Game Theory and the Law

Douglas G. Baird
Robert H. Gertner
Randal C. Picker

Harvard University Press
Cambridge, Massachusetts, and London, England

trons (a tough guy and a wimp) who prefer beer and quiche for breakfast respectively, but who encounter a bully who enjoys fighting wimps but not tough guys. The breakfast order of the tough guy or the wimp turns on the inferences each believes that the bully will draw from the order. Cho and Kreps use this story to motivate refinements that eliminate seemingly implausible solutions to this game (in particular, the solution in which both the tough guy and the wimp order quiche). In the last several years, the *beer-quiche game* has joined the prisoner's dilemma as one of the central paradigms of game theory.

Other important papers on formal models of nonverifiable information include Banks and Sobel (1987), Kohlberg and Mertens (1986), Grossman and Perry (1986), and Farrell (1983). A textbook, such as Fudenberg and Tirole (1991a) or Myerson (1991), is the best place to go for a comprehensive discussion. Gibbons (1992) has clear, simple examples. Our discussion of centralized and decentralized rules was influenced by Farrell (1987) and Bolton and Farrell (1990) as well as by Kaplow (1992) and Posner (1992). The classic model of adverse selection is Akerlof's (1970) analysis of market breakdown in the market for used cars. Signaling models have been applied to a great number of economic settings. Spence (1974) is one of the earliest and contains the model of education as a signal of productivity in the labor market. Applications of signaling models to contract law issues include Aghion and Hermalin (1990), which also models the problem of parental leave legislation.

Information revelation and contract default rules. The information revelation aspects of contractual terms are based on the analysis of Ayres and Gertner (1989). They discuss the problem of both *Hadley v. Baxendale* and *Peevyhouse v. Garland Coal and Mining Co.* Other papers that focus on these issues include Johnston (1990), Bebchuk and Shavell (1991), Ayres and Gertner (1992), and Goetz and Scott (1985). Maute (1993) offers a comprehensive study of *Peevyhouse.*

Signaling, screening, and the costs of information. Screening models have also been applied to a great number of different settings. The classic paper on insurance is by Rothschild and Stiglitz (1976). Stole (1992) examines penalty damages as a screening problem.

CHAPTER 5

Reputation and Repeated Games

To this point, we have examined interactions between individuals as isolated events. Individuals, however, often have repeated dealings with one another. They must take account of how the decisions they make in one interaction will affect what happens in future ones. We can gain an understanding of how people are likely to behave in these situations by examining the strategies that rational players choose in games that consist of a simple game that is repeated many times. We use these models to examine a number of antitrust problems, including tacit collusion, conscious parallelism, and predatory pricing.

Before we look at repeated games, however, we return to the backwards induction solution concept. Backwards induction is useful for analyzing many games, but it rests on assumptions that may generate implausible results when applied to games that are repeated many times. Hence, we must be careful about applying it here. In the next section, we examine backwards induction and its limits more closely by applying it to two commercial transactions, one involving negotiations before a formal contract is signed and the other dealing with an installment sale.

Backwards Induction and Its Limits

Many contract rules establish governing principles for contract negotiations. A good example is the legal rule that conditions the legal enforceability of a contract on the existence of a signed writing evidencing the agreement of the parties. This rule, called the Statute of Frauds, applies to many contracts, including any contract for the sale of land and any contract involving the sale of goods worth more than $500.

The Statute of Frauds does not require that the terms of the contract be set out in any detail. The writing need only show that the parties have entered into a contract; it does not require anything even as rudimentary as a price term.

We can ask whether such a rule is desirable. To be sure, parties often find it in their interest to reduce their contracts to writing, but it might seem that the parties themselves are best positioned to decide whether a written contract brings benefits worth its costs. The rule does make it hard for people to assert that they have a contract with someone with whom they have never dealt. It is difficult, however, to justify the rule on this basis. The risk of liability that a complete stranger poses to a commercial actor seems a small one. After all, a complete stranger will not be able to persuade a finder of fact that a deal existed without being able to provide some evidence of a deal. In the absence of any credible evidence that the stranger ever met the person being sued, it is hard to see how such a lawsuit could be successful.

The Statute of Frauds, however, may reduce the risk of problems that can arise during the course of negotiations that typically take place before a contract is entered. The facts of *Southwest Engineering Co. v. Martin Tractor Co.* illustrate this kind of problem. Southwest was a general contractor that was installing runway lighting at an Air Force base. It needed to buy a stand-by generator and made inquiries with Martin Tractor. A series of negotiations eventually led to a lunch in an airport cafeteria. During this lunch, Martin's employee took out a piece of paper and listed the components of the generator as well as their price. Southwest's employee took a copy of this piece of paper.

The deal fell through a few weeks later and Martin attempted to withdraw "all verbal quotations." Southwest brought suit against Martin, alleging that they had reached an agreement during the lunch. Martin defended its position on the ground that, regardless of whether an agreement was made, it was not legally enforceable because the price quotation that its employee had written on the piece of paper was not sufficient to satisfy the Statute of Frauds. The court ultimately rejected this argument.

In commercial dealings such as the one in *Southwest Engineering*, a series of meetings between the buyer and the seller is commonplace. Parties may reach agreement on many terms, and other terms may be open. In a world in which there is no Statute of Frauds, the moment at which a legally enforceable contract is formed may be unclear. There may be several points during the negotiations at which one party or the other could insist that a deal of some kind had been reached even

when (at least in the view of the other party) none had been reached. *Southwest Engineering* may be such a case. Southwest's employee may have thought that there was a deal, whereas Martin's employee may have thought that details needed to be worked out or that approval of the home office was still necessary.

Before parties sign a formal contract, they may negotiate with each other over an extended period of time. They may discuss possible deals orally. In other cases, a price term may be set down in writing. At every point, however, it may be uncertain whether a contract exists or what its terms are. Consider what might happen if no Statute of Frauds existed. If the parties stop short of signing a written contract, one or the other party may be able to assert that a binding contract nevertheless exists that has terms favorable to that particular party.

We can model this problem as a game in which the players must repeatedly decide whether to continue to negotiate or to exit from the negotiations. The payoff that a party enjoys after one or the other exits turns on whether the negotiations have progressed so far that one or the other can claim that a legally enforceable contract exists. Figure 5.1 presents the simplest version of such a game. Each party has a chance to exit during the negotiations and assert that a deal exists that is favorable to that party but unfavorable to the other. A player who exits enjoys a payoff of \$5 but leaves the other with a payoff of −\$5. If neither party exits and both sign a formal contract, each receives a payoff of \$3.

We can solve the game in Figure 5.1 by using backwards induction. At the start of the game, the buyer must decide whether to ask for a

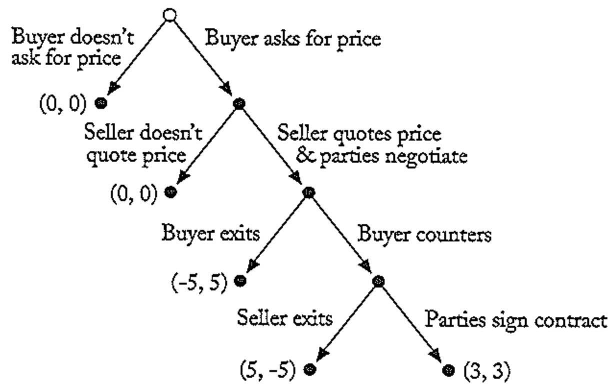

Figure 5.1 No writing requirement. *Payoffs:* Seller, Buyer.

price quote. The buyer cannot decide without forecasting the subsequent play of the game. One useful way to do this is for the buyer to look at how the seller behaves at the last node. At the last node, the seller's best response is to exit rather than to sign the contract. For this reason, the buyer's best response at the next-to-last node is to exit as well. Exit brings a payoff of $5, whereas continuing to negotiate brings the buyer a loss of $5, given that the seller exits if offered the chance. At the previous node, the seller's best response is again to exit. The seller prefers a payoff of $0 to a loss of $5. Because the buyer takes these best responses into account at the start of the game, the buyer is indifferent between never commencing negotiations and asking for a price quote, because in either case the buyer will receive a payoff of $0.

