(a) Suppose that a monopolist can produce good 1 at a constant unit cost of c per unit and cannot engage in any kind of price discrimination. Find its optimal choice of price and quantity. For what values of c will it be true that it chooses to sell to both types of consumers?

(b) Suppose that the monopolist uses a "two-part tariff" where a consumer must pay a lump sum k in order to be able to buy anything at all. A person who has paid the lump sum k can buy as much as he likes at a price of p per unit purchased. Consumers are not able to resell good 1. For $p < 4$, what is the highest amount k that a type A is willing to pay for the privilege of buying at price p? If a type A does pay the lump sum k to buy at price p, how many units will he demand? Describe the function that determines demand for good 1 by type A consumers as a function of p and k. What is the demand function for good 1 by type B consumers? Now describe the function that determines total demand for good 1 by all consumers as a function of p and k.

(c) If the economy consisted only of N type A consumers and no type B consumers, what would be the profit-maximizing choices of p and k?

(d) If $c < 1$, find the values of p and k that maximize the monopolist's profits subject to the constraint that both types of consumers buy from it.

CHAPTER 15

GAME THEORY

Game theory is the study of interacting decision makers. In earlier chapters we studied the theory of optimal decision making by a single agent—a firm or a consumer—in very simple environments. The strategic interactions of the agents were not very complicated. In this chapter we will lay the foundations for a deeper analysis of the behavior of economic agents in more complex environments.

There are many directions from which one could study interacting decision makers. One could examine behavior from the viewpoint of sociology, psychology, biology, etc. Each of these approaches is useful in certain contexts. Game theory emphasizes a study of cold-blooded "rational" decision making, since this is felt to be the most appropriate model for most economic behavior.

Game theory has been widely used in economics in the last decade, and much progress has been made in clarifying the nature of strategic interaction in economic models. Indeed, most economic behavior can be viewed as special cases of game theory, and a sound understanding of game theory is a necessary component of any economist's set of analytical tools.

15.1 Description of a game

There are several ways of describing a game. For our purposes, the **strategic form** and the **extensive form** will be sufficient. Roughly speaking the extensive form provides an "extended" description of a game while the strategic form provides a "reduced" summary of a game.[1] We will first describe the strategic form, reserving the discussion of the extensive form for the section on sequential games.

The strategic form of the game is defined by exhibiting a set of **players**, a set of **strategies**, the choices that each player can make, and a set of **payoffs** that indicate the utility that each player receives if a particular combination of strategies is chosen. For purposes of exposition, we will treat two-person games in this chapter. All of the concepts described below can be easily extended to multiperson contexts.

We assume that the description of the game—the payoffs and the strategies available to the players—are **common knowledge**. That is, each player knows his own payoffs and strategies, and the other player's payoffs and strategies. Furthermore, each player knows that the other player knows this, and so on. We also assume that it is common knowledge that each player is "fully rational." That is, each player can choose an action that maximizes his utility given his subjective beliefs, and that those beliefs are modified when new information arrives according to Bayes' law.

Game theory, by this account, is a generalization of standard, one-person decision theory. How should a rational expected utility maximizer behave in a situation in which his payoff depends on the choices of another rational expected utility maximizer? Obviously, each player will have to consider the problem faced by the other player in order to make a sensible choice. We examine the outcome of this sort of consideration below.

EXAMPLE: Matching pennies

In this game, there are two players, Row and Column. Each player has a coin which he can arrange so that either the head side or the tail side is face-up. Thus, each player has two strategies which we abbreviate as Heads or Tails. Once the strategies are chosen there are payoffs to each player which depend on the choices that both players make.

These choices are made independently, and neither player knows the other's choice when he makes his own choice. We suppose that if both players show heads or both show tails, then Row wins a dollar and Column

[1] The strategic form was originally known as the **normal form** of a game, but this term is not very descriptive and its use has been discouraged in recent years.

The game matrix of matching pennies.

		Player B	
		Heads	Tails
Player A	Heads	1, −1	−1, 1
	Tails	−1, 1	1, −1

loses a dollar. If, on the other hand, one player exhibits heads and the other exhibits tails, then Column wins a dollar and Row looses a dollar.

We can depict the strategic interactions in the following **game matrix**:

The entry in box (Head, Tails) indicates that player Row gets −1 and player Column gets +1 if this particular combination of strategies is chosen. Note that in each entry of this box, the payoff to player Row is just the negative of the payoff to player Column. In other words, this is a **zero-sum game**. In zero-sum games the interests of the players are diametrically opposed and are particularly simple to analyze. However, most games of interest to economists are not zero sum games.

EXAMPLE: The Prisoner's Dilemma

Again we have two players, Row and Column, but now their interests are only partially in conflict. There are two strategies: to Cooperate or to Defect. In the original story, Row and Column were two prisoners who jointly participated in a crime. They could cooperate with each other and refuse to give evidence, or one could defect and implicate the other.

In other applications, cooperate and defect could have different meanings. For example, in a duopoly situation, cooperate could mean "keep charging a high price" and defect could mean "cut your price and steal your competitor's market."

An especially simple description used by Aumann (1987) is the game in which each player can simply announce to a referee: "Give me $1,000," or "Give the other player $3,000." Note that the monetary payments come from a third party, not from either of the players; the Prisoners' Dilemma is a **variable-sum game**.

The players can discuss the game in advance but the actual decisions must be independent. The Cooperate strategy is for each person to announce the $3,000 gift, while the Defect strategy is to take the $1,000 (and run!). Here is the payoff matrix to the Aumann version of the Prisoner's Dilemma, where the units of the payoff are thousands of dollars:

We will discuss this game in more detail below, but we should point out the "dilemma" before proceeding. The problem is that each party has an

Table 15.2 **The Prisoner's Dilemma**

		Player B: Cooperate	Player B: Defect
Player A	Cooperate	3, 3	0, 4
	Defect	4, 0	1, 1

incentive to defect, *regardless* of what he or she believes the other party will do. If I believe that the other person will cooperate and give me a \$3,000 gift, then I will get \$4,000 in total by defecting. On the other hand, if I believe that the other person will defect and just take the \$1,000, then I do better by taking \$1,000 for myself.

EXAMPLE: Cournot duopoly

Consider a simple duopoly game, first analyzed by Cournot (1838). We suppose that there are two firms who produce an identical good at zero cost. Each firm must decide how much output to produce without knowing the production decision of the other duopolist. If the firms produce a total of x units of the good, the market price will be $p(x)$; that is, $p(x)$ is the inverse demand curve facing these two producers.

