An Introduction to Game Theory and the Law

Randal C. Picker
dangelolawlib+randalpicker@gmail.com

Recommended Citation

The Coase Lecture Series

The Coase Lecture series, established in honor of Ronald H. Coase, Clifton R. Musser Professor Emeritus of Economics at the University of Chicago Law School, is intended to provide law students and others with an introduction to important techniques and results in law and economics. The lectures presuppose no background in the subject.

An Introduction to Game Theory and the Law

Randal C. Picker *

I am pleased to have the opportunity to give the third of the three lectures in the Law School's inaugural Coase Lecture Series. I have to confess I am still at a stage in life when I think about how things will look on my resume, and to put down the Coase Lecture, I suspect, adds real luster to it. Nonetheless, we might want to call these lectures something else. My suggestion is "The Bar Stool Lecture Series." That wouldn't sound as distinguished and hence wouldn't do much for my resume, but it more accurately captures what the mission of the talk is. Here is the test for this talk: Given two bar stools and a stack of cocktail napkins, could the ideas in this talk be explained to an intelligent person in a crowded bar with a bank of TVs showing the Bulls and the Blackhawks? If this talk succeeds at that level, I will have accomplished my mission; if it does not, then I will have to consult with the Dean to get a larger budget for field research for my next big talk.

The bar stool test is a test of simplicity, of making an idea accessible to someone who is not a specialist in an area. It is a test that all of Ronald Coase's work that I know passes quite easily. It is the remarkable combination of simplicity and depth, which I guess travel together if you are smart enough, that defines Coase's work. The material that I will discuss today is, I think, fairly simple, though some of it is relatively new. And to give credit where credit is due, some of the work I will describe today is part of a joint effort with Doug Baird and Rob Gertner.

This will be an eight-cocktail-napkin talk: I want to talk about two basic forms for games, the normal form and the extensive form; four ways of predicting the outcomes of games, through dominance arguments, Nash arguments, backwards induction, and forward induction; and two interesting ideas about game theory and the law.

1. Game Theory and Strategic Behavior

As a discipline, law and economics advanced on the strong back of classical microeconomics. Individual decisionmakers maximized utility or profits subject to constraints. These individuals were treated either as pricetakers in competitive settings or pricesetters in monopolies. They were also perfectly informed. A sizable and largely successful academic legal literature grew out of taking first derivatives and ruthlessly applying the discipline of the microeconomist's marginal analysis to a vast array of legal problems.[1]

The last twenty years have seen a major shift in the fundamental methodological tools used by microeconomic theorists. Game theory has emerged to augment the standard, polar approaches of pure competition and monopoly. In a competitive setting, individuals or firms are seen as having no real decisions to make. Prices are given, and individuals and firms are pricetakers. The other production paradigm, monopoly, treats the monopolist as a pricesetter for a given demand curve. In a game-theoretic setting, rational actors need worry about the actions of others—this is the fundamental strategic interdependence that game theory addresses. Other settings lack the back-and-forth quality that characterizes strategic settings.

Game theory sounds like fun—visions of the gamut from Candyland to Monopoly spring to mind. A definition might be useful; as a rough cut, try: *game theory is a set of tools and a language for describing and predicting strategic behavior*. I will discuss in a bit what these tools are and how to apply them, but I want to focus first on the core concept in the definition, strategic behavior. Strategic settings are situations in which one person would like to take into account how a second person will behave in making a decision, and the second person would like to do likewise. Strategic settings typically involve two or more decisionmakers, and the possibility of linking one decision to a second decision, and vice versa.

---

[1] A sample of well-known textbooks and research monographs makes the point. Look at Mitch Polinsky's *Introduction to Law and Economics*, which is now in its second edition; Cooter and Ulen's *Law and Economics* textbook, which came out in 1988; and two torts monographs, Landes and Posner's *The Economic Structure of Tort Law* and Steven Shavell's *Economic Analysis of Accident Law*, both of which were published in 1987.

Consider the airlines industry. Whether Northwest will cut fares may depend on how American and United will respond, and the same, of course, is true for them. Indeed, Northwest recently filed suit against American, claiming that American's introduction of a new pricing schedule was part of a scheme of predatory pricing designed to put Northwest out of business.[2] Oligopolistic industries—airlines, computer microprocessors or operating systems, for example—are natural settings for strategic interactions.