Serious negotiations are never undertaken because of the risk that one or both of the parties to the negotiation will opportunistically claim that a contract in fact exists. As a result, the parties cannot reach the best outcome—that in which a contract is signed and each party receives a payoff of $3. The Statute of Frauds, however, may solve this kind of problem. If a legal action can be brought only if a contract is signed, no player can gain an advantage from exiting. If the payoffs of (0, 0) and (0, 0) replace the payoffs of (−5, 5) and (5, −5), the parties will complete the contract and enter into a mutually beneficial trade.

Under this view, the Statute of Frauds may make parties more willing to begin negotiations with each other. If the Statute of Frauds is to serve this purpose, however, it must require more than evidence that negotiations took place. In *Southwest Engineering* and other Statute of Frauds cases, no one doubts that there were negotiations. A writing requirement, if we have one at all, should be strong enough to provide an answer to the question of whether the negotiations led anywhere. Because the writing in *Southwest Engineering* did not show any more than that the parties had met and negotiated, one can argue that the Statute of Frauds was not satisfied and that the case should have been decided differently.

The extensive form game in Figure 5.1 isolates a principle that may underlie the Statute of Frauds. This perspective may also help us to appreciate why debates over rules such as the Statute of Frauds are so central to the conceptual foundations of contract law. This analysis can also be extended in a straightforward way to other rules governing the formation of contracts, such as those that refuse to recognize agreements that are too indefinite. Unlike warranty and other gap-filling terms, rules such as these set the climate in which negotiations take

place. It is hard for parties to opt out of them, and the stakes involved in getting them right are much higher.

Backwards induction has an intuitive appeal and is easy to apply in many games. Nevertheless, it may lead to solutions that seem implausible. Consider the following game. Seller and Buyer contemplate entering into an agreement. Seller would promise to ship Buyer goods in 100 equal installments and Buyer would promise to pay for each installment after receiving it. Seller has no recourse against Buyer if Buyer refuses to pay, perhaps because legal remedies are too costly. (Seller's only option is to cease any further shipments.) Similarly, Buyer has no recourse if Seller chooses not to perform. It costs Seller $1 to make each installment. Buyer pays $2 for each shipment of goods and the goods have a value to Buyer of $3. The game between Seller and Buyer takes the extensive form game illustrated in Figure 5.2.

To use backwards induction to discover the likely course of play, we

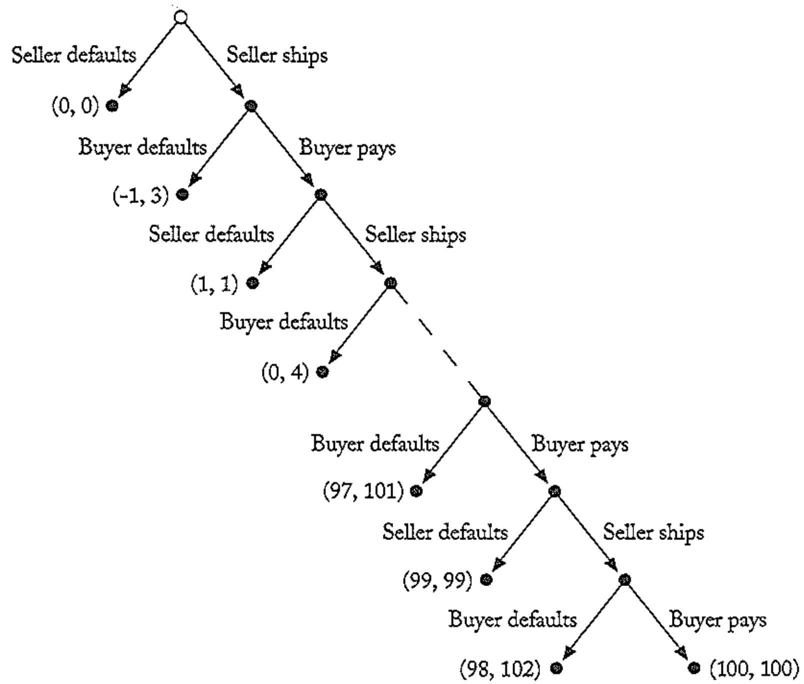

Figure 5.2 Installment sale. *Payoffs:* Seller, Buyer.

again start at the end of the game, with Buyer's decision after it has received the last installment. It can either refuse to pay for the goods and enjoy a payoff of $102 or pay $2 for the last installment and enjoy a payoff of only $100. Buyer is thus better off if it does not pay. Seller anticipates that Buyer will not pay for the last installment. On its last move, Seller therefore faces a payoff of $98 if it ships the goods, but $99 if it does not. Back one more step, Buyer recognizes that Seller will not ship the last installment of goods and Buyer therefore will not pay for the second-to-last shipment of goods. (Buyer enjoys a payoff of $101 if it defaults, instead of a payoff of only $99 if it pays for the goods.) Seller makes a similar calculation on its next-to-last move, and the process continues until we reach Seller's first move—in which it decides not to ship at all. In short, backwards induction suggests that the entire deal will unravel because each party will anticipate that the other will default on its next move.

Common sense tells us that applying backwards induction to this model does not capture how Buyer and Seller would deal with each other. Buyer and Seller might enter into this transaction because some force is at work that we have not taken into account. Buyer, for example, may have trouble buying goods from others if it failed to pay Seller. Moreover, both parties might keep their promise because of the prospect of future dealings with each other. Even apart from reputational issues, however, backwards induction seems to do a poor job of predicting how the game itself might be played.

This game is one that involves many iterations and in which there is only a small cost to a player of deviating from the equilibrium strategy of defaulting at every opportunity. For this reason, the game places unusual weight on the assumption that the payoffs that the parties enjoy are common knowledge. We assume that Buyer and Seller are certain that the game has the payoffs that we show, that each knows that the other player is certain, that each knows that the other player knows, and so forth. They know that if the other did not default and either shipped the goods or paid for them, it would be due to a mistake, a *tremble*, rather than an intentional decision to adopt something other than the equilibrium strategy.

Consider, however, how the solution to the game might change if either player was not certain that a deviation from the equilibrium was simply a mistake that was not likely to be repeated. Focus, for example, on Buyer's move after it receives the first shipment from Seller. It must decide whether to default on its payment obligation and enjoy a benefit of $3 or pay and risk that Seller will default on the next move and leave

it with only $1. If Buyer is not certain that Seller has made a mistake, it must consider how it should respond given that Seller has already acted in a way contrary to the predicted equilibrium.

If Buyer is only slightly unsure about Seller's payoffs, Buyer's best response in a game such as this may be to experiment and cooperate for an additional period. Buyer has little to lose if Seller turns out to have made a mistake, but it has a great deal to gain if Seller's payoffs are in fact different from what they appear to be. Once Buyer is willing to experiment by continuing the game for an additional period, however, Seller has something to gain from making the first installment. Even if Seller's payoffs are exactly what Buyer thinks, Seller may still be better off shipping the goods as long as Buyer harbors some doubts.