If x_i is the production level of firm i, the market price will then be $p(x_1 + x_2)$, and the profits of firm i are given by $\pi_i = p(x_1 + x_2)x_i$. In this game the strategy of firm i is its choice of production level and the payoff to firm i is its profits.

EXAMPLE: Bertrand duopoly

Consider the same setup as in the Cournot game, but now suppose that the strategy of each player is to announce the price at which he would be willing to supply an arbitrary amount of the good in question. In this case the payoff function takes a radically different form. It is plausible to suppose that the consumers will only purchase from the firm with the lowest price, and that they will split evenly between the two firms if they charge the same price. Letting $x(p)$ represent the market demand function, this leads to a payoff to firm 1 of the form:

$$\pi_1(p_1, p_2) = \begin{cases} p_1 x(p_1) & \text{if } p_1 < p_2 \\ p_1 x(p_1)/2 & \text{if } p_1 = p_2 \\ 0 & \text{if } p_1 > p_2. \end{cases}$$

This game has a similar structure to that of the Prisoner's Dilemma. If both players cooperate, they can charge the monopoly price and each reap half of the monopoly profits. But the temptation is always there for one player to cut its price slightly and thereby capture the entire market for itself. But if both players cut price, then they are both worse off.

15.2 Economic modeling of strategic choices

Note that the Cournot game and the Bertrand game have a radically different structure, even though they purport to model the same economic phenomenon—a duopoly. In the Cournot game, the payoff to each firm is a continuous function of its strategic choice; in the Bertrand game, the payoffs are discontinuous functions of the strategies. As might be expected, this leads to quite different equilibria. Which of these models is "right"?

There is little sense to ask which of these is the "right" model in the abstract. The answer is that it depends on what you are trying to model. It is probably more fruitful to ask what considerations are relevant in modeling the set of strategies used by the agents.

One guide is obviously empirical evidence. If observation of OPEC announcements indicates that they attempt to determine production quotas for each member and allow the price to be set on the world oil markets, then presumably it is more sensible to model the strategies of the game as being production levels rather than prices.

Another consideration is that strategies should be something that can be committed to or that are difficult to change once the opponent's behavior is observed. The games described above are "one-shot" games, but the reality that they are supposed to describe takes place in real time. Suppose that I pick a price for my output and then discover that my opponent has set a slightly smaller price. In this case I can quickly revise my own price. Since the strategic variable can be quickly modified once the opponent's play is known, it doesn't make much sense to try to model this sort of interaction in a one-shot game. It seems that a game with multiple stages must be used to capture the full range of strategic behavior possible in a price-setting game of this sort.

On the other hand, suppose that we interpret output in the Cournot game to be "capacity," in the sense that it is an irreversible capital investment capable of producing the indicated amount of output. In this case, once I discover my opponent's production level, it may be very costly to change my own production level. Here capacity/output seems like a natural choice for the strategic variable, even in a one-shot game.

As in most economic modeling, there is an art to choosing a representation of the strategy choices of the game that captures an element of the real strategic iterations, while at the same time leaving the game simple enough to analyze.

15.3 Solution concepts

In many games the nature of the strategic interaction suggests that a player wants to choose a strategy that is not predictable in advance by the other player. Consider, for example, the Matching Pennies game described above. Here it is clear that neither player wants the other player to be able to predict his choice accurately. Thus, it is natural to consider a random strategy of playing heads with some probability p_h and tails with some probability p_t. Such a strategy is called a **mixed strategy**. Strategies in which some choice is made with probability 1 are called **pure strategies**.

If R is the set of pure strategies available to Row, the set of mixed strategies open to Row will be the set of all probability distributions over R, where the probability of playing strategy r in R is p_r. Similarly, p_c will be the probability that Column plays some strategy c. In order to solve the game, we want to find a set of mixed strategies (p_r, p_c) that are, in some sense, in equilibrium. It may be that some of the equilibrium mixed strategies assign probability 1 to some choices, in which case they are interpreted as pure strategies.

The natural starting point in a search for a solution concept is standard decision theory: we assume that each player has some probability beliefs about the strategies that the other player might choose and that each player chooses the strategy that maximizes his expected payoff.

Suppose for example that the payoff to Row is $u_r(r, c)$ if row plays r and Column plays c. We assume that Row has a **subjective probability distribution** over Column's choices which we denote by (π_c); see Chapter 11, page 191, for the fundamentals of the idea of subjective probability. Here π_c is supposed to indicate the probability, as envisioned by Row, that Column will make the choice c. Similarly, Column has some beliefs about Row's behavior that we can denote by (π_r).

We allow each player to play a mixed strategy and denote Row's *actual* mixed strategy by (p_r) and Column's *actual* mixed strategy by (p_c). Since Row makes his choice without knowing Column's choice, Row's probability that a particular outcome (r, c) will occur is $p_r\pi_c$. This is simply the (objective) probability that Row plays r times Row's (subjective) probability that Column plays c. Hence, Row's objective is to choose a probability distribution (p_r) that maximizes

$$\text{Row's expected payoff} = \sum_r \sum_c p_r \pi_c u_r(r, c).$$

Column, on the other hand, wishes to maximize

$$\text{Column's expected payoff} = \sum_c \sum_r p_c \pi_r u_c(r, c).$$

So far we have simply applied a standard decision-theoretic model to this game—each player wants to maximize his or her expected utility given his or her beliefs. Given my beliefs about what the other player might do, I choose the strategy that maximizes my expected utility.

In this model the beliefs that I hold about the other player's strategic choices are exogenous variables. However, now we add an additional twist to the standard decision model and ask what kinds of beliefs are *reasonable* to hold about the other person's behavior? After all, each player in the game knows that the other player is out to maximize his own payoff, and each should use that information in determining what are reasonable beliefs to have about the other player's behavior.

15.4 Nash equilibrium

In game theory we take as given the proposition that each player is out to maximize his own payoff, and, furthermore, that each player *knows* that this is the goal of each other player. Hence, in determining what might be reasonable beliefs for me to hold about what other players might do, I have to ask what *they* might believe about what *I* will do. In the expected payoff formulas given at the end of the last section, Row's behavior—how likely he is to play each of his strategies—is represented by the probability distribution (p_r) and Column's *beliefs* about Row's behavior are represented by the (subjective) probability distribution (π_r).

