

Game Theory

Game Theory

Drew Fudenberg
Jean Tirole

Published in India by



Ane Books Pvt. Ltd.

New Delhi ♦ Chennai ♦ Mumbai
Kolkata ♦ Thiruvananthapuram ♦ Pune ♦ Bengaluru

We begin with a simple, informal example of a game. Rousseau, in his *Discourse on the Origin and Basis of Equality among Men*, comments:

If a group of hunters set out to take a stag, they are fully aware that they would all have to remain faithfully at their posts in order to succeed; but if a hare happens to pass near one of them, there can be no doubt that he pursued it without qualm, and that once he had caught his prey, he cared very little whether or not he had made his companions miss theirs.¹

To make this into a game, we need to fill in a few details. Suppose that there are only two hunters, and that they must decide simultaneously whether to hunt for stag or for hare. If both hunt for stag, they will catch one stag and share it equally. If both hunt for hare, they each will catch one hare. If one hunts for hare while the other tries to take a stag, the former will catch a hare and the latter will catch nothing. Each hunter prefers half a stag to one hare.

This is a simple example of a game. The hunters are the players. Each player has the choice between two strategies: hunt stag and hunt hare. The payoff to their choice is the prey. If, for instance, a stag is worth 4 "utils" and a hare is worth 1, then when both players hunt stag each has a payoff of 2 utils. A player who hunts hare has payoff 1, and a player who hunts stag by himself has payoff 0.

What prediction should one make about the outcome of Rousseau's game? Cooperation—both hunting stag—is an equilibrium, or more precisely a "Nash equilibrium," in that neither player has a unilateral incentive to change his strategy. Therefore, stag hunting seems like a possible outcome of the game. However, Rousseau (and later Waltz (1959)) also warns us that cooperation is by no means a foregone conclusion. If each player believes the other will hunt hare, each is better off hunting hare himself. Thus, the noncooperative outcome—both hunting hare—is also a Nash equilibrium, and without more information about the context of the game and the hunters' expectations it is difficult to know which outcome to predict.

This chapter will give precise definitions of a "game" and a "Nash equilibrium," among other concepts, and explore their properties. There are two nearly equivalent ways of describing games: the *strategic* (or *normal*) form and the *extensive* form.² Section 1.1 develops the idea of the strategic form and of dominated strategies. Section 1.2 defines the solution concept of Nash equilibrium, which is the starting point of most applications of game theory. Section 1.3 offers a first look at the question of when Nash equilibria exist; it is the one place in this chapter where powerful mathematics is used.

1. Quoted by Ordeshook (1986).

2. Historically, the term "normal form" has been standard, but many game theorists now prefer to use "strategic form," as this formulation treats the players' strategies as primitives of the model.

$$\begin{aligned}
 u_1(\sigma_1, \sigma_2) &= \frac{1}{3}(0 \cdot 4 + \frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 6) + \frac{1}{3}(0 \cdot 2 + \frac{1}{2} \cdot 8 + \frac{1}{2} \cdot 3) \\
 &\quad + \frac{1}{3}(0 \cdot 3 + \frac{1}{2} \cdot 9 + \frac{1}{2} \cdot 2) \\
 &= \frac{11}{2}.
 \end{aligned}$$

$$\text{Similarly, } u_2(\sigma_1, \sigma_2) = \frac{27}{6}.$$

1.1.2 Dominated Strategies

Is there an obvious prediction of how the game described in figure 1.1 should be played? Note that, no matter how player 1 plays, R gives player 2 a strictly higher payoff than M does. In formal language, strategy M is *strictly dominated*. Thus, a "rational" player 2 should not play M. Furthermore, if player 1 knows that player 2 will not play M, then U is a better choice than M or D. Finally, if player 2 knows that player 1 knows that player 2 will not play M, then player 2 knows that player 1 will play U, and so player 2 should play L.

The process of elimination described above is called *iterated dominance*, or, more precisely, *iterated strict dominance*.³ In section 2.1 we give a formal definition of iterated strict dominance, as well as an application to an economic example. The reader may worry at this stage that the set of strategies that survive iterated strict dominance depends on the order in which strategies are eliminated, but this is not the case. (The key is that, if strategy s_i is strictly worse than strategy s'_i against all opponents' strategies in some set D , then strategy s_i is strictly worse than strategy s'_i against all opponents' strategies in any subset of D . Exercise 2.1 asks for a formal proof.)

Next, consider the game illustrated in figure 1.2. Here player 1's strategy M is not dominated by U, because M is better than U if player 2 moves R; and M is not dominated by D, because M is better than D when 2 moves L. However, if player 1 plays U with probability $\frac{1}{2}$ and D with probability $\frac{1}{2}$, he is guaranteed an expected payoff of $\frac{1}{2}$ regardless of how player 2 plays, which exceeds the payoff of 0 he receives from M. Hence, a pure strategy

	L	R
U	2,0	-1,0
M	0,0	0,0
D	-1,0	2,0

Figure 1.2

3. Iterated elimination of weakly dominated strategies has been studied by Luce and Raiffa (1957), Fahrquarson (1969), and Moulin (1979).

may be strictly dominated by a mixed strategy even if it is not strictly dominated by any pure strategy.

We will frequently wish to discuss varying the strategy of a single player i while holding the strategies of his opponents fixed. To do so, we let

$$s_{-i} \in S_{-i}$$

denote a strategy selection for all players but i , and write

$$(s'_i, s_{-i})$$

for the profile

$$(s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_I).$$

Similarly, for mixed strategies we let

$$(\sigma'_i, \sigma_{-i}) = (\sigma_1, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_I).$$

Definition 1.1 Pure strategy s_i is *strictly dominated* for player i if there exists $s'_i \in \Sigma_i$ such that

$$u_i(\sigma'_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}. \quad (1.1)$$

The strategy s_i is *weakly dominated* if there exists a σ'_i such that inequality 1.1 holds with weak inequality, and the inequality is strict for at least one s_{-i} .

Note that, for a given s_i , strategy σ'_i satisfies inequality 1.1 for all pure strategies s_{-i} of the opponents if and only if it satisfies inequality 1.1 for all mixed strategies σ_{-i} as well, because player i 's payoff when his opponents play mixed strategies is a convex combination of his payoffs when his opponents play pure strategies.

So far we have considered dominated pure strategies. It is easy to see that a mixed strategy that assigns positive probability to a dominated pure strategy is dominated. However, a mixed strategy may be strictly dominated even though it assigns positive probability only to pure strategies that are not even weakly dominated. Figure 1.3 gives an example. Playing U with probability $\frac{1}{2}$ and M with probability $\frac{1}{2}$ gives expected payoff

	L	R
U	1,3	-2,0
M	-2,0	1,3
D	0,1	0,1

Figure 1.3

	L	R
U	8, 10	-100, 9
D	7, 6	6, 5

Figure 1.4

$-\frac{1}{2}$ regardless of player 2's play and so is strictly dominated by playing D, even though neither U nor M is dominated.

When a game is solvable by iterated strict dominance in the sense that each player is left with a single strategy, as in figure 1.1, the unique strategy profile obtained is an obvious candidate for the prediction of how the game will be played. Although this candidate is often a good prediction, this need not be the case, especially when the payoffs can take on extreme values. When our students have been asked how they would play the game illustrated in figure 1.4, about half have chosen D even though iterated dominance yields (U, L) as the unique solution. The point is that although U is better than D when player 2 is certain not to use the dominated strategy R, D is better than U when there is a 1-percent chance that player 2 plays R. (The same casual empiricism shows that our students in fact do always play L.) If the loss to (U, R) is less extreme, say only -1 , then almost all players 1 choose U, as small fears about R matter less. This example illustrates the role of the assumptions that payoffs and the strategy spaces are common knowledge (as they were in this experiment) and that "rationality," in the sense of not playing a strictly dominated strategy, is common knowledge (as apparently was not the case in this experiment). The point is that the analysis of some games, such as the one illustrated in figure 1.4, is very sensitive to small uncertainties about the behavioral assumptions players make about each other. This kind of "robustness" test—testing how the theory's predictions change with small changes in the model—is an idea that will return in chapters 3, 8, and 11.

At this point we can illustrate a major difference between the analysis of games and the analysis of single-player decisions: In a decision, there is a single decision maker, whose only uncertainty is about the possible moves of "nature," and the decision maker is assumed to have fixed, exogenous beliefs about the probabilities of nature's moves. In a game, there are several decision makers, and the expectations players have about their opponents' play are not exogenous. One implication is that many familiar comparative-statics conclusions from decision theory do not extend once we take into account the way a change in the game may change the actions of *all* players.

Consider for example the game illustrated in figure 1.5. Here player 1's dominant strategy is U, and iterated strict dominance predicts that the

	L	R
U	1, 3	4, 1
D	0, 2	3, 4

Figure 1.5

	L	R
U	-1, 3	2, 1
D	0, 2	3, 4

Figure 1.6

solution is (U, L). Could it help player 1 to change the game and *reduce* his payoffs if U occurs by 2 utils, which would result in the game shown in figure 1.6? Decision theory teaches that such a change would not help, and indeed it would not if we held player 2's action fixed at L. Thus, player 1 would not benefit from this reduction in payoff if it were done without player 2's knowledge. However, if player 1 could arrange for this reduction to occur, and to become known to player 2 before player 2 chose his action, player 1 would indeed benefit, for then player 2 would realize that D is player 1's dominant choice, and player 2 would play R, giving player 1 a payoff of 3 instead of 1.

As we will see, similar observations apply to changes such as decreasing a player's choice set or reducing the quality of his information: Such changes cannot help a player in a fixed decision problem, but in a game they may have beneficial effects on the play of opponents. This is true both when one is making predictions using iterated dominance and when one is studying the equilibria of a game.

1.1.3 Applications of the Elimination of Dominated Strategies

In this subsection we present two classic games in which a *single* round of elimination of dominated strategies reduces the strategy set of each player to a single pure strategy. The first example uses the elimination of strictly dominated strategies, and the second uses the elimination of weakly dominated strategies.