But so is a country road. I have risen for an early-morning walk. I would like to enjoy the view, take in the scenery, and generally ignore the cars going by me. You unfortunately are driving your new Mazda Miata. You want to see how the car handles, to test how it drives through turns and its acceleration. If I knew that you were driving like a maniac, I would want to take that into account in deciding whether to pay much attention to the road. If you knew that I was soaking in the countryside and ignoring the road, you would want to take that into account as well. Our behavioral decisions are intertwined, and we need to take that fact into account when we seek to predict likely outcomes. The legal system should take this into account as well when it establishes antitrust laws for oligopolistic industries or a torts scheme for ordinary accidents.

2. Normal Form Games, Dominant Strategies, and the Hidden Role of Law

A. The Prisoner's Dilemma

The best known bit of game theory is the Prisoner's Dilemma. I will go through the analysis to make clear how much game theory has already crossed over and to establish some terminology, and will then move on to more natural settings. So consider the following "game":

---

[2] See Bridget O'Brian, "Predatory Pricing Issue is Due to be Taken Up in American Air's Trial," *Wall Street Journal,* July 12, 1993, A1.

Prisoner 2

	Silent	Confess
Silent	-2, -2	-10, 0
Confess	0, -10	*-6, -6*

(left labels: **Prisoner 1**, with rows Silent / Confess)

Payoffs: (Prisoner 1, Prisoner 2)

Figure 1: Prisoner's Dilemma

Here is the story that this game is trying to capture. We have two prisoners, or, more generally, two players. They both have committed a serious crime, but the district attorney cannot convict either one of them of 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 for the lesser offense. Each will go to prison for two years. This outcome is represented in the upper left cell.

If, however, one of the prisoners confesses and the other does not, the prisoner who confesses will go free and the other will be tried for the serious crime and given the maximum penalty of ten years in prison. This applies to both prisoners and is represented in the of-diagonal cells. Finally, if both confess, the district attorney will prosecute both for the serious crime, but not ask for the maximum penalty. They will both go to prison for six years. This is the final cell, the lower right cell.

This is a *normal form game*. We have identified the players, our two prisoners; the choices, or strategies, available to them (here, to be silent or confess); and the outcomes associated with the four different strategy pairs. The layout here in the bimatrix is the standard way of representing this normal form game.

Now the solution of the game. Each prisoner wants to minimize time spent behind bars and has no other goal. Moreover, each is indifferent to how much time the other spends in prison. I ignore the possibility of altruism or spite. I also ignore the reputational issues that might arise from being known as a snitch or fear of reprisal for confessing. Finally, the two prisoners have no way of communicating with each other. Each must decide without knowing what the other will do.

This is a game in which each prisoner has a *strictly dominant strategy*. Each is better off confessing regardless of what the other does. One can solve the game by recognizing that each prisoner is likely to reason in the following way: "If the other prisoner has decided to keep silent, I am better off confessing. That way I spend no time behind bars at all, rather than two years. What about the other possibility? If the other prisoner confesses, I am also better off confessing. As bad as serving a six-year sentence might be, serving a ten-year sentence is worse. No matter what the other person does, I am better off confessing. No prison is better than two years and six years is better than ten years." Because both prisoners will likely engage in this reasoning, both are likely to confess.

The outcome—both prisoners confess—seems counterintuitive at first because the prisoners would have been better off if both had remained silent. But this result follows once we assume that we have structured the payoffs correctly. Even if each prisoner erroneously believed that the other was altruistic and would confess, we would still have the same outcome, *given our assumption that the prisoners care only for themselves*. If a prisoner believes (for whatever reason) the other will remain silent, confessing is a way of avoiding prison altogether, the best outcome of all. (Again, if the prisoners care about something in addition to the length of time spent in jail, we have specified the payoffs incorrectly. The premise of the game is that the players are both selfish.) The result is not at all odd once one recognizes that the prisoners lack a means of committing themselves to remaining silent. As long as the two prisoners cannot reach any agreement with each other and as long as their only concern is time spent in prison (and not, let us say, their reputations as finks), their individual interest will lead them to confess, even though they are jointly better off remaining silent.

The power of the Prisoner's Dilemma comes from the incongruence between private benefit and the collective good. Individually rational decisionmaking leads to collective disaster. The Prisoner's Dilemma is thus often seen as one of the main theoretical justifications for government intrusion into private decisionmaking.[3]

---

[3] See, e.g., Cass R. Sunstein, *After the Rights Revolution: Reconceiving the Regulatory State* 49-51 (Harvard Univ. Press, 1990).

Legislation almost appears attractive given the collective disaster that results from individual decisionmaking in the dilemma.

