

INV. .... 89284.  
COL. C. 2810

# Game Theory and the Law

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

Harvard University Press  
Cambridge, Massachusetts, and London, England



## Dynamic Interaction and the Extensive Form Game

### The Extensive Form Game and Backwards Induction

In the last chapter we focused on the normal form game, which contains three elements—the players, the strategy space for each player, and the payoffs of each player for each combination of strategies. In this chapter, we examine the *extensive form game*, which models explicitly the actions that the players take, the sequence in which they take them, and the information they have when they take these actions. Because of its emphasis on actions, the sequence in which those actions are taken, and the information available to the players at each move, the extensive form game is often the appropriate way to model interactions between parties that take place over time.

We begin by using the extensive form game to model a simple debt contract. We go on to examine the problem of market preemption and strategic commitment in contract law. We use these examples once warranty and mitigation in contract law. We use these examples once again to show how problems of strategic behavior can be captured in formal models and to introduce two new solution concepts—*backwards induction* and *subgame perfection*—that are central to the solution of games in which each player takes actions based on what the other player does. The incentives of the lender and the borrower in a straightforward debt contract illustrate the kind of problem we face.

Lender and Debtor consider the following bargain. Lender gives \$10 to Debtor and Debtor promises to pay Lender back in a year's time. The trade is mutually beneficial. Lender has the capital and Debtor has the ability to use it productively. If all goes as planned, Debtor will use the loan to generate additional wealth and will be able to repay Lender in full, with interest. In Chapter 7, we shall examine the class

of cases in which all does not go as planned and Debtor does not have enough to pay Lender back in full. At the moment, however, we assume that only Debtor's willingness to pay Lender back is in doubt.

After Lender makes the loan to Debtor, Debtor may not have sufficient incentive to pay the loan back. If Lender has no ability to call on the state for help in recovering the loan and no sanction to impose on Debtor, it may be in Debtor's self-interest to keep the money. Lender, recognizing that Debtor may not be inclined to repay the loan, might not be willing to make it in the first instance. Lender and Debtor need to create some mechanism that makes it in Debtor's best interest to pay the loan back when it is due.

Both Debtor and Lender are better off if Debtor's self-interest will lead to repayment of the loan. Unless Lender is confident that Debtor will repay the loan, Lender will not lend the money in the first place. Contract law is valuable because it makes it easier for parties to alter their incentives and make their promises *credible*. A promise is not valuable unless its beneficiary believes that it will be kept.

The extensive form game is a useful way to put a formal structure on this problem, and thus to gain a better understanding of it. An extensive form game contains the following elements:

1. The players in the game.
2. When each player can take an action.
3. What choices are available to a player when that player can act.
4. What each player knows about actions that have already been taken (by that player and others) when deciding to take an action.
5. The payoffs to each player that result from each possible combination of actions.

The most common way of illustrating an extensive form game is with an inverted tree diagram. Each possible point in the game is a *node*. By convention, the *initial node* is represented by a hollow circle, and subsequent nodes are shown as filled circles. The *branches* leading away from the node (typically shown as arrows) represent different actions available to that player. Each branch leads in turn to another node. If no branches lead away from that node, then no further actions can be taken should the game reach that point. At such a *terminal node*, there is a payoff for both players. Alternatively, the new node may be another *decision node*, at which a different player or the same player must again

choose an action. Branches leading away from this node once again identify the different actions that are available to a player.

For the kinds of problems that are of interest to us, extensive form games are typically the best vehicles for examining interactions that take place over time, and normal form games are the best for problems of simultaneous decisionmaking. We should note, however, that any problem of strategic behavior can be set out in either form. For example, problems of simultaneous decisionmaking can be captured as extensive form games. The extensive game shows players moving in sequence. We can, however, treat a case in which players move simultaneously in the same way we treat problems in which one player moves after another, but moves without knowledge of what the other player has done.

Figure 2.1 is an extensive form representation of the normal form game that we saw in Figure 1.1. This is the game between the motorist and the pedestrian in which the pedestrian has no right to recover damages from the motorist in the event of an accident. In the extensive form of this game, the players are again the motorist and the pedestrian. The motorist makes the first move, choosing between two actions—either exercising no care or exercising due care. The pedestrian moves second, again deciding whether to take due care or not. The pedestrian is at either of two nodes, the node that arises after the motorist takes no care, or the one after the motorist takes due care. The pedestrian, however, does not know the course of the game up until this point.

We show that the pedestrian cannot distinguish between these two nodes by connecting them with a dotted line. By seeing which nodes are connected in this way, we can determine what each player knows about the actions of the other player at the time that player must move.




Figure 2.1 Regime of no liability (extensive form). Payoffs: Motorist, Pedestrian.

When a player cannot distinguish between different nodes when that player must move, we describe all those nodes as being in the same *information set*. In the extensive form, a strategy for a player specifies the actions a player takes at every information set in the game.

The normal form of a game may have several different representations in the extensive form. For example, we could equally well have shown the pedestrian moving first and the motorist moving second. The solution to the game in either case, of course, is the same as in Chapter 1. This is a game in which the motorist has a dominant strategy (taking no care) and the pedestrian, recognizing this, will take no care as well. The strategy combination in which neither takes care is also the unique Nash equilibrium of the game.

An extensive form representation underscores different aspects of the game. In this case, the extensive form emphasizes that games of simultaneous decisionmaking are games of imperfect information. The players know everything about the game, except the moves that the other players make. The extensive form might have been the most natural one to use if we were looking at a tort problem that had parties acting in sequence. One such problem is the doctrine of last clear chance. This doctrine applies when the motorist can observe the care decision that the pedestrian has made. In such a case, the motorist has the duty to exercise care even if the pedestrian has not.

We can show this as an extensive form game by adding an additional move for the motorist after the pedestrian's move in Figure 2.1. However fast the motorist drove initially, the motorist could see whether the pedestrian was crossing the road carefully and then decide whether to slow down. The decision of whether to slow down would be branches added to what are shown as the terminal nodes in Figure 2.1. The nodes in which the motorist would decide whether to slow down would not be in the same information set (and hence would not be connected with a dotted line) because the motorist would know what action the pedestrian had taken when the time came to make this move. Although the extensive form of this interaction may be easier to analyze, this game could also be represented as a normal form game. There is no longer a one-to-one correspondence between strategies and actions. A single strategy now consists of multiple actions. One such strategy, for example, would be to "exercise no care initially, slow down if the pedestrian exercises no care, do not slow down if the pedestrian exercises care."