We return at the end of the chapter to the question of how uncertainty in the payoffs can affect the play of repeated games. At the moment, however, we want to sound a note of caution. The assumption of common knowledge is a strong one. We must be careful about using it, particularly in games with many iterations and in which deviations from the equilibrium come at low cost and may bring substantial benefits relative to the equilibrium outcome. The test of a concept such as backwards induction is ultimately how well it captures the way people actually behave. Although this solution concept is quite powerful, one cannot apply it or any other solution concept blindly.

Infinitely Repeated Games, Tacit Collusion, and Folk Theorems

When parties interact with each other over time, they often do not know when the interactions will end. The solution to the installment sale game in Figure 5.2 turned on the game's having a definite terminus. In this section, we want to examine games in which there is no definite end. Manufacturers in a small market can earn supracompetitive profits if they simultaneously raise prices and reduce output. The antitrust laws ban explicit cartel agreements, but less clear is how they affect tacit agreements. Tacit agreements may exist when firms in a market increase prices and reduce output without actually entering into an explicit bargain. To understand this issue, we must first examine how tacit agreements could ever come into being and survive for any length of time. A useful way of doing this is to model the pricing and output decisions of those in an industry with few firms as a repeated game of indeterminate length.

In the 1940s, the leading cigarette manufacturers were accused of violations of antitrust laws in the pricing of cigarettes.[1] Pricing in the industry was identical across firms. When one firm changed prices, competitors immediately matched them. The government argued that this pattern allowed the firms to sustain prices at noncompetitive levels. The Supreme Court upheld the jury verdict against the tobacco manufacturers, even though the suit did not allege that there had been any explicit agreement or conspiracy to fix prices.

Subsequent cases, however, have made it clear that the plaintiff in an antitrust case must do more than show parallel behavior among the major firms. Indeed, even conscious parallel behavior falls outside the reach of the antitrust laws. For conduct to run afoul of antitrust laws, some additional "plus factors" must also be at work. In this section, we explore how firms can cooperate on price without explicit agreement and how a variety of practices may facilitate their ability to collude tacitly. These in turn cast light on what one should have to show beyond parallel behavior in order to justify taking action against particular firms.

All models of tacit collusion rely on both the dynamic nature of price competition and the ability of firms to respond to their rivals' pricing decisions. To develop the ideas formally in a simple model, let us begin by assuming that there are two firms in the industry and there are only two prices they can charge: High or Low. In each period they simultaneously decide which price to charge. Profits in that period are given in the bimatrix shown in Figure 5.3.

Joint profits are greatest if both firms choose a high price. If the game is played only once, however, charging a low price is a dominant strategy for each player. If the second firm charges a high price, the first firm earns $6 more by charging a low price rather than a high price. If the second firm charges a low price, the first firm earns $5 more by charging a low price rather than a high price. Thus, no matter what the second firm does, the first firm earns more by charging a low price. Because the game is symmetric, the second firm must also find it in its interest to charge a low price.

		Firm 2	
		High	Low
Firm 1	High	10, 10	0, 16
	Low	16, 0	5, 5

Figure 5.3 Tacit collusion. *Payoffs:* Firm 1, Firm 2.

This game has the same structure as the prisoner's dilemma. Both players adopt the low-price strategy, even though they would be jointly better off if they adopted the high-price strategy. A game in which the firms could choose their price from a continuum is also a game where the unique equilibrium is one in which both firms charge a low price. (The equilibrium is Nash, however, rather than one in which each player has a strategy that is strictly dominant. For any given price above marginal cost, the best response is to undercut it by only a small amount.)

Let us now imagine that the game is repeated a finite number of times. We encounter the same problem we saw in our discussion of the installment sale game represented in Figure 5.2. Each player posits the play of the last period. In the last period, neither has any concern about how an action in this period will affect future play because, by definition, the game will end after this period. In the last period, both act as if they are playing a one-shot game. Each will therefore play its dominant strategy and charge the low price.

Given the outcome in the last period, both players will find it in their interest to charge a low price in the next-to-last period as well. Each player knows that the other will charge a low price in the last period no matter what happens in this game. Hence, both will play in the next-to-last period as if it were a one-shot game as well. They will engage in the same process in the period before the next-to-last one also. The process repeats itself until the players reach the first period. Backwards induction suggests that the players will charge the low price in every period. Mere repetition of the strategic interaction does nothing to allow cooperative pricing.

As in the installment sale game, backwards induction predicts this course of play because of strong assumptions of common knowledge. In this case, the equilibrium turns crucially on the number of repetitions of the game being common knowledge. This assumption is unrealistic. There is no set amount of times that firms in an industry compete with each other, and our model should take this into account.

The finitely repeated game is the appropriate model only if the parties know with certainty exactly when the game will end. For this reason, we focus in this section on infinitely repeated games. A game of infinite length seems equally unrealistic, but, as we shall show, a game of infinite length and a game of uncertain length with a fixed probability of ending after each period have the same structure.

Let us assume that the two firms play the game in Figure 5.3 an infinite number of times. A game that consists of an infinite number of repetitions of another game is a *supergame*. Once we move to infinitely repeated games, we need to understand the relation between payoffs in one game and payoffs in later ones. If a payoff of a fixed amount in the future were worth as much to a player as the same amount today, a player would be indifferent between a strategy which leads to a payoff of $10 each period and a strategy which leads to a payoff of $5 each period. Each strategy yields payoffs that, when summed, are arbitrarily large.

We shall assume that a player values a payoff in a future period less than a payoff in the present period. We account for this by introducing a discount factor, the amount by which the value of a payoff in the next period must be adjusted to reflect its value in the present period. If we have a discount factor of δ, the present value of one dollar earned in the subsequent period is δ dollars.[2]

We can now compare this game to one whose duration is uncertain. The uncertainty could be due to a possible exit by one firm, introduction of a superior product by a competitor, government regulation, or a variety of other factors. Let us assume that the firms believe that the probability that the game will end in any period is $1 - d$, so the probability that it will continue is d. When deciding on a strategy in period 1, a firm discounts the next period's profits by δ times the probability that the game will continue another period, which is d. The firm discounts profits two periods in the future by δ^2 times the probability that it will still be playing, which is d^2.

When we model games of uncertain length, we can use the product of δ and d in every term. More simply, we can continue to use δ, but interpret it as the product of the probability that a game will end in a given period and the amount by which payoffs in each subsequent period are discounted. A game with an uncertain end is equivalent to the infinitely repeated game with a discount factor that is again discounted by a factor equal to the probability that the game will continue in any period. The infinitely repeated game is appropriate as long as there is some uncertainty about when the game will end.

Because there is no last period, we do not have a definite end point from which to start the process of backwards induction. To solve these games, one must describe different strategies and see if any combination of them forms a subgame perfect Nash equilibrium. For example, the strategies "charge a high price in even periods and a low price in odd periods" and "charge a high price in the first period and, in subsequent periods, copy what the other player did on its previous move" are both fully specified strategies in the sense that they tell us how a player moves at every information set in the game. We can then

ask whether such strategies are subgame perfect by looking at all the possible subgames that could arise both on and off the equilibrium path.

