

Game Theory and the Law

Douglas G. Baird

Robert H. Gertner

Randal C. Picker

Harvard University Press

Cambridge, Massachusetts, and London, England

Copyright © 1994 by the President and Fellows
of Harvard College
All rights reserved
Printed in the United States of America
Third printing, 1998

First Harvard University Press paperback edition, 1998

Library of Congress Cataloging-in-Publication Data

Baird, Douglas G., 1953–
Game theory and the law / Douglas G. Baird,
Robert H. Gertner, Randal C. Picker.
p. cm.

Includes bibliographical references and index.

ISBN 0-674-34119-8 (cloth)

ISBN 0-674-34111-2 (pbk.)

1. Law—Methodology. 2. Game theory. I. Gertner,
Robert H. II. Picker, Randal C. III. Title.

K212.B35 1994

340'.1—dc20 94-8625

For
Bruce, Emily, and Lucy,
Leo,
and
Gretchen, Benjamin, and Adam

Contents

Preface xi

Introduction: Understanding Strategic Behavior 1

Bibliographic Notes 4

1 Simultaneous Decisionmaking and the Normal Form Game 6

The Normal Form Game 6

Using Different Games to Compare Legal Regimes 14

The Nash Equilibrium 19

Civil Liability, Accident Law, and Strategic Behavior 24

Legal Rules and the Idea of Strict Dominance 28

Collective Action Problems and the Two-by-Two Game 31

The Problem of Multiple Nash Equilibria 35

Summary 46

Bibliographic Notes 46

2 Dynamic Interaction and the Extensive Form Game 50

The Extensive Form Game and Backwards Induction 50

A Dynamic Model of Preemption and Strategic Commitment 57

Subgame Perfection 63

Summary 75

Bibliographic Notes 77

3 Information Revelation, Disclosure Laws, and Renegotiation 79

Incorporating Beliefs into the Solution Concept 80

The Perfect Bayesian Equilibrium Solution Concept 83

Verifiable Information, Voluntary Disclosure, and the Unraveling Result	89
Disclosure Laws and the Limits of Unraveling	95
Observable Information, Norms, and the Problem of Renegotiation	109
Optimal Incentives and the Need for Renegotiation	112
Limiting the Ability of Parties to Renegotiate	116
Summary	118
Bibliographic Notes	119
4 Signaling, Screening, and Nonverifiable Information	122
Signaling and Screening	122
Modeling Nonverifiable Information	125
Signals and the Effects of Legal Rules	142
Information Revelation and Contract Default Rules	147
Screening and the Role of Legal Rules	153
Summary	156
Bibliographic Notes	157
5 Reputation and Repeated Games	159
Backwards Induction and Its Limits	159
Infinitely Repeated Games, Tacit Collusion, and Folk Theorems	165
Reputation, Predation, and Cooperation	178
Summary	186
Bibliographic Notes	186
6 Collective Action, Embedded Games, and the Limits of Simple Models	188
Collective Action and the Role of Law	189
Embedded Games	191
Understanding the Structure of Large Games	195
Collective Action and Private Information	202
Collective Action Problems in Sequential Decisionmaking	208
Herd Behavior	213
Summary	217
Bibliographic Notes	217
7 Noncooperative Bargaining	219
Modeling the Division of Gains from Trade	219
Legal Rules as Exit Options	224

Bargaining and Corporate Reorganizations	232
Collective Bargaining and Exit Options	237
Summary	241
Bibliographic Notes	241
8 Bargaining and Information	244
Basic Models of the Litigation Process	244
Modeling Separate Trials for Liability and Damages	251
Information and Selection Bias	260
Discovery Rules and Verifiable Information	261
Summary	266
Bibliographic Notes	266
Conclusion: Information and the Limits of Law	268
<i>Notes</i>	275
<i>References</i>	289
<i>Glossary</i>	301
<i>Index</i>	319

Simultaneous Decisionmaking and the Normal Form Game

The Normal Form Game

The simplest strategic problem arises when two individuals interact with each other, and each must decide what to do without knowing what the other is doing. An accident involving a motorist and a pedestrian is such a case. The likelihood of an accident is determined by how carefully the motorist drives and how carefully the pedestrian crosses the street. Each must decide how much care to exercise without knowing how careful the other is. The behavior of the motorist and the pedestrian also depends on the legal regime. Legal scholars have long assumed that motorists will drive more carefully if they are liable for the injuries that the pedestrian suffers in the event of an accident. This observation alone, however, does not tell us how to shape the law of torts, and we must know more about how law affects these simple interactions if we are to understand its effects on more complicated ones.

Much of law and economics scholarship over the past several decades has focused on the intriguing claim that many legal regimes, including all those in Anglo-American tort law, induce the motorist and the pedestrian to act in a way that minimizes the total costs of the accident, costs that include the possibility of injury to the pedestrian as well as the expenses the motorist and the pedestrian incur when they take care to avoid the accident. To draw these conclusions, however, we need to make many assumptions. The pedestrian and the motorist, for example, must know what the legal rule is, and courts must be able to enforce it. Indeed, we cannot draw firm policy prescriptions or choose among possible tort regimes without subjecting all these assumptions to close scrutiny.

Many scholars have undertaken the task of exploring the various assumptions leading to the conclusion that these many different tort regimes are efficient. The debate about which rules work best when certain assumptions are relaxed now fills many volumes. We do not revisit this debate here. Rather, we begin by using the interaction between the motorist and the pedestrian to introduce one of the basic tools of game theory, the normal form game. We then show why many different legal regimes tend, under the same set of assumptions, to induce both parties to act in a way that is mutually beneficial. The model we develop in this chapter allows us to make clear exactly what it means to assume that individuals in the position of the motorist and the pedestrian are rational.

Game theory, like all economic modeling, works by simplifying a given social situation and stepping back from the many details that are irrelevant to the problem at hand. The test of a model is whether it can hone our intuition by illuminating the basic forces that are at work but not plainly visible when we look at an actual case in all its detail. The spirit of the enterprise is to write down the game with the fewest elements that captures the essence of the problem. The use of the word "game" is appropriate because one can reduce the basic elements of complicated social and economic interactions to forms that resemble parlor games.

Our goal in this chapter is to understand the common thread that unites different tort regimes. These regimes range from comparative negligence, in which liability is apportioned between parties according to their relative failure to exercise care, to strict liability with a defense of contributory negligence, in which the motorist is liable to compensate the pedestrian for any injuries unless that pedestrian acted carelessly. To discover what these different tort rules have in common, we can use a model in which the motorist and the pedestrian are each completely informed about everything, except what level of care the other will exercise. They know what it means to act carefully, and they know what consequences the law attaches to any combination of actions. Similarly, we may assume that courts can enforce any given legal regime and that they have all the information they need to do so. For many questions, of course, we cannot make so many simplifying assumptions, but, as we shall see, it is useful to do so here.

We model the interaction between the motorist and the pedestrian by using a traditional game theory model called a *normal form game*, sometimes referred to as the *strategic form* of a game. The normal form game consists of three elements:

1. The *players* in the game.
2. The *strategies* available to the players.
3. The *payoff* each player receives for each possible combination of strategies.

In the accident case, identifying the players is easy, at least if we avoid introducing such complications as whether one or both of the parties is insured. There are only two players, the motorist and the pedestrian.

The next step is to identify the strategies available to the players (or, to use a formal term, the *strategy space* of each of the players) by looking at the options that are open to them. Defining the strategy space is perhaps the most important decision in creating a model in game theory. The range of actions available to a motorist and a pedestrian is broad. Before we even reach the question of how fast the motorist chooses to drive, a choice that lies along a continuum, we face many others, such as the motorist's decision to buy a car, what kind of car to buy, and whether to go on a trip in the first place. Similarly, before choosing how carefully to cross the street, the pedestrian must first decide whether to go on a trip, whether to walk, which route to take, and where to cross the road.

How many of these possibilities are put into the model depends on what we want the model to do. A study of the forces at work in different tort regimes and the assumptions of rationality on which they depend requires only a strategy space in which the players must pick between two actions. In our model, therefore, each of the players faces only a binary choice—either to exercise *due care*, the amount of care that is socially optimal (by driving carefully or crossing the road carefully), or not to exercise care (by, for example, driving too fast or crossing the road without looking).

The last element of the normal form game is the payoff structure. We examine each possible combination of strategies and specify what happens to the pedestrian and the motorist in each case. Tort law is a regime of civil damages. It works by requiring one party to pay another damages under some conditions, but not under others. We can compare different legal regimes by taking games that are the same except for the division of the loss under each combination of strategies.

When the legal rule lets the losses lie with the pedestrian, the payoff to the pedestrian is the amount the pedestrian spends on care plus the expected cost of the accident, that is, the cost of an accident discounted by its likelihood. The likelihood of the accident turns, of course, on the particular combination of strategies that the players have adopted. The

payoff to the motorist is simply the cost of exercising care. When the motorist is obliged to pay damages, the motorist's payoffs are reduced and the pedestrian's are correspondingly increased.

There are several different ways to represent payoffs. The basic idea we wish to convey is that the probability of an accident goes down as investment in care goes up, and the efficacy of one party's investment in care turns on whether the other party invests in care as well. The simplest way to do this is to posit dollar amounts for the costs of taking care and the costs of an injury that reflect these relationships. These amounts are intended to capture only the idea that the probability of an accident decreases as the parties put more effort into being careful, but that at some point the costs of taking additional care do not reduce the likelihood of an accident enough to justify them.

In the model that we build here, exercising care costs each player \$10. An accident, when it occurs, inflicts a \$100 injury on the pedestrian. We assume that the accident is certain to happen unless both players exercise care. (We could, of course, make a less extreme assumption about the need for both parties to take care, but again, the assumption we have made simplifies the problem without compromising our ability to study the effects of different legal regimes.) Finally, we need to make an assumption about the likelihood of an accident in the event that both the motorist and the pedestrian exercise care. We assume that, in this case, there is still a one-in-ten chance of an accident.

In a legal regime in which the motorist is never liable for the accident, if neither exercises care, the motorist enjoys a payoff of \$0 and the pedestrian a payoff of -\$100. If both exercise care, the motorist receives a payoff of -\$10 and the pedestrian a payoff of -\$20. (The pedestrian invests \$10 in care and, assuming that the individual is risk neutral,¹ still faces \$10 in expected accident costs, a one-in-ten chance of a \$100 accident.) If the motorist exercises care and the pedestrian does not, the former receives a payoff of -\$10 (the cost of taking care) and the latter a payoff of -\$100 (the cost of the accident, which is certain to occur unless both take care). Finally, if the motorist does not take care and the pedestrian does, the motorist has a payoff of \$0 and the pedestrian a payoff of -\$110 (the pedestrian invests \$10 in taking care and still suffers a \$100 injury).

