An Introduction to Applicable Game Theory

Robert Gibbons

G ame theory is rampant in economics. Having long ago invaded industrial organization, game-theoretic modeling is now commonplace in interna- tional, labor, macro and public finance, and it is gathering steam in de- velopment and economic history. Nor is economics alone: accounting, finance, law, marketing, political science and sociology are beginning similar experiences. Many modelers use game theory because it allows them to think like an economist when price theory does not apply. That is, game-theoretic models allow economists to study the implications of rationality, self-interest and equilibrium, both in market interactions that are modeled as games (such as where small numbers, hidden information, hidden actions or incomplete contracts are present) and in nonmar- ket interactions (such as between a regulator and a firm, a boss and a worker, and so on).

Many applied economists seem to appreciate that game theory can comple- ment price theory in this way, but nonetheless find game theory more an entry barrier than a useful tool. This paper is addressed to such readers. I try to give clear definitions and intuitive examples of the basic kinds of games and the basic solution concepts. Perhaps more importantly, I try to distill the welter of solution concepts and other jargon into a few basic principles that permeate the literature. Thus, I envision this paper as a tutorial for economists who have brushed up against game theory but have not (yet) read a book on the subject.

The theory is presented in four sections, corresponding to whether the game in question is static or dynamic and to whether it has complete or incomplete

■ Robert Gibbons is the Charles Dyson Professor of Economics and Organizations, Johnson Graduate School of Management, Cornell University, Ithaca, New York, and Research Asso- ciate, National Bureau of Economic Research, Cambridge, Massachusetts.

information. ("Complete information" means that there is no private information: the timing, feasible moves and payoffs of the game are all common knowledge.) We begin with static games with complete information; for these games, we focus on Nash equilibrium as the solution concept. We turn next to dynamic games with complete information, for which we use backward induction as the solution concept. We discuss dynamic games with complete information that have multiple Nash equilibria, and we show how backward induction selects a Nash equilibrium that does not rely on noncredible threats. We then return to the context of static games and introduce private information; for these games we extend the concept of Nash equilibrium to allow for private information and call the resulting solution concept Bayesian Nash equilibrium. Finally, we consider signaling games (the simplest dynamic games with private information) and blend the ideas of backward induction and Bayesian Nash equilibrium to define perfect Bayesian equilibrium.

This outline may seem to suggest that game theory invokes a brand new equilibrium concept for each new class of games, but one theme of this paper is that these equilibrium concepts are very closely linked. As we consider progressively richer games, we progressively strengthen the equilibrium concept to rule out implausible equilibria in the richer games that would survive if we applied equilibrium concepts suitable for simpler games. In each case, the stronger equilibrium concept differs from the weaker concept only for the richer games, not for the simpler games.

Space constraints prevent me from presenting anything other than the basic theory. I omit several natural extensions of the theory; I only hint at the terrific breadth of applications in economics; I say nothing about the growing body of field and experimental evidence; and I do not discuss recent applications outside economics, including fascinating efforts to integrate game theory with behavioral and social-structural elements from other social sciences. To conclude the paper, therefore, I offer a brief guide to further reading.[1]

Static Games with Complete Information

We begin with two-player, simultaneous-move games. (Everything we do for two-player games extends easily to three or more players; we consider sequential-move games below.) The timing of such a game is as follows:

1) Player 1 chooses an action a_1 from a set of feasible actions A_1. Simultaneously, player 2 chooses an action a_2 from a set of feasible actions A_2.

2) After the players choose their actions, they receive payoffs: $u_1(a_1, a_2)$ to player 1 and $u_2(a_1, a_2)$ to player 2.

---

[1] Full disclosure requires me to reveal that I wrote one of the books mentioned in this guide to further reading, so readers should discount my objectivity accordingly. By the gracious consent of the publisher, much of the material presented here is drawn from that book.

Figure 1
An Example of Iterated Elimination of Dominated Strategies

		Player 2		
		Left	Middle	Right
Player 1	Up	1, 0	1, 2	0, 1
	Down	0, 3	0, 1	2, 0

A classic example of a static game with complete information is Cournot's (1838) duopoly model. Other examples include Hotelling's (1929) model of candidates' platform choices in an election, Farber's (1980) model of final-offer arbitration and Grossman and Hart's (1980) model of takeover bids.

Rational Play

Rather than ask how one *should* play a given game, we first ask how one *should not* play the game. Consider the game in Figure 1. Player 1 has two actions, {Up, Down}; player 2 has three, {Left, Middle, Right}. For player 2, playing Right is dominated by playing Middle: if player 1 chooses Up, then Right yields 1 for player 2, whereas Middle yields 2; if 1 chooses Down, then Right yields 0 for 2, whereas Middle yields 1. Thus, a rational player 2 will not play Right.[2]

Now take the argument a step further. If player 1 knows that player 2 is rational, then player 1 can eliminate Right from player 2's action space. That is, if player 1 knows that player 2 is rational, then player 1 can play the game *as if* player 2's only moves were Left and Middle. But in this case, Down is dominated by Up for player 1: if 2 plays Left, then Up is better for 1, and likewise if 2 plays Middle. Thus, if player 1 is rational (and player 1 knows that player 2 is rational, so that player 2's only moves are Left and Middle), then player 1 will not play Down.

Finally, take the argument one last step. If player 2 knows that player 1 is rational, *and* player 2 knows that player 1 knows that player 2 is rational, then player 2 can eliminate Down from player 1's action space, leaving Up as player 1's only move. But in this case, Left is dominated by Middle for player 2, leaving (Up, Middle) as the solution to the game.

This argument shows that some games can be solved by (repeatedly) asking how one should not play the game. This process is called *iterated elimination of dominated strategies.* Although it is based on the appealing idea that rational

---

[2] More generally, action a_1' is *dominated* by action a_1'' for player 1 if, for each action player 2 might choose, 1's payoff is higher from playing a_1'' than from playing a_1'. That is, $u_1(a_1', a_2) < u_1(a_1'', a_2)$ for each action a_2 in 2's action set, A_2. A rational player will not play a dominated action.

Figure 2
A Game without Dominated Strategies to be Eliminated

	L	C	R
T	0, 4	4, 0	5, 3
M	4, 0	0, 4	5, 3
B	3, 5	3, 5	6, 6

players do not play dominated strategies, the process has two drawbacks. First, each step requires a further assumption about what the players know about each other's rationality. Second, the process often produces a very imprecise prediction about the play of the game. Consider the game in Figure 2, for example. In this game there are no dominated strategies to be eliminated. Since all the strategies in the game survive iterated elimination of dominated strategies, the process produces no prediction whatsoever about the play of the game. Thus, asking how one should not play a game sometimes is no help in determining how one should play.

   We turn next to Nash equilibrium—a solution concept that produces much tighter predictions in a very broad class of games. We will see that each of the two games above has a unique Nash equilibrium. In any game, the players' strategies in a Nash equilibrium always survive iterated elimination of dominated strategies; in particular, we will see that (Up, Middle) is the unique Nash equilibrium of the game in Figure 1.