Let us return now to the game involving an extension of credit and the problems that arise when players interact with each other over

time. In this game, Lender and Debtor are the players. Lender and Debtor must each choose between one of two actions. (For Lender, the strategies are "Lend" and "Don't lend." For Debtor, the two strategies are "Pay back if Lender lends" and "Don't pay back if Lender lends.") The problems in this game arise because the parties move in sequence.

The extensive form represented in Figure 2.2a models the game between Lender and Debtor when Lender has no ability to sue Debtor or to use any other mechanism to control Debtor's postborrowing incentives. In this game, Debtor asks to borrow \$100, promising to pay Lender \$105 in a year's time. Lender moves first and decides whether to make the loan. If Lender does not make the loan, both parties receive a payoff of \$0. If Lender does make the loan, the money can be used in Debtor's business and Debtor can earn \$110 over the course of the year. At this point, Debtor must decide whether to pay Lender back. If Debtor repays Lender, both will enjoy a payoff of \$5. (Lender will enjoy \$5 in interest on the loan and Debtor will enjoy the profits that remain.) If Lender makes the loan and Debtor defaults, however, Lender loses \$100 and Debtor gains \$110.

It is in the joint interest of the parties that the loan be made, because when the loan is made, the parties enjoy a joint payoff of \$10 rather than \$0. We can see that, however, if each party acts out of self-interest, the loan will not be made and a mutually beneficial trade will not take place. This conclusion is immediately self-evident in a case such as this, but the process that we use here, backwards induction, can be used in more complicated cases as well.

We start by focusing on the reasoning process of Lender, the player who makes the first move. Lender will determine what Debtor will do when given the move and then reason backwards before deciding




Figure 2.2a Lending without legal enforcement of debt contracts. Payoffs: Lender, Debtor.




Figure 2.2b Lending without legal enforcement of debt contracts (inducted). Payoffs: Lender, Debtor.

whether to make the loan. Hence, we look, as Lender would, at the choice that Debtor faces in the event that the loan is made. Because, by assumption, Lender has already made the loan, Debtor must simply weigh the two possible payoffs, \$110 and \$5. The first dominates the second. Because \$110 is more than \$5, Debtor will choose not to repay the loan if given the choice.

Now that we know how Debtor will move if given the chance, we can truncate the decision tree and eliminate the strategy that we know Debtor will not adopt. The truncated version appears in Figure 2.2b. Once the game tree is truncated in this way, Lender can readily see the consequences of making the loan. Because \$0 is better than a loss of \$100, Lender chooses not to make the loan. Because this outcome is not in the interest of either Lender or Debtor before the fact, they will seek to transform this game into one that has a different solution. A contract can be understood as a mechanism designed to transform this game into one in which the trade goes forward.

A legal rule that allows Lender to call upon the state to enforce a loan in the event of default and make Debtor pay its attorney's fees may make contracting between the parties possible. The legal rule not only prevents Debtor from refusing to repay the loan, but also exposes Debtor to the costs of litigation. Debtor might not contest the lawsuit and incur no costs at all, but let us assume that this debt contract, like most, requires that Debtor reimburse Lender for the latter's litigation costs, which are \$10. The legal rule transforms the game between Lender and Debtor into the game in Figure 2.3.

We can again use backwards induction and look at the decision facing the last player to move. At the time that Debtor must decide whether to repay the loan, Debtor is better off repaying than defaulting. Debtor has no ability to keep the money, and any effort to keep the money will expose Debtor to \$10 in litigation costs, which more than offsets the \$5 profit Debtor would make if the loan were repaid. Because Debtor will be better off repaying the loan, Lender is better off




Figure 2.3 Lending with legal enforcement of debt contracts. Payoffs: Lender, Debtor.

making the loan in the first instance. By changing the payoff to Debtor in the event of default from a gain of \$110 to a loss of \$5, the game becomes one in which Debtor is led by self-interest to repay the loan, and, more important, Lender is led by self-interest to make the loan in the first place.

This example illustrates why it might be in the interests of both parties to allow one to call upon the state to enforce a promise that the other made. It does not, however, tell us that contract law should apply to all agreements; nor does it say anything about the content of contract law. Indeed, it suggests that contract law should be largely an empty vessel. Parties need to be able to ensure that a person who is the victim of a breach has the right to obtain a legally enforceable order that requires the other to pay damages. This ability gives the parties the power at the start of their interaction to alter payoffs arising from particular strategy combinations and to ensure that both parties adopt those strategies that work to their mutual benefit.

Every trade is different, and the constraints that work on parties outside the law are different as well. As long as the agreement between two parties imposes no costs on anyone else, contract law may work best if it simply gives the players the license to perturb the payoffs in a way that advances their mutual interest. Any other method of enforcing contracts may not change the payoffs in a way that guarantees that each party looks out for the interest of the other.

A law that gives Lender the ability to call upon the state to enforce its claim provides parties with a way of transforming a game with a suboptimal equilibrium into another game with an optimal equilibrium. It is important, of course, not to overstate the point, and there are many qualifications that we should make. Most obviously, to say that legally enforceable contracts facilitate mutually beneficial trade is

not to say that the existence of such trade depends on it. Trade can exist and indeed flourish in the absence of legally enforceable contracts. As we shall see in Chapter 5, mechanisms such as reputation can bring about long-term cooperation even if there is no enforceable contract. The prospect of losing future deals both with another party and with people who know that party may be sufficient to ensure that each party performs.

We need to be cautious about the virtues of legal enforceability for a second reason. Making a contract legally enforceable is rarely as easy as the game in Figure 2.3 might suggest. In that game, Lender's money was returned with interest, regardless of whether Debtor defaulted, because we assumed that Lender's collection costs were recoverable in full. When Lender cannot recover such collection costs, the situation becomes more complicated. When the time comes for repayment, Debtor may be able to offer to pay less than what is owed in return for freeing Lender from the costs of litigation. The possibility that Lender may subsequently find it most beneficial to renegotiate the original contract is one of the most important issues in contract law. We shall explore the problem of renegotiation in Chapter 4. At this point, however, we wish to continue examining the extensive form game and, in particular, the solution concept of backwards induction. We do this in the next section by turning to the problem of market preemption and strategic commitment in antitrust.