I say "almost" for two reasons. First, the existence of private failure tells us nothing about whether government decisionmaking enjoys a comparative advantage over private decisionmaking. The Churchill line about democracy—"democracy is the worst form of Government except all those other forms that have been tried from time to time"—may apply here as well. We need to know much about the quality of government decisionmaking before we can summarily abandon private decisionmaking. The second reason for being cautious about relying on a simple game-theoretic model such as the Prisoner's Dilemma to justify legal intervention will require more hardware, so I will return to it at the end of this talk.

B. *An Example from the Law of Torts*

Many legal settings can be represented as normal form games and solved by identifying dominant strategies. Consider an accident on a country road involving a motorist and a pedestrian. The likelihood of an accident turns both on how much care the motorist uses in driving and how much care the pedestrian uses in crossing the street. We do not expect the motorist to drive so slowly that there is never any possibility of hitting a pedestrian. Nor do we insist that the pedestrian cross only when there is no car in sight. We want them both to take sensible precautions. If both act reasonably, the chances of an accident as well as the inconvenience to both parties are minimized. If they could bargain with each other, we would expect that each would agree to act in this way. The problem arises, of course, because the two are strangers and they cannot communicate with each other. The motorist and the pedestrian both recognize that the actions of the other influence what will happen, and that basic fact must be recognized if we are to have a sensible analysis of the situation. Game theory is the right tool for this problem.

To jump right in, consider the following "game":

Motorist

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

Payoffs: (Pedestrian,   Motorist)

Figure 2: No Reallocation Law

Here are the stylized facts that this game is seeking to represent. If an accident takes place between the motorist and the pedestrian, the motorist and her car will not be hurt, but the pedestrian will of course suffer harm. Assume that we can represent the harm to the pedestrian as a dollar amount and set that amount at $100. Both the motorist and the pedestrian decide on how much care to take. Assume that they each choose between taking "no care" and "due care." Representing the decision of how much care to take as a binary choice oversimplifies greatly, but it is the natural place to start. Assume that it costs nothing to exercise "no care" but costs $10 to exercise "due care." "Due care" is really a legal term for a physical level of care. Consistent with the convention, "due care" is the level of care that minimizes the total expected costs of the accident. We also need to know how the care choices relate to the probability of an accident occurring. Assume that the accident is certain to happen unless both the motorist and the pedestrian exercise "due care," but that there is still a one in ten chance of an accident occurring even if both exercise "due care."

So far, we have set out the brute facts of nature: the choices available to our players (the motorist and the pedestrian), or what a game theorist would call the strategies of the players, and the physical consequences associated with those strategies (whether an accident takes place and the resulting harm). To fully specify this game, we need one more item, and it is this item that determines the precise structure of the game set forth above. We need to know the legal rule for allocating the harms of an accident. The problem of strategic behavior that the legal analyst faces is a simple problem of simultaneous decisionmaking. The amount of care that the motorist and pedestrian each take would turn on the amount of care each expects the other to take. The amount of care that each takes will turn in some measure on the legal rule that is in place—when and to

what extent the motorist will have to pay damages to the pedestrian in the event of an accident. The first question for the legal analyst concerns the effect of changes in the legal rule on the behavior of the parties. Start with a rule of no liability, or of letting the parties bear their own losses. In this case, if an accident occurs, the motorist is not harmed and the pedestrian is harmed, and the legal rule of no liability does not reallocate any of the harm by having the motorist pay damages.

We can now explain the game in figure 2 and determine how to solve it. In a legal regime of no liability, a regime in which the motorist was never liable for the accident, the motorist would enjoy a payoff of $0 and the pedestrian a payoff of –$100 if neither exercised care. The cost of "no care" is zero, an accident is certain to happen, and the accident harms the pedestrian to the tune of $100. If both exercised care, the motorist would receive a payoff of –$10 and the pedestrian a payoff of –$20. (The pedestrian invests $10 in care and, assuming the pedestrian is risk neutral, 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 motorist receives a payoff of –$10 (the cost of taking care) and the pedestrian a payoff of –$100 (the cost of the accident, which by assumption is certain to arise 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).

What is the likely outcome of this game? In this model, taking care costs the motorist $10 and provides no benefit to the motorist 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 ("no care") that, in the context of this model, is better for the motorist no matter what choice the pedestrian makes. In the language of game theory, this is a *dominant strategy* (really a *strictly* dominant strategy). In corresponding fashion, a strategy which is always worse than another strategy, again regardless of what the other player does, is a *dominated* strategy. In figure 2, "due care" is a dominated strategy for the motorist. We should predict—as we did in analyzing the Prisoner's Dilemma—that a player will embrace a dominant strategy wherever

possible and will not embrace any strategy that is dominated by another.