A natural consistency requirement is that each player's belief about the other player's choices coincides with the actual choices the other player intends to make. Expectations that are consistent with actual frequencies are sometimes called **rational expectations**. A Nash equilibrium is a certain kind of rational expectations equilibrium. More formally:

Nash equilibrium. *A* **Nash equilibrium** *consists of probability beliefs* (π_r, π_c) *over strategies, and probability of choosing strategies* (p_r, p_c), *such that:*

1) the beliefs are correct: $p_r = \pi_r$ *and* $p_c = \pi_c$ *for all* r *and* c; *and,*

2) each player is choosing (p_r) *and* (p_c) *so as to maximize his expected utility given his beliefs.*

In this definition it is apparent that a Nash equilibrium is an equilibrium in *actions* and *beliefs*. In equilibrium each player correctly foresees how likely the other player is to make various choices, and the beliefs of the two players are mutually consistent.

A more conventional definition of a Nash equilibrium is that it is a pair of mixed strategies (p_r, p_c) such that each agent's choice maximizes his expected utility, given the strategy of the other agent. This is equivalent to the definition we use, but it is misleading since the distinction between the beliefs of the agents and the actions of the agents is blurred. We've tried to be very careful in distinguishing these two concepts.

One particularly interesting special case of a Nash equilibrium is a **Nash equilibrium in pure strategies**, which is simply a Nash equilibrium in which the probability of playing a particular strategy is 1 for each player. That is:

Pure strategies. *A* **Nash equilibrium in pure strategies** *is a pair* (r^*, c^*) *such that* $u_r(r^*, c^*) \geq u_r(r, c^*)$ *for all Row strategies* r, *and* $u_c(r^*, c^*) \geq u_c(r^*, c)$ *for all Column strategies* c.

A Nash equilibrium is a minimal consistency requirement to put on a pair of strategies: if Row believes that Column will play c^*, then Row's best reply is r^* and similarly for Column. No player would find it in his or her interest to deviate unilaterally from a Nash equilibrium strategy.

If a set of strategies is *not* a Nash equilibrium then at least one player is not consistently thinking through the behavior of the other player. That is, one of the players must expect the other player not to act in his own self-interest—contradicting the original hypothesis of the analysis.

An equilibrium concept is often thought of as a "rest point" of some adjustment process. One interpretation of Nash equilibrium is that it is the adjustment process of "thinking through" the incentives of the other player. Row might think: "If I think that Column is going to play some strategy c_1 then the best response for me is to play r_1. But if Column thinks that I will play r_1, then the best thing for him to do is to play some other strategy c_2. But if Column is going to play c_2, then my best response is to play $r_2 \ldots$" and so on. A Nash equilibrium is then a set of beliefs and strategies in which each player's beliefs about what the other player will do are consistent with the other player's actual choice.

Sometimes the "thinking through" adjustment process described in the preceding paragraph is interpreted as an *actual* adjustment process in which each player experiments with different strategies in an attempt to understand the other player's choices. Although it is clear that such experimentation and learning goes on in real-life strategic interaction, this is, strictly speaking, not a valid interpretation of the Nash equilibrium concept. The reason is that if each player knows that the game is going to be repeated some number of times, then each player can plan to base his behavior at time t on observed behavior of the other player up to time t. In this case the correct notion of Nash equilibrium is a sequence of plays that is a best response (in some sense) to a *sequence* of my opponent's plays.

EXAMPLE: Calculating a Nash equilibrium

The following game is known as the "Battle of the Sexes." The story behind the game goes something like this. Rhonda Row and Calvin Column are discussing whether to take microeconomics or macroeconomics this semester. Rhonda gets utility 2 and Calvin gets utility 1 if they both take micro; the payoffs are reversed if they both take macro. If they take different courses, they both get utility 0.

Let us calculate all the Nash equilibria of this game. First, we look for the Nash equilibria in pure strategies. This simply involves a systematic examination of the best responses to various strategy choices. Suppose that Column thinks that Row will play Top. Column gets 1 from playing Left and 0 from playing Right, so Left is Column's best response to Row playing Top. On the other hand, if Column plays Left, then it is easy to see that it is optimal for Row to play Top. This line of reasoning shows that (Top, Left) is a Nash equilibrium. A similar argument shows that (Bottom, Right) is a Nash equilibrium.

The Battle of the Sexes

		Calvin	
		Left (micro)	Right (macro)
Rhonda	Top (micro)	2, 1	0, 0
	Bottom (macro)	0, 0	1, 2

T
15

We can also solve this game systematically by writing the maximization problem that each agent has to solve and examining the first-order conditions. Let (p_t, p_b) be the probabilities with which Row plays Top and Bottom, and define (p_l, p_r) in a similar manner. Then Row's problem is

$$\max_{(p_t, p_b)} p_t[p_l 2 + p_r 0] + p_b[p_l 0 + p_r 1]$$

$$\text{such that } p_t + p_b = 1$$
$$p_t \geq 0$$
$$p_b \geq 0.$$

Let λ, μ_t, and μ_b be the Kuhn-Tucker multipliers on the constraints, so that the Lagrangian takes the form:

$$\mathcal{L} = 2p_t p_l + p_b p_r - \lambda(p_t + p_b - 1) - \mu_t p_t - \mu_b p_b.$$

Differentiating with respect to p_t and p_b, we see that the Kuhn-Tucker conditions for Row are

$$\begin{aligned} 2p_l &= \lambda + \mu_t \\ p_r &= \lambda + \mu_b. \end{aligned} \tag{15.1}$$

Since we already know the pure strategy solutions, we only consider the case where $p_t > 0$ and $p_b > 0$. The complementary slackness conditions then imply that $\mu_t = \mu_b = 0$. Using the fact that $p_t + p_b = 1$, we easily see that Row will find it optimal to play a mixed strategy when $p_l = 1/3$ and $p_r = 2/3$.

Following the same procedure for Column, we find that $p_t = 2/3$ and $p_b = 1/3$. The expected payoff to each player from this mixed strategy can be easily computed by plugging these numbers into the objective function. In this case the expected payoff is 2/3 to each player. Note that each player would prefer either of the pure strategy equilibria to the mixed strategy since the payoffs are higher for each player.

15.5 Interpretation of mixed strategies

It is sometimes difficult to give a behavioral interpretation to the idea of a mixed strategy. For some games, such as Matching Pennies, it is clear that mixed strategies are the only sensible equilibrium. But for other games of economic interest—e.g., a duopoly game—mixed strategies seem unrealistic.

In addition to this unrealistic nature of mixed strategies in some contexts, there is another difficulty on purely logical grounds. Consider again the example of the mixed strategy in the Battle of the Sexes. The mixed strategy equilibrium in this game has the property that if Row is playing his equilibrium mixed strategy, the expected payoff to Column from playing either of his pure strategies must be the same as the expected payoff from playing his equilibrium mixed strategy. The easiest way to see this is to look at the first-order conditions (15.1). Since $2p_l = p_r$, the expected payoff to playing left is the same as the expected payoff to playing right.