#### A Dynamic Model of Preemption and Strategic Commitment

The case of *Federal Trade Commission v. E. I. DuPont de Nemours and Co.* focused on whether DuPont had acted improperly in the market for titanium dioxide,<sup>1</sup> a whitener used mainly in paints and plastics. By the 1970s, DuPont had developed a proprietary manufacturing process that made it the undisputed cost leader in the industry. The company then embarked on a plan to increase its production capacity substantially. Indeed, it planned to increase production by an amount sufficient to enable the company to capture much of the growth in the domestic demand for titanium dioxide through the 1980s. The government brought a complaint against DuPont before the Federal Trade Commission (FTC) in which it asserted, among other things, that the adoption and implementation of these expansion plans were themselves "unfair methods of competition" within the meaning of §5 of the Federal Trade Commission Act.

The threshold question we face is determining exactly how ex-

producing capacity itself could be an "unfair method of competition." Instead of looking at DuPont and the worldwide market for titanium dioxide, we shall look at a problem involving a firm with market power in a small region. Assume that in a small town there is only one cement plant, which we shall call "Incumbent." Trucking in cement from a more distant town is not practicable. For this reason, Incumbent can charge higher prices than cement plants in towns that face competition. Incumbent is very successful and earns profits of \$25 per month. Incumbent has the option of expanding its plant by purchasing land adjacent to it.

Another firm, which we shall call "Entrant," owns land in the town. It must decide between building a competing cement plant or a completely unrelated manufacturing plant that will have no effect on the market for cement. We want to know whether Incumbent will expand its cement plant and whether Entrant will build a second cement plant or a manufacturing plant. The benefits each enjoys from either course of action, of course, depend on what the other does.

We shall attach payoffs to the four possible combinations of decisions by Entrant and Incumbent respectively: (manufacturing, expand); (manufacturing, maintain current size); (cement, expand); and (cement, maintain current size). Entrant earns \$10 if it opens the manufacturing plant. This amount remains constant no matter what Incumbent does. The amount that Entrant earns from opening a second cement plant in town, however, does turn on whether Incumbent expands. Entrant earns \$15 if Incumbent maintains its old size, but it earns only \$5 if Incumbent expands. The profits are much smaller if Incumbent expands because Entrant will have to cut its prices to compete.

Incumbent's payoffs depend on Entrant's decision. If Entrant does not build a competing plant, Incumbent earns only \$25 if it maintains its current size, but \$30 if it expands. If Entrant opens a cement business, Incumbent is better off at its current size, where it earns \$10, rather than expanding, in which case it earns only \$5. In the latter case, not only does Incumbent have a smaller share of the market, but it also incurs the costs of maintaining a larger plant.

So far we have said nothing about the timing of the decisions. Assume that Incumbent cannot make its choice until after Entrant, perhaps because it will take several months to acquire the necessary permits. We can now represent the game by the extensive form in Figure 2.4.

It is easy to use backwards induction to solve this game. When In-




Figure 2.4 Entrant moves first. Payoffs: Entrant, Incumbent.

cumbent is at the left node, it prefers the strategy of expanding. The left-hand node is the one at which Entrant has already decided to go into the manufacturing business. By contrast, Incumbent maintains its current size at the right node, the one at which Entrant opens a second cement plant. Anticipating these decisions, Entrant has a choice between the terminal nodes at the extreme left and right in Figure 2.4. It prefers the one on the extreme right, which gives it a payoff of \$15 rather than \$10. Thus, the equilibrium is for Entrant to open a cement business and for Incumbent to maintain its current size.

The outcome (cement, maintain) does not maximize the two firms' joint profits nor is it the one that Incumbent prefers. Incumbent wants Entrant to believe that it will expand even after Entrant chooses to compete in the cement business. If Entrant believed that Incumbent would in fact expand even after it opened a cement plant, Entrant would be better off building the manufacturing plant (and thereby enjoying a payoff of \$10 rather than \$5). In this case, Incumbent could expand and enjoy profits of \$30 rather than the profits of \$25 it would receive if it maintained its current size.

Entrant, however, knows that Incumbent will not expand once it has opened its cement plant. Moreover, Incumbent cannot threaten to expand regardless of what Entrant does. A threat will be believed only if it is credible. Incumbent's threat to expand after Entrant opens a second cement plant will not be believed because, at that point, Incumbent makes itself worse off by expanding.

In the game set out in Figure 2.4, Incumbent has no way of convincing Entrant that it will expand even if Entrant opens a second cement plant. In many cases, however, Incumbent does in fact have strategies that allow it to commit itself to expanding before Entrant decides whether to enter the manufacturing or the cement business. To put the point another way, the sequence of moves determines the outcome of




Figure 2.5 Incumbent moves first. Payoffs: Incumbent, Entrant.

this game. Incumbent would be better off moving first and therefore will be willing to spend resources to do so.

We shall soon explore how Incumbent might accelerate its decision, but for now let us just assume that Incumbent moves before Entrant. The extensive form of this game is depicted in Figure 2.5. By convention, the payoffs are written in the order in which the players move. Hence, in this game, the payoffs to Incumbent are set out first.

We can use backwards induction to solve this game. If Incumbent chooses to expand its plant, Entrant will enter the manufacturing business ( $10 > 5$ ), whereas if Incumbent chooses to keep its cement plant at its current size, Entrant will open a competing cement plant ( $15 > 10$ ). Incumbent will therefore choose to expand because it predicts that Entrant will open a manufacturing plant in response. This course gives Incumbent \$30. By contrast, if Incumbent decided to maintain its plant at its current size, it would induce Entrant to open a new cement plant, and it would then enjoy profits of only \$10.

As the contrast between these two games suggests, a game's likely course of play is often determined by which party moves first. (In these models, the player who moves first enjoys the advantage, although in some strategic settings moving last is advantageous.) Incumbent therefore wants to commit itself to expanding before Entrant decides whether to enter the cement business.