Nash Equilibrium

   We have just seen that asking how one should not play a given game can shed some light on how one should play. To introduce Nash equilibrium, we take a similarly indirect approach: instead of asking what the solution of a given game is (that is, what all the players should do), we ask what outcomes cannot be the solution. After eliminating some outcomes, we are left with one or more possible solutions. We then discuss which of these possible solutions, if any, deserves further attention. We also consider the possibility that the game has no compelling solution.

   Suppose game theory offers a unique prediction about the play of a particular game. For this predicted solution to be correct, it is necessary that each player be willing to choose the strategy that the theory predicts that individual will play. Thus, each player's predicted strategy must be that player's best response to the predicted strategies of the other players. Such a collection of predicted strategies could be called "strategically stable" or "self-enforcing," because no single player wants to

Figure 3
The Prisoners' Dilemma

Player 2

		L_2	R_2
Player 1	L_1	1, 1	5, 0
	R_1	0, 5	4, 4

deviate from his or her predicted strategy. We will call such a collection of strategies a *Nash equilibrium*.[3]

To relate this definition to the motivation above, suppose game theory offers the actions (a_1^*, a_2^*) as a solution. Saying that (a_1^*, a_2^*) is *not* a Nash equilibrium is equivalent to saying that either a_1^* is not a best response for player 1 to a_2^*, or a_2^* is not a best response for player 2 to a_1^*, or both. Thus, if the theory offers the strategies (a_1^*, a_2^*) as the solution, but these strategies are not a Nash equilibrium, then at least one player will have an incentive to deviate from the theory's prediction, so the prediction seems unlikely to be true.

To see the definition of Nash equilibrium at work, consider the games in Figures 1 and 2. For five of the six strategy pairs in Figure 1, at least one player would want to deviate if that strategy pair were proposed as the solution to the game. Only (Up, Middle) satisfies the mutual-best-response criterion of Nash equilibrium. Likewise, of the nine strategy pairs in Figure 2, only (*B, R*) is "strategically stable" or "self-enforcing." In Figure 2, it happens that the unique Nash equilibrium is efficient: it yields the highest payoffs in the game for both players. In many games, however, the unique Nash equilibrium is not efficient—consider the Prisoners' Dilemma in Figure 3.[4]

Some games have multiple Nash equilibria, such as the Dating Game (or Battle of the Sexes, in antiquated terminology) shown below in Figure 4. The story behind this game is that Chris and Pat will be having dinner together but are currently on their separate ways home from work. Pat is supposed to buy the wine and Chris the main course, but Pat could buy red or white wine and Chris steak or chicken. Both Chris and Pat prefer red wine with steak and white with chicken, but Chris prefers the former combination to the latter and Pat the reverse; that is, the players prefer

---

[3] Formally, in the two-player, simultaneous-move game described above, the actions (a_1^*, a_2^*) are a Nash equilibrium if a_1^* is a best response for player 1 to a_2^*, and a_2^* is a best response for player 2 to a_1^*. That is, a_1^* must satisfy $u_1(a_1^*, a_2^*) \geq u_1(a_1, a_2^*)$ for every a_1 in A_1, and a_2^* must satisfy $u_2(a_1^*, a_2^*) \geq u_2(a_1^*, a_2)$ for every a_2 in A_2.

[4] Another well-known example in which the unique Nash Equilibrium is not efficient is the Cournot duopoly model.

Figure 4
The Dating Game

		Pat	
		Red	White
Chris	Steak	2, 1	0, 0
	Chicken	0, 0	1, 2

to coordinate but disagree about how to do so.[5] In this game, Red Wine and Steak is a Nash equilibrium, as is White Wine and Chicken, but there is no obvious way to decide between these equilibria. When several Nash equilibria are equally compelling, as in the Dating Game, Nash equilibrium loses much of its appeal as a prediction of play. In such settings, which (if any) Nash equilibrium emerges as a convention may depend on accidents of history (Young, 1996).

Other games, such as Matching Pennies in Figure 5, do not have a pair of strategies satisfying the mutual-best-response definition of Nash equilibrium given above. The distinguishing feature of Matching Pennies is that each player would like to outguess the other. Versions of this game also arise in poker, auditing and other settings. In poker, for example, the analogous question is how often to bluff: if player i is known never to bluff, then i's opponents will fold whenever i bids aggressively, thereby making it worthwhile for i to bluff on occasion; on the other hand, bluffing too often is also a losing strategy. Similarly, in auditing, if a subordinate worked diligently, then the boss prefers not to incur the cost of auditing the subordinate, but if the boss is not going to audit, then the subordinate prefers to shirk, and so on.

In any game in which each player would like to outguess the other, there is no pair of strategies satisfying the definition of Nash equilibrium given above. Instead, the solution to such a game necessarily involves uncertainty about what the players will do. To model this uncertainty, we will refer to the actions in a player's action space (A_i) as *pure strategies*, and we will define a *mixed strategy* to be a probability distribution over some or all of the player's pure strategies. A mixed strategy for player i is sometimes described as player i rolling dice to pick a pure strategy, but later in the paper we will offer a much more plausible interpretation based on player j's uncertainty about the strategy player i will choose. Regardless of how one interprets mixed strategies, once the mutual-best-response definition of Nash equilib-

---

[5] I owe the nonsexist, nonheterosexist player names to Matt Rabin. Allison Beezer noted, however, that no amount of Rabin's relabeling could overcome the game's original name, so she suggested the Dating Game. Larry Samuelson suggested the updated choices available to the players.

There are of course many applications of this game, including political groups attempting to establish a constitution, firms attempting to establish an industry standard, and colleagues deciding which days to work at home.

Figure 5
Matching Pennies

Player 2

		Heads	Tails
	Heads	$-1, 1$	$1, -1$
Player 1	Tails	$1, -1$	$-1, 1$

rium is extended to allow mixed as well as pure strategies, then any game with a finite number of players, each of whom has a finite number of pure strategies, has a Nash equilibrium (possibly involving mixed strategies). See Nash's (1950) classic paper for the proof, based on a fixed-point theorem.

Dynamic Games with Complete Information

We turn next to dynamic games, beginning with two-player, sequential-move games. The timing of such a game is as follows:

1) Player 1 chooses an action a_1 from a set of feasible actions A_1.

2) Player 2 observes 1's choice and then chooses an action a_2 from a set of feasible actions A_2.

3) After the players choose their actions, they receive payoffs: $u_1(a_1, a_2)$ to player 1 and $u_2(a_1, a_2)$ to player 2.
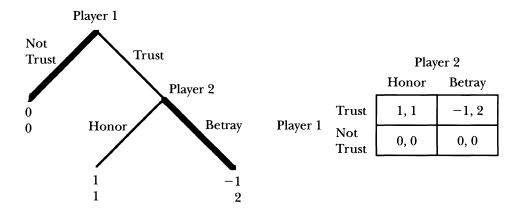
A classic example of a dynamic game with complete information is Stackelberg's (1934) sequential-move version of Cournot duopoly. Other examples include Leontief's (1946) monopoly-union model and Rubinstein's (1982) bargaining model (although the latter may not end after only two moves).