Consider first the following *grim* (or *trigger*) *strategy:* Both firms adopt the strategy of charging a high price in the first period. In all subsequent periods, the firms charge a low price if either firm has ever charged a low price in any previous period, otherwise they charge a high price. First, we will confirm that each firm choosing this strategy forms a Nash equilibrium if the discount factor δ is sufficiently high. Let us consider what Firm 1 gets if it plays its equilibrium strategy. If both firms play their equilibrium strategies, they will each end up charging the high price in each period, since neither initiates a price war. Thus, the present value of profits for Firm 1 is $10/(1 - \delta)$.

At the beginning of each period, the present value of expected profits is equal. The continuation game at the beginning of the tenth period looks identical to the game at the beginning of the first. This game, in other words, is *stationary*. The strategy space of a player, the choices available to a player at any point in the game in which that player moves, remains the same. This stationary structure simplifies the analysis of infinitely repeated games.

We need to ask what will happen if Firm 1 deviates from its strategy and decides to charge a low price at some point, even though both players have charged a high price in all preceding periods. The profit from charging a low price this period is $16. In all subsequent periods, however, Firm 1 will receive $5. (For both firms to charge a low price from then on is a subgame perfect equilibrium.) Thus, the returns from charging a low price are $16 + 5\delta/(1 - \delta)$. Charging a high price when the other player adopts a grim, or trigger, strategy is therefore more profitable than charging a low price as long as

$$\frac{10}{1 - \delta} > 16 + \frac{5\delta}{1 - \delta}.$$

By solving this inequality, we learn that charging a high price is better than charging a low price as long as $\delta > 6/11$. A player receives higher profits today by charging a low price, but the other player responds by charging a low price in all subsequent periods. This punishment brings lower profits in the future that may offset the short-term gains. A player has less incentive to charge a low price as δ rises. Indeed, as δ rises, a player cares more about the profits in future rounds relative to profits in the current round. If $\delta > 6/11$, the strategy combination in which the players both charge a high price until the other charges a low price forms a Nash equilibrium.

Given that Firm 2 is playing this strategy, Firm 1 earns $10/(1 - \delta)$ if it plays its equilibrium strategy of charging a high price as long as Firm 2 has never charged a low price. We have just shown that changing to a strategy in which it charges a low price leaves Firm 1 worse off when $\delta > 6/11$. The only other strategies that Firm 1 could adopt would be ones in which it did not initiate a low price. (Such a strategy would include, for example, one in which a player charges a high price until the other charges a low price and, in that event, charges a low price for five periods and then returns to the high price.) These strategies, however, also yield $10/(1 - \delta)$, given Firm 2's strategy of charging a high price unless Firm 1 charges a low price. All other strategies in which Firm 1 continues to charge a high price offer the same payoff as a trigger strategy. For this reason, a trigger strategy is a best response to the trigger strategy of Firm 2. Hence, a Nash equilibrium exists when both players adopt it.

This strategy combination is also a subgame perfect Nash equilibrium. We can see this by examining the possible subgames and making sure that the strategies form a Nash equilibrium in each of them. There are an infinite number of subgames (one for each possible history at each period), and we cannot check them one by one, but we can nevertheless examine the full spectrum of possible situations. Consider any subgame in which no firm has yet charged a low price. The argument of the preceding paragraph showed that the strategies form a Nash equilibrium in these subgames.

Next consider the subgames in which Firm 1 has charged a low price in some previous period. Firm 2's strategy is to charge a low price in all future periods. From this point on, Firm 2's strategy is independent of Firm 1's actions (it will always charge a low price), so Firm 1 should maximize current profits in each period. In each period, charging a low price is a dominant strategy for Firm 1. The same reasoning applies for Firm 2. Neither has reason to charge a high price when the other is going to charge a low price. Because each firm's strategy is a best response to its rival in these subgames, the strategies form a Nash equilibrium. Having covered all subgames, we have established subgame perfection.

The intuition that explains why these strategies form an equilibrium is simple. If both firms charge a high price initially, but charge a low price forever if the other ever charges a low price, neither player does

better by charging a low price. The short-term gains are insufficient to compensate for the future losses from an infinitely long price war. Given that the best response to a low price is always a low price, a Nash equilibrium exists when both firms charge a low price.

The trigger strategy is not the only equilibrium in the infinitely repeated game. As suggested, for example, it would also be a subgame perfect Nash equilibrium if each firm charged a low price in every period. Neither firm can increase its profits in any period from deviating. A high price lowers current profits and has no effect on future profits. In fact, as we shall soon see, there are an infinite number of subgame perfect Nash equilibria. When δ is sufficiently high, one such strategy is a version of *tit-for-tat*. A firm playing the strict tit-for-tat strategy begins by charging a high price and then charging a low price only if its rival charged a low price in the preceding period. If both players adopt tit-for-tat, each begins with a high price and charges a high price in every period.

Tit-for-tat stands in sharp contrast to trigger strategies in which a single transgression is punished in all subsequent rounds. If a firm deviates from tit-for-tat against a rival playing tit-for-tat, it goes back to a cooperative price, and the punisher returns to the cooperative price in the subsequent period. The firm that deviates suffers one period in which it charges a high price and earns nothing because the other player punishes by charging a low price. In subsequent periods, however, both parties can return to charging a high price. A single defection need not lead to a low payoff in all subsequent periods.

A Nash equilibrium of the supergame exists in which both players adopt tit-for-tat if δ is sufficiently high. There are two types of deviations to consider. First, a firm can charge a low price and never revert to a high price. This firm will enjoy high profits for a single period and then $5 forever. This is equivalent to the payoffs we analyzed earlier. Hence, all that is necessary to prevent this deviation is to ensure that $\delta > 6/11$.

Second, instead of playing the equilibrium strategy, a firm can charge a low price for a single period and then revert to a high price. If Firm 1 follows this strategy, it charges a low price in the first period and Firm 2 charges a high price. In the second period, Firm 1 charges a high price and Firm 2 charges a low price as punishment. The third period is just like the first period because Firm 1 knows that Firm 2 will charge a high price. Thus, if the deviation was profitable in period 1, it will also be profitable in period 3. The game will cycle. Firm 1's profits from following this strategy are $16 in odd periods and $0 in even periods. The discounted value of this deviation is $16/(1 - \delta^2)$. A player is better off playing the equilibrium strategy than this one as long as

$$\frac{10}{1-\delta} > \frac{16}{1-\delta^2}.$$

Hence, the equilibrium strategy is more profitable than alternating between high and low prices as long as $\delta > 6/10$. When δ is greater than $6/10$ (and thus sufficiently large to make deviation to alternating between high and low unattractive), a player's best response is to charge a high price in every period when the other player has adopted tit-for-tat. There are some values of δ (for example, those in which it is larger than $6/11$, but smaller than $6/10$) in which tit-for-tat is not an equilibrium even though infinite punishments for cheating are. For many values of δ, however, both strategies generate an equilibrium when each player adopts them.

We need a slightly different version of tit-for-tat if the strategy combination is to be subgame perfect. Each firm starts with a high price and charges a high price if both cooperated in the previous period or if both defected; otherwise each charges a low price. This formulation of tit-for-tat is necessary to ensure that there is cooperation rather than constant cycling in some subgames that are off the equilibrium path, in particular subgames that begin after one of the players has defected in the previous period.

So far, we have set out three symmetric subgame perfect Nash equilibria of the infinitely repeated game. Indeed, if the discount factor is sufficiently high, any pattern of high prices and low prices can be sustained with the threat to revert to low prices forever if anyone deviates from the pattern. This result can be generalized to all supergames with complete information. That is, any payoff that gives each player at least what that player receives in the single period Nash equilibrium forever, and that is feasible given the payoffs in the game, can be supported as a subgame perfect Nash equilibrium of the supergame. In Figure 5.3, any combination of payoffs that allows each player to average at least \$5 in each round is part of some subgame perfect Nash equilibrium. This result and others like it are known as *folk theorems* because they were part of the received lore of game theory before they appeared in any published papers.

