reach p_2 because then it splits the market and experiences a sizable decrease in profits. Once its price goes above p_2 then firm 1's profits drop further to zero. Indeed the consequence must be that a Nash equilibrium does not exist precisely because firm 1 does not have a "well-behaved" payoff function.

To see this directly, first observe that there cannot be a Nash equilibrium with $p_i \geq p_j > 2$: firm i would want to deviate to some $p \in (2, p_j)$. Second, observe that there cannot be a Nash equilibrium with $p_i \leq p_j < 1$: firm i would want to deviate to any $p > p_j$. Hence if there is a Nash equilibrium, it must have the prices between 1 and 2. Similar to the first observation, within this range we cannot have $p_1 \geq p_2$ (firm

7. The maximization here is for the profit function $v_1 = (100 - p)(p - c_1)$, where $c_1 = 1$.

1 would want to deviate just slightly below p_2), and similar to the second observation, we cannot have $p_2 \le p_1$ (firm 2 would want to deviate to any price above p_1).

This problem of payoff discontinuity is one that we will avoid in the remainder of this book precisely because it will often lead to problems of nonexistence of equilibrium. We need to remember that it is a problem that disappears when we have some discreteness in the actions of players, and if we think that our choice of continuous strategies is one of convenience (to use calculus for optimization), then we may feel comfortable enough ignoring such anomalies. We discuss the existence of Nash equilibria further in Section 6.4.

Remark It is worth pointing out an interesting difference between the Cournot game and the Bertrand game. In the Cournot game the best-response function of each player is downward sloping. That is, the more player j produces, the lower is the best-response quantity of player i. In the Bertrand game, however, for prices between marginal costs (equal to 10 in the leading example) and the monopoly price (equal to 45), the higher the price set by player j, the higher is the best-response price of player i. These differences have received some attention in the literature. Games for which the best response of one player decreases in the choice of the other, like the Cournot game, are called games with **strategic substitutes.** Another example of a game with strategic substitutes is the tragedy of the commons. In contrast, games for which the best response of one player increases in the choice of the other, like the Bertrand game, are called games with **strategic complements.** Another example of a game with strategic complements appears in exercise 5.10 at the end of the chapter. There are several interesting insights to be derived from distinguishing between strategic substitutes and strategic complements. For a nice example see Fudenberg and Tirole (1984).

5.2.5 Political Ideology and Electoral Competition

Given a population of citizens who vote for political candidates, how should candidates position themselves along the political spectrum? One view of the world is that a politician cares only about representing his true beliefs, and that drives the campaign. Another more cynical view is that politicians care only about getting elected and hence will choose a platform that maximizes their chances. This is precisely the view taken in the seminal model introduced by Hotelling (1929).[8]

To consider a simple variant of Hotelling's original model, imagine that there are two politicians, each caring only about being elected. There are 101 citizens, each labeled by an integer $-50, -49, \ldots, 0, \ldots, +49, +50$. Each citizen has political preferences: for simplicity let's call the "-50" citizen the most "left"-leaning citizen and the "$+50$" citizen the most "right"-leaning citizen.

Each candidate i chooses his platform as a policy $a_i \in \{-50, -49, \ldots, 0, \ldots, +49, +50\}$ so that each policy is associated with the citizen for whom this policy is ideal. Each citizen chooses the candidate whose platform is closest to his political preferences. For example, if candidate 1 chooses platform $a_1 = -15$ while candidate 2 chooses platform $a_2 = +22$, then all the citizens at or above $+22$ will surely vote

8. Hotelling's main object of analysis was competition between firms. However, he did also discuss the example of electoral competition, yielding important insights into rational choice–based political science.

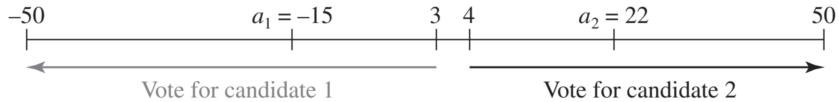

FIGURE 5.4 The Hotelling model of voting behavior.

for candidate 2, all the citizens at or below -15 will surely vote for candidate 1, and those in between will split between the candidates. In particular those citizens at or below 3 will vote for candidate 1, while those at or above 4 will vote for candidate 2. The reason is that -15 is at a distance of 18 away from citizen 3, while $+22$ is at a distance of 19 away. This is shown in Figure 5.4.