The new solution concept in this section is backward induction. We will see that in many dynamic games there are many Nash equilibria, some of which depend on noncredible threats—defined as threats that the threatener would not want to carry out, but will not have to carry out if the threat is believed. Backward induction identifies a Nash equilibrium that does not rely on such threats.

Backward Induction

Consider the Trust Game in Figure 6, in which player 1 first chooses either to Trust or Not Trust player 2. For simplicity, suppose that if player 1 chooses Not Trust then the game ends—1 terminates the relationship. If player 1 chooses to Trust 2, however, then the game continues, and 2 chooses either to Honor or to Betray 1's trust. If player 1 chooses to end the relationship, then both players'

Figure 6
The Trust Game

payoffs are 0. If 1 chooses to Trust 2, then both players' payoffs are 1 if 2 Honors 1's trust, but player 1 receives −1 and player 2 receives 2 if player 2 Betrays 1's trust. All of this is captured by the game tree on the left-hand side of Figure 6. The game begins with a decision node for player 1 and reaches a decision node for player 2 if 1 chooses Trust. At the end of each branch of the tree, player 1's payoff appears above player 2's. The bold branches in the tree will be explained momentarily.
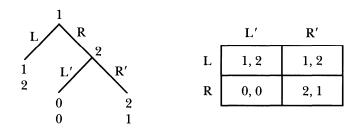
We solve the Trust Game by working backward through the game tree. If player 2 gets to move (that is, if player 1 chooses Trust) then 2 can receive a payoff of 1 by choosing to Honor 1's trust or a payoff of 2 by choosing to Betray 1's trust. Since 2 exceeds 1, player 2 will Betray 1's trust. Knowing this, player 1's initial choice amounts to ending the relationship (and so receiving a payoff of 0) or Trusting player 2 (and so receiving a payoff of −1, after player 2 Betrays 1's trust). Since 0 exceeds −1, player 1 should Not Trust. These arguments are summarized by the bold lines in the game tree.

Thus far, it may appear that simultaneous-move games must be represented in matrix (or "normal") form, as in the previous section, while sequential-move games must be represented using game trees. Similarly, it may appear that we use two different methods to solve these two kinds of games: Nash equilibrium in simultaneous-move games and backward induction in sequential-move games. *These perceptions are not correct.* Either kind of game can be represented in either normal form or a game tree, but for some games it is more convenient to use one than the other. The Trust Game, for example, is represented in normal form on the right-hand side of Figure 6, and using this representation one can verify that the Nash equilibrium is (Not Trust, Betray), just as we found by working backward through the game tree.

The reassurances just offered obscure one subtle point: in some games, there are several Nash equilibria, some of which rely on noncredible threats or promises. Fortunately, the backward-induction solution to a game is always a Nash equilibrium

Figure 7
A Game that Relies on a Noncredible Threat

that does not rely on noncredible threats or promises. As an illustration of a Nash equilibrium that relies on a noncredible threat (but does not satisfy backward induction), consider the game tree and associated normal form in Figure 7. Working backward through this game tree shows that the backward-induction solution is for player 2 to play R' if given the move and for player 1 to play R. But the normal form reveals that there are two Nash equilibria: (R, R') and (L, L'). The second Nash equilibrium exists because player 1's best response to L' by 2 is to end the game by choosing L. But (L, L') relies on the noncredible threat by player 2 to play L' rather than R' if given the move. If player 1 believes 2's threat, then 2 is off the hook because 1 will play L, but 2 would never want to carry out this threat if given the opportunity.

Backward induction can be applied in any finite-horizon game of complete information in which the players move one at a time and all previous moves are common knowledge before the next move is chosen. The method is simple: go to the end of the game and work backward, one move at a time. In dynamic games with simultaneous moves or an infinite horizon, however, we cannot apply this method directly. We turn next to subgame-perfect Nash equilibrium, which extends the spirit of backward induction to such games.

Subgame-Perfect Nash Equilibrium

Subgame-perfect Nash equilibrium is a *refinement* of Nash equilibrium; that is, to be subgame-perfect, the players' strategies must first be a Nash equilibrium and must then fulfill an additional requirement. The point of this additional requirement is, as with backward induction, to rule out Nash equilibria that rely on noncredible threats.

To provide an informal definition of subgame-perfect Nash equilibrium, we return to the motivation for Nash equilibrium—namely, that a unique solution to a game-theoretic problem must satisfy Nash's mutual-best-response requirement. In many dynamic games, the same argument can also be applied to certain pieces of the game, called subgames. A subgame is the piece of an original game that remains to be played beginning at any point at which the complete history of the

play of the game thus far is common knowledge. In the one-shot Trust Game, for example, the history of play is common knowledge after player 1 moves. The piece of the game that then remains is very simple—just one move by player 2.

As a second example of a subgame (and, eventually, of subgame-perfect Nash equilibrium), consider Lazear and Rosen's (1981) model of a tournament. First, the principal chooses two wages—W_H for the winner, W_L for the loser. Second, the two workers observe these wages and then simultaneously choose effort levels. Finally, each worker's output (which equals the worker's effort plus noise) is observed, and the worker with the higher output earns W_H. In this game, the history of play is common knowledge after the principal chooses the wages. The piece of the game that then remains is the effort-choice game between the workers.

Because the workers' effort-choice game has simultaneous moves, we cannot go to the end of the game and work backward *one move at a time*, as with backward induction. (If we go to the end of the game, which worker's move should we analyze first?) Instead, we analyze both workers' moves together. That is, we analyze the entire subgame that remains after the principal sets the wages by solving for the Nash equilibrium in the workers' effort-choice game given arbitrary wages chosen by the principal. Given the workers' equilibrium response to these arbitrary wages, we can then work backward, solving the principal's problem: choose wages that maximize expected profit given the workers' equilibrium response. This process yields the subgame-perfect Nash equilibrium of the tournament game.

There typically are other Nash equilibria of the tournament game that are not subgame-perfect. For example, the principal might pay very high wages because the workers both threaten to shirk if she pays anything less. Solving for the workers' equilibrium response to an arbitrary pair of wages reveals that this threat is not credible. This solution process illustrates Selten's (1965) definition of a subgame-perfect Nash equilibrium: a Nash equilibrium (of the game as whole) is subgame-perfect if the players' strategies constitute a Nash equilibrium in every subgame.[6]

Repeated Games

When people interact over time, threats and promises concerning future behavior may influence current behavior. Repeated games capture this fact of life, and hence have been applied more broadly than any other game-theoretic model (by my armchair count)—not only in virtually every field of economics but also in finance, law, marketing, political science and sociology.