In many repeated games, there are multiple subgame perfect Nash

equilibria, and we again face the familiar problem of choosing among them. We might first consider whether players would be likely to adopt the subgame perfect equilibrium that was *Pareto-optimal*. One might argue that, if one equilibrium gives each party a higher payoff than another, rational parties should never play the second equilibrium. If we further focus on symmetric equilibria, there will be unique predicted payoffs in the game, which may be supported by many equilibrium strategies.

This prediction, however, seems suspect. Pareto-optimal equilibria rely on very severe punishments. The greater the punishment from deviating, the more likely parties are to cooperate. These equilibria are plausible, however, only if all the players believe that the punishments are in fact carried out when someone deviates. The threat of the severe punishment needed to support the Pareto-optimal equilibria may not be credible. Parties may not find it in their self-interest to carry out severe punishments after there has been a deviation.

If parties always want the Pareto-optimal solution, they should adopt the Pareto-optimal equilibrium of the subgame that arises after the deviation. This Pareto-optimal equilibrium, however, would not take into account the need to punish the previous deviation. It would likely provide a less severe level of punishment than is necessary to support the original Pareto-optimal equilibrium. One can argue that parties are likely to adopt a Pareto-optimal equilibrium in the game as a whole only if one can explain why these same parties do not adopt strategies in each subgame that are also Pareto-optimal. If parties consistently gravitate toward the Pareto-optimal equilibrium, they should do it in subgames as well as in the game as a whole.

Another approach is to identify equilibria that are *renegotiation-proof*. Such an equilibrium is one in which the nondeviating party still has the incentive to carry out the punishment should it prove necessary. When the nondeviating party punishes the other party for deviating, the nondeviating party must be willing to incur the costs of inflicting punishment. This way of isolating the strategies that players are likely to adopt raises complications, however. In the subgames that follow any deviation from the equilibrium, the deviating party must willingly choose strategies that yield payoffs that are sufficiently low to make the original deviation unprofitable. One needs to explain why a player who deviates would also be willing to accept low payoffs so that both players can return to the equilibrium path.

None of the existing methods of narrowing the possible equilibria proves completely satisfactory. Nevertheless, the models do show that repetition itself creates the possibility of cooperative behavior. The mechanism that supports cooperation is the threat of future noncooperation for deviations from cooperation. One of the crucial determinants is the discount factor. The higher the discount factor, the more players have to gain from cooperation. If firms can change prices frequently, the appropriate period length is very short and the corresponding discount factor is very high. If firms can change prices only infrequently, as in catalogue sales, the appropriate discount factor will be much lower. What matters, of course, is not whether firms actually change their prices, but whether they have the ability to do so.

When firms in an industry seem to change their prices in unison, we should pay attention to how often prices change. Tacit collusion leading to noncompetitive pricing is much more likely in industries in which firms can adjust prices almost instantly (as they can, for example, in the airline industry) than in other industries (such as the mail-order business) in which they cannot. Firms in the latter industry are more likely to be responding to changes in supply and demand than to a tacit agreement among themselves.

One of the major impediments to tacit collusion is the inability to observe the prices charged by rivals. In many markets, prices are negotiated between buyers and sellers in a way that makes it impossible for rivals to observe each other's prices directly. How does a firm know when a rival has undercut the collusive price? Eventually one may figure it out by a drop in demand or by information from potential customers. If it takes a long time to learn, then the appropriate period length is quite long, and collusion may be difficult to support. Mistakes are also possible, as a firm's demand could drop for reasons other than a rival's secret price cut. In such environments, strategies that lead to cooperation may fare relatively poorly.

Consider what would happen if players were to adopt tit-for-tat when there was uncertainty about the actions of the other players. Assume that a player erroneously believes that the other player defected. The other player may see that player's decision to punish as a defection that triggers punishment in return. Tit-for-tat requires the first player to respond noncooperatively to the second player's response, which in turn generates a noncooperative response from the second player. Even though both players have adopted tit-for-tat, a single mistake could lead them to act noncooperatively forever, or until another mistake is made and a party mistakes punishment for cooperation. Other equilibria that support cooperation suffer from similar problems. For this reason, one might suspect that tacit collusion is more likely in industries

in which pricing information is readily available. A good example is the airline industry, in which price schedules are publicly distributed over a computer network to travel agents and others.

As the number of firms in an industry increases, collusion becomes less likely. By deviating to a lower price and expanding output, one may have a fairly small impact on each of a large number of firms. It becomes less likely that such firms will have an incentive to punish. The detection of cheating also becomes significantly more difficult as the number of firms increases. A firm with many competitors may experience a decrease in market share and have much more trouble determining whether one firm cheated or simply became more efficient.

Even if the other firms can tell that one firm has cheated, it will be hard for them to coordinate the appropriate period of punishment. It might not be in the interest of all the firms to punish the deviating firm by charging a low price, but there may be no way for them to agree on which firm should act as the enforcer. The Fourth Circuit in *Liggett Group, Inc. v. Brown and Williamson Tobacco Corp.*[3] noted exactly this problem:

> Oligopolists might indeed all share an interest in letting one among them discipline another for breaking step and might all be aware that all share this interest. One would conclude, however, that this shared interest would not itself be enough . . . The oligopolist on the sidelines would need to be certain at least that it could trust the discipliner not to expand the low-price segment itself during the fight or after its success.

Changes in the collusive price over time may limit the ability of firms to collude. Let us say that a firm receives private information about changes in demand or cost conditions that imply that the collusive price has dropped. If the firm lowers its price, rivals may think this is an attempt to cheat and respond with punishment. The result may be that firms are unable to change prices as often as is optimal or price wars may break out through mistaken punishments.

One practice that can solve this problem is price leadership, in which price changes are initiated by a single firm in the industry and are followed by all other firms. Differences among firms may also make collusion more difficult, as one firm may wish to cut prices because of lower costs or different demand for its product. Firms will not agree on a single optimal price for the industry, and each firm may compete to induce rivals to accept its preferred price. All of these results have again been captured in models that have the same basic structure as that in Figure 5.3.

Models such as this one suggest that antitrust enforcers pay attention to how rapidly prices change in an industry, how information about prices is conveyed, and how many players there are. We can ask if existing law shares the same focus by looking at the rules governing facilitating devices—industry practices that are thought to make collusion more likely. Specifically, we can look at the scrutiny that antitrust law gives to trade associations. A trade association often collects information about prices, costs, and demand from the member firms. This information may make the detection of cheating a great deal easier. In an oligopoly, resources devoted to monitoring competitors' behavior are a *public good*. The benefits of monitoring accrue to all firms. Thus, individual firms may try to *free ride* on their competitors' monitoring efforts, and less monitoring takes place than if firms could reach explicit agreement. A trade association can get around this problem.

The jurisprudence of trade associations and sharing information among competitors dates to 2 cases in the lumber industry from the 1920s. *American Column and Lumber Co. v. United States* concerned the American Hardwood Manufacturers' Association, a trade association of 400 U.S. hardwood manufacturers.[4] The member firms accounted for 5 percent of the mills in the United States and ⅓ of total production. The association collected and disseminated detailed price and quantity data from its members, inspected member firms, issued a report on changes in market conditions, and held monthly meetings of members. There may have also been recommendations not to increase production. Surveys of members indicated that some felt that the association acted to increase prices.