The outcome is determined by majority rule: if a majority of citizens vote for candidate i then he wins. Since there is an odd number of voters, unless someone is indifferent, one candidate will always win. In the event that a citizen is indifferent between the candidates then the citizen tosses a coin to determine for whom to vote. Assume that our candidates want to win, so that they prefer winning to a tie, and they prefer a tie to losing. Now consider the best response of player i. If player j chooses a policy $a_j > 0$ then by choosing $a_i = a_j$ or $a_i = -a_j$ there will be a tie.[9] By choosing $a_i > a_j$ or $a_i < -a_j$, player i will surely lose, while by choosing $a_i \in [-a_j + 1, a_j - 1]$, player i will win, so any platform in this interval is a best response to $a_j > 0$. Similarly a symmetric argument implies that any platform in the interval $a_i \in [a_j + 1, -a_j - 1]$ is a best response to $a_j < 0$. Observe that the best response to $a_j = 0$ is zero.[10] Thus we can write the best-response correspondence of each player as

$$BR_i(a_j) = \begin{cases} [a_j + 1, -a_j - 1] & \text{if } a_j < 0 \\ 0 & \text{if } a_j = 0 \\ [-a_j + 1, a_j - 1] & \text{if } a_j > 0. \end{cases}$$

From here it is easy to see that there is a unique Nash equilibrium, $a_1 = a_2 = 0$, implying that both candidates position their platforms smack in the middle of the political spectrum! Indeed, as Hotelling (1929, p. 54) wrote, "The competition for votes between the Republican and Democratic parties does not lead to a clear drawing of issues, and adoption of two strongly contrasted positions between which the voter may choose. Instead, each party strives to make its platform as much like the other's as possible." A fine insight in 1929, and one that is echoed frequently today.

This simple example is related to a powerful result known as the **median voter theorem.** It states that if voters are different from one another along a single-dimensional "preference" line, as in Hotelling's model, and if each prefers his own political location, with other platforms being less and less attractive the farther away

9. For example, if candidate 1 chooses $a_1 = -10$ and candidate 2 chooses $a_2 = 10$ then all the "positive" citizens will vote for candidate 2, all the "negative" ones will vote for candidate 1, and citizen 0 is indifferent so he may choose not to vote. In this event there will be a tie. If we force every citizen to vote one way or another then there will be no ties because we assumed an odd number of citizens.

10. Notice that the best response to $a_j = -1$ or $a_j = +1$ is also unique and equal to $a_i = 0$, because in these cases the interval previously identified collapses to 0.

they fall to either side of that location,[11] then the political platform located at the median voter will defeat *any other platform* in a simple majority vote. The theorem was first articulated by Black (1948), and it received prominence in Downs's famous 1957 book. Nevertheless one can see how the seed of the idea had been planted as far back as Hotelling's formalization of spatial competition.

Remark In the more common representation of the Hotelling competition model, the citizens are a continuum of voters, say, given by the interval $A = [\underline{a}, \overline{a}]$, with distribution $F(a)$ to determine the distribution of each political preference. We define the **median voter** as that voter a^m for which $\Pr\{a \leq a^m\} = \frac{1}{2}$. $\left(\text{For example, if } F(\cdot) \text{ is uniform then } a^m = \frac{\underline{a}+\overline{a}}{2}.\right)$ The best-response correspondence is similar to that given earlier, but for all choices of player j that are not equal to a^m player i's best response is an open interval. For example, if $a_j > a^m$, and if $a' < a^m$ is an "opposite" policy that ties with a_j, then the best response of player i is to choose any platform that lies in the open interval (a', a_j). If, however, $a_j = a^m$ then player i's unique best response is to choose $a_i = a^m$, implying that both candidates choosing a^m is the unique Nash equilibrium. You are asked to prove this in exercise 5.15.