This idea by itself, however, tells us only what the motorist 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. Note that this game differs from the Prisoner's Dilemma in this regard, as in that game, both players had a dominant strategy. 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 is likely to 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 idea travels under the name of *iterated dominance* and allows us to solve this game. The pedestrian should understand that the motorist has a dominant strategy—play "no care"—and therefore the pedestrian should play "no care" as well. Given that the motorist plays "no care," the payoff to the pedestrian from playing "due care" is –$110 and that from playing "no care" is –$100. (Recall that the accident is certain to happen unless both players play "due care"; once the motorist will not, the pedestrian is better off by not wasting any money on care.) The pedestrian should play "no care" as well. Neither player exercises care. Note that to reach this solution, we proceeded iteratively: we first identified the strategy that the motorist would play using dominance arguments—this is the first iteration—and we next identified the pedestrian's strategy *given* the motorist's strategy as determined in the first stage of the argument—this is the second iteration. This is the logic of iterated dominance.

This extension of the idea that dominated strategies are not played requires us to make a further assumption about the rationality of the players. Players not only act rationally and do the best they can given their preferences, but they also believe that others act ra-

tionally 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. If we accept this solution concept, we can solve the game in figure 2. The pedestrian will not exercise care because the pedestrian will believe that the motorist will not exercise care and, in that event, the pedestrian, under our assumptions, is better off not exercising care either. We cannot, however, make this prediction as confidently as we can predict the motorist's behavior. The solution to the game turns not only on the motorist acting in a way that advances her self-interest, but also on the pedestrian anticipating that the motorist will in fact act in this way.

You might think that these results are specific to the particular numbers set forth in figure 2. The specific result is, though the result that matters is not. In the example in figure 2, the pedestrian chooses to exercise no care when the motorist exercises no care. That outcome is tied directly to the particular probability function for accidents, which makes it worthless for one player to exercise any care if the other player is exercising no care. In general—meaning for different probability functions for accidents—the pedestrian might choose more or less than "due care." The general result is the result that matters: under a rule of no reallocation of losses and where any harm from the accident will be borne by the pedestrian, the motorist lacks an appropriate incentive to take care. Indeed, as shown above—and this is a general result—exercising "no care" is a dominant strategy.

Thus, play under a rule of no liability puts us far from the social goal of having both players exercise due care. This result in itself is hardly startling. To say that the strategy of taking due care is dominated by another strategy of taking less than due care restates in the language of game theory a familiar insight from law and economics, the insight that in a world without tort law, parties tend to take less than due care because they do not fully internalize the costs of their actions.[4] The motorist enjoys all the benefits of driving fast, but

---

[4] See, e.g., W. Landes and R. Posner, *The Economic Structure of Tort Law* 62 (Cambridge, Harvard Univ. Press, 1987).

does not bear all the costs (the danger of injuring the pedestrian). By capturing the problem of the pedestrian and the motorist in the form of a two-by-two game, however, not only are the incentives of the motorist made manifest, but we can readily understand how a change in the legal rules alters the incentives of the motorist and the pedestrian at the same time.

To see this, consider the legal regime of negligence coupled with contributory negligence. This is the regime that Anglo-American law has embraced for a long time. Under this regime, the pedestrian can recover only if the motorist is negligent and if the pedestrian is not. This rule of law leads to the normal form game set out in figure 3:

Motorist

	No Care	Due Care
No Care	-100, 0	-100, -10
Due Care	-10, -100	*-20, -10*

Pedestrian

Payoffs: (Pedestrian, Motorist)

Figure 3: Negligence with contributory negligence

Compare figure 3 with figure 2. The two figures are identical except in the box in which the pedestrian exercises due care and the motorist fails to do so. In this event, the motorist rather than the pedestrian bears the cost of the accident. The pedestrian bears the cost of the accident whenever the pedestrian fails to exercise care and in the case in which both players exercise care. The legal rule does not change the strategies available to the players or the sum of the payoffs in each box. All that changes is the allocation of the cost of the accident between the parties.

In this game, unlike the game in figure 2, the pedestrian has a dominant strategy. The pedestrian is always better of taking care. The motorist no longer has a dominant strategy. Whether the motorist is better off taking care turns 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 will recognize that the pedestrian will play "due care" and then decide to play "due care." Hence, under this legal regime, both pedestrian and motorist will take due care.