At this point, we have created a normal form game for the interactions between the motorist and the pedestrian in a world in which the pedestrian has no right to recover damages from the motorist in the event of an accident. An important step in modeling such an interaction is to take account of the information that the players possess. In the

game between the motorist and the pedestrian, both know their own payoffs and those of the other player. They also know the strategies available to them and the strategies available to the other player. The only thing they do not know is the strategy the other player actually chooses (that is, they do not know whether the other player chooses to exercise care or not). This is a game of *complete but imperfect information*. If a player were unaware of something other than the strategy choice of another player, such as the payoffs the other player receives, it would be a game of *incomplete information*. It is also possible that both players know everything about the structure of the game and that one player can observe the strategy choice of the other player (as would be the case if the pedestrian could observe the motorist's care decision before determining how much care to exercise). In this situation, we would confront a game in which information is *complete and perfect*.

We can represent a normal form game involving two players who choose among a small number of different strategies with a *bimatrix*. In the bimatrix, each cell contains the payoffs to each player for any given combination of strategies. (This way of illustrating a normal form game is called a "bimatrix" because each cell has two numbers in it; in the ordinary matrix, each cell has only one.) Figure 1.1 illustrates our game using a bimatrix. By convention, the first payoff in each cell is the payoff to the row player, and the second payoff is that to the column player. In this figure, we assume arbitrarily that the pedestrian is the row player and the motorist the column player. Hence, the pedestrian's payoffs are given first in each cell.

The bimatrix is only one way of illustrating a normal form game. (A normal form game consists of players, strategies, and payoffs, regardless of how they are set out.) There are many normal form games that cannot easily be represented in this way. For example, we could create a model in which the motorist and the pedestrian choose from a continuum the amount of care to invest. Such a model would also be a normal

		Motorist	
		No Care	Due Care
Pedestrian	No Care	-100, 0	-100, -10
	Due Care	-110, 0	-20, -10

Figure 1.1 Regime of no liability. Payoffs: Pedestrian, Motorist.

form game, even though the number of strategies available to each player is infinite.

Now that we have reduced the interaction between the motorist and the pedestrian to a normal form game, we must "solve" it. We identify the strategies that the players are likely to adopt and then predict the likely course of play. Games are solved through the use of *solution concepts*, that is, general precepts about how rational parties are likely to choose strategies and about the characteristics of these strategies given the players' goals. Solving a game is the process of identifying which strategies the players are likely to adopt.

We must begin by making a fundamental assumption about how individuals make choices: Individuals are rational in the sense that they consistently prefer outcomes with higher payoffs to those with lower payoffs. We express payoffs in dollars, but this is not necessary. The basic assumption at the heart of this mode of analysis is not that individuals are self-interested profit-maximizers or care only about money, but rather that they act in a way that is sensible for them given their own tastes and predilections. This assumption may not always hold in an individual case, because people at times act in ways that are inconsistent and self-destructive. In general, however, people make the best decisions they can, given their beliefs about what others will do.

Once we assume that the behavior of individuals is rational in this sense, we can identify the strategy the motorist is likely to pick. In this game, taking care costs the motorist \$10 and provides no benefit in return. The motorist always does better by not taking care than by taking care. We can predict the motorist's likely choice of strategy because there is a single strategy—taking no care—that, in the context of this model, is better for the motorist no matter what choice the pedestrian makes. Such a strategy is *strictly dominant*. A *dominant strategy* is a best choice for a player for every possible choice by the other player. One strategy is "dominated by" another strategy when it is never better than that strategy and is sometimes worse. When one strategy is always worse than another, it is "strictly dominated."²

This brings us to our first solution concept: *A player will choose a strictly dominant strategy whenever possible and will not choose any strategy that is strictly dominated by another.* This is the most compelling precept in all of game theory. Few would take issue with the idea that individuals are likely to choose a particular strategy when they can always do better in their own eyes by choosing that strategy than by choosing any other.

This solution concept by itself, however, tells us only what the mo-

torist is likely to do in this model. We cannot use this concept to predict the pedestrian's behavior. Neither of the strategies available to the pedestrian is dominated by the other. It makes sense for the pedestrian not to take care when the motorist does not, but to take care when the motorist does. The pedestrian lacks a dominant strategy because either course of action could be better or worse than the other, depending upon what the motorist does.

To predict the pedestrian's behavior, we need to take the idea that players play dominant strategies one step further. Not only will a player likely adopt a strictly dominant strategy, but a player will predict that the other player will adopt such a strategy and will act accordingly. We can predict, in other words, that the pedestrian will choose a strategy based on the idea that the motorist will not choose a strategy that is strictly dominated by another. This brings us to our second solution concept, that of *iterated dominance*: *A player believes that other players will avoid strictly dominated strategies and acts on that assumption. Moreover, a player believes that other players similarly think that the first player will not play strictly dominated strategies and that they act on this belief. A player also acts on the belief that others assume that the first player believes that others will not play strictly dominated strategies, and so forth ad infinitum.*

This extension of the idea that dominated strategies are not played forces us to make a further assumption about the rationality of the players. We not only act rationally and do the best we can given our preferences, but we also believe that others act rationally as well and do the best they can given their preferences. This solution concept seems plausible if the number of iterations is small. After all, most people act rationally most of the time, and we can choose our own actions in anticipation that they will act this way. In addition, this solution concept relies on the assumptions we have made about the information that the players possess in a way that the strict dominance solution concept does not. With a strictly dominant strategy, a player does not need to know anything about the payoffs to the other player. Indeed, players do not need to know anything about their own payoffs, other than that the dominant strategy provides them with a higher payoff in every instance.

The greater reliance on both rationality and information is worth noting. When a legal rule provides both players with a dominant strategy, a player does not need to know any of the details of the legal rule or any of the effects of the legal rule on the other player. A player needs to know only that one course of action is optimal under all conditions.

As we move to solution concepts such as iterated dominance, we must assume more about what individuals know and more about how they believe others will act. The more we have to make such assumptions, the less certain we can be that our model will accurately predict the way individuals behave.

If we accept the iterated dominance solution concept, we can solve the game in Figure 1.1. The pedestrian will believe that the motorist will not exercise care because not exercising care is a dominant strategy. For this reason, the pedestrian will not exercise care either. Because this solution concept requires stronger assumptions about how individuals behave, however, we cannot predict the pedestrian's behavior as confidently as we can predict the motorist's.

We now face the question of whether we can draw any general conclusions from a model that reduces the problem to so few elements. The model is counterfactual in many respects. It assumes that, when there is no legal rule to shift liability to the motorist, the motorist incurs no costs and suffers no harm when an accident takes place. Most motorists are not indifferent to whether they run people over, quite apart from whether they are held liable. Given this and the many other similar assumptions, we want to know what general conclusions we should draw from the model.

Under our assumptions, once the motorist fails to exercise care, the accident will take place no matter what the pedestrian does. Although badly off by not taking care, the pedestrian is even worse off by taking care. We should not infer, however, that, as a general matter, pedestrians are likely to take no care when motorists fail to take optimal care. The model generated this result only because an accident was certain to occur unless the motorist exercised care. The interactions between investments in care on the part of the motorist and the pedestrian in another model might be quite different. The pedestrian might rationally take *more* rather than *less* care when the pedestrian believes that the motorist will take too little care. The motorist's failure to take care might make it sensible for the pedestrian to be even more vigilant.

The model, however, does contain a robust prediction. In a legal regime of no liability, the motorist will have too little incentive to take care. The motorist will take the optimal amount of care only if one makes implausible assumptions about the way people behave.³ The motorist's strategy of taking the optimal amount of care—the amount of care that minimizes the total costs of accidents—is strictly dominated by a strategy of taking too little care. The amount of care the motorist takes will not be optimal. We would all be better off if we

had a legal rule that induced both the motorist and the pedestrian to exercise due care. A legal regime in which the motorist is never liable contains a built-in bias that is likely to lead either to too many accidents or to unduly costly investments in care by the pedestrian.

This result in itself is hardly startling. To say that the strategy of taking reasonable care is dominated by another strategy of taking less than due care restates a familiar insight: individuals are more likely to be careless in a world in which people are not liable when they act carelessly. In such a world, people do not fully internalize the costs of their actions. The motorist enjoys all the benefits of driving fast but does not bear all the costs, namely, the danger of injuring a pedestrian. When we capture the problem of the pedestrian and the motorist in the form of a two-by-two bimatrix, however, not only are the incentives of the motorist made manifest, but, as we show in the next section, we can see how a change in the legal rules changes the incentives of the motorist and the pedestrian at the same time.

Using Different Games to Compare Legal Regimes

We can now use different variations of the game with the pedestrian and the motorist to compare legal regimes. We have the same players and the same strategies available to the players, but we change the payoffs, taking account of these different legal regimes. The payoff to the pedestrian under any strategy combination, in other words, now includes the expected value of the damage award the pedestrian receives, and the payoff to the motorist is lowered by the expected value of the damage award that must be paid to the pedestrian. The size of the expected damage award will be determined, of course, both by the liability rule for that particular combination of strategies and by the likelihood of an accident when the players choose those strategies.

Instead of a legal regime of no liability, let us go to the opposite extreme. The motorist is liable whenever there is an accident. This is a regime of pure strict liability. The motorist must pay for all of the pedestrian's injuries, regardless of whether either exercises care. This is the opposite of a regime in which there is no liability at all, in the sense that the motorist rather than the pedestrian bears the costs of the accident itself. This game is illustrated in Figure 1.2.

Exercising care still costs each player \$10, and an accident is certain to happen (and cause \$100 in damages) unless both exercise care. As before, there is a one-in-ten chance of an accident even if both exercise care. In this game, however, if neither exercises care, the motorist en-

		Motorist	
		No Care	Due Care
Pedestrian	No Care	0, -100	0, -110
	Due Care	-10, -100	-10, -20

Figure 1.2 Pure strict liability. Payoffs: Pedestrian, Motorist.

joys a payoff of -\$100 and the pedestrian a payoff of \$0. If both exercise care, the motorist receives a payoff of -\$20 and the pedestrian a payoff of -\$10. (The motorist invests \$10 in care and still faces \$10 in expected accident costs, a one-in-ten chance of a \$100 accident.)