5.3 Summary

- Any strategy profile for which players are playing mutual best responses is a Nash equilibrium, making this equilibrium concept self-enforcing.
- If a profile of strategies is the unique survivor of IESDS or is the unique rationalizable profile of strategies then it is a Nash equilibrium.
- If a profile of strategies is a Nash equilibrium then it must survive IESDS and it must be rationalizable, but not every strategy that survives IESDS or that is rationalizable is a Nash equilibrium.
- Nash equilibrium analysis can shed light on phenomena such as the tragedy of the commons and the nature of competition in markets and in politics.

5.4 Exercises

5.1 **Prove Proposition 5.1.**

5.2 **Weak Dominance:** A strategy $s^W \in S$ is a weakly dominant strategy equilibrium if $s_i^W \in S_i$ is a weakly dominant strategy for all $i \in N$, that is, if $v_i(s_i^W, s_{-i}) \geq v_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and for all $s_{-i} \in S_{-i}$. Provide an example of a game for which there is a weakly dominant strategy equilibrium as well as another Nash equilibrium.

5.3 **Nash and IESDS:** Consider a two-player game with m pure strategies for each player that can be represented by an $m \times m$ matrix.

  a. Show that if $m = 2$ and the game has a unique pure-strategy Nash equilibrium then this is the unique strategy profile that survives IESDS.

11. This condition on preferences is called that of "single-peaked" preferences: if we draw the utility function of some voter with his political "bliss point" being $a \in A$, the utility is highest at a and declines in both directions, hence the single peak at a.

b. Show that if $m = 3$ and the game has a unique pure-strategy equilibrium then it may not be the only strategy profile that survives IESDS.

5.4 **Splitting Pizza:** You and a friend are in an Italian restaurant, and the owner offers both of you a free eight-slice pizza under the following condition. Each of you must simultaneously announce how many slices you would like; that is, each player $i \in \{1, 2\}$ names his desired amount of pizza, $0 \le s_i \le 8$. If $s_1 + s_2 \le 8$ then the players get their demands (and the owner eats any leftover slices). If $s_1 + s_2 > 8$, then the players get nothing. Assume that you each care only about how much pizza you individually consume, and the more the better.

a. Write out or graph each player's best-response correspondence.
b. What outcomes can be supported as pure-strategy Nash equilibria?

5.5 **Public Good Contribution:** Three players live in a town, and each can choose to contribute to fund a streetlamp. The value of having the streetlamp is 3 for each player, and the value of not having it is 0. The mayor asks each player to contribute either 1 or nothing. If at least two players contribute then the lamp will be erected. If one player or no players contribute then the lamp will not be erected, in which case any person who contributed will not get his money back.

a. Write out or graph each player's best-response correspondence.
b. What outcomes can be supported as pure-strategy Nash equilibria?

5.6 **Hawk-Dove:** The following game has been widely used in evolutionary biology to understand how fighting and display strategies by animals could coexist in a population. For a typical Hawk-Dove game there are resources to be gained (e.g., food, mates, territories), denoted as v. Each of two players can choose to be aggressive, as Hawk (H), or compromising, as Dove (D). If both players choose H then they split the resources but lose some payoff from injuries, denoted as k. Assume that $k > \frac{v}{2}$. If both choose D then they split the resources but engage in some display of power that carries a display cost d, with $d < \frac{v}{2}$. Finally, if player i chooses H while j chooses D then i gets all the resources while j leaves with no benefits and no costs.

a. Describe this game in a matrix.
b. Assume that $v = 10$, $k = 6$, and $d = 4$. What outcomes can be supported as pure-strategy Nash equilibria?[12]

5.7 **The *n*-Player Tragedy of the Commons:** Suppose there are n players in the tragedy of the commons example in Section 5.2.2.

a. Find the Nash equilibrium of this game. How does n affect the Nash outcome?
b. Find the socially optimal outcome with n players. How does n affect this outcome?
c. How does the Nash equilibrium outcome compare to the socially efficient outcome as n approaches infinity?