Example 1.1: Prisoner's Dilemma

One round of the elimination of strictly dominated strategies gives a unique answer in the famous "prisoner's dilemma" game, depicted in figure 1.7. The story behind the game is that two people are arrested for a crime. The police lack sufficient evidence to convict either suspect and consequently

	C	D
C	1, 1	-1, 2
D	2, -1	0, 0

Figure 1.7

need them to give testimony against each other. The police put each suspect in a different cell to prevent the two suspects from communicating with each other. The police tell each suspect that if he testifies against (doesn't cooperate with) the other, he will be released and will receive a reward for testifying, provided the other suspect does not testify against him. If neither suspect testifies, both will be released on account of insufficient evidence, and no rewards will be paid. If one testifies, the other will go to prison; if both testify, both will go to prison, but they will still collect rewards for testifying. In this game, both players simultaneously choose between two actions. If both players cooperate (C) (do not testify), they get 1 each. If they both play noncooperatively (D, for defect), they obtain 0. If one cooperates and the other does not, the latter is rewarded (gets 2) and the former is punished (gets -1). Although cooperating would give each player a payoff of 1, self-interest leads to an inefficient outcome with payoffs 0. (To readers who feel this outcome is not reasonable, our response is that their intuition probably concerns a different game—perhaps one where players “feel guilty” if they defect, or where they fear that defecting will have bad consequences in the future. If the game is played repeatedly, other outcomes can be equilibria; this is discussed in chapters 4, 5, and 9.)

Many versions of the prisoner's dilemma have appeared in the social sciences. One example is moral hazard in teams. Suppose that there are two workers, $i = 1, 2$, and that each can “work” ($s_i = 1$) or “shirk” ($s_i = 0$). The total output of the team is $4(s_1 + s_2)$ and is shared equally between the two workers. Each worker incurs private cost 3 when working and 0 when shirking. With “work” identified with C and “shirk” with D, the payoff matrix for this moral-hazard-in-teams game is that of figure 1.7, and “work” is a strictly dominated strategy for each worker.

Exercise 1.7 gives another example where strict dominance leads to a unique solution: that of a mechanism for deciding how to pay for a public good.

Example 1.2: Second-Price Auction

A seller has one indivisible unit of an object for sale. There are I potential buyers, or bidders, with valuations $0 \leq v_1 \leq \dots \leq v_I$ for the object, and these valuations are common knowledge. The bidders simultaneously submit bids $s_i \in [0, +\infty)$. The highest bidder wins the object and pays the second bid (i.e., if he wins ($s_i > \max_{j \neq i} s_j$), bidder i has utility $u_i =$

$v_i - \max_{j \neq i} s_j$), and the other bidders pay nothing (and therefore have utility 0). If several bidders bid the highest price, the good is allocated randomly among them. (The exact probability determining the allocation is irrelevant because the winner and the losers have the same surplus, i.e., 0.)

For each player i the strategy of bidding his valuation ($s_i = v_i$) weakly dominates all other strategies. Let $r_i \equiv \max_{j \neq i} s_j$. Suppose first that $s_i > v_i$. If $r_i \geq s_i$, bidder i obtains utility 0, which he would get by bidding v_i . If $r_i \leq v_i$, bidder i obtains utility $v_i - r_i$, which again is what he would get by bidding v_i . If $v_i < r_i < s_i$, then bidder i has utility $v_i - r_i < 0$; if he were to bid v_i , his utility would be 0. The reasoning is similar for $s_i < v_i$: When $r_i \leq s_i$ or $r_i \geq v_i$, the bidder's utility is unchanged when he bids v_i instead of s_i . However, if $s_i < r_i < v_i$, the bidder forgoes a positive utility by underbidding.

Thus, it is reasonable to predict that bidders bid their valuation in the second-price auction. Therefore, bidder i wins and has utility $v_i - v_{i-1}$. Note also that because bidding one's valuation is a dominant strategy, it does not matter whether the bidders have information about one another's valuations. Hence, if bidders know their own valuation but do not know the other bidders' valuations (see chapter 6), it is still a dominant strategy for each bidder to bid his valuation.

1.2 Nash Equilibrium[†]

Unfortunately, many if not most games of economic interest are not solvable by iterated strict dominance. In contrast, the concept of a Nash-equilibrium solution has the advantage of existing in a broad class of games.

1.2.1 Definition of Nash Equilibrium

A Nash equilibrium is a profile of strategies such that each player's strategy is an optimal response to the other players' strategies.

Definition 1.2 A mixed-strategy profile σ^* is a *Nash equilibrium* if, for all players i ,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*) \text{ for all } s_i \in S_i. \quad (1.2)$$

A pure-strategy Nash equilibrium is a pure-strategy profile that satisfies the same conditions. Since expected utilities are “linear in the probabilities,” if a player uses a nondegenerate mixed strategy in a Nash equilibrium (one that puts positive weight on more than one pure strategy) he must be indifferent between all pure strategies to which he assigns positive probability. (This linearity is why, in equation 1.2, it suffices to check that no player has a profitable pure-strategy deviation.)

A Nash equilibrium is *strict* (Harsanyi 1973b) if each player has a unique best response to his rivals' strategies. That is, s_i^* is a strict equi-

librium if and only if it is a Nash equilibrium and, for all i and all $s_i \neq s_i^*$,

$$u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*).$$

By definition, a strict equilibrium is necessarily a pure-strategy equilibrium. Strict equilibria remain strict when the payoff functions are slightly perturbed, as the strict inequalities remain satisfied.^{4,5}

Strict equilibria may seem more compelling than equilibria where players are indifferent between their equilibrium strategy and a nonequilibrium response, as in the latter case we may wonder why players choose to conform to the equilibrium. Also, strict equilibria are robust to various small changes in the nature of the game, as is discussed in chapters 11 and 14. However, strict equilibria need not exist, as is shown by the "matching pennies" game of example 1.6 below: The unique equilibrium of that game is in (nondegenerate) mixed strategies, and no (nondegenerate) mixed-strategy equilibrium can be strict.⁶ (Even pure-strategy equilibria need not be strict; an example is the profile (D, R) in figure 1.18 when $\lambda = 0$.)

To put the idea of Nash equilibrium in perspective, observe that it was implicit in two of the first games to have been studied, namely the Cournot (1838) and Bertrand (1883) models of oligopoly. In the Cournot model, firms simultaneously choose the quantities they will produce, which they then sell at the market-clearing price. (The model does not specify how this price is determined, but it is helpful to think of it being chosen by a Walrasian auctioneer so as to equate total output and demand.) In the Bertrand model, firms simultaneously choose prices and then must produce enough output to meet demand after the price choices become known. In each model, equilibrium is determined by the condition that all firms choose the action that is a best response to the anticipated play of their opponents. It is common practice to speak of the equilibria of these two models as "Cournot equilibrium" and "Bertrand equilibrium," respectively, but it is more helpful to think of them as the *Nash* equilibria of the two different games. We show below that the concepts of "Stackelberg equi-

4. Harsanyi called this "strong" equilibrium; we use the term "strict" to avoid confusion with "strong equilibrium" of Aumann 1959—see note 11.

5. An equilibrium is *quasi-strict* if each pure-strategy best response to one's rivals' strategies belongs to the support of the equilibrium strategy: $\{\sigma_i^*\}_{i \in I}$ is a quasi-strict equilibrium if it is a Nash equilibrium and if, for all i and s_i ,

$$u_i(s_i, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*) \Rightarrow \sigma_i^*(s_i) > 0.$$

The equilibrium in matching pennies is quasi-strict, but some games have equilibria that are not quasi-strict. The game in figure 1.18b for $\lambda = 0$ has two Nash equilibria, (U, L) and (D, R). The equilibrium (U, L) is strict, but the equilibrium (D, R) is not even quasi-strict. Harsanyi (1973b) has shown that, for "almost all games," all equilibria are quasi-strict (that is, the set of all games that possess an equilibrium that is not quasi-strict is a closed set of measure 0 in the Euclidean space of strategic-form payoff vectors).

6. Remember that in a mixed-strategy equilibrium a player must receive the same expected payoff from every pure strategy he assigns positive probability.

librium" and "open-loop equilibrium" are also best thought of as shorthand ways of referring to the equilibria of different games.

Nash equilibria are "consistent" predictions of how the game will be played, in the sense that if all players predict that a particular Nash equilibrium will occur then no player has an incentive to play differently. Thus, a Nash equilibrium, and only a Nash equilibrium, can have the property that the players can predict it, predict that their opponents predict it, and so on. In contrast, a prediction that any fixed non-Nash profile will occur implies that at least one player will make a "mistake," either in his prediction of his opponents' play or (given that prediction) in his optimization of his payoff.

We do not maintain that such mistakes never occur. In fact, they may be likely in some special situations. But predicting them requires that the game theorist know more about the outcome of the game than the participants know. This is why most economic applications of game theory restrict attention to Nash equilibria.

The fact that Nash equilibria pass the test of being consistent predictions does not make them good predictions, and in situations it seems rash to think that a precise prediction is available. By "situations" we mean to draw attention to the fact that the likely outcome of a game depends on more information than is provided by the strategic form. For example, one would like to know how much experience the players have with games of this sort, whether they come from a common culture and thus might share certain expectations about how the game will be played, and so on.

When one round of elimination of strictly dominated strategies yields a unique strategy profile $s^* = (s_1^*, \dots, s_n^*)$, this strategy profile is necessarily a Nash equilibrium (actually the unique Nash equilibrium). This is because any strategy $s_i \neq s_i^*$ is necessarily strictly dominated by s_i^* . In particular,

$$u_i(s_i, s_{-i}^*) < u_i(s_i^*, s_{-i}^*).$$

Thus, s^* is a pure-strategy Nash equilibrium (indeed a strict equilibrium). In particular, not cooperating is the unique Nash equilibrium in the prisoner's dilemma of example 1.1.⁷

We show in section 2.1 that the same property holds for iterated dominance. That is, if a single strategy profile survives iterated deletion of strictly dominated strategies, then it is the unique Nash equilibrium of the game.