A comparison between the two models focuses our attention on the way in which this legal rule works and reveals a counterintuitive insight about the role of law. The only difference between figure 2 and figure 3 is in the box representing the strategy combination in which the pedestrian exercises "due care" and the motorist does not. In figure 2, the payoffs were –$110 and $0 for the pedestrian and the motorist respectively. In figure 3, they are –$10 and –$100. This strategy combination is not the solution to either game: in figure 2, neither player exercises care, while in figure 3 both players exercise care. Yet it is how the negligence/contributory negligence regime reallocates the harm when the pedestrian takes care and the motorist does not—an outcome that is not reached in either game—that completely alters the expected play of the game. Under either liability rule, we would never expect to observe the pedestrian exercising due care and the motorist exercising no care, but it is precisely how the law treats the outcome that will not happen that determines whether the efficient due care–due care outcome occurs. *A legal rule brings about changes through the consequences it attaches to behavior that never happens either when the legal rule is in place or when it is not.*

This model also focuses on a central assumption underlying the Anglo-American rule. To believe that this rule works, we must believe both that the motorist acts rationally and that the motorist believes that the pedestrian acts rationally as well. The motorist will take care in order to avoid liability only if the motorist believes that the pedestrian is similarly motivated to act in a way that tries to avoid bearing the cost of accidents and will take care as well. If the motorist believed that the pedestrian would not take care, the motorist would not take care either. This liability rule turns crucially on the assumption that the motorist believes that the pedestrian will exercise due care.

This explicit game-theoretic approach isolates two features of the law in a useful way. First, it makes clear the rationality assumptions required. We must assume not only that individuals behave rationally, but that individuals expect others to behave rationally as well. Second, 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 alter the behavior of both parties even by changing

outcomes that are never seen under either the new or the old regime.[5]

3. Extensive Form Games and Backwards Induction

Not all games or legal situations can be resolved using dominance arguments. For example, consider the problem of choosing on which side of the road to drive. In this country, we drive on the right hand side, in England, on the left. Think of two players faced with that choice in the absence of a governmental setting:

		Player 2	
		Left	Right
Player 1	Left	4, 3	0, 0
	Right	0, 0	3, 4

Payoffs: (Player 1, Player 2)

Figure 4: Driving Coordination Game (Normal Form)

Player 1 has a slight preference for driving on the left, player 2 for the right, but both care most about making the same decision. (For that reason, this game is often labeled a *coordination game*.) Neither player has a dominant strategy nor is any strategy a dominated strategy. What then is the likely outcome? There is another important approach to solving games, though it will be of only some help here. Consider the following idea: If player 1 knew that player 2 were to play "left," player 1 would play "left" also, and the flipside of that is true as well. The same is true of the combination (right, right): player 1 would play "right" in response to player 2's "right" and player 2 would play "right" in response to player 1's "right." (Left, right) and (right, left) lack this quality: if player 1 chose to play "left" but before committing learned that player 2 was going to play "right," player 1 would abandon "left" and instead play "right." (Left, left) and (right, right) have a stability that the other two outcomes lack. The game

---

[5] For additional analysis of torts issues from the perspective of dominant and dominated strategies, see Daniel Orr, "The Superiority of Comparative Negligence: Another Vote," 20 *J. Legal Stud.* 119 (1991); Tai-Yeong Chung, *Efficiency of Comparative Negligence: A Game Theoretic Analysis*, Mimeo, Department of Economics, Social Science Center, Univ. of Western Ontario, London, Ontario (1992).

theory lingo for this is that both (left, left) and (right, right) are Nash equilibria, Nash coming from the great game theorist John Nash. This game has two pure strategy Nash equilibria. (Pure strategy is more lingo for saying that neither player is playing in a probabilistic fashion.)

In some settings, a game will have a unique Nash equilibrium and it is perhaps understandable that such an equilibrium is considered the most natural outcome to the game. Unfortunately, as in figure 4, many games have multiple Nash equilibria and the games themselves offer no good means for the players to coordinate on those equilibria. As a consequence, if the game in figure 4 were played in an experimental setting, I would expect to see a sizable number of non-Nash (left, right) and (right, left) outcomes. The players would not be happy about this, as this is the worst outcome for them, but the problem with the game is that the players lack any good means for coordinating their choices. Sometimes player 1 would hope that the (left, left) Nash outcome was going to be played while player 2 would be hoping for the (right, right) Nash outcome and that puts the players squarely on (left, right).