5.8 **The *n*-Firm Cournot Model:** Suppose there are n firms in the Cournot oligopoly model. Let q_i denote the quantity produced by firm i, and let

12. In the evolutionary biology literature, the analysis performed is of a very different nature. Instead of considering the Nash equilibrium analysis of a static game, the analysis is a dynamic one in which successful strategies "replicate" in a large population. This analysis is part of a methodology called evolutionary game theory. For more on the subject see Gintis (2000).

$Q = q_i + \cdots + q_n$ denote the aggregate production. Let $P(Q)$ denote the market clearing price (when demand equals Q) and assume that the inverse demand function is given by $P(Q) = a - Q$, where $Q \le a$. Assume that firms have no fixed cost and that the cost of producing quantity q_i is cq_i (all firms have the same marginal cost, and assume that $c < a$).

a. Model this as a normal-form game.
b. What is the Nash (Cournot) equilibrium of the game in which firms choose their quantities simultaneously?
c. What happens to the equilibrium price as n approaches infinity? Is this familiar?

5.9 **Tragedy of the Roommates:** You and your $n - 1$ roommates each have five hours of free time you could spend cleaning your apartment. You all dislike cleaning, but you all like having a clean apartment: each person's payoff is the total hours spent (by everyone) cleaning, minus a number c times the hours spent (individually) cleaning. That is,

$$v_i(s_1, s_2, \ldots, s_n) = -c \cdot s_i + \sum_{j=1}^{n} s_j.$$

Assume everyone chooses simultaneously how much time to spend cleaning.

a. Find the Nash equilibrium if $c < 1$.
b. Find the Nash equilibrium if $c > 1$.
c. Set $n = 5$ and $c = 2$. Is the Nash equilibrium Pareto efficient? If not, can you find an outcome in which everyone is better off than in the Nash equilibrium outcome?

5.10 **Synergies:** Two division managers can invest time and effort in creating a better working relationship. Each invests $e_i \ge 0$, and if both invest more then both are better off, but it is costly for each manager to invest. In particular the payoff function for player i from effort levels (e_i, e_j) is $v_i(e_i, e_j) = (a + e_j)e_i - e_i^2$.

a. What is the best-response correspondence of each player?
b. In what way are the best-response correspondences different from those in the Cournot game? Why?
c. Find the Nash equilibrium of this game and argue that it is unique.

5.11 **Wasteful Shipping Costs:** Consider two countries, A and B, each with a monopolist that owns the only active coal mine in the country. Let firm 1 be the firm located in country A and firm 2 the one in country B. Let q_i^j, $i \in \{1, 2\}$ and $j \in \{A, B\}$ denote the quantity that firm i sells in country j. Consequently let $q_i = q_i^A + q_i^B$ be the total quantity produced by firm $i \in \{1, 2\}$ and let $q^j = q_1^j + q_2^j$ be the total quantity sold in country $j \in \{A, B\}$. The demand for coal in countries A and B is then given respectively by

$$p^j = 90 - q^j, \quad j \in \{A, B\},$$

and the cost of production for each firm is given by

$$c_i(q_i) = 10q_i, \quad i \in \{1, 2\}.$$

a. Assume that the countries do not have a trade agreement and, in fact, that the importation of coal into either country is prohibited.

This implies that $q_2^A = q_1^B = 0$ is set as a political constraint. What quantities q_1^A and q_2^B will both firms produce?

Now assume that the two countries sign a free-trade agreement that allows foreign firms to sell in each country without any tariffs. There are, however, shipping costs. If firm i sells quantity q_i^j in the foreign country (i.e., firm 1 selling in B or firm 2 selling in A) then shipping costs are equal to $10q_i^j$. Assume further that *each firm* chooses a pair of quantities q_i^A, q_i^B simultaneously, $i \in \{1, 2\}$, so that a profile of actions consists of four quantity choices.

b. Model this as a normal-form game and find a Nash equilibrium of the game you described. Is it unique?

Now assume that before the game you described in (b) is played the research department of firm 1 discovers that shipping coal on the existing vessels causes the release of pollutants. If the firm would disclose this report to the World Trade Organization (WTO) then the WTO would prohibit the use of the existing ships. Instead a new shipping technology would be offered that would increase shipping costs to $40q_i^j$ (instead of $10q_i^j$ as given earlier).

c. Would firm 1 be willing to release the information to the WTO? Justify your answer with an equilibrium analysis.