There are many ways of making such a commitment. Although the adjacent land may not become available until much later, Incumbent can contract to buy it now. It can sign contracts for equipment for the new plant, advertise the expansion to customers, and begin to hire and train more employees. These strategies create commitment in two ways. First, some of the policies move the timing forward so that Incumbent's decision precedes Entrant's. Expansion prior to when the capacity is needed is known as *market preemption*. Second, some of these

policies change the payoffs by reducing the relative costs of expansion. Two ingredients are necessary for this second tactic to be successful: Incumbent must bear some of the costs of expansion before Entrant's decision, and these costs must in part be sunk, so that Incumbent will not recover them even if Entrant were to go into the cement business. (As long as Incumbent can get its money back, its threat to expand is not credible and Entrant will find it in its interest to open a second cement plant.)

Let us say that Incumbent can invest money that would otherwise bring a return of \$10 per month on specialized equipment needed only if the plant expands. Incumbent's profits are now \$10 lower if it does not expand. Incumbent has already spent the money that it would not ordinarily spend until after it expanded the plant. Once this investment is sunk, it has no value unless Incumbent expands. The new payoffs are set out in Figure 2.6.

Expanding the plant is a dominant strategy for Incumbent. Entrant anticipates expansion by Incumbent and therefore chooses the manufacturing business, making the outcome (manufacturing, expand). Incumbent spends the same amount on expansion in this game as in the ones set out in Figures 2.4 and 2.5, except that it spends some of the money earlier. Incumbent is better off making a preemptive investment in expansion, even if it costs more, as long as the investment is large enough to make expansion a dominant strategy and the payoff from expanding exceeds \$10. Incumbent is therefore willing to spend up to \$20 more in order to keep Entrant from going into the cement business. This \$20 represents the difference between the profits in the equilibrium in which Incumbent could move first and the equilibrium in which Entrant could move first.

Some types of commitment do not waste resources. If Incumbent




Figure 2.6 Incumbent sinks investment in expansion. Payoffs: Entrant, Incumbent.

merely signed a contract to lease the space, for example, its costs would be sunk, but someone else can use the space before Incumbent's lease begins to run. On the other hand, if Incumbent bought equipment for the larger plant and kept it idle in storage for a time, resources would be wasted. In principle, everyone would be better off if somehow Incumbent did not make such expenditures. To be sure, the money that is being spent belongs to Incumbent, but the private gain to Incumbent exceeds the social gain from this investment. Incumbent is willing to make investments that bring a deadweight social loss as long as the loss is less than the private benefit Incumbent enjoys from keeping Entrant out of the market.

In the model, Incumbent and Entrant are both better off if one expands and the other enters a different business. The model, however, does not show us how this course of action affects consumers. It is possible, for example, that people in the town are better served by one large cement plant than by two smaller ones. The profits might be larger with the expanded plant because of economies of scale, not because of monopoly power. The prices may not be higher if there is only one cement plant rather than two, given the amount of cement that people in the town buy. Even if Incumbent expends resources ensuring this outcome, consumers as a group may not be worse off.

This model shows why DuPont had reason to commit itself to building capacity beyond its current needs. Once it had built the facilities needed to meet projected increases in demand for titanium dioxide, its costs were sunk. If DuPont had not expanded early, other producers might have entered the market. After they entered, it might not have then made sense for DuPont to expand. The company's decision to expand did not affect the strategies available to the other players. They could build new plants if they wished, just as Entrant could still open a second cement plant. By making its investments early, however, DuPont changed the equilibrium strategies of the other players.

We can posit any number of steps that an incumbent might take—from signing a contract to buying equipment—for strategic advantage. The extensive form game we have set out here shows how many kinds of decisions can have strategic consequences. Even an act as simple as buying new equipment may represent a waste of resources, decrease consumer welfare, or both. Because merely changing the order of moves in an extensive form game can lead to a different course of play, any action that affects the sequence in which parties act can, in principle, have anticompetitive implications.

None of this, however, suggests that the scope of antitrust liability

should be expanded or that the FTC was wrong when it refused to rule that DuPont's decision to increase its capacity was an "unfair method of competition." DuPont was, after all, the low-cost producer. The evidence suggested that it was still enjoying economies of scale. DuPont was successful in part because it could produce titanium dioxide more cheaply than anyone else. By committing itself to expand and incurring higher cost producers not to enter the market, DuPont's actions may have prevented others from spending resources inefficiently.

Even in competitive markets efficient firms may be the ones most likely to expand capacity. DuPont's decision may have had some foreseeable anticompetitive effects; indeed, internal memos predicting that the price of titanium dioxide would rise after these steps were taken are consistent with this conclusion. This possibility, however, was not sufficient to persuade the FTC to sanction DuPont and risk shielding less efficient producers from competition. This model suggests how actions, such as overexpansion, can be anticompetitive. Through the use of such models (and more elaborate variations on them), one can identify what actions can produce anticompetitive outcomes and how legal rules can change things for the better.

#### Subgame Perfection

We have been able to solve the preceding extensive form games by using backwards induction. From one perspective, our approach slight the Nash equilibrium solution concept. We could have used it to solve all the extensive form games that we have examined, because any solution found through backwards induction is also a Nash equilibrium. Moreover, many games that cannot be solved with backwards induction can be solved by using the Nash equilibrium concept. Using this concept, however, frequently proves difficult. In many extensive form games, there are multiple Nash equilibria. We have emphasized backwards induction in our discussion thus far because it is often the best way to isolate the one strategy combination that the parties are most likely to adopt.

Backwards induction, however, is not available in those cases in which the last player must move without knowing the other player's previous move. If we turn to the Nash equilibrium solution concept in such a case, we may need some means of identifying those Nash equilibria that are plausible and those that are not. We can get some idea of the problem by looking at a variation of the game involving Lender and Debtor that we examined earlier.



Assume that Lender must engage in extremely costly litigation in the event that Debtor defaults on the loan. Lender will prevail in the end, but the litigation will impose costs on Lender and Debtor of \$125. Hence, instead of both enjoying a \$5 profit from the transaction, each suffers a \$120 loss in the event that Lender makes the loan, Debtor defaults, and Lender sues. (The assumption that litigation costs can exceed the stakes is not implausible, even with loans of substantial size. Litigation is like a contest because the chances of success turn in some measure on which party spends the most. Once one party spends money on litigation, the other party has an incentive to respond by spending a little bit more.) We illustrate this game in Figure 2.7.