If the motorist exercises care and the pedestrian does not, the former receives a payoff of -\$110 (the cost of taking care plus the cost of the accident) and the pedestrian a payoff of \$0. (The pedestrian incurs no costs associated with taking care and is, by assumption, fully compensated for any injury and thus made whole in the event of an accident.) Finally, if the motorist does not take care and the pedestrian does, the motorist has a payoff of -\$100 (the motorist takes no care but is liable for the cost of the accident) and the pedestrian a payoff of -\$10 (the pedestrian invests \$10 in taking care but recovers damages for the injuries arising from the accident).

Before solving this game, note the relationship between the games in Figure 1.1 and Figure 1.2: The sum of the payoffs in each cell in Figure 1.2 is the same as that in the corresponding cell in Figure 1.1. The change in the liability rule, however, reallocates the sum between the pedestrian and the motorist. We can capture the change in the legal rules by changing the payoffs, not by changing the strategies available to the parties themselves. Strategies represent those actions that are physically possible, whereas payoffs tell us the consequences of actions. Because the tort rules we are examining attach consequences to actions, they are reflected in the payoffs, not in the strategies.

We can discover the strategies that the players are likely to adopt by again using the solution concepts of strict dominance and iterated dominance. The pedestrian in this game (rather than the motorist) has a dominant strategy. The pedestrian will not exercise any care, regardless of what the motorist does. Under strict liability, the motorist must fully compensate the pedestrian for any injury. The pedestrian therefore bears none of the costs of the accident and ignores those costs

when deciding whether to exercise care. The pedestrian takes no precautions because precautions are costly and bring no benefit. Exercising no care is a dominant strategy for the pedestrian. We can then invoke the concept of iterated dominance. The motorist recognizes that the pedestrian will act this way and therefore chooses not to exercise care as well. In this regime, the accident is certain to occur, the motorist is in all events liable, and the investment in care brings the motorist no benefit.

This model exposes a weakness of a regime of pure strict liability that parallels the one we saw in a regime of no liability. The pedestrian may have insufficient incentive to exercise care because the motorist bears the costs of the accident. A more elaborate model might provide for damages that do not fully compensate the pedestrian for the cost of the accident. Nevertheless, the general point of the model remains, even if these and other plausible complications are taken into account. The pedestrian will not fully consider the costs of the accident in deciding whether to take care, and hence may take too little care.

There are some accidents—such as airplane crashes—in which the victims have virtually no ability to take precautions. One might also favor a regime of pure strict liability in such cases if there were no way for a court to determine whether a victim took care. One would not, however, want such a regime in a situation in which the victim could take a number of readily visible steps to prevent an accident, and where it was therefore important to have a legal rule that provided the victim with an incentive to take care, just as one would not want a regime of no liability if the injurer needed an incentive to take care. Under a regime of either no liability or strict liability, it is likely to be in the self-interest of at least one of the parties to exercise less than due care.

We turn now to the legal regime of negligence plus contributory negligence, for a long time the prevailing principle of Anglo-American tort law. Under this regime, the pedestrian can recover damages only if the motorist is negligent *and* the pedestrian is not. This rule of law leads to the normal form game set out in Figure 1.3. As before, the legal rule does not change the strategies available to the players or the sum of the payoffs in each cell. All that changes is the allocation of the costs of the accident between the parties.

When we compare Figure 1.3 with Figure 1.1, we see that the two are identical except in the cell in which the pedestrian exercises due care and the motorist fails to do so. In this event, the pedestrian incurs a loss of only \$10, the cost of taking care, and the motorist bears the \$100 cost of the accident. The pedestrian continues to bear the cost of

		Motorist	
		No Care	Due Care
Pedestrian	No Care	-100, 0	-100, -10
	Due Care	-10, -100	-20, -10

Figure 1.3 Negligence with contributory negligence.
Payoffs: Pedestrian, Motorist.

the accident in the other three cases. The first two are cases in which the pedestrian fails to exercise care and the expected cost of the accident is \$100. The third is that case in which both players spend \$10 exercising care and the pedestrian also bears the \$10 expected cost of the accident.

Unlike the game in Figure 1.1, this game is one in which the pedestrian has a dominant strategy. The pedestrian is always better off taking care. The motorist no longer has a dominant strategy. Whether the motorist is better off taking care depends on whether the pedestrian also takes care. If we accept the idea of iterated dominance, however, we can predict the strategy that the motorist will choose. The motorist recognizes that the pedestrian will exercise due care and therefore decides to take due care as well. Hence, under this legal regime, both pedestrian and motorist take due care. Our need to use iterated dominance to solve this game identifies a central assumption underlying this regime. To believe that a rule of negligence coupled with contributory negligence works, we must think that the motorist acts rationally and believes that the pedestrian acts rationally as well.

A comparison between the two models shows how this legal rule works. The only difference between Figure 1.1 and Figure 1.3, as mentioned, is in the cell representing the strategy combination in which the pedestrian exercises due care and the motorist does not. In Figure 1.1, the payoffs were -\$110 and \$0 to the pedestrian and the motorist respectively; in Figure 1.3, they are -\$10 and -\$100. This strategy combination is not the solution to either game, yet changing the payoffs associated with it completely alters the strategies that the parties adopt, and hence the expected play of the game. *The legal rule brings about changes even though it attaches consequences to actions that are never taken, either when the legal rule is in place or when it is not.* We shall return to this idea on a number of occasions. It will prove particularly important

in those contexts in which some parties have information that others do not.

The legal regime in which there is strict liability, subject to a defense of contributory negligence, is set out in Figure 1.4. By injecting contributory negligence into the scheme of things, we make due care a dominant strategy for the pedestrian. Exercising due care results in a payoff to the pedestrian of $-\$10$ instead of $-\$100$. As long as the motorist believes that the pedestrian will play this strictly dominant strategy, the motorist will exercise due care as well, preferring a payoff of $-\$20$ to one of $-\$100$.

The difference between negligence coupled with contributory negligence and strict liability coupled with contributory negligence lies only in the consequences that follow when both players exercise care. In the negligence-based regime, the pedestrian bears the costs of an accident when both players exercise due care, whereas in the strict liability regime, the motorist does. The difference, however, does not affect the solution to the game, because exercising due care never costs a party more than $\$20$, and failure to exercise due care under either regime when the other party does exercise due care leads to a loss of $\$100$.

The comparison between regimes of negligence coupled with contributory negligence and strict liability coupled with contributory negligence in this model makes it easy to understand the well-known insight that both regimes give the two parties incentives to take care. It also unpacks the rationality assumptions that these rules need in order to work well even if we believe that everyone is well informed and that enforcement costs are low. We must assume not only that individuals behave rationally, but also that individuals expect others to behave rationally as well.

This way of looking at the problem reveals one of the important but subtle ways in which a legal rule works. A change in a legal rule can

		Motorist	
		No Care	Due Care
Pedestrian	No Care	-100, 0	-100, -10
	Due Care	-10, -100	-10, -20

Figure 1.4 Strict liability with contributory negligence.
Payoffs: Pedestrian, Motorist.

alter the behavior of both parties even by changing outcomes that are never seen under either the new or the old regime. Similarly, one can make a major alteration, such as changing the identity of the party who bears the loss when both exercise care, without affecting the incentives of either party.

Players choose the strategies that maximize their own payoffs. Hence, players compare their payoffs under a strategy relative to their own payoffs under other strategies. Changing the damages that one player must pay another alters the solution to the game only if the change makes a player's own payoff so much higher (or so much lower) that the player stands to do better (or worse) by choosing that strategy rather than another. A regime of negligence coupled with contributory negligence places the costs of an accident on the pedestrian when both parties exercise care, whereas a regime of strict liability coupled with contributory negligence places them on the motorist. These differences, however, are not large enough to affect the strategy choices of the players. One can change how losses are allocated when both parties exercise due care, as long as the consequences attached to exercising less than due care make that strategy choice less attractive.

The Nash Equilibrium

Regimes of negligence with contributory negligence or strict liability with contributory negligence have a sharp binary character. The pedestrian who falls just short of exercising due care receives nothing. Many have found this outcome normatively troubling and have advocated regimes of comparative negligence, in which both motorist and pedestrian shoulder some of the costs of an accident when both fail to exercise due care. A number of jurisdictions have adopted comparative negligence in their accident law. In this section, we examine comparative negligence regimes and, in the course of analyzing them, discuss another solution concept, the Nash equilibrium. A regime of comparative negligence may be harder to implement than the legal regimes it has replaced. Our focus, however, is again on how comparative negligence regimes differ from others as a matter of first principle. In this section, we examine comparative negligence with an eye to understanding how the incentives of the parties change once both bear some share of the liability when both fail to exercise due care.

The incentives that a comparative negligence regime imposes on parties depend on how liability is allocated when both parties fail to exercise due care. In some jurisdictions, the judge instructs the jury to allo-

cate liability after considering "all the surrounding circumstances"⁴ or "the nature of the conduct of each party and the extent to which each party's conduct caused or contributed to the plaintiff's injury."⁵ Other sharing rules include one that looks at the amount of care that each party took relative to the total amount of care that should have been taken,⁶ and another that reduces damages in the proportion which "the culpable conduct attributable to the claimant . . . bears to the culpable conduct which caused the damages."⁷ The differences among the various sharing rules lead one to ask whether it matters which of these sharing rules one chooses. We can address this question by considering an extreme sharing rule, one in which the person who was the most careless bears a disproportionate share of the costs of the accident.

The players have a choice between exercising no care, some care, and due care. If both exercise no care or if both exercise some care, liability is split equally between them. If one exercises no care and the other exercises some care, however, the party who exercised no care bears a disproportionately large share and the party who exercised some care bears only a small portion of the costs of the accident. An accident imposes costs of \$100 and is certain to happen unless both parties exercise due care.

Due care costs each player \$3. Some care costs each party \$1. When both exercise due care, the chances of an accident drop to one in fifty, and the expected costs of the accident to the pedestrian, the party who bears the costs when both exercise due care, is \$2. If both parties exercise no care, they each bear the \$50 loss. If both exercise some care, they each bear a \$50 loss from the accident and \$1 from the cost of the care they did exercise. If one exercises no care and the other exercises some care, the first bears \$99 of the cost of the accident, while the other, more careful person bears only \$1 in liability for the accident, plus the \$1 cost of taking care that the player already incurred. The game is illustrated in Figure 1.5.

The incentives of the players are not so clear-cut under this comparative negligence regime as they were under the ones we examined earlier. Neither player has a strictly dominant strategy. In the comparative negligence regime set out in Figure 1.5, some care is usually worse than due care, but sometimes it is better than due care (in those cases in which the other player exercises no care). Neither due care nor some care dominates the other. Nevertheless, most people looking at this game have the intuition that both pedestrian and motorist will exercise due care.