Subject to the Churchill caveat, legal intervention might again be appropriate. To get at this and to introduce another form for representing games, suppose, for example, the government gave the first person the right to set the rules of the road. This game could be represented in the following way:
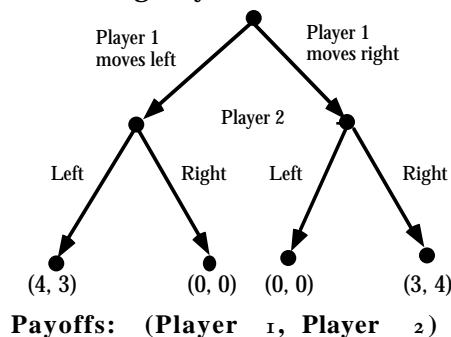

Payoffs: (Player 1, Player 2)

Figure 5: Driving Sequential Game (Extensive Form)

This game represents the players' choices through something akin to a decision tree. This representation is known as the *extensive form* of a game. Figure 5 differs from a decision tree in that it represents de-

cisions by two players, but the basic idea is the same. Pursuant to governmental edict, player 1 chooses first, player 2 second, and each still chooses between "left" and "right." In this game, player 2 observes player 1's choice, which is the essential difference between this game and our prior game in figure 4.

This game can be solved using another solution technique, *backwards induction*. If player 1 moves "left," player 2 will choose between "left," with a payoff of 3, and "right" with a payoff of 0. Player 2 would clearly play "left." If player 1 moves "right," player 2 will choose between "left," with a payoff of 0, and "right" with a payoff of 4, and hence will choose "right." Player 1 thus faces moving "left," and receiving 4 and moving "right" and receiving 3, and hence would move "left." Legislation changing the sequence of moves turns a simultaneous decisionmaking game into a sequential game and establishes a clear outcome. The indeterminacy of the simultaneous game is eliminated. Note that the government allocation of the right to move first has distributional consequences. In this game, player 1 receives 4 and player 2 gets 3. If the right to move first were allocated to player 2, player 2 would get 4 and player 1 would receive 3.

Standard setting, such as establishing the rules of the road, is a conventional use of governmental power. The games in figures 4 and 5 should make clear the possible benefits associated with these activities.

4. Embedded Games: Caveat Legislator

I started the analysis with the Prisoner's Dilemma, as it is easily the best-known game and is most often invoked in defense of legal intervention. Such an analysis often does little more than to suggest that a particular situation has the form of the dilemma and then to claim that intervention would be appropriate. This may be a serious mistake. Whether a Prisoner's Dilemma creates problems depends on the larger structure in which the game exists. Put differently, a small game, such as the Prisoner's Dilemma, may arise in a much larger game. The very existence of the Prisoner's Dilemma in the large game may have beneficial, rather than negative, consequences. A simple example should make this clear. Consider the games set forth in figure 6:

	Player 2	
	Left	Right
Up	2, 2	3, 1.5
Down	1.5, 3	2.5, 2.5

Player 1 (row labels)

Prisoner's Dilemma

	Player 2	
	Left	Right
Up	6, 1.5	0, 0
Down	0, 0	1, 3

Player 1 (row labels)

Coordination Game

Payoffs: (Player 1, Player 2)

Figure 6

Figure 6 illustrates a Prisoner's Dilemma and a coordination game. (I have changed the payoffs from the prior versions of these games, but that is irrelevant here.) In the first game in figure 6, player 1 will play "up," as that is his dominant strategy. (If player 2 were to play "left," player 1 gets a payoff of 2 from "up" and a payoff of 1.5 from "down;" if player 2 were to play "right," player 1 would get a payoff of 3 from "up" and of 2.5 from "down;" "up" is therefore a dominant strategy.) Players 1 and 2 are in symmetric positions in the first game, so player 2 has a dominant strategy of "left." Both players have dominant strategies, resulting in the payoff of (2, 2), which is worse than (2.5, 2.5) from (right, right).

Game 2 in figure 6 is a coordination game, meaning here, as before, that the game has two pure strategy Nash equilibria. The strategy combination (up, left) is one equilibrium: if player 1 were to play "up," player 2 would want to play "left," as that results in a payoff of 1.5 rather than the payoff of 0 obtained by playing "right." And if player 2 were to play "left," player 1 would prefer "up" and 6 to "down" and 0. Thus, (up, left) forms a Nash equilibrium. A similar analysis holds for (down, right). As before in figure 4, game theory offers us little basis for choosing between these two equilibria.

That's where the Prisoner's Dilemma comes in; it will take us two steps to get there. Start with the game set forth in figure 7:

Mixed Form

Payoffs: (Player 1, Player 2)

Figure 7: Embedded Coordination Game

I have embedded the coordination game from figure 6 into a larger game. In this game, player 1 makes an initial move in which player 1 has a chance to decide between taking a certain payoff of 2 or playing a coordination game. If the coordination game is played, player 2 knows that player 1 has elected to forego the certain payoff of 2 and has instead chosen to play the coordination game with player 2. This coordination game is identical to that in figure 6. In that game, players 1 and 2 move simultaneously, and, most importantly, neither can observe the choice of the other.