We can solve the game in Figure 2.7 using backwards induction. Lender decides whether to lend money by asking what course can be taken after Debtor defaults. At this point, Lender does nothing. Lender prefers a payoff of  $-\$100$  to a payoff of  $-\$120$ . Lender therefore infers that Debtor defaults when the loan has been made because Debtor knows that Lender will not sue after a default. Lender, believing that Debtor will default when a loan is made, refuses to make the loan in the first place.

We can also use the Nash equilibrium concept to solve this game. We do this by again looking at possible strategy combinations and asking whether each player is playing a best response given the strategy of the other. Recall that a strategy for a player spells out what that player does at each information set, even when it is an information set that will not be reached because of an action which that player has taken previously. In other words, even if Lender does not make the loan in




Figure 2.7 Debt contract with costly enforcement. Payoffs: Lender, Debtor.

the first place, Lender's strategy still reveals what Lender would do in the event that Debtor defaulted.

The logic of defining strategies in this way will become apparent as we proceed. It rests on the idea that we cannot analyze whether Lender's decision to lend the money makes sense unless we know what Lender would do when Debtor defaults. We implicitly used this notion when we solved the game in Figure 2.7 using backwards induction. (Lender adopts the strategy of refusing to make the loan and does not sue when Debtor defaults, whereas Debtor adopts the strategy in which Debtor defaults after the loan is made.)

This combination of strategies is also a Nash equilibrium. Lender's best response to Debtor's strategy of defaulting is not to lend the money in the first place. Lender is better off not making the loan and receiving a payoff of \$0 than making the loan. When Lender makes the loan, Debtor defaults. When Debtor defaults, Lender does not sue. This leaves Lender with a loss of \$120. Given that Lender's strategy in this proposed equilibrium is not to lend the money, Debtor's decision to default is a best response. Debtor cannot do any better by changing this strategy.

The Nash equilibrium concept, however, does not identify a unique solution to this game. It is also a Nash equilibrium for Lender to make the loan and to sue in the event of default, and for Debtor to pay the loan back. Lender is better off making the loan than not, given that Debtor repays. A payoff of \$5 is better than a payoff of \$0. Similarly, Debtor is better off repaying than defaulting, given that Lender sues. Debtor prefers a payoff of \$5 to a payoff of  $-\$120$ . A moment's reflection, however, will reveal that parties are unlikely to adopt these strategies. The solution rests upon Lender's threat to sue in the event that Debtor defaults and this threat is not credible. After Debtor defaults, it is no longer in Lender's self-interest to sue.

The Nash equilibrium solution concept can be refined to eliminate this second equilibrium. This equilibrium does not take into account Lender's incentives when Debtor actually defaults. The equilibrium rests on the idea that Lender sues when the game takes that course, but it does not seem plausible that Lender sues when Debtor defaults. As we have set the payoffs (which ignore, for example, reputational effects that might follow from a failure to sue), it is not in Lender's self-interest to sue.

A threat is ineffective unless the person who makes it will actually carry it out if called upon to do so. In a game such as the one in Figure 2.7, Lender will carry out a threat to sue Debtor to recover the loan

only if, at the time Lender decides whether to sue, suing brings greater benefits to Lender than doing nothing. Any solution is suspect if it rests, directly or indirectly, on the assumption that Lender would bring a suit even though its costs would exceed the amount of the recovery.

We need a way to incorporate this idea into the Nash equilibrium concept. To put our ambition in game-theoretic language, we need a refinement of the Nash equilibrium solution concept so that we can exclude strategy combinations that, even though they are Nash, make implausible assumptions about the actions that the players would take, but do not in fact take under the proposed equilibrium. These are actions that are *off the equilibrium path*. This refinement of the Nash solution concept is known as *subgame perfection*.

A Nash equilibrium is subgame perfect if the players' strategies constitute a Nash equilibrium in every *subgame* of a game. A subgame is a move or a set of moves of an extensive form game that can be viewed in isolation. More formally, a subgame of a game in the extensive form is any part of a game that meets the following three conditions:

1. It begins at a decision node that is in an information set by itself.
2. It includes all the decision nodes and terminal nodes that follow it in the game and no others.
3. No nodes belong to an information set that includes nodes that do not also follow the decision node that begins the subgame.

The game in Figure 2.7 has two subgames. The first begins with Debtor's choice, the second with Lender's decision to sue or not to sue. The first subgame has two Nash equilibria. In the first, Debtor chooses to repay the loan and Lender sues when Debtor defaults; in the second, Debtor defaults and Lender does not sue. Focusing on this subgame does not allow us to eliminate any of the Nash equilibria in the whole game. But subgame perfection requires that the candidate solution be Nash in every subgame, so consider the subgame that begins when Lender chooses whether to sue. This subgame has only one Nash equilibrium: Lender does not sue.

To be "subgame perfect," a solution to the entire game set out in Figure 2.7 must be one in which Lender never adopts a strategy in which Lender would sue were Debtor to default. The proposed solution in which Lender makes the loan and Debtor repays it depends on Lender's suing if Debtor were to default. Hence, it cannot be a subgame

perfect equilibrium. The only equilibrium that is subgame perfect is precisely the equilibrium that is reached through backwards induction.

As we have seen, the game in Figure 2.7 could be solved through a straightforward application of backwards induction. Backwards induction, however, is not available in games that have information sets with more than one node. Backwards induction depends on knowing the last move of the game. This may not be possible if a player does not know what moves have already been taken when the time comes to move. Subgame perfection is a useful refinement of the Nash equilibrium concept. Like backwards induction, it eliminates those solutions that rest upon threats that are not credible, as well as other implausible actions off the equilibrium path; unlike backwards induction, however, it is still available when information sets contain more than one node.

There are three complications that one might introduce into our model of the interaction between Lender and Debtor. The first two give formal content to the ideas of vengeance and retaliation. First, Lender may not make decisions entirely on the basis of monetary costs. To the extent that Lender suffers psychological harm or loss of esteem among peers from failing to recover a loan, we need to alter the payoff structure of the game accordingly. Although litigation has a dollar amount that exceeds the size of the loan, these other costs associated with failing to recover the loan may make litigation the course that brings Lender the greatest benefits. This model, like those in Chapter 1, posits hard out-of-pocket costs as the only component of the payoffs, but nothing in the structure of the model requires this.