Although each would be better off exercising only some care if the other exercised no care, neither the motorist nor the pedestrian expects

		Motorist		
		No Care	Some Care	Due Care
Pedestrian	No Care	-50, -50	-99, -2	-100, -3
	Some Care	-2, -99	-51, -51	-101, -3
	Due Care	-3, -100	-3, -101	-5, -3

Figure 1.5 Comparative negligence (skewed sharing rule).
Payoffs: Pedestrian, Motorist.

the other to exercise no care. The strategy combination in which one motorist exercises no care and the other exercises some care is not a likely course of play. While one player (the one who exercises some care) favors this strategy combination over all the rest, the other player (the one who exercises no care) does not, and therefore should be expected to choose a different strategy.

There are two formal ways of capturing this intuition and solving this game. The first idea is an application of the dominance ideas that we have already developed. No care on the part of the pedestrian and the motorist is dominated by due care. Because neither will play a strictly dominated strategy, we can reduce the game to a simple two-by-two game in which the only strategies on the part of each player are due care and some care. At this point, due care is a strictly dominant strategy for both motorist and pedestrian.

The second way we could solve this game is based on the following principle: *The combination of strategies that players are likely to choose is one in which no player could do better by choosing a different strategy given the strategy the other chooses. The strategy of each player must be a best response to the strategies of the other.* The solution concept based on this principle is known as a *Nash equilibrium*.⁸ Introduced by John Nash in 1950, the Nash equilibrium has emerged as a central—probably the central—solution concept of game theory.

As applied to this game, this principle tells us that a strategy of some care on the part of either the pedestrian or the motorist and no care on the part of another is not likely to be the combination of strategies that the players adopt. Given that one of the players has adopted a policy of some care, the other player is better off using due care rather than no care.

The Nash equilibrium has loomed larger than dominance solvability

because it can be usefully applied to more games of interest to economists. If the successive elimination of dominated strategies leads to a unique outcome, that outcome is also the unique Nash equilibrium of that game. A game that we cannot solve through the successive elimination of dominated strategies, however, often has a Nash equilibrium.

The ability of the Nash equilibrium concept to solve additional games makes it both more powerful and more controversial than solution concepts based on the idea that players do not choose dominated strategies. One can point to games in which the unique Nash equilibrium may not be the combination of strategies that players would in fact adopt.⁹ Moreover, the Nash solution concept often does not identify a unique solution to a game. When there are multiple Nash equilibria, we may not be able to identify one of these as that which the players are likely to choose. Indeed, when there are multiple Nash equilibria, there is no guarantee that the outcome of the game is going to be a Nash equilibrium. Each player, for example, might adopt a strategy that is part of a different Nash equilibrium, and the combination of strategies might not be Nash. Nevertheless, the Nash solution concept is often useful in the context of a game such as this one and many others that we shall examine. We therefore focus on its formal definition more closely.

In a two-person game, a pair of strategies will form a Nash equilibrium when each player cannot do better given the strategy the other player has adopted. A Nash equilibrium, in other words, is a pair of strategies such that each is a best response to the other. To test whether a strategy combination forms a Nash equilibrium, let us call the strategy for the first player x^* and the strategy for the second player y^* . Now we need to ask whether, given that the second player will play y^* , the first player can do strictly better by switching to some strategy other than x^* . Similarly, we need to ask whether, given that the first player will play x^* , the second player can do strictly better by switching to some strategy other than y^* . If there is no better strategy for the first player than x^* in response to the second player's y^* , and if there is no better strategy for the second player than y^* in response to x^* , then this pair is a Nash equilibrium for the game.

Virtually all games of interest to us have at least one Nash equilibrium. More important, a strategy combination that is not a Nash equilibrium is unlikely to be the solution to the game. We can see this by assuming for a moment the opposite—that a particular strategy combination that is not Nash is the solution to the game. If such a combination is the solution to a game, both players should be able to identify

this fact beforehand. If the strategy is not Nash, it follows, by definition, that one of the players is choosing a strategy that is not a best response given what the other player is doing. Put yourself in the position of the player whose strategy is not a best response. Why should you choose the strategy that is asserted to be part of the solution to the game? Given what the other player is supposed to do in this purported solution, you can do better. You are not acting rationally if you pick a strategy that does not maximize your own payoff.

If we return to the game that models a comparative negligence regime with a sharing rule that skews damages toward the party who was the most careless, we can see that only the strategy combination in which both players exercise due care is a Nash equilibrium. The strategy of due care for one player is always the best response when the other player exercises due care. If the players adopted any other combination of strategies, at least one of them would be choosing a strategy that was not a best response to the other, and that player could receive a higher payoff by switching to a different strategy.

Assume, for example, that the pedestrian exercised some care and the motorist exercised no care. The pedestrian has no incentive to exercise due care, given the motorist's strategy. The pedestrian prefers a payoff of $-\$2$ to one of $-\$3$. The motorist's strategy of no care, however, is not a best response to the pedestrian's strategy of some care. The motorist is better off exercising some care (and enjoying a payoff of $-\$51$ rather than $-\$99$) or exercising due care (and enjoying a payoff of $-\$3$). The strategy combination in which both players exercise due care (and enjoy payoffs of $-\$5$ and $-\$3$ respectively) is the only combination in which neither player has an incentive to change strategy—or "deviate"—given the strategy of the other.

In this model, notwithstanding the extreme sharing rule, each player has the correct incentive. The model suggests that a comparative negligence regime is likely to induce both parties to exercise due care, independent of the particular sharing rule that a comparative negligence regime adopts. As others have shown,¹⁰ even when the strategy space of each player is continuous, the only Nash equilibrium (and the only combination of strategies that survives the repeated elimination of dominated strategies) is the strategy combination in which both players exercise due care. As long as one accepts the Nash solution concept or the concept of iterated elimination of dominated strategies as a good prediction of the strategies that players will adopt, a comparative negligence rule gives players the right incentives in this simple environment, regardless of how the sharing rule is itself defined.

Civil Liability, Accident Law, and Strategic Behavior

The common law influences behavior by allowing injured individuals to bring actions for civil damages under specified circumstances. Such a legal regime stands in marked contrast to a regulatory regime that prescribes certain courses of conduct or subjects particular actions to criminal sanctions. A civil damages rule, seen through a game-theoretic lens, is simply a rule that reallocates the payoffs between players for each combination of strategies. The amount of wealth in each cell of the bimatrix remains the same, but it is distributed differently.

The power of a civil damages rule to affect the behavior of the players should not be underestimated. In games of complete but imperfect information, an infinite number of civil damages rules exists such that any outcome (including that which is the social optimum) is the only one that survives the iterated elimination of dominated strategies.¹¹ For this reason, it should come as no surprise that a number of different regimes (including the common law regime of negligence coupled with contributory negligence) provide players with the correct incentives as a matter of first principle. The common law is "efficient" in the limited sense that it gives the parties the correct incentives under a number of strong assumptions, but so do many other legal regimes. The interesting question is whether common threads unite these different rules of civil liability beyond the fact that they all induce parties to act efficiently under the same set of assumptions.

Regimes of negligence with contributory negligence, strict liability with contributory negligence, and comparative negligence all share three features:

1. The legal regimes are regimes of compensatory damages. Parties always bear their own cost of care, and the legal rules never require an injurer to pay more than is necessary to compensate the victim for the injury.
2. An injurer who exercises at least due care pays no damages whenever the victim does not exercise at least due care, and, in parallel fashion, a victim is fully compensated for any injuries suffered whenever the victim exercises at least due care and the injurer does not.
3. When both the injurer and the victim exercise at least due care, the costs of the accident are borne by one or the other or divided between them in some fixed proportion.

There are a number of other legal regimes that share these features as well. These include, for example, a legal regime in which the injurer

is always liable when negligent and there is no defense of contributory negligence. They also include a regime of strict liability coupled with comparative negligence, in which the injurer is liable for the accident if both the injurer and the victim exercise care, but the losses are shared when both do not. There are other regimes that we do not see—such as those in which the losses are divided evenly between the parties when both exercise care—that also share these three features.

These three common characteristics are quite general. Regimes with radically different distributional consequences all share them. Nevertheless, they are themselves sufficient to ensure that both the injurer and the victim take due care. The proof of this proposition using the Nash equilibrium concept is the easiest to show, and we can set it out quickly.

Note first that excess care can never be part of a Nash equilibrium. If the other player exercises less than due care, a player avoids liability completely by playing due care. Excess care just creates costs and provides no additional benefit for this player. The costs of the accident have already been shifted to the first player. Alternatively, if the other player exercises due care, excess care cannot be a best response for a player, given how we have constructed our definition of due care. (If it were a best response, the due care-due care strategy combination could not be the social optimum.) We can therefore restrict our focus to strategies of due care or too little care.

The strategy combination in which both take care is a Nash equilibrium. Note that there are two generic allocations of liability when both players exercise due care: either one player bears the full costs of the accident, or the costs are shared. Consider the incentives of a player who does not bear the full costs of the accident. For this player, deviating from due care means bearing the full costs of the accident. For the deviation to be sensible, the gains from lowering the private cost of care must exceed, not only the extra expected costs of the accident, but also the additional fraction of those costs now borne by this player. This cannot happen, given how due care is defined.

A player is always better off exercising due care and bearing only a part of the expected costs of the accident than exercising less than due care and bearing all the costs of the accident. The increased costs of taking care are necessarily less than the reduction in the player's share of the accident costs. The reduction in the accident costs alone more than offsets the added costs of exercising due care rather than some lesser level of care. Any player who does not bear the full costs of the accident—namely, both players when costs are shared and the player bearing none of the costs when one player bears all of the

costs—has due care as a best response when the other player exercises due care.

Consider finally a player bearing the full costs of the accident. This player cannot shift liability through the choice of strategy and thus just cares about minimizing the costs of the accident. Given that the other player is playing due care, the remaining social costs—all of which are borne by the first player—are minimized by selecting due care. Again, this follows from our definition of due care.

We must also ask whether any other combination of strategies could be a Nash equilibrium if a legal rule had these features. Consider any strategy combination in which one party is exercising due care and the other is not. In this event, the party who is not exercising due care bears all the costs of the accident itself and the costs of care which that party is taking. Thus, this party is better off deviating and exercising due care. Even if the party still bears all these costs, exercising due care leaves the party better off. The reduction in the expected accident costs necessarily offsets the costs of additional care. When one party exercises due care, the best response for the second party is always to exercise due care as well.