But this is no accident. It must always be the case that for any mixed strategy equilibrium, if one party believes that the other player will play the equilibrium mixed strategy, then he is indifferent as to whether he plays his equilibrium mixed strategy, or any pure strategy that is part of his mixed strategy. The logic is straightforward: if some pure strategy that is part of the equilibrium mixed strategy had a higher expected payoff than some other component of the equilibrium mixed strategy, then it would pay to increase the frequency with which one played the strategy with the higher expected payoff. But if all of the pure strategies that are played with positive probability in a mixed strategy have the same expected payoff, this must also be the expected payoff of the mixed strategy. And this in turn implies that the agent is indifferent as to which pure strategy he plays or whether he plays the mixed strategy. This "degeneracy" arises since the expected utility function is linear in probabilities. One would like there to be some more compelling reason to "enforce" the mixed strategy outcome.

In some settings this may not pose a serious problem. Suppose that you are part of a large group of people who meet each other randomly and play Matching Pennies once with each opponent. Suppose that initially everyone is playing the unique Nash equilibrium in mixed strategies of $(\frac{1}{2}, \frac{1}{2})$. Eventually some of the players tire of playing the mixed strategy and decide to play heads or tails all of the time. If the number of people who decide to play Heads all the time equals the number who decide to play Tails, then nothing significant has changed in any agent's choice problem: each agent would still rationally believe that his opponent has a 50:50 chance of playing Heads or Tails.

In this way each member of the population is playing a pure strategy, but in a given game the players have no way of knowing which pure strategy their opponent is playing. This interpretation of mixed strategy probabilities as being population frequencies is common in modeling animal behavior.

Another way to interpret mixed strategy equilibria is to consider a given individual's choice of whether to play Heads or play Tails in a one-shot game. This choice may be thought to depend on idiosyncratic factors that cannot be determined by opponents. Suppose for example that one calls Heads if you're in a "heads mood" and one calls tails if you're in a "tails mood." You may be able to observe your own mood, but your opponent cannot. Hence, from the viewpoint of each player, the other person's strategy is random, even though one's own strategy is deterministic. What matters about a player's mixed strategy is the uncertainty it creates in the *other* players of the game.

15.6 Repeated games

We indicated above that it was not appropriate to expect that the outcome of a repeated game with the same players as simply being a repetition of the one-shot game. This is because the strategy space of the repeated game is much larger: each player can determine his or her choice at some point as a function of the entire *history* of the game up until that point. Since my opponent can modify his behavior based on my history of choices, I must take this influence into account when making my own choices.

Let us analyze this in the context of the simple Prisoner's Dilemma game described earlier. Here it is in the "long-run" interest of both players to try to get to the (Cooperate, Cooperate) solution. So it might be sensible for one player to try to "signal" to the other that he is willing to "be nice" and play cooperate on the first move of the game. It is in the short-run interest of the other player to Defect, of course, but is this really in his long-run interest? He might reason that if he defects, the other player may lose patience and simply play Defect himself from then on. Thus, the second player might lose in the long run from playing the short-run optimal strategy. What lies behind this reasoning is the fact that a move that I make now may have repercussions in the future—the other player's future choices may depend on my current choices.

Let us ask whether the strategy of (Cooperate, Cooperate) can be a Nash equilibrium of the repeated Prisoner's Dilemma. First we consider the case of where each player knows that the game will be repeated a fixed number of times. Consider the reasoning of the players just before the last round of play. Each reasons that, at this point, they are playing a one-shot game. Since there is no future left on the last move, the standard logic for Nash equilibrium applies and both parties Defect.

Now consider the move before the last. Here it seems that it might pay each of the players to cooperate in order to signal that they are "nice guys" who will cooperate again in the next and final move. But we've just seen that when the next move comes around, each player will want to play Defect. Hence there is no advantage to cooperating on the next to the last move—as long as both players believe that the other player will Defect on the *final* move, there is no advantage to try to influence future behavior by being nice on the penultimate move. The same logic of **backwards induction** works for two moves before the end, and so on. In a repeated Prisoner's Dilemma with a known number of repetitions, the Nash equilibrium is to Defect in every round.

The situation is quite different in a repeated game with an infinite number of repetitions. In this case, at each stage it is known that the game will be repeated at least one more time and therefore there will be some (potential) benefits to cooperation. Let's see how this works in the case of the Prisoner's Dilemma.

Consider a game that consists of an infinite number of repetitions of the Prisoner's Dilemma described earlier. The strategies in this repeated game are sequences of functions that indicate whether each player will Cooperate or Defect at a particular stage as a function of the history of the game up to that stage. The payoffs in the repeated game are the discounted sums of the payoffs at each stage; that is, if a player gets a payoff at time t of u_t, his payoff in the repeated game is taken to be $\sum_{t=1}^{\infty} u_t/(1+r)^t$, where r is the discount rate.

I claim that as long as the discount rate is not too high there exists a Nash equilibrium pair of strategies such that each player finds it in his interest to cooperate at each stage. In fact, it is easy to exhibit an explicit example of such strategies. Consider the following strategy: "Cooperate on the current move unless the other player defected on the last move. If the other player defected on the last move, then Defect forever." This is sometimes called a **punishment strategy**, for obvious reasons: if a player defects, he will be punished forever with a low payoff.

To show that a pair of punishment strategies constitutes a Nash equilibrium, we simply have to show that if one player plays the punishment strategy the other player can do no better than playing the punishment strategy. Suppose that the players have cooperated up until move T and consider what would happen if a player decided to Defect on this move. Using the numbers from the Prisoner's Dilemma example on page 261, he would get an immediate payoff of 4, but he would also doom himself to an infinite stream of payments of 1. The discounted value of such a stream of payments is $1/r$, so his total expected payoff from Defecting is $4 + 1/r$.

On the other hand, his expected payoff from continuing to cooperate is $3/r$. Continuing to cooperate is preferred as long as $3/r > 4 + 1/r$, which reduces to requiring that $r < 1/2$. As long as this condition is satisfied, the punishment strategy forms a Nash equilibrium: if one party plays the punishment strategy, the other party will also want to play it, and neither party can gain by unilaterally deviating from this choice.

This construction is quite robust. Essentially the same argument works for any payoffs that exceed the payoffs from (Defect, Defect). A famous result known as the Folk Theorem asserts precisely this: in a repeated Prisoner's Dilemma any payoff larger than the payoff received if both parties consistently defect can be supported as a Nash equilibrium. The proof is more-or-less along the lines of the construction given above.