The second complication requires us to change the structure of the model. Lender may lend money to many different borrowers. To model the behavior of Lender, we may not be able to focus on a one-time game. Instead, we might want to imagine this game as part of a larger game in which Lender engages in the same transaction repeatedly. To the extent that this is the case, we must ask whether Lender will pursue litigation because of the effect of a failure to sue on interactions with future creditors. We show how reputation can be modeled as a repeated game in Chapter 5.

The last complication arises from the way in which the strategy space in this model is limited. This model does not allow for the possibility of settlement. We explore in Chapter 8 how, once one introduces the possibility of settlement, Debtor might be willing to settle in the face of Lender's threat to sue, making the threat once again credible. The virtue of putting a formal structure on the behavior of Lender is that

it makes us confront the question of whether a threat is credible and, if it is, why.

In the rest of this section, we show how subgame perfection illuminates a problem in contract law that often arises when goods are sold and then fail to work as expected. The buyer asserts that the seller failed to deliver the goods that the seller promised. The seller, by contrast, will argue that the goods were as promised, but that the buyer did not use them properly. Disputes arise because much depends on both the seller and the buyer exercising care in making or using the goods that may not be visible to the other unless there is expensive litigation.

We can illustrate the problem by examining several variations on *General Foods v. Valley Lea Dairy*, a commercial law dispute typical of many that arise every year. In 1978, General Foods bought 40,000 pounds of dry milk from a dairy cooperative. The dairy delivered the milk in 9 separate lots. Although it had no explicit contractual obligation to do so, General Foods tested each lot and found that 1 was contaminated with salmonella. This discovery led General Foods to retest the 8 other lots. When no further evidence of salmonella was found, it used them in its milk chocolate.

At this point, General Foods ran a third round of tests. Because the milk was once again in liquid form, the test was more accurate than the earlier ones. Before the results of this test were available, however, General Foods sold the chocolate to several candy makers. The third round of tests eventually revealed that the 8 lots were tainted. All the candy made from the chocolate had to be destroyed. General Foods reimbursed the buyers of the chocolate for the losses they suffered and in turn demanded that the dairy make it whole. General Foods, however, was unsuccessful in its efforts to recover civil damages from the dairy that had sold it the contaminated milk.<sup>2</sup> Much in this case turned on its peculiar facts. Our interest is in the general way in which this problem should be approached.

When food products such as milk are sold, there is always the risk that they are contaminated. The parties, however, want to confront the danger that salmonella presents in a sensible way. The problem is similar to the ones we examined in the last chapter. Each party should use due care in manufacturing and processing the milk. The contract they write should ensure that both parties have the incentive to test in a way that is cost effective, and, when they do not reach an explicit agreement, the gap-filling rules of contract law should try to give them the right incentive. If the laws do not do this job, we run a risk analogous

to the one we saw in the example of the debt contract at the outset of this chapter. Rather than entering into a contract and taking advantage of the skills of another, a company such as General Foods may bring the operation inside the firm, where it will do the task less efficiently than the dairy could if it had the right set of incentives. If a legal regime can give the dairy the right incentives, everyone could be made better off.

Consider a variation on the facts of *General Foods*. A food processor that uses dry milk in various products it sells must decide whether to process its own dry milk or to buy it from a dairy. If it processes the milk itself, the firm can cheaply monitor whether the milk is contaminated. The workers who inspect its other food processing operations at the same site can inspect the initial drying of the milk with little added cost. A dairy can dry the milk more cheaply in its more specialized plant located some distance away. If the dairy dries the milk, however, the food processor must trust the dairy to hire people to monitor the initial drying process and run the first set of tests.

The dairy may perform two kinds of tests—either high or low. The high test is more expensive but also more accurate. The food processor runs a second set of tests after the dry milk is delivered. It has a choice of three kinds of tests: high, medium, or low. As before, the more expensive tests are more accurate. We can illustrate this problem as the extensive form game in Figure 2.8a.

Processor moves first and decides whether to buy the milk on the market. Dairy then decides on the kind of testing it will do, and then




Figure 2.8a Sale without warranty. Payoffs: Processor, Dairy.

Processor decides on its tests. Even when there is no enforceable contract between it and Processor, Dairy incurs some added costs from using the cheaper test because of the loss to its reputation and potential lawsuits from third parties. Whether these costs are worth incurring turns on what kinds of tests Processor runs. As in the game in Figure 2.1, we draw a dashed line between the two possible decision nodes of Processor because they are both in the same information set. At the time when it must make its decision, Processor does not know whether Dairy has used the high or the low test and hence does not know whether it is at the decision node on the left or the right.

We cannot reason backwards move by move as we could in the games involving Lender and Debtor. If Processor knew that it was at the decision node that arises after Dairy decides to test high, it would test medium. If Processor finds itself at the right node, however, the outcome is different. If Dairy tests low, Processor is better off testing high. We cannot reason backwards from Processor's last move because it depends on what Dairy did in its previous move.

We can, however, solve this game by a different kind of backwards reasoning. We can isolate one part of the game—the subgame that begins after Processor decides to have Dairy process the milk—and solve it before solving the game as a whole. If we ignore Processor's initial move (the decision of whether to process its own milk or buy it from the dairy), we have a game that can stand independently. The single node at which Dairy must decide on the kind of test it wants to use could be an initial node of a free-standing game. We can solve the larger game by solving this subgame and, with that solution in hand, use backwards induction to solve the game as a whole.

The subgame that begins with Dairy's first move is one that we can easily capture in the normal form, because, after that point, each player must make its move without knowing what the other has done. We can illustrate the normal form of this game with a three-by-two bi-matrix, shown in Figure 2.8b.

There are two solution concepts that we can use to find the likely course of play in this subgame. First, we can note that this subgame has only one Nash equilibrium, the strategy combination in which Dairy tests low and Processor tests high. If Dairy is going to test low, Processor is better off testing high and enjoying a return of \$0 or -\$5. If Processor is going to test high, Dairy, of course, has no incentive to deviate either. It prefers receiving \$5 from testing low to \$4 from testing high. This subgame has no other Nash equilibria.

		Dairy's Level of Care	
		High	Low
Processor's Care	High	6, 4	2, 5
	Medium	10, 2	0, 4
Low	5, -2	-5, -5	

Figure 2.8b Sale without warranty (subgame after Processor decides to buy). Payoffs: Processor, Dairy.