Consider finally the possibility that both players exercise less than due care. Either player could deviate, play due care, and incur only the costs of due care. We need to ask whether one party or another will have an incentive to play such a strategy. If so, the strategy of less than due care cannot be a best response for that player. In order for a strategy combination in which both players exercise less than due care to be a Nash equilibrium, two conditions must hold simultaneously. First, the injurer's share of the liability plus the injurer's cost of taking care must be less than the injurer's cost of taking due care. (The injurer's cost of taking due care is the relevant value for comparison. Given that the victim is exercising less than due care, the injurer can avoid liability completely by taking due care.) Second, the part of the injury that remains uncompensated plus the victim's cost of taking care must be less than the victim's cost of taking due care.

If both conditions hold at the same time, the costs to both the injurer and the victim together in this strategy combination are less than the costs to both of taking due care. In other words, the costs of the injurer's and the victim's taking care plus the costs of the accident—the social costs of the accident—must be less than the cost to the victim and the injurer of taking due care. This, however, cannot be true because, by definition, when the victim and the injurer exercise due care they minimize the total social costs of the accident. The costs of taking due care

can never exceed the total social costs of an accident under any other combination of strategies. For this reason, one player would always prefer to exercise due care rather than less than due care in a strategy combination in which the other player was exercising less than due care as well.

We have now ruled out the possibility that any strategy combination in which one party does not exercise due care can be a Nash equilibrium. Therefore, the only Nash equilibrium in a game of complete but imperfect information in which the applicable legal regime satisfies these three conditions is the strategy combination in which both players exercise due care. Seen from this perspective, the various legal regimes that govern torts under Anglo-American rule are different variations on the same basic principle. Rules such as negligence, negligence coupled with contributory negligence, comparative negligence, or strict liability coupled with contributory negligence all share three very general attributes, which are themselves sufficient to bring about efficient outcomes in games of complete but imperfect information.

Under the strong assumptions we have been making, all the Anglo-American tort regimes induce parties to behave in a way that is socially optimal. They are all compensatory damage regimes in which a party never bears the costs of the accident if that party takes due care and the other does not. As long as a rule has these features, parties will have the right set of incentives. The rules have dramatically different distributional consequences, but these variations themselves do not give parties an incentive to behave differently.

Three general observations can be drawn from this examination of the common principle that links these different regimes. First, all these rules work in the same way and all depend on the same assumptions about the rationality of both injurers and victims. They require us to assume not only that individuals act rationally, but that they expect others to do so as well. Second, because these rules provide parties with the same set of incentives, choosing among the different rules requires us to examine all those things that are assumed away in this environment, such as whether a rule is likely to lead to more litigation or whether a court is more likely to make errors in enforcing a particular rule. We also cannot ignore the informational demands that the rules place on the parties. Parties do not need to know the particular content of the legal rule as long as they know that it is in their interest to exercise due care; all the rules, however, depend on at least one of the parties knowing what constitutes due care in any particular context.

Finally, this approach to the problem naturally leads to asking what

other kinds of rules are possible. In the next section, we ask whether rules exist under which exercising due care is a strictly dominant strategy for both sides. If such a rule can be fashioned, we would not have to assume that parties expect each other to behave rationally. Each party would have the incentive to take care, regardless of what that party thought the other would do. We would not, of course, necessarily want to embrace such a rule if it existed, because it might come at too great a cost. Nevertheless, the first step in understanding how legal rules work is understanding what assumptions are essential to the enterprise.

Legal Rules and the Idea of Strict Dominance

In this section, we ask whether it is possible to state a rule of civil damages such that both players always find it in their interest to exercise due care, regardless of what each player expects the other to do. We begin by asking whether a regime can have this feature if it shares the same premise as those we have examined so far—legal regimes of compensatory damages in which a player who exercises due care never bears the costs of the injury if the other fails to exercise due care as well. We then ask whether other rules exist that are not built on this principle.

There is an intuitive way to describe the basic feature that we should see in a compensatory damages regime in which due care strictly dominates less than due care. Let us return to our example with the pedestrian and the motorist. The rule should ensure that the motorist and the pedestrian are always rewarded for the investments in care that they make, no matter what the other does. Hence, we want to make sure that both are better off for every dollar of additional care that they invest until they have invested in the optimal amount of care. In other words, their expected liability should go down by at least a dollar for each additional dollar they invest in taking care.

The game in Figure 1.5 proved difficult precisely because it did not have this feature. The pedestrian who invested only \$1 in care was exposed to only \$1 of liability when the motorist invested nothing. The pedestrian who invested in due care had to spend \$2 more but would reduce the expected liability by only \$1. When the motorist takes no care, the added costs to the pedestrian of taking due care instead of some care are greater than the benefits; hence, the pedestrian has no incentive to do so. Similarly, in regimes of negligence or strict liability coupled with contributory negligence, the motorist's investment in

care brings no benefits to the motorist when the pedestrian exercises no care. A rule that ensures that both parties have an incentive to exercise due care no matter what the other does requires that the private benefits to a party from taking care always equal or exceed the private costs of taking care.

We can specify a sharing rule in a comparative negligence regime in which the costs that a party faces in taking care always correspond with the benefits that party receives in the way of reduced liability: A party who fails to exercise due care should bear the liability in proportion to the amount that party failed to spend on due care relative to the amount both parties fell short of exercising due care. Let us return to the example in Figure 1.5. Consider how this rule would allocate liability in the event that the pedestrian exercised some care (spending \$1 instead of \$3) and the motorist exercised no care (instead of spending \$3). In this case, the pedestrian should have spent \$2 more, and both parties together should have spent \$5 more (\$2 from the pedestrian and \$3 from the motorist). Hence, the pedestrian should bear $\frac{2}{5}$, or 40 percent, of the liability for the accident. This rule does not allocate liability when both parties exercise due care. As the earlier discussion of negligence and strict liability suggested, the allocation of liability when both parties exercise due care does not affect the strategies that players adopt in a compensatory damage regime.

As stated, this sharing rule ensures only that exercising due care on the part of each party dominates all strategy combinations in which both parties exercise less than due care. The possibility that a player could exercise excessive care therefore must be taken into account. Once we do this, however, we discover that there is no compensatory damages rule of the type we have been discussing in which exercising due care is a strictly dominant strategy for both sides. Let us assume that the pedestrian is rational, but believes that the motorist is not and that the motorist will take excessive care. Excessive care reduces the likelihood of an accident. This may in turn lead a pedestrian who believes that the motorist will exercise excessive care to take less than due care. (Indeed, if the motorist actually did exercise excessive care, we might want the pedestrian to exercise less than due care.)

To modify our comparative negligence rule so that parties always find it in their interest to take due care no matter what they believe others will do, we need to modify this sharing rule to provide that a person also shoulders some of the liability when that individual exercises too much care. Such a rule is counterintuitive because the party who does this already bears the social costs of taking too much care.

The rule's justification lies in making it more likely that parties will take due care, not in the way it parcels out liability when one party acts contrary to self-interest. The obstacles that stand in the way of implementing this rule, including the difficulties of ascertaining due care in any case, are both obvious and substantial. In equilibrium, both parties would exercise due care. But the distribution of liability in combinations of strategies that are not part of the equilibrium are counter-intuitive.

Consider the case in which the pedestrian crosses the street carelessly and is injured, even though the motorist exercised not merely due care, but excessive care. Under this rule, the pedestrian could sue the motorist and obtain a partial recovery of damages to the extent that the motorist was more careful than was socially optimal. Such an outcome seems wrong for two reasons. First, the motorist is punished even though the motorist bears all the costs of driving too carefully. It seems strange to force a party to pay damages to someone else when that party already bears the costs of departing from the social optimum. Second, given that the pedestrian did not exercise care, we may be better off if the motorist exercises excessive care. We can justify this allocation of damages only because we do not expect the motorist ever to exercise excessive care.

There is a simple rule that makes playing due care a strictly dominant strategy for both parties. The key is to relax the assumptions that damages be compensatory and that a party bears only the costs of taking care when that party exercises due care and the other does not. Consider the following regime:

1. Whenever both parties fail to exercise due care, and exercise instead too much or too little care, each party must bear some cost. (This requirement is usually trivial. If a party exercises any care or suffers some of the injury and has no right to recover damages from the other party, that cost is sufficient. The rule does require that the injurer pay some amount in damages in the event that the injurer exercises no care, even though this amount can be quite small.)
2. When one party exercises due care and the other does not, the latter must compensate the former for any injury suffered and must also reimburse the first party for the costs of taking care. (In other words, if the injurer fails to exercise due care, but the victim does, the injurer must compensate the victim not only for the injury, but also for the costs of care that the victim

incurred. Similarly, if the injurer exercises care and the victim does not, the victim must not only suffer the costs of the injury, but also compensate the injurer for the care taken.)

3. When both parties exercise due care, there is some rule allocating costs between them.

This rule is one in which due care is a strictly dominant strategy for both sides. Exercising due care strictly dominates any other strategy when the other player is not exercising due care. In all of these cases, exercising due care costs nothing and doing anything else costs something. Similarly, exercising due care dominates all other strategies when the other player exercises due care as well. Even if one party bears the full costs of the accident when exercising due care, that party is better off exercising due care than doing anything else. When the other player exercises due care, doing so must be the best response for the player. This player bears all the costs of the accident by exercising something other than due care. Hence, a player always has the incentive to exercise due care. Exercising due care minimizes the costs that the player faces, regardless of how much of the costs of the accident the player bears.

This rule has the same knife-edge characteristic that we see in rules incorporating negligence or contributory negligence. Such a rule can work only if we are confident that injurers, victims, and the courts can all identify the due care standard. Many other factors, such as the cost of legal error, need to be taken into account. It is, however, possible to create legal regimes in which parties must focus only on their own actions and do not have to take into account what others are likely to do. These regimes have the virtue of making fewer assumptions about individual rationality. They also should make us skeptical of relying on analyses of legal regimes that depend heavily on the same or similar assumptions and invoke the Nash equilibrium solution concept. If these assumptions hold, we do not even need the Nash equilibrium solution concept. In games of complete but imperfect information, the socially optimal outcome can be implemented with civil damages in strictly dominant strategies.

Collective Action Problems and the Two-by-Two Game

Unlike the game involving the motorist and the pedestrian, in which just two people were involved, many of the problems of strategic behavior facing a legal analyst are problems of collective action in which

many individuals are involved. Nevertheless, these interactions often can also usefully be reduced to two-person games. Consider a problem that can arise in areas that are subject to flooding. At common law, flood waters are regarded as a "common enemy," and individual landowners have a right to build levees to keep flood waters off their land. This legal regime, however, creates a serious problem. Building a levee in one place increases the threat of flooding elsewhere. The response of individuals who are on the other side of the river or are upstream or downstream is to build new levees or increase the height of those that are already in place. In the end, investments in levees may not bring the landowners benefits commensurate with their costs relative to where they would be if no levees were built at all.¹²

In the game involving the motorist and the pedestrian, both players made their care decisions at the same moment in time. The game-theoretic problems involving simultaneous decisionmaking extend to a broader class of cases, however. They include any situation in which the players must act without knowing what the other player has done. Moreover, when enough people are involved so that negotiations between them are costly, the decision of each person may have little effect on the decisions of others. One may know what others do but have little ability to influence them. For such interactions, a simultaneous-move game may again be a useful model.