5.12 **Asymmetric Bertrand:** Consider the Bertrand game with $c_1(q_1) = q_1$ and $c_2(q_2) = 2q_2$ and demand equal to $p = 100 - q$, in which firms must choose prices in increments of \$0.01. We have seen in Section 5.2.4 that one possible Nash equilibrium is $(p_1^*, p_2^*) = (1.99, 2.00)$.

a. Show that there are other Nash equilibria for this game.
b. How many Nash equilibria does this game have?

5.13 **Comparative Economics:** Two high-tech firms (1 and 2) are considering a joint venture. Each firm i can invest in a novel technology and can choose a level of investment $x_i \in [0, 5]$ at a cost of $c_i(x_i) = x_i^2/4$ (think of x_i as how many hours to train employees or how much capital to spend for R&D labs). The revenue of each firm depends on both its investment and the other firm's investment. In particular if firms i and j choose x_i and x_j, respectively, then the **gross revenue** to firm i is

$$R(x_i, x_j) = \begin{cases} 0 & \text{if } x_i < 1 \\ 2 & \text{if } x_i \geq 1 \quad \text{and } x_j < 2 \\ x_i \cdot x_j & \text{if } x_i \geq 1 \quad \text{and } x_j \geq 2. \end{cases}$$

a. Write down mathematically and draw the **profit function** (gross revenue minus costs) of firm i as a function of x_i for three cases: (i) $x_j < 2$, (ii) $x_j = 2$, and (iii) $x_j = 4$.
b. What is the best-response function of firm i?
c. It turns out that there are two *identical* pairs of such firms; that is, the description applies to both pairs. One pair is in Russia, where coordination is hard to achieve and businesspeople are very cautious, and the other pair is in Germany, where coordination is common and

businesspeople expect their partners to go the extra mile. You learn that the Russian firms are earning significantly lower profits than the German firms, despite the fact that their technologies are identical. Can you use Nash equilibrium analysis to shed light on this dilemma? If so, be precise and use your previous analysis to do so.

5.14 **Negative Ad Campaigns:** Each one of two political parties can choose to buy time on commercial radio shows to broadcast negative ad campaigns against its rival. These choices are made simultaneously. Government regulations forbid a party from buying more than 2 hours of negative campaign time, so that each party cannot choose an amount of negative campaigning above 2 hours. Given a pair of choices (a_1, a_2), the payoff of party i is given by the following function: $v_i(a_1, a_2) = a_i - 2a_j + a_i a_j - (a_i)^2$.

a. What is the normal-form representation of this game?
b. What is the best-response function for each party?
c. What is the pure-strategy Nash equilibrium? Is it unique?
d. If the parties could sign a binding agreement on how much to campaign, what levels would they choose?

5.15 **Hotelling's Continuous Model:** Consider Hotelling's model, in which citizens are a continuum of voters on the interval $A = [-a, a]$, with uniform distribution $U(a)$.

a. What is the best response of candidate i if candidate j is choosing $a_j > 0$?
b. Show that the unique Nash equilibrium is $a_1 = a_2 = 0$.
c. Show that for a general distribution $F(\cdot)$ over $[-a, a]$ the unique Nash equilibrium is where each candidate chooses the policy associated with the median voter.

5.16 **Hotelling's Price Competition:** Imagine a continuum of potential buyers, located on the line segment [0, 1], with uniform distribution. (Hence the "mass" or quantity of buyers in the interval $[a, b]$ is equal to $b - a$.) Imagine two firms, players 1 and 2, who are located at each end of the interval (player 1 at the 0 point and player 2 at the 1 point). Each player i can choose its price p_i, and each customer goes to the vendor who offers him the highest value. However, price alone does not determine the value; distance is important as well. In particular each buyer who buys the product from player i has a net value of $v - p_i - d_i$, where d_i is the distance between the buyer and vendor i and represents the transportation costs of buying from vendor i. Thus buyer $a \in [0, 1]$ buys from 1 and not 2 if $v - p_1 - d_1 > v - p_2 - d_2$ and if buying is better than getting zero. (Here $d_1 = a$ and $d_2 = 1 - a$. The buying choice would be reversed if the inequality were reversed.) Finally, assume that the cost of production is zero.