In this section, we analyze the infinitely repeated Trust Game, borrowed from Kreps's (1990a) analysis of corporate culture. All previous outcomes are known before the next period's Trust Game is played. Both players share the interest rate r per period.[7] Consider the following "trigger" strategies:

---

[6] Any finite game has a subgame-perfect Nash equilibrium, possibly involving mixed strategies, because each subgame is itself a finite game and hence has a Nash equilibrium.
[7] The interest rate r can be interpreted as reflecting both the rate of time preference and the probability that the current period will be the last, so that the "infinitely repeated" game ends at a random date.

Player 1: In the first period, play Trust. Thereafter, if all moves in all previous periods have been Trust and Honor, play Trust; otherwise, play Not Trust.

Player 2: If given the move this period, play Honor if all moves in all previous periods have been Trust and Honor; otherwise, play Betray.

Recall that in the one-shot version of the Trust Game, backward induction yields (Not Trust, Betray), with payoffs of (0, 0). Given the trigger strategies stated above for the repeated game, this backward-induction outcome of the stage game will be the "punishment" outcome if cooperation collapses in the repeated game. Under these trigger strategies, the payoffs from "cooperation" are (1, 1), but cooperation creates an incentive for "defection," at least for player 2: if player 1 chooses Trust, player 2's one-period payoff would be maximized by choosing to Betray, producing payoffs of (−1, 2). Thus, player 2 will cooperate if the present value of the payoffs from cooperation (1 in each period) exceeds the present value of the payoffs from detection followed by punishment (2 immediately, but 0 thereafter). The former present value exceeds the latter if the interest rate is sufficiently small (here, $r \le 1$).[8]

What about player 1? Suppose player 2 is playing his strategy given above. Because player 1 moves first, she has no chance to defect, in the sense of cheating while player 2 attempts to cooperate. The only possible deviation for player 1 is to play Not Trust, in which case player 2 does not get the move that period. But 2's strategy then specifies that any future Trusts will be met with Betrayal. Thus, by playing Not Trust, player 1 gets 0 this period and 0 thereafter (because playing Not Trust forever after is 1's best response to 2's anticipated Betrayal of Trust). So if player 2 is playing his strategy given above, then it is optimal for player 1 to play hers. Thus, if the interest rate is sufficiently small, then the trigger strategies stated above are a Nash equilibrium of the repeated game.[9]

The general point is that cooperation is prone to defection—otherwise we should call it something else, such as a happy alignment of the players' self-interests. But in some circumstances, defection can be met with punishment, in which case a potential defector must weigh the present value of continued cooperation against the short-term gain from defection followed by the long-term loss from punishment. If the players are sufficiently patient (that is, the interest rate is sufficiently small), then cooperation can occur in an equilibrium of the repeated game when it cannot in the one-shot game.

---

[8] If player 1 is playing her strategy given above, then it is a best response for player 2 to play his strategy if $\{1 + (1/r)\}1 \ge 2 + (1/r) \cdot 0$, or $r \le 1$. More generally, if a player's payoffs (per period) are C from cooperation, D from defection and P from punishment, then the player has an incentive to cooperate if $\{1 + (1/r)\}C \ge D + (1/r)P$, or $r \le (C - P)/(D - C)$.

[9] In fact, this Nash equilibrium of the repeated game is subgame-perfect.

Static Games with Incomplete Information

We turn next to games with *incomplete information,* also called *Bayesian games.* In a game of complete information, the players' payoff functions are common knowledge, whereas in a game of incomplete information at least one player is uncertain about another player's payoff function. One common example of a static game of incomplete information is a sealed-bid auction: each bidder knows his or her own valuation for the good being sold, but does not know any other bidder's valuation; bids are submitted in sealed envelopes, so the players' moves are effectively simultaneous. Most economically interesting Bayesian games are dynamic, however, because the existence of private information leads naturally to attempts by informed parties to communicate (or mislead) and to attempts by uninformed parties to learn and respond.

We first use the idea of incomplete information to provide a new interpretation for mixed-strategy Nash equilibria in games with *complete* information—an interpretation of player i's mixed strategy in terms of player j's uncertainty about i's action, rather than in terms of actual randomization on i's part. Using this simple model as a template, we then define a static Bayesian game and a Bayesian Nash equilibrium. Reassuringly, we will see that a Bayesian Nash equilibrium is simply a Nash equilibrium in a Bayesian game: the players' strategies must be best responses to each other.

Mixed Strategies Reinterpreted

Recall that in the Dating Game discussed earlier, there are two pure-strategy Nash equilibria: (Steak, Red Wine) and (Chicken, White Wine). There is also a mixed-strategy Nash equilibrium, in which Chris chooses Steak with probability $\frac{2}{3}$ and Chicken with probability $\frac{1}{3}$, and Pat chooses White Wine with probability $\frac{2}{3}$ and Red Wine with probability $\frac{1}{3}$. To verify that these mixed strategies constitute a Nash equilibrium, check that given Pat's strategy, Chris is indifferent between the pure strategies of Steak and Chicken and so also indifferent among all probability distributions over these pure strategies. Thus, the mixed strategy specified for Chris is one of a continuum of best responses to Pat's strategy. The same is true for Pat, so the two mixed strategies are a Nash equilibrium.

Now suppose that, although they have known each other for quite some time, Chris and Pat are not quite sure of each other's payoffs, as shown in Figure 8. Chris's payoff from Steak with Red Wine is now $2 + t_c$, where t_c is privately known by Chris; Pat's payoff from Chicken with White Wine is now $2 + t_p$, where t_p is privately known by Pat; and t_c and t_p are independent draws from a uniform distribution on $[0, x]$. The choice of a uniform distribution is only for convenience, but we do have in mind that the values of t_c and t_p only slightly perturb the payoffs in the original game, so think of x as small. All the other payoffs are the same as in the original complete-information game.

We will construct a pure-strategy Bayesian Nash equilibrium of this incomplete-information version of the Dating Game in which Chris chooses Steak if t_c exceeds

Figure 8
The Dating Game with Incomplete Information

		Pat	
		Red	White
Chris	Steak	$2 + t_c, 1$	$0, 0$
	Chicken	$0, 0$	$1, 2 + t_p$

a critical value, c, and chooses Chicken otherwise, and Pat chooses White Wine if t_p exceeds a critical value, p, and chooses Red Wine otherwise. In such an equilibrium, Chris chooses Steak with probability $(x - c)/x$, and Pat chooses White Wine with probability $(x - p)/x$. (For example, if the critical value c is nearly x, then the probability that t_c will exceed c is almost zero.) We will show that as the incomplete information disappears—that is, as x approaches zero—the players' behavior in this pure-strategy Bayesian Nash equilibrium of the incomplete-information game approaches their behavior in the mixed-strategy Nash equilibrium in the original complete-information game. That is, both $(x - c)/x$ and $(x - p)/x$ approach $\frac{2}{3}$ as x approaches zero.