We can set out the essence of this problem with flooding by imagining that there are only two landowners, each of whom must independently decide whether to build a levee. We illustrate this game in Figure 1.6. If neither builds a levee, each will experience some flooding and suffer a loss of \$4. A levee costs each landowner \$2, but it eliminates the flooding problem only if the other does not build a levee. If one landowner builds a levee and the other does not, the landowner without a levee suffers a large flood and a \$10 loss. If both build levees,

		Landowner 2	
		Don't Build	Build
Landowner 1	Don't Build	-4, -4	-10, -2
	Build	-2, -10	-5, -5

Figure 1.6 Levee collection action game.
Payoffs: Landowner 1, Landowner 2.

they suffer \$3 in flood damage. This amount is less than when neither builds a levee, but the landowners are worse off because the cost of the levee exceeds the amount saved from reducing the amount of flooding.

The two-by-two game that captures collective action problems like the one in Figure 1.6 is commonly called a *prisoner's dilemma*. The name comes from the story that was first told in the 1950s to illustrate the following strategic interaction: Two criminals are arrested. They both have committed a serious crime, but the district attorney cannot convict either of them for this crime without extracting at least one confession. The district attorney can, however, convict them both on a lesser offense without the cooperation of either. The district attorney tells each prisoner that if neither confesses, they will both be convicted of the lesser offense. Each will go to prison for two years. If, however, one of the prisoners confesses and the other does not, the former will go free and the latter will be tried for the serious crime and given the maximum penalty of ten years in prison. If both confess, the district attorney will prosecute them for the serious crime but will not ask for the maximum penalty. They will both go to prison for six years.

Each prisoner wants only to minimize time spent behind bars and has no other goal. Moreover, each is indifferent to how much time the other spends in prison. Finally, the two prisoners have no way of reaching an agreement with each other. Figure 1.7 reduces this story to a normal form game.

The games illustrated in Figures 1.6 and 1.7 have the same structure. Each landowner and each prisoner has a strictly dominant strategy—build a levee or confess. If the other landowner does not build a levee, the first can reduce flooding costs from \$4 to \$2 by building a levee. (There is no flood damage and the levee costs \$2.) If the other landowner does build a levee, building a levee reduces losses from \$10 to \$5. (When the other landowner builds a levee and the first does not, the first landowner who does not build a levee incurs \$10 in flood dam-

		Prisoner 2	
		Silent	Confess
Prisoner 1	Silent	-2, -2	-10, 0
	Confess	0, -10	-6, -6

Figure 1.7 Prisoner's dilemma. Payoffs: Prisoner 1, Prisoner 2.

age. When the first landowner does build a levee at a cost of \$2, flood damage drops to \$3 for a total cost of \$5.) Either way, a landowner is better off building a levee. Similarly, a prisoner is much better off confessing than remaining silent if the other prisoner is going to confess. Six years in prison is preferable to ten years. A prisoner is even better off confessing if the other remains silent. By confessing, the prisoner can avoid prison altogether. No matter what the other prisoner does, a prisoner is better off confessing.

One would much rather not incur the cost of building a levee and suffer from a moderate flood than spend money on a levee and suffer from only slightly less flooding. Similarly, a prisoner would much rather spend two years in prison than six. These outcomes, however, are possible only when the players can reach a binding agreement. In both games each player has a strictly dominant strategy, and the strategy combination the players choose leaves them both worse off than they would be if they could cooperate with each other.

Collective action problems that fit the paradigm of the prisoner's dilemma present a possible case for legal intervention. For example, the government might have the expertise to build a system of levees that would minimize the costs of flooding to all the landowners as a group. As one court put it: "[T]he only adequate method of preventing this result was the unification of the individualistic and antagonistic efforts of the land owners on the opposite sides of the river into one comprehensive co-ordinating plan looking toward the flood control of the river in its entirety."¹³

This kind of problem is also generally known as a *tragedy of the commons*, named for the problem that arises when shepherds who share a common pasture overgraze it. Each shepherd does not incur all the costs of adding an additional sheep to the flock. Each additional sheep reduces the amount of grass available for the other sheep. The benefits to a single shepherd from grazing an additional sheep on the common pasture may be greater than the harm to the other sheep in that shepherd's flock, but smaller than the harm to all the sheep that graze there. Each shepherd enjoys all the benefits of grazing an additional sheep, but the harm to all the other sheep is borne by the shepherds as a group. Moreover, there are so many shepherds that the cost of reaching a consensual bargain among all of them is prohibitive. Hence, the shepherds collectively graze too many sheep on the common pasture.

The existence of transaction costs makes simultaneous decisionmaking an appropriate model for talking about this kind of problem. The model we used reduced the collective action problem to its barest ele-

ments. If we were interested in other questions (such as how the total payoffs in equilibrium change relative to the social optimum as the number of landowners changes), we would need to develop a more elaborate model in which there are many players.

The Problem of Multiple Nash Equilibria

We can illustrate the power of the two-person, two-by-two game by looking at another problem involving flooding and levees. In the previous example, we confronted landowners on opposite sides of the river. A different kind of problem can arise with landowners on the same side of the river. It may be in the interest of a landowner to build a levee and maintain it only if adjacent landowners build levees and maintain them as well. If any levee is improperly maintained, all landowners suffer damage in the event of a flood. A game that captures this problem is set out in Figure 1.8.

Maintaining a levee in this game costs \$4, and a flood brings damages of \$6. If both landowners maintain the levee, there is no flood, but both incur the \$4 cost of maintaining the levee. If neither maintains the levee, they save money on maintenance but suffer flood damage of \$6. If one maintains the levee and the other does not, the first suffers flood damage of \$6 and incurs maintenance costs of \$4, for a total loss of \$10. The second suffers \$6 from flood damage but incurs no maintenance costs.

Like many two-person, two-by-two games, this also fits within a well-known paradigm with a story attached to it. This game, known as the *stag hunt*, has two players who each have only two strategies: There are two hunters. Each must decide whether to hunt hare or stag. A hunter can catch a hare alone, but will catch a stag only if the other

		Landowner 2	
		Maintain	Don't Maintain
Landowner 1	Maintain	-4, -4	-10, -6
	Don't Maintain	-6, -10	-6, -6

Figure 1.8 Levee coordination game.
Payoffs: Landowner 1, Landowner 2.

		Hunter 2	
		Stag	Hare
Hunter 1	Stag	10, 10	0, 8
	Hare	8, 0	8, 8

Figure 1.9 Stag hunt. Payoffs: Hunter 1, Hunter 2.

hunter is also pursuing it. Sharing in half a stag, however, is better than catching a single hare.

The bimatrix takes the form shown in Figure 1.9. In this game, the strategy that each player adopts is good or bad depending on what the other does. If the first hunter were certain that the second would hunt stag, the first hunter would also decide to hunt stag. If the second hunter were going to hunt hare, however, the first hunter would hunt hare as well. The hunters' interests do not conflict. Each prefers to hunt stag, but only if the other does—and neither can be certain that the other will. Stag hunting will take place only if each is assured that the other will hunt stag.¹⁴

Solving either our second levee game or the stag hunt game introduces a new complication. The two landowners are each best off if both maintain the levee. The two hunters are best off if both hunt stag. The strategy combination in which both maintain the levee—or hunt stag—is a Nash equilibrium of this game. If the other landowner is maintaining the levee, the first landowner's best response is to maintain the levee as well. The first landowner should prefer a payoff of $-\$4$ to a payoff of $-\$6$. The second landowner is in a perfectly symmetrical position. We can engage in the same analysis for the stag hunt.

We cannot, however, be confident that both landowners will maintain the levee or that both hunters will hunt stag because games that have this structure have more than one Nash equilibrium. Consider the strategy combination in which neither landowner maintains the levee or in which both hunters hunt hare. If the other landowner is not maintaining the levee, the first landowner's best response is not to maintain the levee either. If one hunter is hunting hare, the other hunter's best response is to hunt hare as well. Once the other landowner is not going to maintain the levee, the first landowner is going to suffer from a flood whether or not the first landowner maintains the levee. The first landowner would rather suffer a loss of $\$6$ than a loss of $\$10$.

Similarly, once one hunter is not going to hunt stag, the other hunter will receive a payoff of $\$0$ from hunting stag. The hunter would rather take the $\$8$ payoff from hunting hare.

Making matters more complicated, games of this type have a third Nash equilibrium. So far, we have restricted our attention to "pure" strategy equilibria. A *pure strategy equilibrium* is a Nash equilibrium in which the equilibrium strategies are played with certainty, or with probability one. When the Nash equilibrium involves only strategies that are played with certainty, we have a pure strategy equilibrium. The alternative to a pure strategy equilibrium is a *mixed strategy equilibrium*, in which, in equilibrium, each player adopts a strategy that randomizes among a number of pure strategies.

An example of a mixed strategy would arise if one landowner randomly decided to maintain or not maintain the levee with equal probability. This particular mixed strategy, however, is not part of a Nash equilibrium. To see this, we need to discover the other landowner's best response to this strategy. The other landowner would calculate the expected payoffs from each of the pure strategies of maintaining and not maintaining the levee. A landowner always receives a payoff of $-\$6$ when that landowner does not maintain the levee. We now must examine the payoff to this landowner from maintaining the levee when the first pursues this mixed strategy. By maintaining the levee (and incurring a cost of $\$4$ in all cases and flood damage of $\$6$ in half), this landowner receives an expected payoff of $-\$7$.¹⁵ Hence, a landowner's best response to this mixed strategy is not to maintain the levee.

We now know that, if the first landowner would maintain the levee half the time, the second landowner's best response would be to not maintain it at all. This strategy combination can be a Nash equilibrium only if maintaining the levee or not with equal likelihood is a best response to not maintaining it at all. It is not. When the other landowner does not maintain, this mixed strategy brings an expected cost of $\$8$. (A loss of $\$6$ half the time and a loss of $\$10$ the other half.) The first landowner could do better by playing the strategy of not maintaining the levee with certainty and enjoy a payoff of $-\$6$.