Conversely, any Nash-equilibrium strategy profile must put weight only on strategies that are not strictly dominated (or, more generally, do not survive iterated deletion of strictly dominated strategies), because a player

7. The same reasoning shows that if there exists a single strategy profile surviving one round of deletion of weakly dominated strategies, this strategy profile is a Nash equilibrium. So, bidding one's valuation in the second-price auction (example 1.2) is a Nash equilibrium.

could increase his payoff by replacing a dominated strategy with one that dominates it. However, Nash equilibria may assign positive probability to weakly dominated strategies.

1.2.2 Examples of Pure-Strategy Equilibria

Example 1.3: Cournot Competition

We remind the reader of the Cournot model of a duopoly producing a homogeneous good. The strategies are quantities. Firm 1 and firm 2 simultaneously choose their respective output levels, q_i , from feasible sets $Q_i = [0, \infty)$, say. They sell their output at the market-clearing price $p(q)$, where $q = q_1 + q_2$. Firm i 's cost of production is $c_i(q_i)$, and firm i 's total profit is then

$$u_i(q_1, q_2) = q_i p(q) - c_i(q_i).$$

The feasible sets Q_i and the payoff functions u_i determine the strategic form of the game. The "Cournot reaction functions" $r_1: Q_2 \rightarrow Q_1$ and $r_2: Q_1 \rightarrow Q_2$ specify each firm's optimal output for each fixed output level of its opponent. If the u_i are differentiable and strictly concave, and the appropriate boundary conditions are satisfied,⁸ we can solve for these reaction functions using the first-order conditions. For example, $r_2(\cdot)$ satisfies

$$p(q_1 + r_2(q_1)) + p'(q_1 + r_2(q_1))r_2(q_1) - c_2'(r_2(q_1)) = 0. \quad (1.3)$$

The intersections (if any exist) of the two reaction functions r_1 and r_2 are the Nash equilibria of the Cournot game: Neither firm can gain by a change in output, given the output level of its opponent.

For instance, for linear demand ($p(q) = \max(0, 1 - q)$) and symmetric, linear cost ($c_i(q_i) = cq_i$, where $0 \leq c \leq 1$), firm 2's reaction function, given by equation 1.3, is (over the relevant range)

$$r_2(q_1) = (1 - q_1 - c)/2.$$

By symmetry, firm 1's reaction function is

$$r_1(q_2) = (1 - q_2 - c)/2.$$

The Nash equilibrium satisfies $q_2^* = r_2(q_1^*)$ and $q_1^* = r_1(q_2^*)$ or $q_1^* = q_2^* = (1 - c)/3$.

Example 1.4: Hotelling Competition

Consider Hotelling's (1929) model of differentiation on the line. A linear city of length 1 lies on the abscissa of a line, and consumers are uniformly

8. The "appropriate boundary conditions" refer to sufficient conditions for the optimal reaction of each firm to be in the interior of the feasible set Q_i . For example, if all positive outputs are feasible ($Q_i = [0, +\infty)$), it suffices that $p(q) - c_i'(0) > 0$ for all q (which, in general, implies that $c_i'(0) = 0$) for $r_i(q_i)$ to be strictly positive for all q_i , and $\lim_{q \rightarrow \infty} p(q) + p'(q)q - c_i'(q) < 0$ for $r_i(q_i)$ to be finite for all q_i .

distributed with density 1 along this interval. There are two stores (firms) located at the two extremes of the city, which sell the same physical product. Firm 1 is at $x = 0$, firm 2 at $x = 1$. The unit cost of each store is c . Consumers incur a transportation cost t per unit of distance. They have unit demands and buy one unit if and only if the minimum generalized price (price plus transportation cost) for the two stores does not exceed some large number \bar{s} . If prices are "not too high," the demand for firm 1 is equal to the number of consumers who find it cheaper to buy from firm 1. Letting p_i denote the price of firm i , the demand for firm 1 is given by

$$D_1(p_1, p_2) = x,$$

where

$$p_1 + tx = p_2 + t(1 - x)$$

or

$$D_1(p_1, p_2) = \frac{p_2 - p_1 + t}{2t}$$

and

$$D_2(p_1, p_2) = 1 - D_1(p_1, p_2).$$

Suppose that prices are chosen simultaneously. A Nash equilibrium is a profile (p_1^*, p_2^*) such that, for each player i ,

$$p_i^* \in \arg \max_{p_i} \{(p_i - c)D_i(p_i, p_j^*)\}.$$

For instance, firm 2's reaction curve, $r_2(p_1)$, is given (in the relevant range) by

$$D_2(p_1, r_2(p_1)) + [r_2(p_1) - c] \frac{\partial D_2}{\partial p_2}(p_1, r_2(p_1)) = 0.$$

In our example, the Nash equilibrium is given by $p_1^* = p_2^* = c + t$ (and the above analysis is valid as long as $c + 3t/2 \leq \bar{s}$).

Example 1.5: Majority Voting

There are three players, 1, 2, and 3, and three alternatives, A, B, and C. Players vote simultaneously for an alternative; abstaining is not allowed. Thus, the strategy spaces are $S_i = \{A, B, C\}$. The alternative with the most votes wins; if no alternative receives a majority, then alternative A is selected. The payoff functions are

$$u_1(A) = u_2(B) = u_3(C) = 2,$$

$$u_1(B) = u_2(C) = u_3(A) = 1,$$

cost of inspection (h) must equal the expected wage savings (xw). Hence, $y = g/w$ and $x = h/w$ (both x and y belong to $(0, 1)$).

1.2.4 Multiple Nash Equilibria, Focal Points, and Pareto Optimality

Many games have several Nash equilibria. When this is the case, the assumption that a Nash equilibrium is played relies on there being some mechanism or process that leads all the players to expect the same equilibrium.

One well-known example of a game with multiple equilibria is the "battle of the sexes," illustrated by figure 1.10a. The story that goes with the name "battle of the sexes" is that the two players wish to go to an event together, but disagree about whether to go to a football game or the ballet. Each player gets a utility of 2 if both go to his or her preferred event, a utility of 1 if both go to the other's preferred event, and 0 if the two are unable to agree and stay home or go out individually. Figure 1.10b displays a closely related game that goes by the names of "chicken" and "hawk-dove." (Chapter 4 discusses a related dynamic game that is also called "chicken.") One version of the story here is that the two players meet at a one-lane bridge and each must choose whether to cross or to wait for the other. If both play T (for "tough"), they crash in the middle of the bridge and get -1 each; if both play W (for "weak"), they wait and get 0; if one player chooses T and the other chooses W, then the tough player crosses first, receiving 2, and the weak one receives 1. In the bridge-crossing story, the term "chicken" is used in the colloquial sense of "coward." (Evolutionary biologists call this game "hawk-dove," because they interpret strategy T as "hawk-like" and strategy W as "dove-like.")

Though the different payoff matrices in figures 1.10a and 1.10b describe different sorts of situations, the two games are very similar. Each of them has three equilibria: two in pure strategies, with payoffs $(2, 1)$ and $(1, 2)$, and

9. Building on this result, one can compute the optimal contract, i.e., the w that maximizes the principal's expected payoff

$$v(1-x) - w(1-xy) - hy = v(1-h/w) - w.$$

The optimal wage is thus $w = \sqrt{hv}$ (assuming $\sqrt{hv} > g$). Note that the principal would be better off if he could "commit" to an inspection level. To see this, consider the different game in which the principal plays first and chooses a probability y of inspection, and the agent, after observing y , chooses whether to shirk. For a given w ($> g$), the principal can choose $y = g/w + \epsilon$, where ϵ is positive and arbitrarily small. The agent then works with probability 1, and the principal has (approximately) payoff

$$v - w - hg/w > v(1-h/w) - w.$$

Technically, commitment eliminates the constraint $xw \geq h$ (i.e., that it is *ex post* worthwhile to inspect). (It is crucial that the principal is committed to inspecting with probability y . If the "toss of the coin" determining inspection is not public, the principal has an *ex post* incentive not to inspect, as he knows that the agent works.) This reasoning will become familiar in chapter 3. See chapters 5 and 10 for discussions of how repeated play might make the commitment credible whereas it would not be if the game was played only once.

	B	F
F	0,0	2,1
B	1,2	0,0

a

	T	W
T	-1,-1	2,1
W	1,2	0,0

b

Figure 1.10

one that is mixed. In the battle of the sexes, the mixed equilibrium is that player 1 plays F with probability $\frac{2}{3}$ (and B with probability $\frac{1}{3}$) and player 2 plays B with probability $\frac{2}{3}$ (and F with probability $\frac{1}{3}$). To obtain these probabilities, we solve out the conditions that the players be indifferent between their two pure strategies. So, if x and y denote the probabilities that player 1 plays F and player 2 plays B, respectively, player 1's indifference between F and B is equivalent to

$$0 \cdot y + 2 \cdot (1-y) = 1 \cdot y + 0 \cdot (1-y),$$

or

$$y = \frac{2}{3}.$$

Similarly, for player 2 to be indifferent between B and F it must be the case that

$$0 \cdot x + 2 \cdot (1-x) = 1 \cdot x + 0 \cdot (1-x),$$

or

$$x = \frac{2}{3}.$$

In the chicken game of figure 1.10b, the mixed-strategy equilibrium has players 1 and 2 play tough with probability $\frac{1}{2}$.

If the two players have not played the battle of the sexes before, it is hard to see just what the right prediction might be, because there is no obvious way for the players to coordinate their expectations. In this case we would not be surprised to see the outcome (B, F). (We would still be surprised if (B, F) turned out to be the "right" prediction, i.e., if it occurred almost every time.) However, Schelling's (1960) theory of "focal points" suggests that in some "real-life" situations players may be able to coordinate on a particular equilibrium by using information that is abstracted away by the strategic form. For example, the names of the strategies

	L	R
U	0, 0, 10	-5, -5, 0
D	-5, -5, 0	1, 1, -5
	A	

	L	R
U	-2, -2, 0	-5, -5, 0
D	-5, -5, 0	-1, -1, 5
	B	

Figure 1.12

player 3 chooses matrices. (Harsanyi and Selten (1988) give a closely related example where player 3 moves before players 1 and 2.) This game has two pure-strategy Nash equilibria, (U, L, A) and (D, R, B), and an equilibrium in mixed strategies. Bernheim, Peleg, and Whinston do not consider mixed strategies, so we will temporarily restrict our attention to pure ones. The equilibrium (U, L, A) Pareto-dominates (D, R, B). Is (U, L, A) then the obvious focal point? Imagine that this was the expected solution, and hold player 3's choice fixed. This induces a two-player game between players 1 and 2. In this two-player game, (D, R) is the Pareto-dominant equilibrium! Thus, if players 1 and 2 expect that player 3 will play A, and if they can coordinate their play on their Pareto-preferred equilibrium in matrix A, they should do so, upsetting the "good" equilibrium (U, L, A).