Suppose that Pat will play the strategy described above for the incomplete-information game. Chris can then compute that Pat chooses White Wine with probability $(x - p)/x$ and Red with probability p/x, so Chris's expected payoffs from choosing Steak and from choosing Chicken are $p(2 + t_c)/x$ and $(x - p)/x$, respectively. Thus, Chris's best response to Pat's strategy has the form described above: choosing Steak has the higher expected payoff if and only if $t_c \geq (x - 3p)/p \equiv c$. Similarly, given Chris's strategy, Pat can compute that Chris chooses Steak with probability $(x - c)/x$ and Chicken with probability c/x, so Pat's expected payoffs from choosing White Wine and from choosing Red Wine are $c(2 + t_p)/x$ and $(x - c)/x$, respectively. Thus, choosing White Wine has the higher expected payoff if and only if $t_p \geq (x - 3c)/c \equiv p$.

We have now shown that Chris's strategy (namely, Steak if and only if $t_c \geq c$) and Pat's strategy (namely, White Wine if and only if $t_p \geq p$) are best responses to each other if and only if $(x - 3p)/p = c$ and $(x - 3c)/c = p$. Solving these two equations for p and c shows that the probability that Chris chooses Steak, namely, $(x - c)/x$, and the probability that Pat chooses White Wine, namely, $(x - p)/x$, are equal. This probability approaches $\frac{2}{3}$ as x approaches zero (by application of l'Hopital's rule). Thus, as the incomplete information disappears, the players' behavior in this pure-strategy Bayesian Nash equilibrium of the incomplete-information game approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information.

Harsanyi (1973) showed that this result is quite general: a mixed-strategy Nash

equilibrium in a game of complete information can (almost always) be interpreted as a pure-strategy Bayesian Nash equilibrium in a closely related game with a little bit of incomplete information. Put more evocatively, the crucial feature of a mixed-strategy Nash equilibrium is not that player j chooses a strategy randomly, but rather that player i is uncertain about player j's choice; this uncertainty can arise either because of randomization or (more plausibly) because of a little incomplete information.

Static Bayesian Games and Bayesian Nash Equilibrium

Recall from the first section that in a two-player, simultaneous-move game of complete information, first the players simultaneously choose actions (player i chooses a_i from the feasible set A_i) and then payoffs $u_i(a_i, a_j)$ are received. To describe a two-player, simultaneous-move game of incomplete information, the first step is to represent the idea that each player knows his or her own payoff function but may be uncertain about the other player's payoff function. Let player i's possible payoff functions be represented by $u_i(a_i, a_j; t_i)$, where t_i is called player i's *type* and belongs to a set of possible types (or *type space*) T_i. Each type t_i corresponds to a different payoff function that player i might have. In an auction, for example, a player's payoff depends not only on all the players' bids (that is, the players' actions a_i and a_j) but also on the player's own valuation for the good being auctioned (that is, the player's type t_i).

Given this definition of a player's type, saying that player i knows his or her own payoff function is equivalent to saying that player i knows his or her type. Likewise, saying that player i may be uncertain about player j's payoff function is equivalent to saying that player i may be uncertain about player j's type t_j. (In an auction, player i may be uncertain about player j's valuation for the good.) We use the probability distribution $p(t_j|t_i)$ to denote player i's belief about player j's type, t_j, given player i's knowledge of her own type, t_i. For notational simplicity we assume (as in most of the literature) that the players' types are independent, in which case $p(t_j \mid t_i)$ does not depend on t_i, so we can write player i's belief as $p(t_j)$.[10]

Joining these new concepts of types and beliefs with the familiar elements of a static game of complete information yields a *static Bayesian game*, as first defined by Harsanyi (1967, 1968a,b). The timing of a two-player static Bayesian game is as follows:

1) Nature draws a type vector $t = (t_1, t_2)$, where t_i is independently drawn from the probability distribution $p(t_i)$ over player i's set of possible types T_i.

2) Nature reveals t_i to player i but not to player j.

---

[10] As an example of correlated types, imagine that two firms are racing to develop a new technology. Each firm's chance of success depends in part on how difficult the technology is to develop, which is not known. Each firm knows only whether it has succeeded, not whether the other has. If firm 1 has succeeded, however, then it is more likely that the technology is easy to develop and so also more likely that firm 2 has succeeded. Thus, firm 1's belief about firm 2's type depends on firm 1's knowledge of its own type.

3) The players simultaneously choose actions, player i choosing a_i from the feasible set A_i.

4) Payoffs $u_i(a_i, a_j; t_i)$ are received by each player.[11]

It may be helpful to check that the Dating Game with incomplete information described above is a simple example of this abstract definition of a static Bayesian game.

We now need to define an equilibrium concept for static Bayesian games. To do so, we must first define the players' strategy spaces in such a game, after which we will define a Bayesian Nash equilibrium to be a pair of strategies such that each player's strategy is a best response to the other player's strategy. That is, given the appropriate definition of a strategy in a static Bayesian game, the appropriate definition of equilibrium (now called *Bayesian Nash equilibrium*) is just the familiar definition from Nash.[12]

A strategy in a static Bayesian game is an action rule, not just an action. More formally, a (pure) strategy for player i specifies a feasible action (a_i) for *each* of player i's possible types (t_i). In the Dating Game with incomplete information, for example, Chris's strategy was a rule specifying Chris's action for each possible value of t_c: Steak if t_c exceeds a critical value, c, and Chicken otherwise. Similarly, in an auction, a bidder's strategy is a rule specifying the player's bid for each possible valuation the bidder might have for the good.

In a static Bayesian game, player 1's strategy is a best response to player 2's if, for each of player 1's types, the action specified by 1's action rule for that type maximizes 1's expected payoff, given 1's belief about 2's type and given 2's action rule. In the Bayesian Nash equilibrium we constructed in the Dating Game, for example, there was no incentive for Chris to change even one action by one type, given Chris's belief about Pat's type and given Pat's action rule (namely, choose White Wine if t_p exceeds a critical value, p, and choose Red Wine otherwise). Likewise, in a Bayesian Nash equilibrium of a two-bidder auction, bidder 1 has no incentive to change even one bid by one valuation-type, given bidder 1's belief about bidder 2's type and given bidder 2's bidding rule.[13]

---

[11] There are games in which one player has private information not only about his or her own payoff function but also about another player's payoff function. As an example, consider an asymmetric-information Cournot model in which costs are common knowledge, but one firm knows the level of demand and the other does not. Since the level of demand affects both players' payoff functions, the informed firm's type enters the uninformed firm's payoff function. To allow for such information structures, the payoff functions in a Bayesian game can be written as $u_i(a_i, a_j; t_i, t_j)$.

[12] Given the close connection between Nash equilibrium and Bayesian Nash equilibrium, it should not be surprising that a Bayesian Nash equilibrium exists in any finite Bayesian game.