EXAMPLE: Maintaining a cartel

Consider a simple repeated duopoly which yields profits (π_c, π_c) if both firms choose to play a Cournot game and (π_j, π_j) if both firms produce the level of output that maximizes their joint profits—that is, they act as a **cartel**. It is well-known that the levels of output that maximize joint profits are typically not Nash equilibria in a single-period game—each producer has an incentive to dump extra output if he believes that the other producer will keep his output constant. However, as long as the discount rate is not too high, the joint profit-maximizing solution will be a Nash equilibrium of the repeated game. The appropriate punishment strategy is for each firm to produce the cartel output unless the other firm deviates, in which case it will produce the Cournot output forever. An argument similar to the Prisoner's Dilemma argument shows that this is a Nash equilibrium.

15.7 Refinements of Nash equilibrium

The Nash equilibrium concept seems like a reasonable definition of an equilibrium of a game. As with any equilibrium concept, there are two questions of immediate interest: 1) will a Nash equilibrium generally exist; and 2) will the Nash equilibrium be unique?

Existence, luckily, is not a problem. Nash (1950) showed that with a finite number of agents and a finite number of pure strategies, an equilibrium will always exist. It may, of course, be an equilibrium involving mixed strategies.

Uniqueness, however, is very unlikely to occur in general. We have already seen that there may be several Nash equilibria to a game. Game theorists have invested a substantial amount of effort into discovering further criteria that can be used to choose among Nash equilibria. These criteria are known as **refinements** of the concept of Nash equilibrium, and we will investigate a few of them below.

15.8 Dominant strategies

Let r_1 and r_2 be two of Row's strategies. We say that r_1 **strictly dominates** r_2 for Row if the payoff from strategy r_1 is strictly larger than the payoff for r_2 no matter what choice Column makes. The strategy r_1 **weakly dominates** r_2 if the payoff from r_1 is at least as large for all choices Column might make and strictly larger for some choice.

A **dominant strategy equilibrium** is a choice of strategies by each player such that each strategy (weakly) dominates every other strategy available to that player.

One particularly interesting game that has a dominant strategy equilibrium is the Prisoner's Dilemma in which the dominant strategy equilibrium is (Defect, Defect). If I belief that the other agent will Cooperate, then it is to my advantage to Defect; and if I believe that the other agent will Defect, it is still to my advantage to Defect.

Clearly, a dominant strategy equilibrium is a Nash equilibrium, but not all Nash equilibria are dominant strategy equilibria. A dominant strategy equilibrium, should one exist, is an especially compelling solution to the game, since there is a unique optimal choice for each player.

15.9 Elimination of dominated strategies

When there is no dominant strategy equilibrium, we have to resort to the idea of a Nash equilibrium. But typically there will be more than one Nash equilibrium. Our problem then is to try to eliminate some of the Nash equilibria as being "unreasonable."

One sensible belief to have about players' behavior is that it would be unreasonable for them to play strategies that are are dominated by other strategies. This suggests that when given a game, we should first eliminate all strategies that are dominated and then calculate the Nash equilibria of

the remaining game. This procedure is called **elimination of dominated strategies**; it can sometimes result in a significant reduction in the number of Nash equilibria.

For example consider the game

A game with dominated strategies

		Player B: Left	Player B: Right
Player A	Top	2, 2	0, 2
	Bottom	2, 0	1, 1

Note that there are two pure strategy Nash equilibria, (Top, Left) and (Bottom, Right). However, the strategy Right weakly dominates the strategy Left for the Column player. If the Row agent assumes that Column will never play his dominated strategy, the only equilibrium for the game is (Bottom, Right).

Elimination of *strictly* dominated strategies is generally agreed to be an acceptable procedure to simplify the analysis of a game. Elimination of *weakly* dominated strategies is more problematic; there are examples in which eliminating weakly dominated strategies appears to change the strategic nature of the game in a significant way.

15.10 Sequential games

The games described so far in this chapter have all had a very simple dynamic structure: they were either one-shot games or a repeated sequence of one-shot games. They also had a very simple information structure: each player in the game knew the other player's payoffs and available strategies, but did not know in advance the other player's actual choice of strategies. Another way to say this is that up until now we have restricted our attention to games with **simultaneous moves**.

But many games of interest do not have this structure. In many situations at least *some* of the choices are made sequentially, and one player may know the other player's choice *before* he has to make his own choice. The analysis of such games is of considerable interest to economists since many economic games have this structure: a monopolist gets to observe consumer demand behavior before it produces output, or a duopolist may

Table 15.5 **The payoff matrix of a simultaneous-move game.**

		Player B Left	Player B Right
Player A	Top	1,9	1,9
	Bottom	0,0	2,1

observe his opponent's capital investment before making its own output decisions, etc. The analysis of such games requires some new concepts.

Consider for example, the following simple game.

It is easy to verify that there are two pure strategy Nash equilibria in this game, (Top, Left) and (Bottom, Right). Implicit in this description of this game is the idea that both players make their choices simultaneously, without knowledge of the choice that the other player has made. But suppose that we consider the game in which Row must choose first, and Column gets to make his choice after observing Row's behavior.

In order to describe such a sequential game it is necessary to introduce a new tool, the **game tree**. This is simply a diagram that indicates the choices that each player can make at each point in time. The payoffs to each player are indicated at the "leaves" of the tree, as in Figure 15.1. This game tree is part of the a description of the game in **extensive form**.

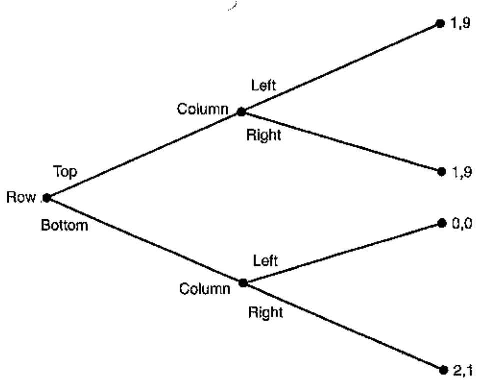

Figure 15.1 **A game tree.** This illustrates the payoffs to the previous game when Row gets to move first.

The nice thing about the tree diagram of the game is that it indicates the dynamic structure of the game—that some choices are made before others. A choice in the game corresponds to the choice of a branch of the tree. Once a choice has been made, the players are in a **subgame** consisting of the strategies and payoffs available to them from then on.