The Court found that the detailed disclosure of information was inconsistent with competition: "This is not the conduct of competitors, but is so clearly that of men united in an agreement, express or implied, to act together and pursue a common purpose under a common guide that, if it did not stand confessed a combination to restrict production and increase prices . . . that conclusion must inevitably have been inferred from the facts which were proved." In dissent, Justice Holmes and Justice Brandeis just as strongly argued that the sharing of information was consistent with competition.

Four years later the Court considered the same issues again. *Maple Flooring Manufacturers' Association v. United States* involved a twenty-two-member trade association that had 70 percent of the maple, beech, and birch flooring market. The association shared information about costs, freight rates, price, and quantity. It also had meetings of members to discuss industry conditions. The Court held that there was no

violation of the antitrust laws.[5] Speaking quite generally, Justice Stone echoed the views of Brandeis and Holmes:

> We decide only that trade associations or combinations of persons or corporations which openly and fairly gather and disseminate information as to the cost of their product, the volume of production, the actual price which the product has brought in past transactions, stocks of merchandise on hand, approximate cost of transportation from the principal point of shipment to the points of consumption as did these defendants and who, as they did, meet and discuss such information and statistics without however reaching or attempting to reach any agreement or any concerted action with respect to prices or production or restraining competition, do not thereby engage in unlawful restraint of commerce.

In both these cases, the Court focused narrowly on the information that the associations gathered. This is not sufficient. Information about other firms in an industry can be put to both competitive and anticompetitive uses. Firms in a competitive market need to know what their rivals are doing. They must set their prices not by their own costs but by the prevailing market price. If one firm has developed an imaginative or more effective marketing strategy, consumers benefit when other firms are free to learn about it and imitate it. Information about rivals' prices, however, is also an essential element of successful collusion. To discover whether a trade association is promoting or hindering competition, one should focus not on information in the abstract, but on whether the structure of the particular industry is one in which tacit collusion is likely. If it is, one should look with suspicion on devices that make information about price readily available.

An important modern case on facilitating devices was brought by the FTC against the producers of ethyl chloride, a lead-based antiknock gasoline additive. The firms in the industry appeared to be very profitable despite the homogeneity of the product and a declining market. The FTC claimed that a number of industry practices facilitated cooperative pricing. The practices included public preannouncement of price changes, uniform delivered pricing, and most-favored nation clauses (if prices are decreased, recent purchasers must be rebated the difference).

These practices may in fact make tacit collusion more likely. Preannouncement of price changes allows rivals to react prior to any effect of price change on market shares. This reduction in reaction time implies that the appropriate period length is very short, so discount factors are high and collusion is easier; there is little ability to steal market share before competitors can react. Uniform delivered pricing, a practice in which all buyers pay a single price independent of transportation costs, can make it easier for firms to observe rivals' prices. It may therefore make it easier to respond to a price cut. Most-favored nation clauses reduce the incentive to cut prices because the price-cutter bears the additional cost of the rebates to previous customers. The reduction in gains from price cutting may make cooperative pricing easier to sustain.[6]

The court ruled that these practices were not violations of the antitrust laws because the existence of these practices was insufficient to show that an agreement had been reached among the firms. The implication is that, absent explicit agreements among competitors, facilitating practices are not illegal and the mere existence of these practices is insufficient evidence to prove agreement. As we have seen, however, an anticompetitive equilibrium can emerge even in the absence of explicit agreement. The justification for a legal rule in such cases then rests ultimately not on the absence of any ability on the part of firms to engage in tacit collusion, but rather on our inability to do much about it.

Price is merely one attribute among many. We might limit the ability of firms to change price, but we must also place other limits on these firms for the legal rule to be effective. In the airline industry, for example, one would have to confront the availability of different kinds of fares on each flight as well as the number of flights. One could limit the ability of airlines to respond to fare changes quickly by curtailing the use of computer reservation systems, but these systems have many benefits for consumers.

Parties who engage in tacit collusion are behaving quite differently from firms that enter into explicit cartels. The managers of firms that engage in tacit collusion may not even know what they are doing. They may, for example, use base-point pricing (basing shipping costs from a single source, even if it is not the source from which the goods will actually be shipped) because it is common practice; they may not recognize that the practice helps to support an anticompetitive equilibrium.

Reputation, Predation, and Cooperation

In a static environment, small firms can benefit when one of their competitors is a large monopolist with market power. The monopolist deliberately reduces output in order to keep price above the competitive level. The small firms may benefit because they too can charge the supracompetitive price. Things, however, are not always static. A frequent claim of small firms in an industry is that the large firm, instead of keeping prices high, will deliberately cut prices below cost for a time until the smaller firms are driven out of business, at which time it will return to charging the monopoly price. This practice is known as *predatory pricing*. For predatory pricing to be worthwhile, of course, the lower profits the predator earns during the predation stage must be more than made up by the increased profits after the prey has exited.

Many have argued that courts should be extremely reluctant to find that large firms have engaged in predatory conduct. What small firms see as predation may simply be the response of large firms to competition. Predatory pricing can work only if the implicit threat of the large firm (to continue low prices until the smaller firms leave) is credible. Because the larger firm sells more goods at the lower price than the small firm, predatory pricing is necessarily more costly to it. The threat to predate therefore may not be credible. A number of scholars have found only weak evidence for predation in those cases in which it has been alleged and have concluded that it may not make sense for the law to try to deal with it.[7] In their view, the benefits to consumers from confronting the problem of predatory pricing do not come close to matching the costs of litigating predation cases. Moreover, firms may charge high prices and leave consumers worse off merely because they risk being held liable for predation if they charge the competitive price.

To understand the legal debate, one must first understand how it might be rational for one firm to engage in predation. One of the best-known predation cases involved the practices of the American Tobacco Company.[8] The Supreme Court found that American Tobacco would enter a market of a much smaller firm, lower its prices dramatically—sometimes incurring substantial losses—persuade the smaller firm to merge with it, and then shut that firm down. By showing its willingness to incur such losses, it was argued, American Tobacco could persuade existing firms to sell out at artificially low prices and deter others from entering the tobacco industry in the first place. We want to ask whether one can explain how American Tobacco's threat to lower prices was credible, given that it was facing different small firms over time.

We can examine the problem of reputation-based predation by taking an entry and pricing game and repeating it over time. In each of the one-shot games, there are two players—Incumbent and Entrant. Initially, Entrant must decide whether to enter or stay out. If there is entry, Incumbent must decide whether to accommodate entry or predate. The payoffs are shown in the extensive form representation of Figure 5.4.

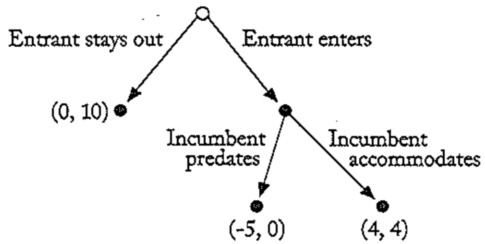

Figure 5.4 Predation and rational incumbent. *Payoffs:* Entrant, Incumbent.