Now consider how players 1 and 2 should reason. Player 2 decides whether to play "left" or "right" only after observing that player 1 has moved "right." Player 2 does not know whether player 1 moved "up" or "down," but player 2 should not expect player 1 ever to move "down" after having moved "right." Moving "down" is dominated by any strategy in which player 1 moves "left." Player 1's maximum payoff of 1 in the game that follows after playing "right" followed by "down" is dominated by the payoff from playing "left." Hence, if player 1 moves "right," player 1 should follow that move by moving "up." Were player 1 to do otherwise, player 1 would have adopted a dominated strategy. Believing that others would not play dominated strategies, player 2 will play "left" in response to player 1's initial move of "right." Because player 2 believes player 1 will move "up" after moving "right," player 2 ensures a payoff of 1.5 rather than 0 by moving "up." Player 1, recognizing that player 2 will move "left," will play the strategy of moving "right" and "up" and enjoy a payoff of 6, rather than one in which player 1 moves "left" and enjoys a payoff of only 2. Even though this coordination game standing alone does

not have a unique solution, it does have one when it is part of a larger game.[6]

Now take the next step. Replace the solitary payoff of (2, 2) with our Prisoner's Dilemma game from figure 6:


Payoffs: (Player 1, Player 2)

Figure 8: Embedded Prisoner's Dilemma and Coordination Games

In this game, player 1 moves "left" or "right" first, and this move is observed by player 2. If player 1 moves "left," the Prisoner's Dilemma game is played. If player 2 moves "right," the coordination game is played.

How should this game be solved? In the same way we solved the game in figure 7. In the Prisoner's Dilemma, each player has a dominant strategy and a payoff of (2,2) should result. If player 1 were to play "left," he would obtain 2. That payoff is better than any payoff that can result by playing "right" followed by "down." Hence, player 1 would follow "right" only with "up." Player 2 should understand this and play "left" following player 1's initial "right." This would result in a payoff of 6 to player 1. Player 1 should therefore play "right" followed by "up" and player 2 should play "left." This results in payoffs of 6 and 1.5, for a total of 7.5, the maximum available on these particular (and cooked) numbers.

Step back and note what has happened. We started with two games in figure 6, the Prisoner's Dilemma and a coordination game. Taking either of these as freestanding games would suggest that legal intervention might be appropriate. The Prisoner's Dilemma

---

[6] This is an example of *forward induction*. For an introduction, see Drew Fudenberg and Jean Tirole, *Game Theory*, §11.3 (MIT Press, 1991).

plays out inefficiently, and the existence of multiple equilibria in the coordination game means we can have little confidence of an efficient outcome there. Yet bring these two games together in a single larger game, and private decisionmaking leads to an efficient outcome. The very existence of the Prisoner's Dilemma makes it possible to coordinate on a particular Nash equilibrium in the coordination game.

The punch line here is that game structure matters, and often matters a lot. Identification of the game itself is of great importance. Misidentification usually occurs when the small, freestanding game is viewed as *the* game. A modeler who focused on the interaction captured in the Prisoner's Dilemma in figure 8 rather than the entire game would be misled. It is a mistake to suggest that a Prisoner's Dilemma may arise in a particular context and to use that to justify legal intervention. The larger game structure must be understood, as these rather stylized games suggest. The counterintuitive (at least to me) suggestion of figure 8 is that the existence of a scenario in which a Prisoner's Dilemma game might arise actually *helps* the players to achieve the best outcome.

All of this should introduce a level of caution into willy-nilly invocations of the Prisoner's Dilemma as a basis for legislation. More generally, it is critical to understand the context in which a particular game occurs and the extent to which it is embedded in a larger game.[7] Understanding that may make it clear that the very form of the game is up for grabs. For example, the dominant theoretical justification for bankruptcy is that creditors of the failing firm face a collective action problem akin to that in the Prisoner's Dilemma.[8] (This is often called the common pool problem.) One solution is a government-created collective procedure, the modern bankruptcy proceeding. Nonetheless, to accept that the creditors of the firm must play the financial equivalent of the Prisoner's Dilemma is a mistake. Together with the debtor, the creditors have an interest in

---

[7] For a similar point in a political science context, see George Tsebelis, *Nested Games: Rational Choice in Comparative Politics* 7 (Univ. of California Press, 1990).