It is straightforward to calculate the Nash equilibria in each of the possible subgames, particularly in this case since the example is so simple. If Row chooses top, he effectively chooses the very simple subgame in which Column has the only remaining move. Column is indifferent between his two moves, so that Row will definitely end up with a payoff of 1 if he chooses Top.

If Row chooses Bottom, it will be optimal for Column to choose Right, which gives a payoff of 2 to Row. Since 2 is larger than 1, Row is clearly better off choosing Bottom than Top. Hence the sensible equilibrium for this game is (Bottom, Right). This is, of course, one of the Nash equilibria in the simultaneous-move game. If Column announces that he will choose Right, then Row's optimal response is Bottom, and if Row announces that he will choose Bottom then Column's optimal response is Right.

But what happened to the *other* equilibrium, (Top, Left)? If Row believes that Column will choose Left, then his optimal choice is certainly to choose Top. But why should Row believe that Column will actually choose Left? Once Row chooses Bottom, the optimal choice in the resulting subgame is for Column to choose Right. A choice of Left at this point is not an equilibrium choice in the relevant subgame.

In this example, only one of the two Nash equilibria satisfies the condition that it is not only an overall equilibrium, *but also an equilibrium in each of the subgames.* A Nash equilibrium with this property is known as a **subgame perfect** equilibrium.

It is quite easy to calculate subgame-perfect equilibria, at least in the kind of games that we have been examining. One simply does a "backwards induction" starting at the last move of the game. The player who has the last move has a simple optimization problem, with no strategic ramifications, so this is an easy problem to solve. The player who makes the second to the last move can look ahead to see how the player with the last move will respond to his choices, and so on. The mode of analysis is similar to that of dynamic programming; see Chapter 19, page 359. Once the game has been understood through this backwards induction, the agents play it going forwards.[2]

The extensive form of the game is also capable of modeling situations where some of the moves are sequential and some are simultaneous. The necessary concept is that of an **information set**. The information set of

[2] Compare to Kierkegaard (1938): "It is perfectly true, as philosophers say, that life must be understood backwards. But they forget the other proposition, that it must be lived forwards." [465]

an agent is the set of all nodes of the tree that cannot be differentiated by the agent. For example, the simultaneous-move game depicted at the beginning of this section can be represented by the game tree in Figure 15.2. In this figure, the shaded area indicates that Column cannot differentiate which of these decisions Row made at the time when Column must make his own decision. Hence, it is just as if the choices are made simultaneously.

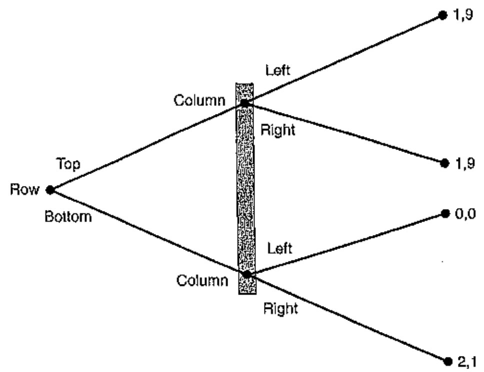

Figure 15.2 **Information set.** This is the extensive form to the original simultaneous-move game. The shaded information set indicates that Column is not aware of which choice Row made when he makes his own choice.

Thus the extensive form of a game can be used to model everything in the strategic form *plus* information about the sequence of choices and information sets. In this sense the extensive form is a more powerful concept than the strategic form, since it contains more detailed information about the strategic interactions of the agents. It is the presence of this additional information that helps to eliminate some of the Nash equilibria as "unreasonable."

EXAMPLE: A simple bargaining model

Two players, A and B, have \$1 to divide between them. They agree to spend at most three days negotiating over the division. The first day, A will make an offer, the next day B comes back with a counteroffer, and on the third day A gets to make one final offer. If they cannot reach an agreement in three days, both players get zero.

A and B differ in their degree of impatience: A discounts payoffs in the future at a rate of α per day, and B discounts payoffs at a rate of β per day. Finally, we assume that if a player is indifferent between two offers, he will accept the one that is most preferred by his opponent. This idea is that the opponent could offer some arbitrarily small amount that would make the player strictly prefer one choice, and that this assumption allows us to approximate such an "arbitrarily small amount" by zero. It turns out that there is a unique subgame perfect equilibrium of this bargaining game.[3]

As suggested above, we start our analysis at the end of the game, right before the last day. At this point A can make a take-it-or-leave-it offer to B. Clearly, the optimal thing for A to do at this point is to offer B the smallest possible amount that he would accept, which, by assumption, is zero. So if the game actually lasts three days, A would get 1 and B would get zero (i.e., an arbitrarily small amount).

Now go back to the previous move, when B gets to propose a division. At this point B should realize that A can guarantee himself 1 on the next move by simply rejecting B's offer. A dollar next period is worth α to A this period, so any offer less than α would be sure to be rejected. B certainly prefers $1-\alpha$ now to zero next period, so he should rationally offer α to A, which A will then accept. So if the game ends on the second move, A gets α and B gets $1-\alpha$.

Now move to the first day. At this point A gets to make the offer and he realizes that B can get $1-\alpha$ if he simply waits until the second day. Hence A must offer a payoff that has at least this present value to B in order to avoid delay. Thus he offers $\beta(1-\alpha)$ to B. B finds this (just) acceptable and the game ends. The final outcome is that the game ends on the first move with A receiving $1-\beta(1-\alpha)$ and B receiving $\beta(1-\alpha)$.

Figure 15.3A illustrates this process for the case where $\alpha = \beta < 1$. The outermost diagonal line shows the possible payoff patterns on the first day, namely all payoffs of the form $x_A + x_B = 1$. The next diagonal line moving towards the origin shows the present value of the payoffs if the game ends in the second period: $x_A + x_B = \alpha$. The diagonal line closest to the origin shows the present value of the payoffs if the game ends in the third period; this equation for this line is $x_A + x_B = \alpha^2$. The right angled path depicts the minimum acceptable divisions each period, leading up to the final subgame perfect equilibrium. Figure 15.3B shows how the same process looks with more stages in the negotiation.

It is natural to let the horizon go to infinity and ask what happens in the infinite game. It turns out that the subgame perfect equilibrium division

[3] This is a simplified version of the Rubinstein-Ståhl bargaining model; see the references at the end of the chapter for more detailed information.