a. What is the best-response function of each player?
b. Assume that $v = 1$. What is the Nash equilibrium? Is it unique?
c. Now assume that the transportation costs are $\frac{1}{2}d_i$, so that a buyer buys from 1 if and only if $v - p_1 - \frac{1}{2}d_1 > v - p_2 - \frac{1}{2}d_2$. Write down the best-response function of each player and solve for the Nash equilibrium.
d. Following your analysis in (c), imagine that transportation costs are αd_i, with $\alpha \in [0, 1]$. What happens to the Nash equilibrium as $\alpha \to 0$? What is the intuition for this result?

5.17 **To Vote or Not to Vote:** Two candidates, D and R, are running for mayor in a town with n residents. A total of $0 < d < n$ residents support candidate D, while the remainder, $r = n - d$, support candidate R. The value for each resident for having his candidate win is 4, for having him tie is 2, and for having him lose is 0. Going to vote costs each resident 1.

a. Let $n = 2$ and $d = 1$. Write down this game as a matrix and solve for the Nash equilibrium.
b. Let $n > 2$ be an even number and let $d = r = \frac{n}{2}$. Find all the Nash equilibria.
c. Assume now that the cost of voting is equal to 3. How does your answer to (a) and (b) change?

5.18 **Political Campaigning:** Two candidates are competing in a political race. Each candidate i can spend $s_i \geq 0$ on ads that reach out to voters, which in turn increases the probability that candidate i wins the race. Given a pair of spending choices (s_1, s_2), the probability that candidate i wins is given by $\frac{s_i}{s_1+s_2}$. If neither spends any resources then each wins with probability $\frac{1}{2}$. Each candidate values winning at a payoff of $v > 0$, and the cost of spending s_i is just s_i.

a. Given two spend levels (s_1, s_2), write down the expected payoff of a candidate i.
b. What is the function that represents each player's best-response function?
c. Find the unique Nash equilibrium.
d. What happens to the Nash equilibrium spending levels if v increases?
e. What happens to the Nash equilibrium levels if player 1 still values winning at v but player 2 values winning at kv, where $k > 1$?

6

Mixed Strategies

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

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

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

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

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

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

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

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

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

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

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

6.1 Strategies, Beliefs, and Expected Payoffs

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

6.1.1 Finite Strategy Sets

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

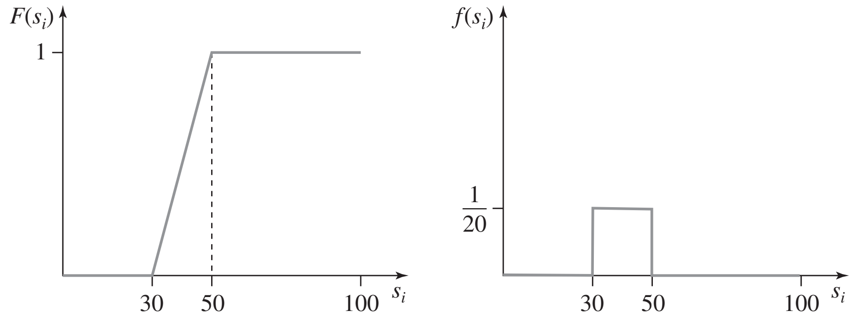

FIGURE 6.1 A continuous mixed strategy in the Cournot game.

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

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

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

6.1.2 Continuous Strategy Sets

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

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

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

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

These two functions are depicted in Figure 6.1.

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

6.1.3 Beliefs and Mixed Strategies

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

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

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

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

6.1.4 Expected Payoffs

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

For more on this topic see Section 19.4.4.