We can also solve this game through the successive elimination of dominated strategies. Processor will never play low because the strategies of high and medium dominate it. Dairy will observe that, so long as Processor never moves low, it is better off moving low rather than high; hence, Dairy will move low. Processor, recognizing that Dairy will move low, will realize that it is better off moving high. Thus, the solution to the subgame is one in which Processor moves high and Dairy moves low.

Once we determine how the subgame is likely to be played if the players were to reach it, the solution to the game as a whole becomes clear. Processor knows that it will enjoy an expected payoff of only \$2 if it decides to buy the dry milk from Dairy, but that it will enjoy a payoff of \$5 if it makes the dry milk inside the firm. Because Processor will find it in its self-interest to play high whenever it decides to buy milk from Dairy, we must reject any proposed solution that rests upon Processor's adopting some different strategy if the game proceeds along that path. The game as a whole has several Nash equilibria, but of these only one seems plausible, and this strategy combination is the unique subgame perfect Nash equilibrium of the game. This equilibrium is the one in which Processor makes its own dry milk, Dairy tests low—when given the chance to move—and Processor tests high when it buys milk from Dairy.

Up to this point, we have been considering the problem between Processor and Dairy as it would exist if the milk were sold without a warranty. The outcome that has the greatest joint payoff for both players is that in which Processor buys the milk and tests medium, while Dairy tests high. Unless the parties take actions to perturb the payoffs, however, they will end up far from this outcome. One way to perturb

		Dairy's Level of Care	
		High	Low
Processor's Care	High	6, 4	5, 2
	Medium	10, 2	3, 1
Low	5, -2	-2, -8	

Figure 2.8c Sale with warranty (subgame, with contract).  
Payoffs: Processor, Dairy.

the payoffs is to provide a schedule of damages that Dairy owes Processor when the milk proves to be tainted. Contract damages do not change the joint payoffs, but they change the way they are distributed between the parties. Contract damages work in the same way as tort damages.

As we showed in the last chapter, there are many schedules of damages in a simultaneous move game that induce both parties to adopt the course that is most mutually beneficial. One such schedule gives us the subgame that begins after Processor decides to buy the milk, under which Dairy must pay damages of \$3 any time it tests low.<sup>3</sup> The payoff transformations brought about by this damages schedule are shown in Figure 2.8c.

The unique Nash equilibrium of this subgame is one in which Processor moves medium and Dairy moves high. When Dairy moves high, Processor's best response is to play medium and enjoy a payoff of \$10, rather than play either high or low and receive a payoff of \$6 and \$5 respectively. Dairy will test high if Processor tests medium, preferring a payoff of \$2 to a payoff of \$1. Hence, this strategy combination is Nash.

We can also check that no other combination is Nash. As noted, Processor's best response to Dairy's strategy of high is medium. Any other combination in which Dairy tests high cannot be a Nash equilibrium because Processor's strategy would not be a best response. If Dairy plays low, Processor's best response is high. The strategy combination of high on the part of Processor and low on the part of Dairy is not Nash because Dairy's best response to Processor's strategy of high is to play high as well.<sup>4</sup>

Given that in the subgame the strategy combination in which Dairy

tests high and Processor tests medium is the only Nash equilibrium, we know that the only subgame perfect equilibrium to the game as a whole is the one in which Processor and Dairy adopt these strategies. If Dairy tests high when given the chance, Processor's best response in the game as a whole is to buy milk from Dairy rather than process the milk itself. Given the actions that Dairy will take in the subgame, Processor enjoys a payoff now of \$10 if it buys milk from Dairy, rather than a payoff of \$5 if it makes the milk itself.

This game has other Nash equilibria, but they rest on implausible assumptions about play off the equilibrium path. Consider the following combination of strategies: Dairy adopts a strategy of low in the event that Processor buys Dairy's milk. Processor in turn manufactures its own milk and tests high when it buys milk from Dairy. This equilibrium is Nash. Given that Dairy is going to test low, Processor's best response is to make its own dry milk. When Processor makes its own milk, it earns \$5, but if it buys from Dairy and Dairy tests low, Processor can earn no more under this schedule of damages than \$5. Dairy has no incentive to switch from a strategy of testing low. Because Processor is going to make its own milk, Dairy receives the same payoff of \$0 regardless of how it tests when it sells the milk. Each player adopts a best response given the strategy of the other; hence, this combination of strategies is a Nash equilibrium.

This equilibrium, however, is not a plausible solution to the game. A combination of strategies in which Processor makes its own milk and Dairy tests low is Nash only because, in equilibrium, Dairy does not actually have a chance to move. Dairy would not in fact test low if Processor were to buy the milk from it. Under the new payoffs, testing low is no longer in Dairy's interest. Testing low, after all, is a dominated strategy and a player will not play a dominated strategy if given a chance to move. Moreover, the other player will act on that assumption.

Our solution to the game must take into account the incentives of the parties at every possible decision point. The technique of focusing on the play in each subgame to solve the game as a whole is useful because it eliminates any Nash equilibrium that rests upon players' taking implausible actions that are off the equilibrium path. Once we find that a particular combination of strategies is the unique subgame perfect Nash equilibrium, we have solved the game. For this reason, we can predict that, in this game, Processor should adopt a strategy of buying from Dairy and testing medium. Dairy should adopt a strategy of testing high in the event that it sells the milk. There are no other

Nash equilibria in which the players behave rationally both on and off the equilibrium path.

In this example, the parties wrote a contract that changed the payoffs in three cells of the payoff bimatrix. Most contracts do not have elaborate schedules of damages. Moreover, it is not possible for the background legal rule to provide such details for every possible commercial transaction. As we saw in the last chapter, however, it is possible to state in general terms obligations to pay damages that give everyone the right incentives. For this reason, the parties may be able to write a contract in which they adopt a general rule that transforms the payoff structure in a way that ensures that both parties act with the interests of the other in mind.

A crucial question—and one upon which we focus in the next two chapters—is how to formulate a general rule that can work even in a world in which information is hard to come by, both for the parties and for the courts called upon to enforce the rule. In the game involving Dairy and Processor, for example, we would want to explore whether we could induce optimal behavior with a schedule of damages that was tied to whether the milk was tainted, not to the kind of test that Dairy ran. It is much easier for a court to determine the former than the latter.