This may suggest how we find a combination of mixed strategies that is a Nash equilibrium. A player will be willing to randomize between two pure strategies only if that player is indifferent as to which of the strategies is played. A landowner plays the pure strategy of maintaining the levee if the payoff from it exceeds that from not maintaining the levee; the landowner plays the pure strategy of not main-

taining it if the payoff from this strategy exceeds that from maintaining it. Hence, a landowner is willing to play a mixed strategy only if the payoffs from the two pure strategies are equivalent.

To understand how this works, return to Figure 1.8. We can see that, unless the first landowner is likely to maintain the levee, the second is better off not maintaining it. We want, however, to talk about this more precisely. Let p_1 be the first landowner's probability of maintaining the levee. There is a corresponding probability of not maintaining of $1 - p_1$. Let p_2 do the same for the second landowner. For the first landowner's given mixed strategy $(p_1, 1 - p_1)$, the second landowner will be indifferent between maintaining and not maintaining the levee if the expected payoffs from the two strategies are the same. The second landowner's expected payoff from not maintaining the levee is $-\$6$, independent of the first landowner's strategy. If the second landowner maintains the levee, the second landowner's payoff is $-\$4$ when the first landowner maintains it and $-\$10$ when the first landowner does not.

We determine the second landowner's expected payoff from maintaining the levee for any probability of maintaining or not maintaining it on the part of the first landowner by adding $-4 \times p_1$ and $-10 \times (1 - p_1)$. This amount is greater or less than the second landowner's $-\$6$ payoff from not maintaining, depending on the value of p_1 . The second landowner's expected payoff from maintaining is equal to the payoff from not maintaining only when p_1 , the first landowner's probability of maintaining the levee, has a certain value. This value is $\frac{2}{3}$.¹⁶ The first landowner has to be twice as likely to maintain as not, or the second landowner will not be willing to give up a certain loss of $\$6$ in exchange for the possibility of losing only $\$4$, but risking a possible loss of $\$10$.

When the first landowner adopts the mixed strategy of maintaining with $\frac{2}{3}$ probability, the second landowner is indifferent between maintaining and not maintaining. Given the first landowner's mixed strategy, any strategy the second landowner adopts, including any mixed strategy, is a best response. This game is symmetrical; hence, when the second landowner maintains with the levee $\frac{2}{3}$ probability, the first landowner is indifferent between maintaining and not maintaining or playing any mixed strategy. Any of these is again a best response to this mixed strategy.

When both landowners adopt the mixed strategy of maintaining the levee with $\frac{2}{3}$ probability, each is choosing a best response given the strategy of the other. The other landowner's decision to adopt this

mixed strategy makes any strategy, including this mixed strategy, a best response. Therefore, this combination of mixed strategies is a Nash equilibrium. If one landowner were to adopt anything other than this mixed strategy in response to this mixed strategy on the part of the other, however, we could not have a Nash equilibrium. The first player's strategy would be a best response to the other player's mixed strategy, but the other player's mixed strategy would not be a best response to the strategy of the first.

Thus, the game in Figure 1.8 has three Nash equilibria, two in pure strategies and a third in mixed strategies. When a game has several Nash equilibria, it is not immediately self-evident how we should predict the strategies that the players will adopt. If there are ways to identify the one Nash equilibrium that individuals are likely to play and others that they are not, we may still be able to take advantage of the Nash equilibrium concept even when a game has multiple Nash equilibria. For this reason, much of the work in game theory over the last decade has focused on the question of whether we can isolate the kind of Nash equilibria that rational individuals are likely to play. We rely on these *refinements* of the Nash equilibrium concept when we examine a number of different legal rules in later chapters. At some point, however, we have to confront the limits of game theory. Although parties are likely to choose strategies that form a Nash equilibrium whenever a game has a predictable outcome, not all games have predictable outcomes.

One way of choosing among different Nash equilibria is to examine the different equilibria and ask whether any of them is especially prominent. Such a strategy combination is a *focal point*. It is also called a *Schelling point*, after Thomas Schelling, who examined this idea in an important early work on game theory.¹⁷ The classical illustration of a focal point comes from experiments run several decades ago, in which a group of individuals were given the following thought experiment: You and another person must meet in New York on a particular day. You have no way of communicating with each other beforehand, however. You must therefore choose a time and location and hope against hope that the other person chooses the same time and spot.

This game has an infinite number of Nash equilibria. Every time and every location is a Nash equilibrium. Given that one player is in the middle of some block at some time during the day, the other player is better off being there at that time than waiting at any other place at any other time. Notwithstanding the infinite number of equilibria, however, the majority of those who participated in these experiments

adopted the same strategy: they waited at noon at the information booth at Grand Central Station.

Those who engage in the same thought experiment today might not choose Grand Central Station. Grand Central Station no longer has the prominence it once had. In any given group, however, some focal point might exist. (Indeed, among game theorists familiar with the experiment, Grand Central Station may remain a focal point.) Returning to the game in Figure 1.8, one can argue that maintaining the levee is a focal point both because it is the outcome that brings the greatest benefit to the parties and because neither party is better off in any of the other Nash equilibria.

Experimental work on coordination games, however, suggests that players do not necessarily choose the Nash equilibrium that is in the individual interests of the parties and in their joint interest as well. Consider the game set out in Figure 1.10. There are two pure strategy Nash equilibria in this game—the strategy combination in which Player 1 plays up and Player 2 plays left, and that in which Player 1 plays middle and Player 2 plays center. Experiments suggest that individuals are overwhelmingly likely to choose the strategy combination of up and left, even though it leaves both players worse off than the combination middle and center.¹⁸

Players might adopt the Nash equilibrium that was in their joint interest if there were some possibility of preplay communication between the parties even if they had no way to reach a binding agreement. If two landowners each told the other that they were going to maintain their levees, each one might believe the other, because neither has anything to gain by persuading the other to adopt a strategy that is Nash and then deviating from it. If that other person actually adopts the

		Player 2		
		Left	Center	Right
Player 1	Up	350, 350	350, 250	1000, 0
	Middle	250, 350	550, 550	0, 0
	Down	0, 1000	0, 0	600, 600

Figure 1.10 Coordination game experiment. *Payoffs:* Player 1, Player 2.
Source: This game is taken from Cooper, Dejong, Forsythe, and Ross (1990).

strategy, the landowner's position cannot be improved by doing something else. Introducing these ideas of preplay communication, however, makes sense only if the parties can, in fact, communicate with each other. Moreover, such communication may not be effective if each player prefers a different Nash equilibrium and announces that preference to the other.

The danger that parties might not settle on the outcome that is in their joint interest even when it is a Nash equilibrium may provide a justification for a legal regime that changes the payoffs. For example, a legal rule that requires landowners to maintain levees once they build them would make the strategy of maintaining a levee a strictly dominant one for both landowners.¹⁹ Landowners are made whole in the event that they maintain the levee and others do not, but they pay damages if they fail to maintain it when others do. Hence, maintaining the levee becomes a strictly dominant strategy for both landowners. We show the changes brought by a legal rule that requires landowners to maintain levees in Figure 1.11.

Parties who interact with each other can face other kinds of coordination problems as well. These can also be captured in two-by-two games. One such example is the problem of driving on one side of the road or the other. One type of driver might prefer the left-hand side of the road and the other the right-hand side, but each would rather drive on the less-favored side of the road if everyone else drove on that side as well. We see such problems of coordination in many places. In Chapter 6, we look at this problem in the context of the emergence of standards in an industry with several different firms. Legal rules may affect whether firms adopt a common standard and whether the one they adopt is the one that makes everyone better off.

The two-by-two game that captures this kind of problem is illustrated in Figure 1.12 and is typically called the *battle of the sexes*. It ac-

		Landowner 2	
		Maintain	Don't Maintain
Landowner 1	Maintain	-4, -4	-4, -12
	Don't Maintain	-12, -4	-6, -6

Figure 1.11 Levee coordination game (with legal duty to maintain).
Payoffs: Landowner 1, Landowner 2.

		Spouse 2	
		Fight	Opera
Spouse 1	Fight	8, 4	3, 3
	Opera	2, 2	4, 8

Figure 1.12 Battle of the sexes. Payoffs: Spouse 1, Spouse 2.

quired this name (obviously many years ago) because the story usually told to exemplify it was about a conflict between a couple who wanted to spend the evening together but had different preferences about whether to go to a fight or to an opera. Both would rather be with the other at the event they did not like rather than go alone to the event they preferred, but the first choice of both would be to go with their spouse to their favored event. Neither, however, is able to communicate with the other. Each must guess what the other will do.

It is a Nash equilibrium for both to go to the fight, for both to go to the opera, or for each to randomize between the two. In coordination games such as this, both players want to coordinate their actions, but each player wants a different outcome. To craft a legal rule that brings about cooperation in such cases, one must not only evaluate whether an outcome is efficient, but also weigh the competing interests of the players.

To this point, we have examined games where either there was a single Nash equilibrium in which parties adopt pure strategies, or there were multiple Nash equilibria. There are also games in which the only Nash equilibrium is one in which both players adopt mixed strategies. The prototypical game of this kind is *matching pennies*, illustrated in Figure 1.13.

		Player 2	
		Heads	Tails
Player 1	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

Figure 1.13 Matching pennies. Payoffs: Player 1, Player 2.

Each player chooses heads or tails. The first player wins both pennies if both choose heads or both choose tails. The second player wins if one chooses heads and the other tails. This game is one in which there are no Nash equilibria in pure strategies. Given any combination of pure strategies, one player is always better off changing to the other. The only Nash equilibrium is one involving mixed strategies. Matching pennies is a classic *zero-sum game*, in which any gains to one player come at the expense of the other.

This paradigm most naturally applies to problems of law enforcement. Criminals have a powerful incentive to avoid acting in predictable ways. The same holds true for the police. Even if a law enforcement agency believes that a particular corner in a city is often used for drug deals, permanently placing a police officer on the corner will merely displace the deals elsewhere. If instead a random strategy is played, the police officer has a better chance of catching someone. Playing any strategy with certainty is likely to be foolish in the case of law enforcement—or indeed in any situation in which one party monitors another. The outcome of such games may well be a mixed-strategy equilibrium.

There is one other kind of strategic interaction worth noting. We explore it as a standard problem in negotiation. In negotiations one acquires an edge by being able to commit oneself to a particular strategy. If two players are splitting a dollar, the first player will acquire an enormous advantage if that player can make an offer that gives the other player almost nothing and at the same time make a credible commitment never to change the offer. The other player then has a choice between taking the little the first player offers immediately or rejecting it and never receiving anything. A problem can arise, however, if both players make such a commitment. A familiar problem in labor negotiations illustrates this point.