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

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

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

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

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

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

6.2 Mixed-Strategy Nash Equilibrium

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

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

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

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

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

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

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

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

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

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

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

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

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

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

6.2.1 Example: Matching Pennies

Consider the Matching Pennies game,

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

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

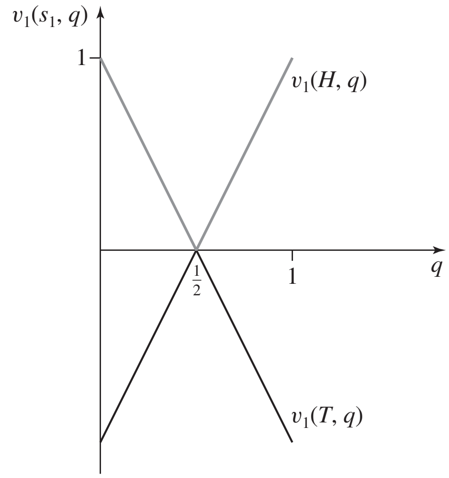

FIGURE 6.2 Expected payoffs for player 1 in the Matching Pennies game.

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

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

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

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

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

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

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

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

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

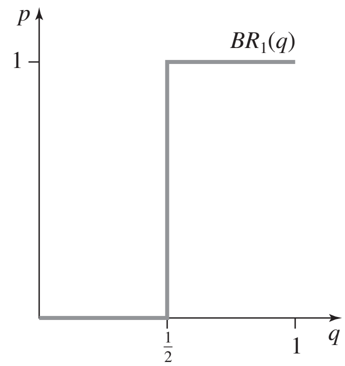

FIGURE 6.3 Player 1's best-response correspondences in the Matching Pennies game.

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

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

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

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

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

and this implies that player 2's best response is

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

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

We know from proposition 6.1 that when player 1 is mixing between H and T, both with positive probability, then it must be the case that his payoffs from H and from T are identical. This, it turns out, imposes a restriction on the behavior of player 2, given by the choice of q. Player 1 is willing to mix between H and T if and only if $v_1(H, q) = v_1(T, q)$, which will hold if and only if $q = \frac{1}{2}$. This is the way in which *the indifference of player 1 imposes a restriction on player 2:* only when player 2 is playing $q = \frac{1}{2}$ will player 1 be willing to mix between his actions H and T. Similarly player 2 is willing to mix between H and T only when $v_2(p, H) = v_2(p, T)$, which is true only when $p = \frac{1}{2}$. We have come to the conclusion of our quest for a Nash equilibrium in this game. We can see that there is indeed a pair of mixed strategies that form a Nash equilibrium, and these are precisely when $(p, q) = \left(\frac{1}{2}, \frac{1}{2}\right)$.

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

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

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

6.2.2 Example: Rock-Paper-Scissors

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

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

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

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

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

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

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

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

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

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

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

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

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

6.2.3 Multiple Equilibria: Pure and Mixed

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

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

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

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

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

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

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

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

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

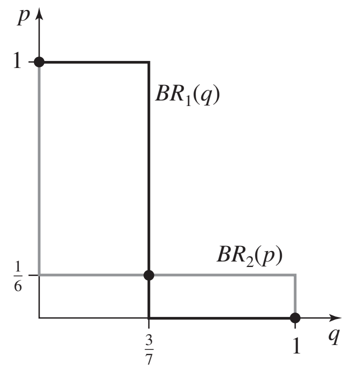

FIGURE 6.4 Best-response correspondences and Nash equilibria.

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

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

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

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

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

6.3 IESDS and Rationalizability Revisited

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

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

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

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

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

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

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

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

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

	L	C	R		$\left(0, \frac{1}{2}, \frac{1}{2}\right)$
U	5, 1	1, 4	1, 0		2
M	3, 2	0, 0	3, 5	Player 2's expected payoff from mixing C and $R \Rightarrow$	2.5
D	4, 3	4, 4	0, 3		3.5

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

		Player 2	
		C	R
	U	1, 4	1, 0
Player 1	M	0, 0	3, 5
	D	4, 4	0, 3
	$\left(0, \frac{1}{2}, \frac{1}{2}\right)$	2	1.5

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

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

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

This last 2×2 game cannot be further reduced.