In response to this example, Bernheim, Peleg, and Whinston propose the idea of a coalition-proof equilibrium, as a way of extending the idea of coordinating on the Pareto-dominant equilibrium to games with more than two players.¹¹

To summarize our remarks on multiple equilibria: Although some games have focal points that are natural predictions, game theory lacks a general and convincing argument that a Nash outcome will occur.¹² However, equilibrium analysis has proved useful to economists, and we will focus attention on equilibrium in this book. (Chapter 2 discusses the "rationalizability" notion of Bernheim and Pearce, which investigates the predictions

11. The definition of a coalition-proof equilibrium proceeds by induction on coalition size. First one requires that no one-player coalition can deviate, i.e., that the given strategies are a Nash equilibrium. Then one requires that no two-player coalition can deviate, given that once such a deviation has "occurred" either of the deviating players (but none of the others) is free to deviate again. That is, the two-player deviations must be Nash equilibria of the two-player game induced by holding the strategies of the others fixed. And one proceeds in this way up to the coalition of all players. Clearly (U, L, A) in figure 1.12 is not coalition-proof; brief inspection shows that (D, R, B) is.

Coalition-proof equilibrium is a weakening of Aumann's (1959) "strong equilibrium," which requires that no subset of players, taking the actions of others as given, can jointly deviate in a way that increases the payoffs of all its members. Since this requirement applies to the grand coalition of all players, strong equilibria must be Pareto efficient, unlike coalition-proof equilibria. No strong equilibrium exists in the game of figure 1.12.

12. Aumann (1987) argues that the "Harsanyi doctrine," according to which all players' beliefs must be consistent with Bayesian updating from a common prior, implies that Bayesian rational players must predict a "correlated equilibrium" (a generalization of Nash equilibrium defined in section 2.2).

one can make without invoking equilibrium. As we will see, rationalizability is closely linked to the notion of iterated strict dominance.)

1.2.5 Nash Equilibrium as the Result of Learning or Evolution

To this point we have motivated the solution concepts of dominance, iterated dominance, and Nash equilibrium by supposing that players make their predictions of their opponents' play by introspection and deduction, using their knowledge of the opponents' payoffs, the knowledge that the opponents are rational, the knowledge that each player knows that the others know these things, and so on through the infinite regress implied by "common knowledge."

An alternative approach to introspection for explaining how players predict the behavior of their opponents is to suppose that players extrapolate from their past observations of play in "similar games," either with their current opponents or with "similar"¹³ ones. At the end of this subsection we will discuss how introspection and extrapolation differ in the nature of their assumptions about the players' information about one another.

The idea of using learning-type adjustment processes to explain equilibrium goes back to Cournot, who proposed a process that might lead the players to play the Cournot-Nash equilibrium outputs. In the Cournot adjustment process, players take turns setting their outputs, and each player's chosen output is a best response to the output his opponent chose the period before. Thus, if player 1 moves first in period 0, and chooses q_1^0 , then player 2's output in period 1 is $q_2^1 = r_2(q_1^0)$, where r_2 is the Cournot reaction function defined in example 1.3. Continuing to iterate the process,

$$q_1^2 = r_1(q_2^1) = r_1(r_2(q_1^0)),$$

and so on. This process may settle down to a steady state where the output levels are constant, but it need not do so. If the process does converge to (q_1^*, q_2^*) , then $q_2^* = r_2(q_1^*)$ and $q_1^* = r_1(q_2^*)$, so the steady state is a Nash equilibrium.

If the process converges to a particular steady state for all initial quantities sufficiently close to it, we say that the steady state is *asymptotically stable*. As an example of an asymptotically stable equilibrium, consider the Cournot game where $p(q) = 1 - q$, $c_i(q_i) = 0$, and the feasible sets are $Q_i = [0, 1]$. The reaction curves for this game are $r_i(q_j) = (1 - q_j)/2$, and the unique Nash equilibrium is at the intersection of the reaction curves, which is the point $A = (\frac{1}{3}, \frac{1}{3})$. Figure 1.13 displays the path of the Cournot adjust-

13. Of course the distinction between introspection and extrapolation is not absolute. One might suppose that introspection leads to the idea that extrapolation is likely to work, or conversely that past experience has shown that introspection is likely to make the correct prediction.




Figure 1.13

ment or *tâtonnement* process for the initial condition $q_1^0 = \frac{1}{2}$. The process converges to the Nash equilibrium from every starting point; that is, the Nash equilibrium is globally stable.

Now suppose that the cost and demand functions yield reaction curves as in figure 1.14 (we spare the reader the derivation of such reaction functions from a specification of cost and demand functions). The reaction functions in figure 1.14 intersect at three points, B, C, and D, all of which are Nash equilibria. Now, however, the intermediate Nash equilibrium, C, is not stable, as the adjustment process converges either to B or to D unless it starts at exactly C.

Comparing figures 1.13 and 1.14 may suggest that the question of asymptotic stability is related to the relative slopes of the reaction functions, and this is indeed the case. If the payoff functions are twice continuously differentiable, the slope of firm i 's reaction function is

$$\frac{dr_i}{dq_j} = -\frac{\partial^2 u_i}{\partial q_i \partial q_j} / \frac{\partial^2 u_i}{\partial q_i^2},$$

and a sufficient condition for an equilibrium to be asymptotically stable is that

$$\left| \frac{dr_1}{dq_2} \right| \left| \frac{dr_2}{dq_1} \right| < 1$$

or

$$\frac{\partial^2 u_1}{\partial q_1 \partial q_2} \frac{\partial^2 u_2}{\partial q_1 \partial q_2} < \frac{\partial^2 u_1}{\partial q_1^2} \frac{\partial^2 u_2}{\partial q_2^2}$$

in an open neighborhood of the Nash equilibrium.




Figure 1.14

Technical aside The condition for asymptotic stability when firms react *simultaneously*, instead of alternatively, to their opponent's most recent outputs is the same as the one just described. To see this, suppose that both players simultaneously adjust their quantities each period by choosing a best response to their opponent's output in the previous period. View this as a dynamic process

$$q^t = (q_1^t, q_2^t) = (r_1(q_2^{t-1}), r_2(q_1^{t-1})) \equiv f(q^{t-1}).$$

From the study of dynamical systems (Hirsch and Smale 1974), we know that a fixed point q^* of f is asymptotically stable in this process if all the eigenvalues of $\partial f(q^*)$ have real parts whose absolute value is less than 1. The condition on the slopes of the reaction functions is exactly sufficient to imply that this eigenvalue condition is satisfied. Classic references on the stability of the Cournot adjustment process include Fisher 1961, Hahn 1962, Seade 1980, and Dixit 1986; see Moulin 1986 for a discussion of more recent work and of subtleties that arise with more than two players.

One way to interpret Cournot's adjustment process with either alternating or simultaneous adjustment is that in each period the player making a move expects that his opponent's output in the future will be the same as it is now. Since output in fact changes every period, it may seem more plausible that players base their forecasts on the average value of their opponent's past play, which suggests the alternative dynamic process

$$q_i^t = r_i \left(\sum_{s=0}^{t-1} q_i^s / t \right).$$

This alternative has the added value of converging under a broader set

	L	M	R
U	0,0	4,5	5,4
M	5,4	0,0	4,5
D	4,5	5,4	0,0

Figure 1.15

of assumptions, which makes it more useful as a tool for computing equilibria.¹⁴

However, even when players do respond to the past averages of their opponents' play, the adjustment process need not converge, especially once we move away from games with one-dimensional strategy spaces and concave payoffs. The first example of cycles in this context is due to Shapley (1964), who considered the game illustrated here in figure 1.15.

Suppose first that, in each period, each player chooses a best response to the action his opponent played the period before. If play starts at the point (M, L), it will proceed to trace out the cycle (M, L), (M, R), (U, R), (U, M), (D, M), (D, L), (M, L). If instead players take turns reacting to one another's previous action, then once again play switches from one point to the next each period. If players respond to their opponents' average play, the play cycles increasingly (in fact, geometrically) slowly but never converges: Once (M, L) is played, (M, R) occurs for the next two periods, then player 1 switches to U; (U, R) occurs for the next four periods, then player 2 switches to M; after eight periods of (U, M), player 1 switches to D; and so on.

Thus, even assuming that behavior follows an adjustment process does not imply that play must converge to a Nash equilibrium. And the adjustment processes are not compelling as a description of players' behavior. One problem with all the processes we have discussed so far is that the players ignore the way that their current action will influence their opponent's action in the next period. That is, the adjustment process itself may not be an equilibrium of the "repeated game," where players know they face one another repeatedly.¹⁵ It might seem natural that if the same two players face each other repeatedly they would come to recognize the dynamic effect of their choices. (Note that the effect is smaller if players react to past averages.)

14. For a detailed study of convergence when Cournot oligopolists respond to averages, see Thorlund-Petersen 1990.

15. If firms have perfect foresight, they choose their output taking into account its effect on their rival's future reaction. On this, see exercise 13.2. The Cournot tâtonnement process can be viewed as a special case of the perfect-foresight model where the firms have discount factor 0.

A related defense of Nash equilibrium supposes that there is a large group of players who are matched at random and asked to play a specific game. The players are not allowed to communicate or even to know who their opponents are. At each round, each player chooses a strategy, observes the strategy chosen by his opponent, and receives the corresponding payoff. If there are a great many players then a pair of players who are matched today are unlikely to meet again, and players have no reason to worry about how their current choice will affect the play of their future opponents. Thus, in each period the players should tend to play the strategy that maximizes that period's expected payoff. (We say "tend to play" to allow for the possibility that players may occasionally "experiment" with other choices.)