The problem that parties face when they enter into their initial contract is similar to the one facing a lawmaker who is devising tort rules and trying to minimize the social costs of accidents. If we make strong assumptions about the rationality of the players and the information available to both the players and the courts, many different rules are possible that transform the payoff structure from what it would be if the buyer had no cause of action against the seller and that also give each party the right set of incentives. A lawmaker, however, needs to design rules that handle those cases in which the parties do not explicitly take a particular contingency into account.

The Uniform Commercial Code governs a contract for the sale of goods and applies to cases such as *Valley Lea*. Two of its provisions (both of which the parties can waive if they choose) are relevant for us. The first provision is the implied warranty of merchantability. A merchant seller such as Dairy promises that its goods pass without objection in the trade under the contract description. If they do not, the seller must make the buyer whole. The second is the provision that requires the buyer to mitigate losses. A seller is liable only for those losses that “could not reasonably be prevented.”<sup>5</sup> Under this rule,

Dairy is liable only for the losses Processor suffers if Processor exercises reasonable care and takes those steps that are cost effective.

As our discussion from the last chapter suggests, such a rule gives the players the incentives to exercise care, provided that we make a number of assumptions about such things as what losses could not be “reasonably” prevented. We should not, however, end the inquiry there. To say that this rule is efficient under this set of assumptions is very far from saying that it is the only rule that is efficient, given these assumptions. Recall that any extensive form game can be represented as a normal form game; and, as we saw in the last chapter, an infinite number of civil damages rules exist that give the players the right incentives. In the last chapter, we worried about whether a rule existed that made the fewest rationality assumptions. This concern, however, should not be our only one. Perhaps of even greater importance is whether a particular rule is informationally parsimonious. We may prefer a simple rule—one that parties can understand and that we can enforce—that gives us results that come close to the optimum, even if it would not work as well under ideal conditions.

The common law embraces the idea that parties are liable for expectation damages, subject to a duty to mitigate on the part of the innocent party. This rule itself is not perfect,<sup>6</sup> but even this rule, which requires a court to determine how much a party would have received if the contract had been performed, may be too informationally demanding. When parties draft their own damage provisions, they often opt for ones that are considerably more simple.<sup>7</sup> They often provide for fixed damages in the event of breach, and many items that would fall within the scope of expectation damages, such as consequential damages, are routinely excluded. The game with Processor and Dairy in this chapter showed how a simple damage schedule might work. In the next chapter, we turn to the considerably more difficult problem of showing how we can explicitly take account of informational problems in fashioning legal rules.

### Summary

In this chapter, we have shown how we could capture as extensive form games situations in which individuals interact with each other over time. As in the last chapter, we have shown how we could model legal rules providing for civil damages as transformations of the payoffs. The payoff to one player under each strategy combination rises

by the same amount that another player's falls. Once again, legal rules often ensured that the likely solution to the game was one that was in the joint interests of the players by changing payoffs that were off the equilibrium path. The dynamic character of these games, however, introduced a new problem.

In the tort problems that we examined in the last chapter, what mattered was that players internalize the cost of the decisions they make. When the other player makes decisions at the same time, little else matters. When players interact over time, however, what matters is not so much that a player internalizes all the costs of a decision as that a player always has an incentive to make decisions that keep the game on the path leading to the socially desirable outcome. The structure that the extensive form game imposes on problems in contract law and other problems involving long-term relationships between parties provides an opportunity to look at such doctrines as the parole evidence rule, conditions, and the Statute of Frauds, all of which have the effect of attaching consequences to actions that take place off the equilibrium path of the game.

The structure of the extensive form game and the idea of subgame perfection give us a way to distinguish credible threats from ones that are not. As we saw in the game involving Incumbent and Entrant, regulatory regimes such as antitrust must be sensitive to the frequently subtle ways in which parties can act so as to change the incentives they and others will have at subsequent points in time—and thus convert noncredible threats into credible ones. In addition, this structure provides a way of looking at regulatory regimes in general, and it focuses on the dynamic consequences of imposing a new regulatory regime.

We can illustrate this problem of dynamic consistency by considering a weakness in a one-time tax amnesty law that is coupled with the promise that another such law will never be passed. The ambition of the law is to force tax evaders to come forward, but also to convince people that tax evaders in the future will not be treated so generously. Such a law may not work because the promise never to institute another amnesty (essential to deterring people from future tax evasion) may not be credible. A regulatory regime, like the structure of a contract between private parties, must be dynamically consistent over time.

A policy in which one offers amnesty on only a single occasion can make sense only if one can explain why offering amnesty again at some later time would not appear at least as attractive. One must look not only at the play of the game as a whole, but at the play of the subgame

that arises after the tax amnesty takes effect, people once again evade taxes, and the legislature decides whether to keep its promise not to provide another amnesty. The initial resolve of the legislature is credible only if it will be in the interest of the legislature to keep that resolve at a later point in time. If the legislature is not able to tie its hands, the promise never to have another amnesty will not be credible.

The same problem can arise in the context of litigation. In a Chapter 11 reorganization, for example, the debtor often requests an extension of the period in which the debtor has the exclusive right to propose a plan of reorganization. The court often grants the debtor's request for an extension, but asserts at the same time that it will not grant any additional extensions. Creditors are hurt by the passage of time because of the time value of money, but, except for the passage of time, nothing may change if no plan of reorganization is agreed upon before the period expires again. The court may find itself in exactly the same position when the debtor requests another extension of the exclusivity period. In such a case, the court's initial statement that it will not make another extension may not be believed. To model extensions of the exclusivity period, we must take the problem of dynamic consistency into account.

#### Bibliographic Notes

*The extensive form game.* The concepts of the extensive form game and subgame perfection are the basic elements of modern game theory. Kreps (1990b) provides a good formal introduction to the extensive form game.

*The extensive form game and the debt contract.* The idea that parties write contracts and devise mechanisms to ensure that each has the incentive to cooperate with the other is well known. Kronman (1985) is a good example. Kreps (1990b) as well as Fudenberg and Tirole (1991a) explores the potential weaknesses of backwards induction and subgame perfection as solution concepts.

*Market preemption and strategic commitment.* The literature on the value of commitment starts with the classic book by Schelling (1960). Applications to entry and capacity decisions are contained in Dixit (1979) and Spence (1977). The titanium dioxide industry is discussed in Chemaawat (1984).