Parties to labor negotiations are not permitted to adopt a fixed policy of never deviating from the first offer they make.²⁰ To see why this legal rule may be sensible, consider the following situation. Each party to a labor negotiation can either commit in advance to making a single offer (and never being able to deviate from it) or to entering into a series of negotiations in which the parties engage in the ordinary process of exchanging offers and counteroffers.

We set out such a game in Figure 1.14. If the employer and the union each decide to bargain, there is no strike and each receives a payoff of \$5. If one of them commits and the other does not, there is again no strike. Instead of an even division of the \$10 surplus, however, the

	Union	
	Bargain	Commit
Bargain	5, 5	3, 7
Commit	7, 3	2, 2

Figure 1.14 Collective bargaining. Payoffs: Employer, Union.

party who commits enjoys a payoff of \$7 and the other enjoys a payoff of \$3. If both commit themselves, there is a strike and each enjoys a payoff of only \$2.

This is another game in which there are three Nash equilibria. Two are pure strategy equilibria in which one commits oneself to a fixed offer and the other is prepared to bargain. The third is a mixed strategy equilibrium in which, under these numbers, the employer and the union bargain a third of the time and commit themselves two-thirds of the time. Neither of the pure strategy equilibria seems to be a focal point because the combined payoffs are identical and the two players have strong and perfectly opposite preferences. Hence, this game may not have a predictable outcome. It is possible that the union and the employer will adopt either the mixed strategy equilibrium or some other combination of strategies in which both players sometimes commit themselves and a strike arises.

This two-by-two game differs from the battle of the sexes because the pure strategy Nash equilibria involve strategy combinations in which each player does the opposite of what the other player does. The general type of two-by-two game that captures this problem is known as *chicken*. The story behind this game comes again from an earlier time. Two teenagers drive cars headlong at each other. A driver gains stature when that driver drives headlong and the other swerves. Both drivers die, however, if neither swerves. Each player's highest payoff comes when that player drives head on and the other swerves; the second highest comes when that player swerves and the other player swerves as well; and the third highest comes when that player swerves and the other drives. The lowest payoff results when both drive. This game has multiple Nash equilibria, but unlike the stag hunt, the assurance game, or the battle of the sexes, the pure strategy equilibria are ones in which each player adopts a different action (that is, one swerves and the other drives).

We need to be cautious, however, about drawing any firm conclusions from our collective bargaining game. For example, it may not make sense to model the process as one in which the union and the employer decide independently of each other whether to commit themselves to an offer that cannot subsequently be changed. We may need to take account of the dynamic aspects of their interaction. The employer and the union may be able to negotiate with each other before either one makes an offer that cannot subsequently be reversed. To be useful, a model may have to incorporate these negotiations as well.

At this point, it is useful to acknowledge the strengths as well as the limits of using two-by-two normal form games to understand strategic behavior. Paradigmatic games such as chicken, the battle of the sexes, matching pennies, and the prisoner's dilemma can provide useful benchmarks. It may seem unduly limiting to begin with a game in which there are only two players and two strategies, but such simple games can capture the dynamics of many interactions. When we can use these games, the forces that are at work are readily apparent, and it is easy to understand the effects of different legal rules.

Nevertheless, one should use these paradigms with caution. First, one always wants to examine the strategic elements in a given situation and avoid being drawn too quickly to a well-known paradigm such as the prisoner's dilemma. Such paradigms can become Procrustean beds, and, by rushing to one or another too quickly, one may miss important parts of a problem. It is better to capture the problem in normal form and then look for the appropriate paradigm, rather than to shoehorn the problem into one at the outset. Taking advantage of a two-by-two game also requires an understanding of its limits. The prisoner's dilemma, for example, captures the basic feature of collective action and common pool problems, but a model with more elements will reveal details that the prisoner's dilemma does not. If one is interested in the dynamics of a particular collective action problem, the prisoner's dilemma may not be useful.

One must also guard against looking at interactions between players in isolation. A problem that may look like a prisoner's dilemma or some other simple two-by-two game may be part of a much larger game. One cannot assume that, once embedded in a larger game, the play of the smaller game will be the same. Moreover, many interactions between individuals are inherently dynamic. People deal with each other over time and make decisions in response to what the other does. Two-by-two games that model simultaneous decisionmaking are not useful vehicles for analyzing such problems.

Summary

The two-by-two bimatrix, and hence the familiar games that take this form, are well suited to analyzing the way legal rules affect the behavior of players when each must make decisions without knowing what the other will do. We have shown in this chapter how they provide a useful way to understand how different tort regimes operate and the strategic problems underlying specific issues in property law, labor law, and elsewhere.

We must compare two or more games in order to understand the effects of different legal regimes. The fewer the elements of each game, the easier it is to understand how they are different and how players in them might act differently. We should resist adding complications unless we are satisfied that they are necessary, for they tend to obscure the basic forces at work. The test of a model is not whether it is "realistic," but whether it sheds light on the problem at hand.

If a problem does not involve strategic behavior, we should not bring the tools of game theory to bear upon it. Similarly, when we encounter a problem of strategic behavior, we must be sure that the tool that we use is the appropriate one. Most important, we must first ensure that the problem to be analyzed dictates the tools that are used; second, we must use those tools that are best suited to the problem, rather than the ones that are the most accessible. Elegance and power are definite virtues of the two-by-two bimatrix, but these virtues may also lead to its being used in contexts for which it is unsuited, or for which other, more technically difficult tools are more suited. We begin to develop these in the next chapter.

Bibliographic Notes

The assumptions of game theory. As we emphasize in our discussion, game theory shares its basic premises with classical economics. For an elegant exposition of the basic principles of economics, see Becker (1971). Varian (1992) and Deaton and Muellbauer (1980) carefully explore different assumptions about preferences and choice. A good axiomatic discussion of decision theory is Kreps (1988). Kreps (1990b) introduces economics within a game-theoretic context. Elster (1986) provides an eloquent discussion of the limits of rationality and the theory of choice.

With the assumptions of game theory in hand, we can build a structure that does cast light on legal problems. None of this, however, is

to suggest that these assumptions are trivial or unproblematic. For systematic criticism of the basic assumptions of economics, see Thaler (1991). Kahneman, Knetsch, and Thaler (1991) provides a general discussion of anomalies and the way in which the assumptions of economics and experiments appear to diverge. Kahneman, Knetsch, and Thaler (1990) looks specifically at the Coase theorem. For criticisms of the von Neumann-Morgenstern expected utility theory, see Hampton (1992).

Dominant strategies. A general discussion of dominance solvability and elimination of weakly dominated strategies may be found in Kreps (1990b, pp. 417–421). For a standard proof of the existence of a Nash equilibrium in the class of games we have considered in this chapter, see Friedman (1990, pp. 63–77). Several notes of caution should be made about our reliance on dominance arguments in this chapter. There is much debate over the significance of different solution concepts and their refinements. Although dominance arguments are arguably among the least controversial, they are not entirely free from controversy either. Nozick (1985) provides examples of games in which playing a dominant strategy leads to what might be considered an unreasonable result; see also Myerson (1991, pp. 192–195). In a similar vein, Cooper, DeJong, Forsythe, and Ross (1990) suggests that dominated strategies should not be entirely discounted in experimental situations.

Tort law and game theory. The torts literature is enormous, but for our purposes, three works stand out as benchmarks. Brown (1973) is generally credited with initiating the formal analysis of torts with his explicitly game-theoretic approach. Brown used elements of noncooperative game theory to explore different liability rules by invoking Nash solutions (although not under that name) rather than appealing to dominance solvability. Explicit use of game theory, however, has been limited. Landes and Posner (1987) and Shavell (1987) reveal the accumulated understanding of two decades' worth of economic analysis of torts. Neither text makes overt use of game-theoretic concepts. Landes and Posner makes no reference to formal solution concepts, and, although Shavell does use "equilibrium" (and even "Nash equilibrium"—once), he generally avoids such terms. Arguments made in these books, however, do rely implicitly upon the ideas of game theory. For example, Landes and Posner uses a dominance argument—without labeling it as such—in explaining why a defense of contributory

negligence is unnecessary for the negligence rule to achieve the due care outcome. (See Landes and Posner (1987), pp. 75–76.) Landes and Posner (1981) provides a useful introduction to the economic theory of torts.

More recent works on torts have started to return to the explicit use of game theory. Examples include Orr (1991) and Chung (1992). Both argue for a comparative negligence standard on the basis of dominant strategy for both parties. The sharing rule for comparative negligence cases that we introduce in the text—and variations of it—have been mentioned or advocated in a variety of articles. Rea (1987) recognizes that it produces the correct marginal incentives; Orr (1991), Chung (1992), Sobelsohn (1985), and Schwartz (1978) also discuss the rule.

A large literature exists on the merits of different legal regimes. The case for comparative negligence is presented in Cooter and Ulen (1986), Calabresi and Hirschhoff (1972), and Diamond (1974). There is also a large collection of works that examines tort standards from empirical and historical perspectives. See Curran (1992), Hammitt, Carroll, and Relles (1985), or White (1989).

The growth area in the recent torts literature has been the inclusion of friction in the standard accident models, either by relaxing the assumptions of perfect information or by making the legal system costly and/or unpredictable. On the cost of litigation, see Ordover (1978), Ordover (1981), Hylton (1990), and Hylton (1992). For an analysis of how errors by courts affect different legal regimes, see Hylton (1990), Polinsky (1987), Friedman (1992), and Rubinfeld (1987). Finally, Kornhauser and Revesz (1991) discusses the relative merits of negligence and strict liability in the context of an environmental tort.

The two-by-two game. The prisoner's dilemma was first discovered and recounted in its modern form in 1950 by scientists at RAND. As a matter of notation, we adopt the convention of referring to the game as the "prisoner's" dilemma, rather than the plural "prisoners' dilemma," because the emphasis, in our view, should be upon the individual and the choices which that individual faces. The dilemma is one that each prisoner confronts separately.

The history of the prisoner's dilemma and the social and political context in which it evolved are recounted in Poundstone (1992). Using an iterated prisoner's dilemma, Axelrod (1984b) examines issues of cooperation in the absence of legal institutions, while Axelrod (1984a) tackles that and other topics such as evolutionary stability and cooperation in WWI trenches. The general literature on the prisoner's dilemma

is large. Some of the works that are most relevant given our perspective include Wiley (1988), Kreps, Milgrom, Roberts, and Wilson (1982), and Leinfellner (1986). Gibbons (1992) shows how the tragedy of the commons can be modeled as a normal form game with n -players and a continuous strategy space.

The stag hunt game can be traced back to an informal discussion of the problem by Jean-Jacques Rousseau. The strategic effects of Boulwareism are well known; see McMillan (1992). We discuss some additional issues in labor law in Chapters 3 and 7.