The next step is to specify how players adjust their expectations about their opponents' play in light of their experience. Many different specifications are possible, and, as with the Cournot process, the adjustment process need not converge to a stable distribution. However, if players observe their opponents' strategies at the end of each round, and players eventually receive a great many observations, then one natural specification is that each player's expectations about the play of his opponents converges to the probability distribution corresponding to the sample average of play he has observed in the past. In this case, if the system converges to a steady state, the steady state must be a Nash equilibrium.¹⁶

Caution The assumption that players observe one another's strategies at the end of each round makes sense in games like the Cournot competition where strategies correspond to uncommitted choices of actions. In the general extensive-form games we introduce in chapter 3, strategies are contingent plans, and the observed outcome of play need not reveal the action a player would have used in a contingency that did not arise (Fudenberg and Kreps 1988).

The idea of a large population of players can also be used to provide an alternative interpretation of mixed strategies and mixed-strategy equilibria. Instead of supposing that individual players randomize among several strategies, a mixed strategy can be viewed as describing a situation in which different fractions of the population play different pure strategies. Once again a Nash equilibrium in mixed strategies requires that all pure strategies that receive positive probability are equally good responses, since if one pure strategy did better than the other we would expect more and more of the players to learn this and switch their play to the strategy with the higher payoff.

16. Recent papers on the explanation of Nash equilibrium as the result of learning include Gul 1989, Milgrom and Roberts 1989, and Nyarko 1989.

The large-population model of adjustment to Nash equilibrium has yet another application: It can be used to discuss the adjustment of population fractions by *evolution* as opposed to learning. In theoretical biology, Maynard Smith and Price (1973) pioneered the idea that animals are genetically programmed to play different pure strategies, and that the genes whose strategies are more successful will have higher reproductive fitness. Thus, the population fractions of strategies whose payoff against the current distribution of opponents' play is relatively high will tend to grow at a faster rate, and, any stable steady state must be a Nash equilibrium. (Non-Nash profiles can be unstable steady states, and not all Nash equilibria are locally stable.) It is interesting to note that there is an extensive literature applying game theory to questions of animal behavior and of the determination of the relative frequency of male and female offspring. (Maynard Smith 1982 is the classic reference.)

More recently, some economists and political scientists have argued that evolution can be taken as a metaphor for learning, and that evolutionary stability should be used more broadly in economics. Work in this area includes Axelrod's (1984) study of evolutionary stability in the repeated prisoner's dilemma game we discuss in chapter 4 and Sugden's (1986) study of how evolutionary stability can be used to ask which equilibria are more likely to become focal points in Schelling's sense.

To conclude this section we compare the informational assumptions of deductive and extrapolative explanations of Nash equilibrium and iterated strict dominance. The deductive justification of the iterated deletion of strictly dominated strategies requires that players are rational and know the payoff functions of all players, that they know their opponents are rational and know the payoff functions, that they know the opponents know, and so on for as many steps as it takes for the iterative process to terminate. In contrast, if players play one another repeatedly, then, even if players do not know their opponents' payoffs, they will eventually learn that the opponents do not play certain strategies, and the dynamics of the learning system will replicate the iterative deletion process. And for an extrapolative justification of Nash equilibrium, it suffices that players know their own payoffs, that play eventually converges to a steady state, and that if play does converge all players eventually learn their opponents' steady-state strategies. Players need not have any information about the payoff functions or information of their opponents.

Of course, the reduction in the informational requirements is made possible by the additional hypotheses of the learning story: Players must have enough experience to learn how their opponents play, and play must converge to a steady state. Moreover, we must suppose either that there is a large population of players who are randomly matched, or that, even though the same players meet one another repeatedly, they ignore

any dynamic links between their play today and their opponents' play tomorrow.

1.3 Existence and Properties of Nash Equilibria (technical)¹¹

We now tackle the question of the existence of Nash equilibria. Although some of the material in this section is technical, it is quite important for those who wish to read the formal game-theory literature. However, the section can be skipped in a first reading by those who are pressed for time and have little interest in technical detail.

1.3.1 Existence of a Mixed-Strategy Equilibrium

Theorem 1.1 (Nash 1950b) Every finite strategic-form game has a mixed-strategy equilibrium.

Remark Remember that a pure-strategy equilibrium is an equilibrium in degenerate mixed strategies. The theorem does not assert the existence of an equilibrium with nondegenerate mixing.

Proof Since this is the archetypal existence proof in game theory, we will go through it in detail. The idea of the proof is to apply Kakutani's fixed-point theorem to the players' "reaction correspondences." Player i 's reaction correspondence, r_i , maps each strategy profile σ to the set of mixed strategies that maximize player i 's payoff when his opponents play σ_{-i} . (Although r_i depends only on σ_{-i} and not on σ_i , we write it as a function of the strategies of all players, because later we will look for a fixed point in the space Σ of strategy profiles.) This is the natural generalization of the Cournot reaction function we defined above. Define the correspondence $r: \Sigma \rightrightarrows \Sigma$ to be the Cartesian product of the r_i . A fixed point of r is a σ such that $\sigma \in r(\sigma)$, so that, for each player, $\sigma_i \in r_i(\sigma)$. Thus, a fixed point of r is a Nash equilibrium.

From Kakutani's theorem, the following are sufficient conditions for $r: \Sigma \rightrightarrows \Sigma$ to have a fixed point:

- (1) Σ is a compact,¹⁷ convex,¹⁸ nonempty subset of a (finite-dimensional) Euclidean space.
- (2) $r(\sigma)$ is nonempty for all σ .
- (3) $r(\sigma)$ is convex for all σ .

17. A subset X of a Euclidean space is compact if any sequence in X has a subsequence that converges to a limit point in X . The definition of compactness for more general topological spaces uses the notion of "cover," which is a collection of open sets whose union includes the set X . X is compact if any cover has a finite subcover.

18. A set X in a linear vector space is convex if, for any x and x' belonging to X and any $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)x'$ belongs to X .

	L	R
U	1, 1	0, 0
D	0, 0	$\lambda, 2$

a

	L	R
U	1, 1	0, 0
D	0, 0	λ, λ

b

Figure 1.18

sider two sequences $\lambda^n \rightarrow \lambda$ and $\sigma^n \rightarrow \sigma$ such that $\sigma^n \in r(\sigma^n)$ and $\sigma \notin r(\sigma)$. That is, σ^n is a Nash equilibrium of $G(\lambda^n)$, but σ is not a Nash equilibrium of $G(\lambda)$. Then there is a player i and a δ_i that does strictly better than σ_i against σ_{-i} . Since payoffs are continuous in λ , for any λ^n near λ and any σ_i^n near σ_i , δ_i is a strictly better response to σ_{-i}^n than σ_i^n is—a contradiction.

It is important to note that this does not mean that the correspondence $E(\cdot)$ is continuous. Loosely speaking, a closed graph (plus compactness) implies that the set of equilibria cannot shrink in passing to the limit. If σ^n are Nash equilibria of $G(\lambda^n)$ and $\lambda^n \rightarrow \lambda$, then σ^n has a limit point $\sigma \in E(\lambda)$. However, $E(\lambda)$ can contain additional equilibria that are not limits of equilibria of "nearby" games. Thus, $E(\cdot)$ is not lower hemi-continuous, and hence is not continuous. We illustrate this with the two games in figure 1.18. In both of these games, (U, L) is the unique Nash equilibrium if $\lambda < 0$, while for $\lambda > 0$ there are three equilibria (U, L), (D, R), and an equilibrium in mixed strategies. While the equilibrium correspondence has a closed graph in both games, the two games have very different sets of equilibria at the point $\lambda = 0$.

First consider the game illustrated in figure 1.18a. For $\lambda > 0$, there are two pure-strategy equilibria and a unique equilibrium with nondegenerate mixing, as each player can be indifferent between his two choices only if the other player randomizes. If we let p denote the probability of U and q denote the probability of L, a simple computation shows that the unique mixed-strategy equilibrium is

$$(p, q) = \left(\frac{2}{3}, \frac{\lambda}{1 + \lambda} \right).$$

As required by a closed graph, the profiles $(p, q) = (1, 1)$, $(0, 0)$, and $(\frac{2}{3}, 0)$ are all Nash equilibria at $\lambda = 0$. There are also additional equilibria for $\lambda = 0$ that are not limits of equilibria for any sequence $\lambda^n \rightarrow 0$, namely $(p, 0)$ for any $p \in [0, \frac{2}{3}]$. When $\lambda = 0$, player 1 is willing to randomize even if player 2 plays R with probability 1, and so long as the probability of U is not too large player 2 is still willing to play R. This illustrates how the equilibrium correspondence can fail to be lower hemi-continuous.

In the game of figure 1.18b, the equilibria for $\lambda > 0$ are $(1, 1)$, $(0, 0)$, and $(\lambda/(1 + \lambda), \lambda/(1 + \lambda))$, whereas for $\lambda = 0$ there are only two equilibria:




Figure 1.19




Figure 1.20

$(1, 1)$ and $(0, 0)$. (To see this, note that if p is greater than 0 then player 2 will set $q = 1$, and so p must equal 1, and $(1, 1)$ is the only equilibrium with $q > 0$.)

At first sight a decrease in the number of equilibria might appear to violate the closed-graph property, but this is not the case: For λ positive but small, the mixed-strategy equilibrium $(\lambda/(1 + \lambda), \lambda/(1 + \lambda))$ is very close to the pure-strategy equilibrium $(0, 0)$. Figures 1.19 and 1.20 display the equilibrium correspondences of these two games. More precisely, for each λ we display the set of p such that (p, q) is an equilibrium of $N(\lambda)$ for some q ; this allows us to give a two-dimensional diagram.

Inspection of the diagrams reveals that each of these games has an odd number of Nash equilibria everywhere except $\lambda = 0$. Chapter 12 explains that this observation is generally true: If the strategy spaces are held fixed, there is an odd number of Nash equilibria for "almost all" payoff functions.