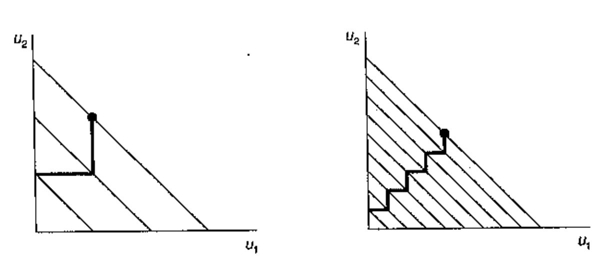

Figure 15.3 **A bargaining game.** The heavy line connects together the equilibrium outcomes in the subgames. The point on the outermost line is the subgame-perfect equilibrium.

is

$$\text{payoff to } A = \frac{1-\beta}{1-\alpha\beta}$$
$$\text{payoff to } B = \frac{\beta(1-\alpha)}{1-\alpha\beta}.$$

Note that if $\alpha = 1$ and $\beta < 1$, then player A receives the entire payoff, in accord with the principal expressed in the Gospels: "Let patience have her [subgame] perfect work." (James 1:4).

15.11 Repeated games and subgame perfection

The idea of subgame perfection eliminates Nash equilibria that involve players threatening actions that are not credible—i.e., they are not in the interest of the players to carry out. For example, the Punishment Strategy described earlier is not a subgame perfect equilibrium. If one player actually deviates from the (Cooperate, Cooperate) path, then it is not necessarily in the interest of the other player to actually defect *forever* in response. It may seem reasonable to punish the other player for defection to some degree, but punishing forever seems extreme.

A somewhat less harsh strategy is known as Tit-for-Tat: start out cooperating on the first play and on subsequent plays do whatever your opponent did on the previous play. In this strategy, a player is punished for defection, but he is only punished once. In this sense Tit-for-Tat is a "forgiving" strategy.

Although the punishment strategy is not subgame perfect for the repeated Prisoner's Dilemma, there are strategies that can support the cooperative solution that *are* subgame perfect. These strategies are not easy to describe, but they have the character of the West Point honor code: each player agrees to punish the other for defecting, *and* also punish the other for failing to punish another player for defecting, and so on. The fact that you will be punished if you don't punish a defector is what makes it subgame perfect to carry out the punishments.

Unfortunately, the same sort of strategies can support many other outcomes in the repeated Prisoner's Dilemma. The **Folk Theorem** asserts that essentially all distributions of utility in a repeated one-shot game can be equilibria of the repeated game.

This excess supply of equilibria is troubling. In general, the larger the strategy space, the more equilibria there will be, since there will be more ways for players to "threaten" retaliation for defecting from a given set of strategies. In order to eliminate the "undesirable" equilibria, we need to find some criterion for eliminating strategies. A natural criterion is to eliminate strategies that are "too complex." Although some progress has been made in this direction, the idea of complexity is an elusive one, and it has been hard to come up with an entirely satisfactory definition.

15.12 Games with incomplete information

Up until now we have been investigating games of complete information. In particular, each agent has been assumed to know the payoffs of the other player, and each player knows that the other agent knows this, etc. In many situations, this is not an appropriate assumption. If one agent doesn't know the payoffs of the other agent, then the Nash equilibrium doesn't make much sense. However, there is a way of looking at games of incomplete information due to Harsanyi (1967) that allows for a systematic analysis of their properties.

The key to the Harsanyi approach is to subsume all of the uncertainty that one agent may have about another into a variable known as the agent's **type**. For example, one agent may be uncertain about another agent's valuation of some good, about his or her risk aversion and so on. Each type of player is regarded as a different player and each agent has some prior probability distribution defined over the different types of agents.

A **Bayes-Nash equilibrium** of this game is then a set of strategies for each type of player that maximizes the expected value of each type of player, given the strategies pursued by the other players. This is essentially the same definition as in the definition of Nash equilibrium, except for the additional uncertainty involved about the type of the other player. Each player knows that the other player is chosen from a set of possible types, but doesn't know exactly which one he is playing. Note in order to have a complete description of an equilibrium we must have a list of strategies for *all* types of players, not just the actual types in a particular situation, since each individual player doesn't know the actual types of the other players and has to consider all possibilities.

In a simultaneous-move game, this definition of equilibrium is adequate. In a sequential game it is reasonable to allow the players to update their beliefs about the types of the other players based on the actions they have observed. Normally, we assume that this updating is done in a manner consistent with Bayes' rule.[4] Thus, if one player observes the other choosing some strategy s, the first player should revise his beliefs about what type the other player is by determining how likely it is that s would be chosen by the various types.

EXAMPLE: A sealed-bid auction

Consider a simple sealed-bid auction for an item in which there are two bidders. Each player makes an independent bid without knowing the other player's bid and the item will be awarded to the person with the highest bid. Each bidder knows his own valuation of the item being auctioned, v, but neither knows the other's valuation. However, each player believes that the other person's valuation of the item is uniformly distributed between 0 and 1. (And each person *knows* that each person believes this, etc.)

In this game, the type of the player is simply his valuation. Therefore, a Bayes-Nash equilibrium to this game will be a *function*, $b(v)$, that indicates the optimal bid, b, for a player of type v. Given the symmetric nature of the game, we look for an equilibrium where each player follows an identical strategy.

It is natural to guess that the function $b(v)$ is strictly increasing; i.e., higher valuations lead to higher bids. Therefore, we can let $V(b)$ be its inverse function so that $V(b)$ gives us the valuation of someone who bids b. When one player bids some particular b, his probability of winning is the probability that the other player's bid is less than b. But this is simply the probability that the other player's valuation is less than $V(b)$. Since v is uniformly distributed between 0 and 1, the probability that the other player's valuation is less than $V(b)$ is $V(b)$.

Hence, if a player bids b when his valuation is v, his expected payoff is

$$(v - b)V(b) + 0[1 - V(b)].$$

The first term is the expected consumer's surplus if he has the lowest bid; the second term is the zero surplus he receives if he is outbid. The optimal bid must maximize this expression, so

$$(v - b)V'(b) - V(b) = 0.$$

For each value of v, this equation determines the optimal bid for the player as a function of v. Since $V(b)$ is by hypothesis the function that describes the relationship between the optimal bid and the valuation, we must have

$$(V(b) - b)V'(b) \equiv V(b)$$

[4] See Chapter 11, page 191 for a discussion of Bayes' rule.

for all b.

The solution to this differential equation is

$$V(b) = b + \sqrt{b^2 + 2C},$$

where C is a constant of integration. (Check this!) In order to determine this constant of integration we note that when $v = 0$ we must have $b = 0$, since the optimal bid when the valuation is zero must be 0. Substituting this into the solution to the differential equation gives us

$$0 = 0 + \sqrt{2C},$$

which implies $C = 0$. It follows that $V(b) = 2b$, or $b = v/2$, is a Bayes-Nash equilibrium for the simple auction. That is, it is a Bayes-Nash equilibrium for each player to bid half of his valuation.