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

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

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

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

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

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

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

Hence for two-player games the set of strategies that survive IESDS is the same as the set of strategies that are rationalizable. Proving this is not that simple and is beyond the scope of this text. The eager and interested reader is encouraged to read Chapter 2 of Fudenberg and Tirole (1991), and the daring reader can refer to the original research papers by Bernheim (1984) and Pearce (1984), which simultaneously introduced the concept of rationalizability.[7]

6.4 Nash's Existence Theorem

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

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

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

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

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

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

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

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

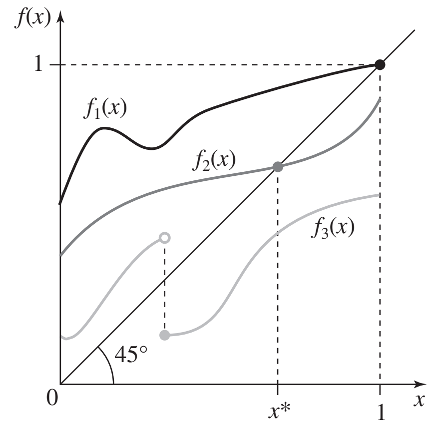

FIGURE 6.5 Brouwer's fixed-point theorem.

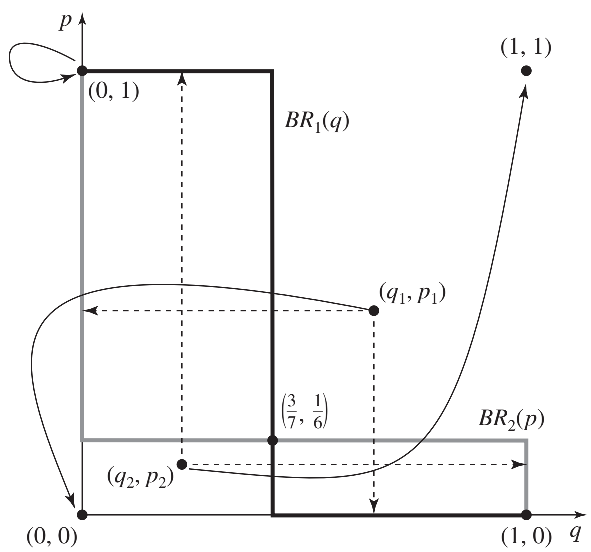

FIGURE 6.6 Mapping mixed strategies using the best-response correspondence.

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

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

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

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

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

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

and

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

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

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

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

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

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

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

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

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

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

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

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

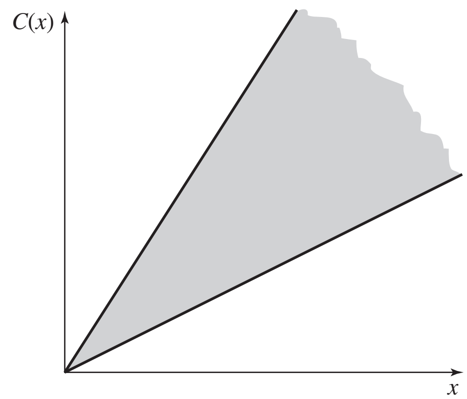

FIGURE 6.7 A correspondence with a closed graph.

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

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

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

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

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

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

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

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

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

6.5 Summary

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

6.6 Exercises

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

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

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

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

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

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

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

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

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

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

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

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

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

6.7 **Grad School Competition:** Two students sign up to prepare an honors thesis with a professor. Each can invest time in his own project: either no time, one week, or two weeks (these are the only three options). The cost of time is 0 for no time, and each week costs 1 unit of payoff. The more time a student puts in the better his work will be, so that if one student puts in more time than the other there will be a clear "leader." If they put in the same amount of time then their thesis projects will have the same quality. The professor, however, will give out only one grade of A. If there is a clear leader then he will get the A, while if they are equally good then the professor will toss a fair coin to decide who gets the A. The other student will get a B. Since both wish to continue on to graduate school, a grade of A is worth 3 while a grade of B is worth 0.

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

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

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

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

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

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

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

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

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

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

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