Finally, note that in figures 1.18a and 1.18b, although (D, R) is not a Nash equilibrium for $\lambda < 0$, it is an " ϵ -Nash equilibrium" in the sense of Radner (1980) if $\epsilon \geq |\lambda|$: Each player's maximum gain to deviation is less than ϵ . More generally, an equilibrium of a given game will be an ϵ -Nash equilibrium for games "nearby"—a point developed and exploited by Fudenberg and Levine (1983, 1986), whose results are discussed in chapter 4.

1.3.3 Existence of Nash Equilibrium in Infinite Games with Continuous Payoffs

Economists often use models of games with an uncountable number of actions (as in the Cournot game of example 1.3 and the Hotelling game of example 1.4). Some might argue that prices or quantities are "really" infinitely divisible, while others might argue that "reality" is discrete and the continuum is a mathematical abstraction, but it is often easier to work with a continuum of actions rather than a large finite grid. Moreover, as Dasgupta and Maskin (1986) argue, when the continuum game does not have a Nash equilibrium, the equilibria corresponding to fine, discrete grids (whose existence was proved in subsection 1.3.1) could be very sensitive to exactly which finite grid is specified: If there were equilibria of the finite-grid version of the game that were fairly insensitive to the choice of the grid, one could take a sequence of finer and finer grids "converging" to the continuum, and the limit of a convergent subsequence of the discrete-space equilibria would be a continuum equilibrium under appropriate continuity assumptions. (To put it another way, one can pick equilibria of the discrete-grid version of the game that do not fluctuate with the grid if the continuum game has an equilibrium.)

Theorem 1.2 (Debreu 1952; Glicksberg 1952; Fan 1952) Consider a strategic-form game whose strategy spaces S_i are nonempty compact convex subsets of an Euclidean space. If the payoff functions u_i are continuous in s and quasi-concave in s_i , there exists a pure-strategy Nash equilibrium.²²

Proof The proof is very similar to that of Nash's theorem: We verify that continuous payoffs imply nonempty, closed-graph reaction correspondences, and that quasi-concavity in players' own actions implies that the reaction correspondences are convex-valued. ■

Note that Nash's theorem is a special case of this theorem. The set of mixed strategies over a finite set of actions, being a simplex, is a compact, convex subset of an Euclidean space; the payoffs are polynomial, and therefore quasi-concave, in the player's own mixed strategy.

If the payoff functions are not continuous, the reaction correspondences can fail to have a closed graph and/or fail to be nonempty. The latter problem arises because discontinuous functions need not attain a maximum, as for example the function $f(x) = -|x|$, $x \neq 0$, $f(0) = -1$. To see how the reaction correspondence may fail to have a closed graph even when optimal reactions always exist, consider the following two-player game:

$$S_1 = S_2 = [0, 1],$$

$$u_1(s_1, s_2) = -(s_1 - s_2)^2,$$

²² It is interesting to note that Debreu (1952) used a generalization of theorem 1.2 to prove that competitive equilibria exist when consumers have quasi-convex preferences.

$$u_2(s_1, s_2) = \begin{cases} -(s_1 - s_2 - \frac{1}{3})^2, & s_1 \geq \frac{1}{3} \\ -(s_1 - s_2 + \frac{1}{3})^2, & s_1 < \frac{1}{3}. \end{cases}$$

Here each player's payoff is strictly concave in his own strategy, and a best response exists (and is unique) for each strategy of the opponent. However, the game does not have a pure-strategy equilibrium: Player 1's reaction function is $r_1(s_2) = s_2$, while player 2's reaction function is $r_2(s_1) = s_1 - \frac{1}{3}$ for $s_1 \geq \frac{1}{3}$, $r_2(s_1) = s_1 + \frac{1}{3}$ for $s_1 < \frac{1}{3}$, and these reaction functions do not intersect.

Quasi-concavity is hard to satisfy in some contexts. For example, in the Cournot game the quasi-concavity of payoffs requires strong conditions on the second derivatives of the price and cost functions. Of course, Nash equilibria can exist even when the conditions of the existence theorems are not satisfied, as these conditions are sufficient but not necessary. However, in the Cournot case Roberts and Sonnenschein (1976) show that pure-strategy Cournot equilibria can fail to exist with "nice" preferences and technologies.

The absence of a pure-strategy equilibrium in some games should not be surprising, since pure-strategy equilibria need not exist in finite games, and these games can be approximated by games with real-valued action spaces but nonconcave payoffs. Figure 1.21 depicts the payoffs of player 1, who chooses an action s_1 in the interval $[\underline{s}_1, \bar{s}_1]$. Payoff function u_1 is continuous in s but not quasi-concave in s_1 . This game is "almost" a game where player 1 has two actions, s_1' and s_1'' . Suppose the same holds for player 2. Then the game is similar to a game with two actions per player, and we know (from "matching pennies," for instance) that such games may have no pure-strategy equilibrium.

When payoffs are continuous (but not necessarily quasi-concave), mixed strategies can be used to obtain convex-valued reactions, as in the following theorem.

Theorem 1.3 (Glicksberg 1952) Consider a strategic-form game whose strategy spaces S_i are nonempty compact subsets of a metric space. If the payoff functions u_i are continuous then there exists a Nash equilibrium in mixed strategies.




Figure 1.21

Here the mixed strategies are the (Borel) probability measures over the pure strategies, which we endow with the topology of weak convergence.²³ Once more, the proof applies a fixed-point theorem to the reaction correspondences. As we remarked above, the introduction of mixed strategies again makes the strategy spaces convex, the payoffs linear in own strategy and continuous in all strategies (when payoffs are continuous functions of the pure strategies, they are continuous in the mixed strategies as well²⁴), and the reaction correspondences convex-valued. With infinitely many pure strategies, the space of mixed strategies is infinite-dimensional, so a more powerful fixed-point theorem than Kakulani's is required. Alternatively, one can approximate the strategy spaces by a sequence of finite grids. From Nash's theorem, each grid has a mixed-strategy equilibrium. One then argues that, since the space of probability measures is weakly compact, the sequence of these discrete equilibria has an accumulation point. Since the payoffs are continuous, it is easy to verify that the limit point is an equilibrium.

We have already seen that pure-strategy equilibria need not exist when payoffs are discontinuous. There are many examples to show that in this case mixed-strategy equilibria may fail to exist as well. (The oldest such example we know of is given in Sion and Wolfe 1957—see exercise 2.2 below.) Note: The Glicksberg theorem used above fails because when the pure-strategy payoffs are discontinuous the mixed-strategy payoffs are discontinuous too. Thus, as before, best responses may fail to exist for some of the opponents' strategies. Section 12.2 discusses the existence of mixed-strategy equilibria in discontinuous games and conditions that guarantee the existence of pure-strategy equilibria.

Exercises

Exercise 1.1* This exercise asks you to work through the characterization of all the Nash equilibria of general two-player games in which each player has two actions (i.e., 2×2 matrix games). This process is time consuming but straightforward and is recommended to the student who is unfamiliar with the mechanics of determining Nash equilibria.

Let the game be as illustrated in figure 1.22.

The pure-strategy Nash equilibria are easily found by testing each cell of the matrix; e.g., (U, L) is a Nash equilibrium if and only if $a \geq e$ and $b \geq d$.

23. Fix a compact metric space A . A sequence of measures μ^n on A converges "weakly" to a limit μ if $\int f d\mu^n \rightarrow \int f d\mu$ for every real-valued continuous function f on A . The set of probability measures on A endowed with the topology of weak convergence is compact.
24. This is an immediate consequence of the definition of convergence we gave in note 23.

	L	R
U	a, b	c, d
D	e, f	g, h

Figure 1.22

	L	R
U	1, -1	3, 0
D	4, 2	0, -1

Figure 1.23

To determine the mixed-strategy equilibria requires more work. Let x be the probability player 1 plays U and let y be the probability player 2 plays L. We provide an outline, which the student should complete:

- (i) Compute each player's reaction correspondence as a function of his opponent's randomizing probability.
- (ii) For which parameters is player i indifferent between his two strategies regardless of the play of his opponent?
- (iii) For which parameters does player i have a strictly dominant strategy?
- (iv) Show that if neither player has a strictly dominant strategy, and the game has a unique equilibrium, the equilibrium must be in mixed strategies.
- (v) Consider the particular example illustrated in figure 1.23.
 - (a) Derive the best-response correspondences graphically by plotting player i 's payoff to his two pure strategies as a function of his opponent's mixed strategy.
 - (b) Plot the two reaction correspondences in the (x, y) space. What are the Nash equilibria?

Exercise 1.2* Find all the equilibria of the voting game of example 1.5.

Exercise 1.3 (Nash demand game)* Consider the problem of dividing a pie between two players. If we let x and y denote player 1's and player 2's payoffs, the vector (x, y) is feasible if and only if $x \geq x_0$, $y \geq y_0$, and $g(x, y) \leq 1$, where g is a differentiable function with $\partial g / \partial x > 0$ and $\partial g / \partial y > 0$ (for instance, $g(x, y) = x + y$). Assume that the feasible set is convex. The point (x_0, y_0) will be called the *status quo*. Nash (1950a) proposed axioms which implied that the "right" way to divide the pie is the allocation (x^*, y^*) that maximizes the product of the differences from the status quo $(x - x_0)(y - y_0)$ subject to the feasibility constraint $g(x, y) \leq 1$. In his 1953

paper, Nash looked for a game that would give this axiomatic bargaining solution as a Nash equilibrium.

(a) Suppose that both players simultaneously formulate demands x and y . If (x, y) is feasible, each player gets what he demanded. If (x, y) is infeasible, player 1 gets x_0 and player 2 gets y_0 . Show that there exists a continuum of pure-strategy equilibria, and, more precisely, that any efficient division (x, y) (i.e., feasible and satisfying $g(x, y) = 1$) is a pure-strategy-equilibrium outcome.

(b)** Consider Binmore's (1981) version of the Nash "modified demand game." The feasible set is defined by $x \geq x_0$, $y \geq y_0$, and $g(x, y) \leq z$, where z has cumulative distribution F on $[z, \bar{z}]$ (suppose that $\forall z$, the feasible set is nonempty). The players do not know the realization of z before making demands. The allocation is made as previously, after the demands are made and z is realized. Derive the Nash-equilibrium conditions. Show that when F converges to a mass point at 1, any Nash equilibrium converges to the axiomatic bargaining solution.