The way that we arrived at the solution to this game is reasonably standard. Essentially, we guessed that the optimal bidding function was invertible and then derived the differential equation that it must satisfy. As it turned out, the resulting bid function had the desired property. One weakness of this approach is that it only exhibits one particular equilibrium to the Bayesian game—there could in principle be many others.

As it happens, in this particular game, the solution that we calculated is unique, but this need not happen in general. In particular, in games of incomplete information it may well pay for some players to try to hide their true type. For example, one type may try to play the same strategy as some other type. In this situation the function relating type to strategy is not invertible and the analysis is much more complex.

15.13 Discussion of Bayes-Nash equilibrium

The idea of Bayes-Nash equilibrium is an ingenious one, but perhaps it is too ingenious. The problem is that the reasoning involved in computing Bayes-Nash equilibria is often very involved. Although it is perhaps not unreasonable that purely *rational* players would play according to the Bayes-Nash theory, there is considerable doubt about whether real players are able to make the necessary calculations.

In addition, there is a problem with the predictions of the model. The choice that each player makes depends crucially on his beliefs about the distribution of various types in the population. Different beliefs about the frequency of different types leads to different optimal behavior. Since we generally don't observe players beliefs about the prevalence of various types of players, we typically won't be able to check the predictions of the model. Ledyard (1986) has shown that essentially *any* pattern of play is a Bayes-Nash equilibrium for some pattern of beliefs.

Nash equilibrium, in its original formulation, puts a consistency requirement on the beliefs of the agents—only those beliefs compatible with maximizing behavior were allowed. But as soon as we allow there to be many types of players with different utility functions, this idea loses much of its force. Nearly any pattern of behavior can be consistent with some pattern of beliefs.

Notes

The concept of Nash equilibrium comes from Nash (1951). The concept of Bayesian equilibrium is due to Harsanyi (1967). More detailed treatments of the simple bargaining model may be found in Binmore & Dagupta (1987).

This chapter is just a bare-bones introduction to game theory; most students will want to study this subject in more detail. Luckily several fine treatments have recently become available that provide a more rigorous and detailed treatment. For review articles see Aumann (1987), Myerson (1986), and Tirole (1988). For book-length treatments see the works by Kreps (1990), Binmore (1991), Myerson (1991), Rasmusen (1989), and Fudenberg & Tirole (1991).

Exercises

15.1. Calculate all the Nash equilibria in the game of Matching Pennies.

15.2. In a finitely repeated Prisoner's Dilemma game we showed that it was a Nash equilibrium to defect every round. Show that, in fact, this is the dominant strategy equilibrium.

15.3. What are the Nash equilibria of the following game after one eliminates dominated strategies?

Player A	Player B: Left		Right
Top	3, 3	0, 3	0, 0
	3, 0	2, 2	0, 2
Bottom	0, 0	2, 0	1, 1

15.4. Calculate the expected payoff to each player in the simple auction game described in the text if each player follows the Bayes-Nash equilibrium strategy, conditional on his value v.

15.5. Consider the game matrix given here.

Player A	Player B: Left	Right
Top	a, b	c, d
Bottom	e, f	g, h

(*a*) If (top, left) is a dominant strategy equilibrium, then what inequalities must hold among $a, \ldots, h$?

(*b*) If (top, left) is a Nash equilibrium, then which of the above inequalities must be satisfied?

(*c*) If (top, left) is a dominant strategy equilibrium, must it be a Nash equilibrium?

15.6. Two California teenagers Bill and Ted are playing Chicken. Bill drives his hot rod south down a one-lane road, and Ted drives his hot rod north along the same road. Each has two strategies: Stay or Swerve. If one player chooses Swerve he looses face; if both Swerve, they both lose face. However, if both choose Stay, they are both killed. The payoff matrix for Chicken looks like this:

Player A	Player B: Left	Right
Top	-3,-3	2,0
Bottom	0,2	1,1

(*a*) Find all pure strategy equilibria.

(*b*) Find all mixed strategy equilibria.

(*c*) What is the probability that both teenagers will survive?

15.7. In a repeated, symmetric duopoly the payoff to both firms is π_j if they produce the level of output that maximizes their joint profits and π_c if they produce the Cournot level of output. The maximum payoff that one player can get if the other chooses the joint profit maximizing output is π_d. The discount rate is r. The players adopt the punishment strategy of reverting to the Cournot game if either player defects from the joint profit-maximizing strategy. How large can r be?

15.8. Consider the game shown below:

Player A	Player B: Left		Right
Top	1, 0	1, 2	2,-1
	1, 1	1, 0	0,-1
Bottom	-3,-3	-3,-3	-3,-3

(*a*) Which of Row's strategies is strictly dominated no matter what Column does?

(*b*) Which of Row's strategies is weakly dominated?

(*c*) Which of Column's strategies is strictly dominated no matter what Row does?

(*d*) If we eliminate Column's dominated strategies, are any of Row's strategies weakly dominated?

15.9. Consider the following coordination game

		Player B: Left	Player B: Right
Player A	Top	2,2	-1,-1
	Bottom	-1,-1	1,1

(*a*) Calculate all the pure strategy equilibria of this game.

(*b*) Do any of the pure strategy equilibria dominate any of the others?

(*c*) Suppose that Row moves first and commits to either Top or Bottom. Are the strategies you described above still Nash equilibria?

(*d*) What are the subgame perfect equilibria of this game?

CHAPTER 16

OLIGOPOLY

Oligopoly is the study of market interactions with a small number of firms. The modern study of this subject is grounded almost entirely in the theory of games discussed in the last chapter. This is, of course, a very natural development. Many of the earlier *ad hoc* specifications of strategic market interactions have been substantially clarified by using the concepts of game theory. In this chapter we will investigate oligopoly theory primarily, though not exclusively, from this perspective.

16.1 Cournot equilibrium

We begin with the classic model of **Cournot equilibrium**, already mentioned as an example in the last chapter. Consider two firms which produce a homogeneous product with output levels y_1 and y_2, and thus an aggregate output of $Y = y_1 + y_2$. The market price associated with this output (the inverse demand function) is taken to be $p(Y) \equiv p(y_1 + y_2)$. Firm i has a cost function given by $c_i(y_i)$ for $i = 1, 2$.

Firm 1's maximization problem is

$$\max_{y_1} \pi_1(y_1, y_2) = p(y_1 + y_2)y_1 - c_1(y_1).$$