We can solve the one-shot game through backwards induction. Once Entrant enters, Incumbent faces a payoff of \$4 from accommodating and \$0 from predating. Hence, Entrant believes at the start that Incumbent accommodates when the former enters. Because Entrant earns nothing by staying out and accommodating brings a payoff of \$4, Entrant enters. When this game is repeated a finite number of times, there is a unique subgame perfect Nash equilibrium—Entrant enters and Incumbent accommodates. Incumbent would like to threaten predation to deter entry but the threat is not credible. In the last market, Incumbent accommodates in the event of entry because there are no other markets from which entrants might be deterred. Knowing that Incumbent will accommodate, Entrant decides to enter. In the second-to-last market, Incumbent accommodates in the event of entry because it knows that Entrant will enter in the next period regardless of what it does. Entrant, knowing this, will enter in the second-to-last market. This process of unraveling continues until the first market is reached. One can also imagine that Incumbent faces sequential potential entry in ten separate but identical markets where it is a monopolist. This repeated game, however, has the same outcome. The unique subgame perfect Nash equilibrium is entry and accommodation in each market.

If the game is of infinite or uncertain length and Incumbent faces the same entrant in every period, however, there are an infinite number of equilibria. If Incumbent faces a different entrant in every period, however, the only subgame perfect equilibrium is one in which there is accommodation in every period. Incumbent can punish any given entrant only after all that entrant's costs are sunk. Any given entrant is unmoved by Incumbent's threat to predate in future periods, and Incumbent always accommodates whenever it focuses exclusively on the one-shot game.

At this point, however, we return to the problem that we encountered in the installment sale game at the beginning of the chapter. The solution to the games of finite and indefinite length turns on the payoffs' being common knowledge. The outcome may be altogether different if Entrant is only slightly unsure about Incumbent's payoffs from predating or accommodating. At this point, we want to show how we can build a model that captures this element of uncertainty.

Let us assume that there is a small chance that the payoff to Incumbent from accommodating is not \$4, but −\$1. Incumbent, in other words, is one of two types. Most incumbents are "rational"[9] and they evaluate only the monetary returns from accommodating. A few incumbents, however, are "aggressive." They suffer a loss of face if they fail to carry out a threat and the profits they would earn from accommodating are not enough to make up for it. We set out the one-shot game in Figure 5.5.

In this one-shot game, Incumbent predates when Entrant enters. Incumbent prefers a payoff of \$0 to one of −\$1. For this reason, Entrant does not enter. Entrant prefers the \$0 payoff from staying out to the −\$5 payoff from entering. Incumbent's threat in this game is credible because Incumbent is moved by self-interest to carry it out when the occasion arises.

At this point, we want to consider the repeated game that arises when Entrant is not certain whether Incumbent is the rational type who receives the payoffs shown in Figure 5.4 or the aggressive type who receives the payoffs shown in Figure 5.5. Entrants in later markets observe the previous actions of Incumbent. They can infer that if Incumbent ever accommodated, Incumbent must be rational, but they cannot be certain when Incumbent predates whether Incumbent is aggressive or Incumbent is rational but is mimicking the actions of the aggressive type in order to deter entry. We illustrate the extensive form of the one-shot game in Figure 5.6.

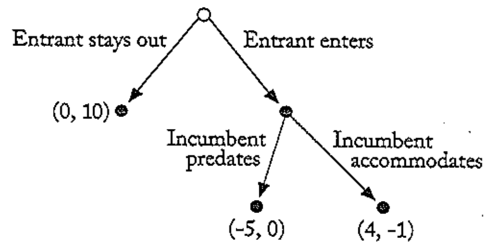

Figure 5.5 Predation and aggressive incumbent.
Payoffs: Entrant, Incumbent.

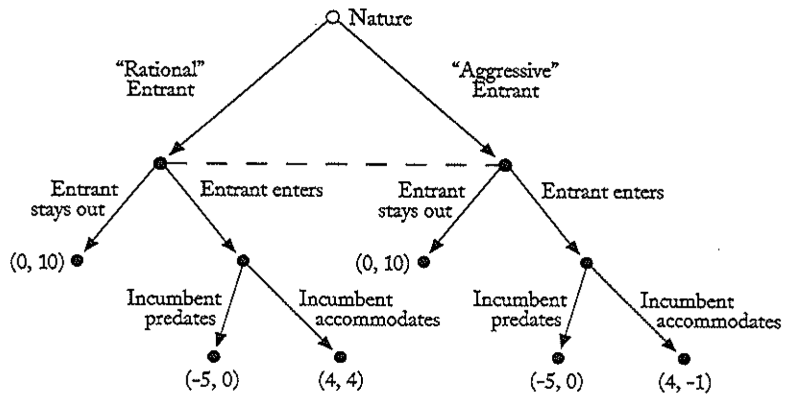

Figure 5.6 Predation with rational and aggressive incumbents. *Payoffs:* Entrant, Incumbent.

The equilibrium in the last period is easy to solve. The rational incumbent accommodates when Entrant enters and the aggressive incumbent predates. Entrant enters when it believes that the probability that Incumbent is rational is greater than 5/9 and stays out otherwise. The rational incumbent has no reason to predate. The rational incumbent, however, may choose to predate in early rounds in order to deter future entry by making potential entrants think that it may be an aggressive incumbent.

Entrant updates beliefs about Incumbent's type by looking at how Incumbent responded to entry in previous rounds. If Incumbent responded to entry in an early round by accommodating, future entrants will know with certainty that Incumbent is rational. From then until the end, backwards induction suggests that Incumbent will receive only $4 each period. If Incumbent responded to entry in early rounds by predating, however, future entrants will remain uncertain about whether Incumbent is rational.

If there are many periods left, it may be part of an equilibrium for the rational incumbent to predate. The equilibrium works in the following way: If Entrant expects both types to predate, Entrant will choose to stay out. We also need to know, however, about what happens off the equilibrium path. We must ask what happens if Entrant deviates and enters anyway. If Incumbent accommodates, future entrants know that Incumbent is rational, so they will enter. If Incumbent predates,

future entrants are still in the dark. If it was worthwhile for both types of incumbents to predate in the preceding period, it may still pay for them to predate during the next period. If this is the case, Entrant will stay out. Therefore, by predating in the preceding period, the rational incumbent may deter future entry.

When there is some uncertainty in the payoffs, predation is a sensible strategy for both types of incumbents. The rational incumbent pools with the aggressive incumbent. The inability of entrants to distinguish the different incumbents deters entry. A key element of this argument is that a rational incumbent need not bear the costs of predation in the early periods. In equilibrium, there is no entry. If, however, there should be entry in the early periods, the rational incumbent will predate to maintain a reputation for aggressive behavior that deters future entry. One can extend this same idea to the infinitely repeated game.

This game is one way to model reputation. Incumbents who fail to predate reveal their type. They are willing to predate in order to convince future entrants that they are aggressive rather than rational. Indeed, predating is the way in which a reputation as an aggressive player is established. There are two important characteristics of equilibrium in the reputation game. First, the reputation equilibrium is unique, even if the game is infinitely repeated. We do not have the problem we encountered in other infinitely repeated games of multiple equilibria, only some of which lead to cooperative pricing. In this model, a reputation for aggressive responses to entry is the only equilibrium. Second, as we add periods, the fraction of periods in which entry is deterred approaches one, so the final periods where reputation does not work become economically unimportant.

The reputation-based model of predation suggests the possibility of successful, rational, predatory pricing. In addition to the reputation model, there are also signaling models of predation, in which an incumbent charges a low price to signal its low cost to an entrant or a potential entrant. The idea is to convince the entrant that competing with the incumbent will not be profitable, so the entrant should exit or not bother to enter. Another type of predation model is based upon the limited financial resources of the entrant rather than the entrant's cost disadvantage. If financial markets are imperfect, making it difficult for a firm to get outside financing, a protracted price war may force the financially weak firm out of business. This is sometimes known as the deep-pockets model of predatory pricing. The existence of three consistent models of rational predatory pricing—reputation, signaling, and deep pockets—does not mean that predatory pricing is rife. The contribution to antitrust policy from these models is to indicate what the necessary assumptions are to make predatory pricing work.