Exercise 1.4 (Edgeworth duopoly)** There are two identical firms producing a homogeneous good whose demand curve is $q = 100 - p$. Firms simultaneously choose prices. Each firm has a capacity constraint of K . If the firms choose the same price they share the market equally. If the prices are unequal, $p_1 < p_2$, the low-price firm, i , sells $\min(100 - p_i, K)$ and the high-price firm, j , sells $\min[\max(0, 100 - p_j - K), K]$. (There are many possible rationing rules, depending on the distribution of consumers' preferences and on how consumers are allocated to firms. If the aggregate demand represents a group of consumers each of whom buys one unit if the price p_i is less than his reservation price of r , and buys no units otherwise, and the consumer's reservation prices are uniformly distributed on $[0, 100]$, the above rationing rule says that the high-value consumers are allowed to purchase at price p_i before lower-value consumers are.) The cost of production is 10 per unit.

(a) Show that firm 1's payoff function is

$$u_1(p_1, p_2) = \begin{cases} (p_1 - 10)\min(100 - p_1, K), & p_1 < p_2 \\ (p_1 - 10)\min(50 - p_1/2, K), & p_1 = p_2 \\ (p_1 - 10)\min(100 - K - p_1, K), & p_1 > p_2, p_1 < 100 - K \\ 0, & \text{otherwise.} \end{cases}$$

(b) Suppose $30 < K < 45$. (Note that these inequalities are strict.) Show that this game does not have a pure-strategy Nash equilibrium by proving the following sequence of claims:

(i) If (p_1, p_2) is a pure-strategy Nash equilibrium, then $p_1 = p_2$. (Hint: If $p_1 \neq p_2$, then the higher-price firm has customers (Why?) and so the

lower-price firm's capacity constraint is strictly binding. What happens if this firm charges a slightly higher price?)

(ii) If (p, p) is a pure-strategy Nash equilibrium, then $p > 10$.

(iii) If (p, p) is a pure-strategy Nash equilibrium, then p satisfies $p \leq 100 - 2K$.

(iv) If (p, p) is a pure-strategy Nash equilibrium, then $p = 100 - 2K$. (Hint: If $p < 100 - 2K$, is a deviation to a price between p and $100 - 2K$ profitable for either firm?)

(v) Since $K > 30$, there exists $\delta > 0$ such that a price of $100 - 2K + \delta$ earns a firm a higher profit than $100 - 2K$ when the other firm charges $100 - 2K$.

Note: The Edgeworth duopoly game does satisfy the assumptions of theorem 1.3 (restrict prices to the set $[0, 100]$) and so has a mixed-strategy equilibrium.

Exercise 1.5 (final-offer arbitration)* Farber (1980) proposes the following model of final-offer arbitration. There are three players: a management ($i = 1$), a union ($i = 2$), and an arbitrator ($i = 3$). The arbitrator must choose a settlement $t \in \mathbb{R}$ from the two offers, $s_1 \in \mathbb{R}$ and $s_2 \in \mathbb{R}$, made by the management and the union respectively. The arbitrator has exogenously given preferences $v_0 = -(t - s_0)^2$. That is, he would like to be as close to his "bliss point," s_0 , as possible. The management and the union don't know the arbitrator's bliss point; they know only that it is drawn from the distribution P with continuous, positive density p on $[s_0, \bar{s}_0]$. The management and the union choose their offers simultaneously. Their objective functions are $u_1 = -t$ and $u_2 = +t$, respectively.

Derive and interpret the first-order conditions for a Nash equilibrium. Show that the two offers are equally likely to be chosen by the arbitrator.

Exercise 1.6** Show that the two-player game illustrated in figure 1.24 has a unique equilibrium. (Hint: Show that it has a unique pure-strategy equilibrium; then show that player 1, say, cannot put positive weight on both U and M; then show that player 1, say, cannot put positive weight on both U and D, but not on M, for instance.)

	L	M	R
U	1, -2	-2, 1	0, 0
M	-2, 1	1, -2	0, 0
D	0, 0	0, 0	1, 1

Figure 1.24

Exercise 1.7 (public good)* Consider an economy with I consumers with "quasi-linear" utility functions,

$$u_i = V_i(x, \theta_i) + t_i,$$

where t_i is consumer i 's income, x is a public decision (for instance, the quantity of a public good), $V_i(x, \theta_i)$ is consumer i 's gross surplus for decision x , and θ_i is a utility parameter. The monetary cost of decision x is $C(x)$.

The socially efficient decision is

$$x^*(\theta_1, \dots, \theta_I) \in \arg \max_x \left\{ \sum_{i=1}^I V_i(x, \theta_i) - C(x) \right\}.$$

Assume (i) that the maximand in this program is strictly concave and (ii) that for all θ_{-i} , θ_i , and θ'_i ,

$$\theta'_i \neq \theta_i \Rightarrow x^*(\theta_{-i}, \theta'_i) \neq x^*(\theta_{-i}, \theta_i).$$

Condition ii says that the optimal decision is responsive to the utility parameter of each consumer. (Condition i is satisfied if x belongs to \mathbb{R} , V_i is strictly concave in x , and C is strictly convex in x . Furthermore, if θ_i belongs to an interval of \mathbb{R} , V_i and C are twice differentiable, $\partial V_i / \partial x \partial \theta_i > 0$ or < 0 , and x^* is an interior solution, then x^* is strictly increasing or strictly decreasing in θ_i , so that condition (ii) is satisfied as well.)

Now consider the following "demand-revelation game": Consumers are asked to announce their utility parameters simultaneously. A pure strategy for consumer i is thus an announcement $\hat{\theta}_i$ of his parameter ($\hat{\theta}_i$ may differ from the true parameter θ_i). The realized decision is the optimal one for the announced parameters $x^*(\hat{\theta}_1, \dots, \hat{\theta}_I)$, and consumer i receives a transfer from a "social planner" equal to

$$t_i(\hat{\theta}_1, \dots, \hat{\theta}_I) = K_i + \sum_{j \neq i} V_j(x^*(\hat{\theta}_1, \dots, \hat{\theta}_I), \theta_j) - C(x^*(\hat{\theta}_1, \dots, \hat{\theta}_I)),$$

when K_i is a constant.

Show that telling the truth is dominant, in that any report $\hat{\theta}_i \neq \theta_i$ is strictly dominated by the truthful report $\hat{\theta}_i = \theta_i$.

Because each player has a dominant strategy, it does not matter whether he knows the other players' utility parameters. Hence, even if the players do not know one another's payoffs (see chapter 6), it is still rational for them to tell the truth. This property of the dominant-strategy demand-revelation mechanism (called the *Groves mechanism*) makes it particularly interesting in a situation in which a consumer's utility parameter is known only to that consumer.

Exercise 1.8* Consider the following model of bank runs, which is due to Diamond and Dybvig (1983). There are three periods ($t = 0, 1, 2$). There are many consumers—a continuum of them, for simplicity. All consumers are *ex ante* identical. At date 0, they deposit their entire wealth, \$1, in a bank.

The bank invests in projects that yield \$ R each if the money is invested for two periods, where $R > 1$. However, if a project is interrupted after one period, it yields only \$1 (it breaks even). Each consumer "dies" (or "needs money immediately") at the end of date 1 with probability x , and lives for two periods with probability $1 - x$. He learns which one obtains at the beginning of date 1. A consumer's utility is $u(c_1)$ if he dies in period 1 and $u(c_1 + c_2)$ if he dies in period 2, where $u' > 0$, $u'' < 0$, and c_1 and c_2 are the consumptions in periods 1 and 2.

An optimal insurance contract (c_1^*, c_2^*) maximizes a consumer's *ex ante* or expected utility. The consumer receives c_1^* if he dies at date 1, and otherwise consumes nothing at date 1 and receives c_2^* at date 2. The contract satisfies $xc_1^* + (1-x)c_2^*/R = 1$ (the bank breaks even) and $u'(c_1^*) = R u'(c_2^*)$ (equality between the marginal rates of substitution). Note that $1 < c_1^* < c_2^*$. The issue is whether the bank can implement this optimal insurance scheme if it is unable to observe who needs money at the end of the first period. Suppose that the bank offers to pay $r_1 = c_1^*$ to consumers who want to withdraw their money in period 1. If $f \in [0, 1]$ is the fraction of consumers who withdraw at date 1, each withdrawing consumer gets r_1 if $f r_1 \leq 1$, and gets $1/f$ if $f r_1 > 1$. Similarly, consumers who do not withdraw at date 1 receive $\max\{0, R(1 - r_1 f)/(1 - f)\}$ in period 2.

- Show that it is a Nash equilibrium for each consumer to withdraw at date 1 if and only if he "dies" at that date.
- Show that another Nash equilibrium exhibits a bank run ($f = 1$).
- Compare with the stag hunt.

Exercise 1.9* Suppose $p(q) = a - bq$ in the Cournot duopoly game of example 1.3.

- Check that the second-order and boundary conditions for equation (1.3) are satisfied. Compute the Nash equilibrium.
- Now suppose there are I identical firms, which all have cost function $c_i(q_i) = cq_i$. Compute the limit of the Nash equilibria as $I \rightarrow \infty$. Comment.

Exercise 1.10* Suppose there are I farmers, each of whom has the right to graze cows on the village common. The amount of milk a cow produces depends on the total number of cows, N , grazing on the green. The revenue produced by n_i cows is $n_i v(N)$ for $N < \bar{N}$, and $v(N) \equiv 0$ for $N \geq \bar{N}$, where $v(0) > 0$, $v' < 0$, and $v'' \leq 0$. Each cow costs c , and cows are perfectly divisible. Suppose $v(0) > c$. Farmers simultaneously decide how many cows to purchase; all purchased cows will graze on the common.

- Write this as a game in strategic form.
- Find the Nash equilibrium, and compare it against the social optimum.
- Discuss the relationship between this game and the Cournot oligopoly model.

(This exercise, constructed by R. Gibbons, is based on a discussion in Hume 1739.)