[13] It may seem strange to define equilibrium in terms of action rules. In an auction, for example, why can't a bidder simply consider what bid to make given her actual valuation? Why does it matter what bids she would have made given other valuations? To see through this puzzle, note that for bidder 1 to compute an optimal bid, bidder 1 needs a conjecture about bidder 2's entire bidding rule. And to determine whether even one bid from this rule is optimal, bidder 2 would need a conjecture about bidder 1's entire bidding rule. Akin to a rational expectations equilibrium, these conjectured bidding rules must be correct in a Bayesian Nash equilibrium.

Dynamic Games with Incomplete Information

As noted earlier, the existence of private information leads naturally to at-
tempts by informed parties to communicate (or to mislead) and to attempts by
uninformed parties to learn and respond. The simplest model of such attempts is
a signaling game: there are two players—one with private information, the other
without; and there are two stages in the game—a signal sent by the informed party,
followed by a response taken by the uninformed party. In Spence's (1973) classic
model, for example, the informed party is a worker with private information about
his or her productive ability, the uninformed party is a potential employer (or a
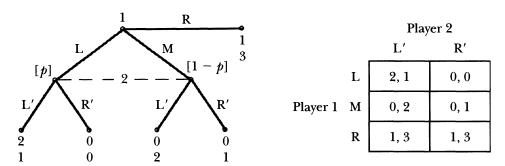market of same), the signal is education, and the response is a wage offer.

Richer dynamic Bayesian games allow for reputations to be developed, main-
tained or milked. In the first such analysis, Kreps, Milgrom, Roberts and Wilson
(1982) showed that a finitely repeated prisoners' dilemma that begins with a little
bit of (the right kind of) private information can have equilibrium cooperation in
all but the last few periods. In contrast, a backward-induction argument shows that
equilibrium cooperation cannot occur in any round of a finitely repeated prisoners'
dilemma under complete information, because knowing that cooperation will break
down in the last round causes it to break down in the next-to-last round, and so on
back to the first round. Signaling games, reputation games and other dynamic
Bayesian games (like bargaining games) have been very widely applied in many
fields of economics and in accounting, finance, law, marketing and political science.
For example, see Benabou and Laroque (1992) on insiders and gurus in financial
markets, Cramton and Tracy (1992) on strikes and Rogoff (1989) on monetary
policy.

Perfect Bayesian Equilibrium

To analyze dynamic Bayesian games, we introduce a fourth equilibrium con-
cept: perfect Bayesian equilibrium. The crucial new feature of perfect Bayesian
equilibrium is due to Kreps and Wilson (1982): beliefs are elevated to the level of
importance of strategies in the definition of equilibrium. That is, the definition of
equilibrium no longer consists of just a strategy for each player but now also in-
cludes a belief for each player whenever the player has the move but is uncertain
about the history of prior play.[14] The advantage of making the players' beliefs an
explicit part of the equilibrium is that, just as we previously insisted that the players
choose credible (that is, subgame-perfect) strategies, we can now also insist that
they hold reasonable beliefs.

---

[14] Kreps and Wilson (1982) formalize this perspective on equilibrium by defining *sequential equilibrium,*
an equilibrium concept that is equivalent to perfect Bayesian equilibrium in many economic applications
but in some cases is slightly stronger. Sequential equilibrium is more complicated to define and to apply
than perfect Bayesian equilibrium, so most authors now use the latter. Kreps and Wilson show that any
finite game (with or without private information) has a sequential equilibrium, so the same can be said
for perfect Bayesian equilibrium.

Figure 9
Why Players' Beliefs are as Important as their Strategies

To illustrate why the players' beliefs are as important as their strategies, consider the example in Figure 9. (This example shows that perfect Bayesian equilibrium refines subgame-perfect Nash equilibrium; we return to dynamic Bayesian games in the next subsection.) First, player 1 chooses among three actions: L, M and R. If player 1 chooses R then the game ends without a move by player 2. If player 1 chooses either L or M then player 2 learns that R was not chosen (but not which of L or M was chosen) and then chooses between two actions, L' and R', after which the game ends. (The dashed line connecting player 2's two decision nodes in the game tree on the left of Figure 9 indicates that if player 2 gets the move, player 2 does not know which node has been reached—that is, whether player 1 has chosen L or M. The probabilities p and $1 - p$ attached to player 2's decision nodes will be explained below.) Payoffs are given in the game tree.

The normal-form representation of this game on the right-hand side of Figure 9 reveals that there are two pure-strategy Nash equilibria: (L, L') and (R, R'). We first ask whether these Nash equilibria are subgame-perfect. Because a subgame is defined to begin when the history of prior play is common knowledge, there are no subgames in the game tree above. (After player 1's decision node at the beginning of the game, there is no point at which the complete history of play is common knowledge: the only other nodes are player 2's, and if these nodes are reached, then player 2 does not know whether the previous play was L or M.) If a game has no subgames, then the requirement of subgame-perfection—namely, that the players' strategies constitute a Nash equilibrium on every subgame—is trivially satisfied. Thus, in any game that has no subgames the definition of subgame-perfect Nash equilibrium is equivalent to the definition of Nash equilibrium, so in this example both (L, L') and (R, R') are subgame-perfect Nash equilibria. Nonetheless, (R, R') clearly depends on a noncredible threat: if player 2 gets the move, then playing L' dominates playing R', so player 1 should not be induced to play R by 2's threat to play R' if given the move.

One way to strengthen the equilibrium concept so as to rule out the sub-game-perfect Nash equilibrium (R, R') is to impose two requirements.

Requirement 1: Whenever a player has the move and is uncertain about the history of prior play, the player must have a *belief* over the set of feasible histories of play.

Requirement 2: Given their beliefs, the players' strategies must be *sequentially rational.* That is, whenever a player has the move, the player's action (and the player's strategy from then on) must be optimal given the player's belief at that point (and the other players' strategies from then on).

In the example above, Requirement 1 implies that if player 2 gets the move, then player 2 must have a belief about whether player 1 has played L or M. This belief is represented by the probabilities p and $1 - p$ attached to the relevant nodes in the game tree. Given player 2's belief, the expected payoff from playing R' is $p \cdot 0 + (1 - p) \cdot 1 = 1 - p$, while the expected payoff from playing L' is $p \cdot 1 + (1 - p) \cdot 2 = 2 - p$. Since $2 - p > 1 - p$ for any value of p, Requirement 2 prevents player 2 from choosing R'. Thus, simply requiring that each player have a belief and act optimally given this belief suffices to eliminate the implausible equilibrium (R, R') in this example.

What about the other subgame-perfect Nash equilibrium, (L, L')? Requirement 1 dictates that player 2 have a belief but does not specify what it should be. In the spirit of rational expectations, however, player 2's belief in this equilibrium should be $p = 1$. We state this idea a bit more formally as

Requirement 3: Where possible, beliefs should be determined by Bayes' rule from the players' equilibrium strategies.

We give other examples of Requirement 3 below.