We have also developed the predation model in order to make a more general point. Models such as this one relax the assumption of common knowledge by introducing a little uncertainty into the payoffs. These show how patterns of behavior based on reputation might emerge in contexts in which folk theorems suggest that many equilibria exist. If only a few players have different payoff structures, the course of the repeated game for everyone may become radically different.

We have created a repeated game in which the problem was one of predation, but we could as easily have made the same point with a variation on the installment sale game in Figure 5.2. In that case, injecting uncertainty in the payoffs suggests how socially desirable cooperation might emerge. Even if only a few players value their reputations, it may be in the interest of everyone else to mimic them. Thus, in many situations, it may be possible to have cooperation for a long time as the unique outcome to the game even if it would be in the interest of most players to defect in the one-shot game. The problem is not explaining why most people would cooperate in the one-shot game, but rather why a few do.

Because the equilibrium may depend on the behavior of a few players who are unusual, rather than the players who are typical, it is worth spending a moment focusing on these players. The individuals who support these equilibria cannot simply be those who place a special value on the long-term. What matters is their payoff in the one-shot game. Their payoffs in the individual games must be such that, when they are called upon to carry out a threat or to cooperate, they will find it in their self-interest to do so. As we saw in Chapter 2, players may be able to make sunk investments that change the payoffs of the one-shot game. More generally, individuals may be able to change their payoffs in these repeated games by making investments now that will alter their own preferences in the future.

The preferences that an individual has will affect the payoffs that that individual receives in the one-shot game, and individuals can deliberately alter their preferences over time. Individual merchants, for example, might choose at the outset of their careers to enter a social circle whose members are those with whom they are likely to do business. Such a social circle might ostracize anyone thought to engage in sharp practices. A person might develop community ties and adopt a social life (such as by joining a church and other community groups) that would make it costly for that person to break a promise made to another merchant, even if that person were not subject to any legal damages.

To alter one's payoffs in the future, however, one must do more than feign attachment to a particular church or social group. One must actually form these attachments and suffer real costs from noncooperation. Some of these might be visible and some might not be. To the extent that these ties are verifiable, a particular individual can capture the entire benefit; to the extent that they are not, however, long-term relationships may be possible for all merchants. As long as some have genuinely committed to a particular social circle or a particular church group, others may find it in their interest to mimic them.

We can also imagine other ways in which players might be able to change the payoffs they receive in one-shot games. The owners of a firm, for example, might take advantage of the separation of ownership and control in many firms. They could hire managers who do not incur the costs of carrying out a threat. Various golden parachutes and other kinds of arrangements with managers may undercut their incentives to work hard, but they may also make managers more willing to carry out a threat. The profits of the firm would fall if the threat ever had to be carried out, but the compensation package may make the managers less concerned about profits and more concerned about preserving market share and gross revenues.

Once it is in the interest of a few managers to predate even when it is not in the interests of the owners, it becomes in the interest of other managers to mimic them (at least in early periods), even if their sole concern is profits. Here again, what matters is whether some firms are structured in such a way that the managers do not act in the interest of the firm when the time comes to predate and that these managers cannot be distinguished from others.

Legal rules can affect these cases of repeated interaction with small uncertainty in the payoffs in two different ways. First, legal rules can affect whether some people can make investments that change the way they play the one-shot game. For example, a legal rule that made it harder to exclude an individual from a social circle might have the unintended effect of reducing the informal sanctions for sharp practices and the uncertainty in the payoff from long-term cooperation. Without at least the possibility of such sanctions, long-term cooperation may become less likely. Second, legal rules may make it harder or easier for other players to mimic those whose self-interest leads to cooperation. Disclosure laws about the relationship between managers and firms may make it easier for players to identify the payoffs their competitors face and thus make it harder for ordinary players to be confused with those who enjoy a special set of payoffs.

Summary

Repeated games are inherently hard to analyze. In the first instance, we must be more cautious in making the assumption of common knowledge, and we must further ensure that the model reflects such things as the uncertainty in the length of the game and in the payoffs that the players receive.

When games are of uncertain length, folk theorems tell us that the mere possibility of repeated interaction does not ensure that cooperative behavior will emerge. Nevertheless, it takes only some uncertainty in the payoffs of the one-shot game for an equilibrium to emerge in the game that is in the interest of an individual player, and perhaps in the interest of society as a whole.

Legal rules designed to change the equilibrium of a repeated game may be ineffective. Players may have the ability to change the structure of the game. More subtly, legal rules may affect the equilibrium of a repeated game by changing the payoffs to a few players in the one-shot game. They may also make it harder for the mine-run of players to mimic those who have unusual payoffs.

Bibliographic Notes

Antitrust, neoclassical economics, and game theory. Both the traditional and the game-theoretic literatures on the economics of antitrust are extensive. One of the primary focuses of the field of industrial organization is to understand the effects and appropriate application of antitrust laws and other governmental regulation of competition. A large percentage of the applications of game theory has been in the field of industrial organization. Textbooks include Carlton and Perloff (1990), Scherer and Ross (1990), and Tirole (1988). Tirole's book emphasizes the game-theoretic aspects of modern industrial organization. Another useful reference is the collection of survey articles edited by Schmalensee and Willig (1989).

The limits of backwards induction. The idea that backwards induction leads to unraveling back to the first move is called the *chain-store paradox* because it was developed in the context of predatory pricing with an incumbent firm that owned a chain of stores in many towns and that faced a potential entrant in each town; see Selten (1978). The question faced was whether predation would be rational—backwards induction suggests that it would not be.

Infinitely repeated games, tacit collusion, and the folk theorem. An excellent discussion of the various folk theorem results is contained in Fudenberg and Tirole (1991a). The literature on the repeated prisoner's dilemma is quite large. A natural place to start is Axelrod's book (1984a). Stigler (1964) was the first to draw attention to the issue of detection and punishment as essential factors in the ability to support tacit collusion in a non-game-theoretic model. The idea was captured in a game-theoretic model by Green and Porter (1984). The model has been extended by Abreu, Pearce, and Stacchetti (1986). Also see Tirole (1988) for a simple and very accessible version of the model. Literature on the selection of equilibria in infinitely repeated games based on renegotiation includes Pearce (1988) and Farrell and Maskin (1989). Rotemberg and Saloner (1986) is an example of a paper that uses a selection criterion based on the most-collusive equilibrium to generate comparative static results on the ability of firms to collude.

Reputation and predation. The reputation model discussed in the text is based on Milgrom and Roberts (1982). The other pathbreaking reputation papers are Kreps and Wilson (1982) and Kreps, Milgrom, Roberts, and Wilson (1982). The facts of predation cases reported at the appellate level, however, do not provide evidence for this sort of predation. Nor is it easy to muster evidence that golden parachutes and the like are designed to give managers preferences that make the threat of predation credible; see Lott and Opler (1992). Ellickson (1991) discusses the interaction between legal and extralegal sanctions. For discussions of how individuals can alter their future preferences, see Becker (1993) and Sunstein (1993).

In addition to the textbook discussions, an excellent survey of models of predation is Ordover and Saloner (1989). The extent to which the theoretical possibility of predation translates into actual predatory behavior is a subject of controversy. See Easterbrook (1981) for a discussion of many of the issues.