[8] See Thomas H. Jackson, *The Logic and Limits of Bankruptcy Law* 10 (Harvard Univ. Press, 1986).

taking actions ahead of time to mitigate the possible harms of the dilemma. Security interests can be understood as one important way of completely avoiding the dilemma.[9] Again, the point here is that we must understand the context in which a game would otherwise take place. The game in figure 8 makes this point—I hope—in relatively stark fashion.

5. Conclusion

This talk sketches out some of the basic ideas of game theory. There is a standard language for representing situations, giving rise to the normal form and extensive form games, and ways to discuss solutions, such as dominant strategy solutions, Nash equilibria, backwards induction, and forward induction. I hope that I have suggested a number of ways in which these ideas help us generate counterintuitive insights about legal problems. The central lesson of the torts example is that a legal rule brings about changes through the consequences it attaches to behavior that never happens either when the legal rule is in place or when it is not. I found that surprising. I found even more surprising the notion that having a Prisoner's Dilemma handy might actually help solve collective action problems, rather than create them, and that this should make us cautious in relying on the Prisoner's Dilemma to justify legal intervention. I would have found it difficult to reach either of these points without using game theory, though there very well may be other routes.

I return to where I started. The bar stool test demands simplicity. The work of Ronald Coase, and a lecture worthy of his name, demands both simplicity and depth. I hope that the ideas set forth here at least come close on both scores. Nonetheless, if I have failed, I accept no blame and instead place it squarely on the shoulders of Dean Geoffrey Stone. Any failings must reflect the fact that I spent too little time in bars in preparing this talk and that in turn can be attributed to the measly research budget for it. Notwithstanding this, I am prepared to move forward and undertake more research and we can begin at the reception that immediately follows.

---

[9] See Randal C. Picker, "Security Interests, Misbehavior, and Common Pools," 59 *U. Chi. L. Rev.* 645 (1992).

Readers with comments should address them to:

Randal C. Picker
Professor of Law
The Law School
The University of Chicago
1111 East 60th Street
Chicago, Illinois 60637

CHICAGO WORKING PAPERS IN LAW AND ECONOMICS
(SECOND SERIES)

1. William M. Landes, Copyright Protection of Letters, Diaries and Other Unpublished Works: An Economic Approach (July 1991).

2. Richard A. Epstein, The Path to *The T. J. Hooper*: The Theory and History of Custom in the Law of Tort (August 1991).

3. Cass R. Sunstein, On Property and Constitutionalism (September 1991).

4. Richard A. Posner, Blackmail, Privacy, and Freedom of Contract (February 1992).

5. Randal C. Picker, Security Interests, Misbehavior, and Common Pools (February 1992).

6. Tomas J. Philipson & Richard A. Posner, Optimal Regulation of AIDS (April 1992).

7. Douglas G. Baird, Revisiting Auctions in Chapter 11 (April 1992).

8. William M. Landes, Sequential versus Unitary Trials: An Economic Analysis (July 1992).

9. William M. Landes & Richard A. Posner, The Influence of Economics on Law: A Quantitative Study (August 1992).

10. Alan O. Sykes, The Welfare Economics of Immigration Law: A Theoretical Survey With An Analysis of U.S. Policy (September 1992).

11. Douglas G. Baird, 1992 Katz Lecture: Reconstructing Contracts (November 1992).

12. Gary S. Becker, The Economic Way of Looking at Life (January 1993).

13. J. Mark Ramseyer, Credibly Committing to Efficiency Wages: Cotton Spinning Cartels in Imperial Japan (March 1993).

14. Cass R. Sunstein, Endogenous Preferences, Environmental Law (April 1993).

15. Richard A. Posner, What Do Judges and Justices Maximize? (The Same Thing Everyone Else Does) (April 1993).

16. Lucian Arye Bebchuk and Randal C. Picker, Bankruptcy Rules, Managerial Entrenchment, and Firm-Specific Human Capital (August 1993).

17. J. Mark Ramseyer, Explicit Reasons for Implicit Contracts: The Legal Logic to the Japanese Main Bank System (August 1993).

18. William M. Landes and Richard A. Posner, The Economics of Anticipatory Adjudication (September 1993).

19. Kenneth W. Dam, The Economic Underpinnings of Patent Law (September 1993).

20. Alan O. Sykes, An Introduction to Regression Analysis (October 1993).

21. Richard A. Epstein, The Ubiquity of the Benefit Principle (March 1994).

22. Randal C. Picker, An Introduction to Game Theory and the Law, (June 1994).