In simple economic applications, including the signaling games discussed below, Requirements 1 through 3 constitute the definition of *perfect Bayesian equilibrium.* In richer applications, more requirements need to be imposed to eliminate implausible equilibria.[15]
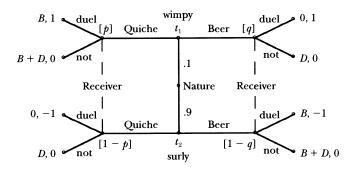
Signaling Games

We now return to dynamic Bayesian games, where we will apply perfect Bayesian equilibrium. For simplicity, we restrict attention to (finite) signaling games, which have the following timing:

1) Nature draws a type t_i for the Sender from a set of feasible types $T = \{t_1, \ldots, t_I\}$ according to a probability distribution $p(t_i)$.

---

[15] To give a sense of the issues not addressed by Requirements 1 through 3, suppose players 2 and 3 have observed the same events, and then both observe a deviation from the equilibrium by player 1. Should players 2 and 3 hold the same belief about earlier unobserved moves by player 1? Fudenberg and Tirole (1991a) give a formal definition of perfect Bayesian equilibrium for a broad class of dynamic Bayesian games and provide conditions under which their perfect Bayesian equilibrium is equivalent to Kreps and Wilson's (1982) sequential equilibrium.

Figure 10
The Beer and Quiche Signaling Game

2) The Sender observes t_i and then chooses a message m_j from a set of feasible messages $M = \{m_1, \ldots, m_J\}$.

3) The Receiver observes m_j (but not t_i) and then chooses an action a_k from a set of feasible actions $A = \{a_1, \ldots, a_K\}$.

4) Payoffs are given by $U_S(t_i, m_j, a_k)$ and $U_R(t_i, m_j, a_k)$.

In Cho and Kreps's (1987) "Beer and Quiche" signaling game, shown in Figure 10, the type, message and action spaces (T, M and A, respectively) all have only two elements. While most game trees start at the top, a signaling game starts in the middle, with a move by Nature that determines the Sender's type: here $t_1 = $ "wimpy" (with probability .1) or $t_2 = $ "surly" (with probability .9).[16] Both Sender types then have the same choice of messages—Quiche or Beer (as alternative breakfasts). The Receiver observes the message but not the type. (As above, the dashed line connecting two of the Receiver's two decision nodes indicates that the Receiver knows that one of the nodes in this "information set" was reached, but does not know which node—that is, the Receiver observes the Sender's breakfast but not his type.) Finally, following each message, the Receiver chooses between two actions—to duel or not to duel with the Sender.

The qualitative features of the payoffs are that the wimpy type would prefer to have quiche for breakfast, the surly type would prefer to have beer, both types would prefer not to duel with the Receiver, and the Receiver would prefer to duel with the wimpy type but not to duel with the surly type. Specifically, the preferred breakfast is worth $B > 0$ for both sender types, avoiding a duel is worth $D > 0$ for both Sender types, and the payoff from a duel with the wimpy (respectively, surly) type is 1 (respectively, -1) for the Receiver; all other payoffs were zero.

The point of a signaling game is that the Sender's message may convey

---

[16] Readers over the age of 35 may recognize that the labels in this game were inspired by *Real Men Don't Eat Quiche*, a highly visible book when this example was conceived.

information to the Receiver. We call the Sender's strategy *separating* if each type sends a different message. In Beer and Quiche, for example, the strategy [Quiche if wimpy, Beer if surly] is a separating strategy for the Sender. At the other extreme, the Sender's strategy is called *pooling* if each type sends the same message. In a model with more than two types there are also *partially pooling* (or *semiseparating*) strategies in which all the types in a given set of types send the same message, but different sets of types send different messages. Perfect Bayesian equilibria involving such strategies for the Sender are also called separating, pooling, and so on.

If $B > D$, then the Sender's strategy [Quiche if wimpy, Beer if surly] and the Receiver's strategy [duel after Quiche, no duel after Beer], together with the beliefs $p = 1$ and $q = 0$ satisfy Requirements 1 through 3 and so are a perfect Bayesian equilibrium of the Beer and Quiche signaling game. Put more evocatively, when $B > D$, having the preferred breakfast is more important than avoiding a duel, so each Sender type chooses its preferred breakfast, thereby signaling its type; signaling this information works against the wimpy type (because it induces the Receiver to duel), but this consideration is outweighed by the importance of getting the preferred breakfast.

We can also ask whether Beer and Quiche has other perfect Bayesian equilibria. The three other pure strategies the Sender could play are [Quiche if wimpy, Quiche if surly], [Beer if wimpy, Quiche if surly] and [Beer if wimpy, Beer if surly]. When $B > D$, the lowest payoff the wimpy Sender-type could receive from playing Quiche (B) exceeds the highest available from playing Beer (D), so the wimpy type will not play Beer, leaving [Quiche if wimpy, Quiche if surly] as the only other strategy the Sender might play. Analogously, the lowest payoff the surly Sender-type could receive from playing Beer (B) exceeds the highest available from playing Quiche (D), so the surly type will not play Quiche. Thus, the separating perfect Bayesian equilibrium derived above is the unique perfect Bayesian equilibrium of the Beer and Quiche signaling game when $B > D$.

What about when $B < D$? Now there is no separating perfect Bayesian equilibrium.[17] But there are two pooling perfect Bayesian equilibria. It is straightforward to show that when $B < D$, the Sender's strategy [Beer if wimpy, Beer if surly] and the Receiver's strategy [duel after Quiche, no duel after Beer], together with the beliefs $p = 1$ and $q = .1$ satisfy Requirements 1 through 3. (In fact, any $p \geq .5$ would work as well.) This pooling equilibrium makes sense (just as the separating equilibrium above made sense when $B > D$): the surly type gets its preferred breakfast and avoids a duel; because $B < D$, the wimpy type now prefers to hide behind the high prior probability of the surly type (.9, which dissuades the Receiver from dueling without further information) rather than have its preferred breakfast.

There is also another pooling equilibrium: when $B < D$, the Sender's strategy [Quiche if wimpy, Quiche if surly] and the Receiver's strategy [no duel after

---

[17] To see why, work out what the Receiver would do if, say, the wimpy Sender-type chose Quiche and the surly choose Beer, and then work out whether these Sender-types would in fact make these choices, given the response just calculated for the Receiver.

Quiche, duel after Beer], together with the beliefs $p = .1$ and $q = 1$ satisfy Requirements 1 through 3. (In fact, any $q \geq .5$ would work as well.) Cho and Kreps argue that the Receiver's belief in this equilibrium is counterintuitive. Their "Intuitive Criterion" refines perfect Bayesian equilibrium by putting additional restrictions on beliefs (beyond Requirement 3) that rule out this pooling equilibrium (but not the previous pooling equilibrium, in which both types choose Beer).

Further Reading

I hope this paper has clearly defined the four major classes of games and their solution concepts, as well as sketched the motivation for and connections among these concepts. This may be enough to allow some applied economists to grapple with game-theoretic work in their own research areas, but I hope to have interested at least a few readers in more than this introduction.