Exercise 1.11** We mentioned that theorem 1.3, which concerns the existence of a mixed-strategy Nash equilibrium when strategy spaces are nonempty, compact subsets of a metric space (\mathbb{R}^n , say) and when the payoff functions are continuous, can also be proved by taking a sequence of discrete approximations of the strategy spaces that "converge" to it. Go through the steps of the proof as carefully as you can.

Here is a sketch of the proof: Each discrete grid has a mixed-strategy equilibrium. By compactness, the sequence of discrete-grid equilibria has an accumulation point. Argue that this limit must be an equilibrium of the limit game with a continuum of actions. (This relies on the discrete grids becoming increasingly good approximations and the payoffs being continuous.)

Exercise 1.12* Consider a simultaneous-move auction in which two players simultaneously choose bids, which must be in nonnegative integer multiples of one cent. The higher bidder wins a dollar bill. If the bids are equal, neither player receives the dollar. Each player must pay his own bid, whether or not he wins the dollar. (The loser pays too.) Each player's utility is simply his net winnings; that is, the players are risk neutral. Construct a symmetric mixed-strategy equilibrium in which every bid less than 1.00 has a positive probability.

References

- Aumann, R. 1959. Acceptable points in general cooperative n -person games. In *Contributions to the Theory of Games IV*. Princeton University Press.
- Aumann, R. 1987. Correlated equilibrium as an extension of Bayesian rationality. *Econometrica* 55: 1–18.
- Aumann, R. 1990. Communication need not lead to Nash equilibrium. Mimeo, Hebrew University of Jerusalem.
- Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books.
- Berge, C. 1963. *Topological Spaces*. Macmillan.
- Bernheim, D. 1984. Rationalizable strategic behavior. *Econometrica* 52: 1007–1028.
- Bernheim, D., D. Peleg, and M. Whinston. 1987. Coalition-proof Nash equilibria. I: Concepts. *Journal of Economic Theory* 42: 1–12.
- Bernheim, D., and M. Whinston. 1987. Coalition-proof Nash equilibria. II: Applications. *Journal of Economic Theory* 42: 13–22.
- Bertrand, J. 1883. Théorie mathématique de la richesse sociale. *Journal des Savants* 499–508.
- Binmore, K. 1981. Nash bargaining theory II. London School of Economics.
- Cournot, A. 1838. *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. English edition: *Researches into the Mathematical Principles of the Theory of Wealth*, ed. N. Bacon (Macmillan, 1897).
- Dasgupta, P., and E. Maskin. 1986. The existence of equilibrium in discontinuous economic games. I: Theory. *Review of Economic Studies* 53: 1–26.

- Debreu, D. 1952. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences* 38: 886–893.
- Diamond, D., and P. Dybvig. 1983. Bank runs, deposit insurance and liquidity. *Journal of Political Economy* 91: 401–419.
- Diamond, P. 1982. Aggregate demand in search equilibrium. *Journal of Political Economy* 90: 881–894.
- Dixit, A. 1986. Comparative statics for oligopoly. *International Economic Review* 27: 107–122.
- Fahrquarson, R. 1969. *Theory of Voting*. Yale University Press.
- Fan, K. 1952. Fixed point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences* 38: 121–126.
- Farber, H. 1980. An analysis of final-offer arbitration. *Journal of Conflict Resolution* 35: 683–705.
- Farrell, J., and G. Saloner. 1985. Standardization, compatibility, and innovation. *Rand Journal of Economics* 16: 70–83.
- Fisher, F. 1961. The stability of the Cournot oligopoly solution: The effects of speed of adjustment and increasing marginal costs. *Review of Economic Studies* 28: 125–135.
- Fudenberg, D., and D. Kreps. 1988. A theory of learning, experimentation, and equilibrium in games. Mimeo, Stanford Graduate School of Business.
- Fudenberg, D., and D. Levine. 1983. Subgame-perfect equilibria of finite and infinite horizon games. *Journal of Economic Theory* 31: 251–268.
- Fudenberg, D., and D. Levine. 1986. Limit games and limit equilibria. *Journal of Economic Theory* 38: 261–279.
- Glicksberg, I. L. 1952. A further generalization of the Kakutani fixed point theorem with application to Nash equilibrium points. *Proceedings of the National Academy of Sciences* 38: 170–174.
- Green, J., and W. Heller. 1981. Mathematical analysis and convexity with applications to economics. In *Handbook of Mathematical Economics*, volume 1, ed. K. Arrow and M. Intriligator. North-Holland.
- Gul, F. 1989. Rational strategic behavior and the notion of equilibrium. Mimeo, Stanford Graduate School of Business.
- Hahn, F. 1962. The stability of the Cournot oligopoly solution. *Review of Economic Studies* 29: 929–931.
- Harsanyi, J. 1973a. Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points. *International Journal of Game Theory* 1: 1–23.
- Harsanyi, J. 1973b. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory* 2: 235–250.
- Harsanyi, J., and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press.
- Hirsch, M., and S. Smale. 1974. *Differential Equations, Dynamical Systems and Linear Algebra*. Academic Press.
- Hotelling, H. 1929. Stability in competition. *Economic Journal* 39: 41–57.
- Hume, D. 1739. *A Treatise on Human Nature* (Everyman edition: J. M. Dent, 1952).
- Kuhn, H. 1953. Extensive games and the problem of information. *Annals of Mathematics Studies*, No. 28. Princeton University Press.
- Luce, D., and H. Raiffa. 1957. *Games and Decisions*. Wiley.
- Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge University Press.
- Maynard Smith, J., and G. R. Price. 1973. The logic of animal conflicts. *Nature* 246: 15–18.
- Milgrom, P., and J. Roberts. 1989. Adaptive and sophisticated learning in repeated normal forms. Mimeo, Stanford University.
- Moulin, H. 1979. Dominance solvable voting schemes. *Econometrica* 47: 1337–1353.

- Moulin, H. 1984. Dominance solvability and Cournot stability. *Mathematical Social Sciences* 7: 83-102.
- Moulin, H. 1986. *Game Theory for the Social Sciences*. New York University Press.
- Nash, J. 1950a. The bargaining problem. *Econometrica* 18: 155-162.
- Nash, J. 1950b. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences* 36: 48-49.
- Nash, J. 1953. Two-person cooperative games. *Econometrica* 21: 128-140.
- Novshek, W. 1985. On the existence of Cournot equilibrium. *Review of Economic Studies* 52: 85-98.
- Nyarko, Y. 1989. Bayesian learning in games. Mimeo, New York University.
- Ordeshook, P. 1986. *Game Theory and Political Theory: An Introduction*. Cambridge University Press.
- Radner, R. 1980. Collusive behavior in non-cooperative epsilon equilibria of oligopolies with long but finite lives. *Journal of Economic Theory* 22: 136-154.
- Roberts, J., and H. Sonnenschein. 1976. On the existence of Cournot equilibrium without concave profit functions. *Journal of Economic Theory* 13: 112-117.
- Schelling, T. 1960. *The Strategy of Conflict*. Harvard University Press.
- Scade, J. 1980. The stability of Cournot revisited. *Journal of Economic Theory* 23: 15-17.
- Shapley, L. 1964. Some topics in two-person games. In *Contributions to the Theory of Games* (Princeton Annals of Mathematical Studies, no. 52).
- Sion, M., and P. Wolfe. 1957. On a game without a value. In *Contributions to the Theory of Games, Volume III* (Princeton Annals of Mathematical Studies, no. 39).
- Sugden, R. 1986. *The Economic of Rights, Cooperation, and Welfare*. Blackwell.
- Thorlund-Petersen, L. 1990. Iterative computation of Cournot equilibrium. *Games and Economic Behavior* 2: 61-95.
- von Neumann, J. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Annalen* 100: 295-320.
- Waltz, K. 1959. *Man, the State, and War*. Columbia University Press.

Most economic applications of game theory use the concept of Nash equilibrium or one of the more restrictive "equilibrium refinements" we introduce in later chapters. However, as we warned in chapter 1, in some situations the Nash concept seems too demanding. Thus, it is interesting to know what predictions one can make without assuming that a Nash equilibrium will occur. Section 2.1 presents the notions of iterated strict dominance and rationalizability, which derive predictions using only the assumptions that the structure of the game (i.e., the strategy spaces and the payoffs) and the rationality of the players are common knowledge. As we will see, these two notions are closely related, as rationalizability is essentially the contrapositive of iterated strict dominance.

Section 2.2 introduces the idea of a correlated equilibrium, which extends the Nash concept by supposing that players can build a "correlating device" that sends each of them a private signal before they choose their strategy.

2.1 Iterated Strict Dominance and Rationalizability[†]

We introduced iterated strict dominance informally at the beginning of chapter 1. We will now define it formally, derive some of its properties, and apply it to the Cournot model. We will then define rationalizability and relate the two concepts. As throughout, we restrict our attention to finite games except where we explicitly indicate otherwise.

2.1.1 Iterated Strict Dominance: Definition and Properties

Definition 2.1 The process of iterated deletion of strictly dominated strategies proceeds as follows: Set $S_i^0 \equiv S_i$ and $\Sigma_i^0 \equiv \Sigma_i$. Now define S_i^n recursively by

$$S_i^n = \{s_i \in S_i^{n-1} \mid \text{there is no } \sigma_i \in \Sigma_i^{n-1} \text{ such that } u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}^{n-1}\}$$

and define

$$\Sigma_i^n = \{\sigma_i \in \Sigma_i \mid \sigma_i(s_i) > 0 \text{ only if } s_i \in S_i^n\}.$$

Set

$$S_i^\infty = \bigcap_{n=0}^{\infty} S_i^n.$$

S_i^∞ is the set of player i 's pure strategies that survive iterated deletion of strictly dominated strategies. Set Σ_i^∞ to be all mixed strategies σ_i such that there is no σ_i' with $u_i(\sigma_i', s_{-i}) > u_i(\sigma_i, s_{-i})$ for all $s_{-i} \in S_{-i}^\infty$. This is the set of player i 's mixed strategies that survive iterated strict dominance.

In words, S_i^∞ is the set of player i 's strategies that are not strictly dominated when players $j \neq i$ are constrained to play strategies in S_j^{n-1} and Σ_i^n