An economist seeking further reading on game theory has the luxury of a great deal of choice—at least eight new books, as well as at least two earlier texts, one now in its second edition. (I apologize for excluding several other books written either for or by noneconomists, as well as any books by and for economists that have escaped my attention.) These ten books are Binmore (1992), Dixit and Nalebuff (1991), Friedman (1990), Fudenberg and Tirole (1991b), Gibbons (1992), Kreps (1990b), McMillan (1992), Myerson (1991), Osborne and Rubinstein (1994) and Rasmussen (1989). These books are all excellent, but I think it fair to say that different readers will find different books appropriate, depending on the reader's background and goals. At the risk of offending my fellow authors, let me hazard some characterizations and suggestions.

Roughly speaking, some books emphasize theory, others economic applications, and still others "the real world." Given a book's emphasis, there is then a question regarding its level. I see Binmore, Friedman, Fudenberg-Tirole, Kreps, Myerson and Osborne-Rubinstein as books that emphasize theory. If I were trying to transform a bright undergraduate into a game theorist (as distinct from an applied modeler), I would start with either or both of Binmore and Kreps, and then proceed to any or all of Friedman, Fudenberg-Tirole, Myerson and Osborne-Rubinstein. In contrast, I see Gibbons and Rasmussen (and, to some extent, McMillan) as books that emphasize economic applications. Each is accessible to a bright undergraduate, but could also provide the initial doctoral training for an applied modeler and perhaps the full doctoral training for an applied economist wishing to consume (rather than construct) applied models. The next step for those who wish to construct such models might be to sample from Fudenberg-Tirole, as the most applications oriented of the advanced theory books. Finally, I see Dixit-Nalebuff and McMillan as books that emphasize the real world (McMillan being more closely tied to applications from the economics literature). These are the texts to use to teach an undergraduate (or an MBA) to think strategically, although for this purpose one should also read the collected works of Thomas Schelling. These books would also

be useful additions to the training of an applied modeler, in the hope that the student would learn to keep his or her eye on the empirical ball.

All of this further reading is for economists seeking a deeper treatment of the theory. I wish I could offer analogous recommendations for those seeking further reading on the many ways game theory has been used to build new theoretical models, both inside and outside economics; this will have to await a future survey. More importantly, I eagerly await the first thorough assessment of how game-theoretic models in economics have fared when confronted with field data of the kind commonly used to assess price-theoretic models. For an important step in a related direction, see Roth and Kagel's (1995) excellent *Handbook of Experimental Economics*, which describes laboratory evidence pertaining to many game-theoretic models.

References

Benabou, Roland, and Guy Laroque, "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*, August 1992, *107,* 921–58.

Binmore, Ken, *Fun and Games: A Text on Game Theory.* Lexington, Mass.: D. C. Heath & Co, 1992.

Cho, In-Koo, and David Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics,* May 1987, *102,* 179–222.

Cournot, A., *Recherches sur les Principes Mathematiques de la Theorie des Richesses.* 1838. English edition, Bacon, N., ed., *Researches into the Mathematical Principles of the Theory of Wealth.* New York: Macmillan, 1897.

Cramton, Peter, and Joseph Tracy, "Strikes and Holdouts in Wage Bargaining: Theory and Data," *American Economic Review,* March 1992, *82,* 100–21.

Dixit, Avinash, and Barry Nalebuff, *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life.* New York: Norton, 1991.

Farber, Henry, "An Analysis of Final-Offer Arbitration," *Journal of Conflict Resolution,* December 1980, *35,* 683–705.

Friedman, James, *Game Theory with Applications to Economics.* 2nd ed., Oxford: Oxford University Press, 1990.

Fudenberg, Drew, and Jean Tirole, "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory,* April 1991a, *53,* 236–60.

Fudenberg, Drew, and Jean Tirole, *Game Theory.* Cambridge, Mass.: Massachusetts Institute of Technology Press, 1991b.

Gibbons, Robert, *Game Theory for Applied Economists.* Princeton, N.J.: Princeton University Press, 1992.

Grossman, Sanford, and Oliver Hart, "Takeover Bids, the Free-Rider Problem, and the Theory of the Corporation," *Bell Journal of Economics,* Spring 1980, *11,* 42–64.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': I. The Basic Model," *Management Science,* November 1967, *14,* 159–82.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': II. Bayesian Equilibrium Points," *Management Science,* January 1968a, *14,* 320–34.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': III. The Basic Probability Distribution of the Game," *Management Science,* March 1968b, *14,* 486–502.

Harsanyi, John, "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory,* 1973, *2*:1, 1–23.

Hotelling, Harold, "Stability in Competition," *Economic Journal,* March 1929, *39,* 41–57.

Kreps, David, "Corporate Culture and Economic Theory." In Alt, J., and K. Shepsle, eds., *Perspectives on Positive Political Economy.* Cambridge: Cambridge University Press, 1990a, pp. 90–143.

Kreps, David, *Game Theory and Economic Modeling.* Oxford: Oxford University Press, 1990b.

Kreps, David, and Robert Wilson, "Sequential Equilibrium," *Econometrica,* July 1982, *50,* 863–94.

Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson, "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory,* August 1982, *27,* 245–52.

Lazear, Edward, and Sherwin Rosen, "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy,* October 1981, *89,* 841–64.

Leontief, Wessily, "The Pure Theory of the Guaranteed Annual Wage Contract," *Journal of Political Economy,* February 1946, *54,* 76–9.

McMillan, John, *Games, Strategies, and Managers.* Oxford: Oxford University Press, 1992.

Myerson, Roger, *Game Theory: Analysis of Conflict.* Cambridge, Mass.: Harvard University Press, 1991.

Nash, John, "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences,* 1950, *36,* 48–9.

Osborne, Martin, and Ariel Rubinstein, *A Course in Game Theory.* Cambridge, Mass.: Massachusetts Institute of Technology Press, 1994.

Rasmussen, Eric, *Games and Information: An Introduction to Game Theory.* New York: Basil Blackwell, 1989.

Rogoff, Kenneth, "Reputation, Coordination, and Monetary Policy." In Barro, R., ed., *Modern Business Cycle Theory.* Cambridge, Mass.: Harvard University Press, 1989, pp. 236–64.

Roth, Alvin, and John Kagel, *Handbook of Experimental Economics.* Princeton, N.J.: Princeton University Press, 1995.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica,* January 1982, *50,* 97–109.

Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit," *Zeitschrift für Gesamte Staatswissenschaft,* 1965, *121,* 301–24.

Spence, A. Michael, "Job Market Signaling," *Quarterly Journal of Economics,* August 1973, *87,* 355–74.

von Stackelberg, H., *Marktform und Gleichgewicht.* Vienna: Julius Springer, 1934.

Young, H. Peyton, "The Economics of Convention," *Journal of Economic Perspectives,* Spring 1996, *10,* 105–22.