

The Game-Theoretic Approach to Modeling Strategic Behavior

Stephen Martin
FEE/F&O
University of Amsterdam
Roetersstraat 11
1018 WB Amsterdam
The Netherlands
smart@fee.uva.nl

June 5, 2001

Abstract

With minor changes (principally, the discussion of trembling-hand equilibrium), this document is the appendix to the first edition of *Advanced Industrial Economics*
GTA2001.tex

Contents

1	Introduction	3
2	Modeling interactive behavior	3
3	Formal representations of a game	5
3.1	The extensive form	5
3.1.1	Information concepts	8
3.2	The strategic form	13
4	Solution concepts	17
4.1	Noncooperative equilibrium – pure strategies	17
4.2	Noncooperative equilibrium – mixed strategies	20
4.3	Subgame-perfect equilibrium	21
4.4	Trembling-hand perfect equilibrium	24
4.4.1	Noncooperative equilibria for the unperturbed game . .	25
4.4.2	The perturbed game	27
4.5	Sequential equilibrium	31
4.5.1	The consistency of beliefs in sequential equilibrium . .	32
4.5.2	Sequential equilibrium for the game of Figure 10	36
4.5.3	Sequential equilibrium for the game of Figure 9	36
5	Supergames	37
6	Experimental games	41
7	Conclusion	43
8	Problems	45
9	References	48

The first player to make a bad move loses the game.

Japanese proverb.

1 Introduction

Game theory suffers, literally, from a bad name. Game theory, at least as applied to industrial economics, is no more the theory of games than imaginary numbers are imaginary. Any electrical engineer will testify that the analysis of imaginary numbers has important real-world implications, and any industrial economist will testify that game theory has important real-world implications. It might be more appropriate to label the discipline that is called game theory “interactive decision theory” (Aumann, 1987) or (at least as regards applications to industrial economics) “theory of strategic rivalry.” But names of disciplines, like names of people, are inherited. It is better to try to understand the discipline rather than to try to rename it. Here we consider the elements of game theory that have been most useful for industrial economics.¹

2 Modeling interactive behavior

Suppose that the Marx brothers meet for a poker game. What would we need to do to formally analyze the game?²

To begin with, we would have to identify the players – Groucho, Chico, Harpo, and Karl. We would have to specify the rules of the game: aces are high, deuces are low, and so on. Sometimes particular players can select the rules of the game (if the dealer wishes, he can make one-eyed jacks wild). The circumstances under which this would be possible, and the universe of rules from which the dealer could select, would have to be indicated. The rules of the game include the procedure for selecting the dealer and the order of play – first the player to the immediate left of the dealer, next the player to the immediate left of the second player, etc. We would have to specify the choices or options available to each player – to pass, to draw cards, to ante,

¹See Friedman (1977; 1983b, ch. 9; 1986); Aumann (1985a); Shubik (1981); Tirole (1983, 1988); Fisher (1989) versus Shapiro (1989); Fudenberg and Tirole (1989); Rasmusen (1989); Kreps (1990a, b); Sutton (1990); and Roth (1991). See also the following entries in Eatwell et al. (1987): cooperative equilibrium; cooperative games; cores; game theory; games with incomplete information; Nash equilibrium; large economies; non-cooperative games; oligopoly and game theory; prisoner’s dilemma; repeated games; Shapley value.

²For a capsule summary of the rules of poker, see Binmore (1990, p. 96).

to call, to fold, and so on. We would have to specify the way the winnings or losses of each player – the payoffs – depend on the strategies of all players.

We would have to specify what each player knows. Each player knows who the other players are, and each player knows the rules. Each player knows how the dealer is determined, and each player knows his own cards. At any point in a hand of poker, each player knows what all preceding players have done in that hand. We would have to specify the payoff or reward to each player – what he or she would get, as a function of his or her strategy and of the strategies of all other players.

This is quite a lot, but we are not done yet. What we have described is perhaps sufficient to characterize a single hand of poker. But an evening of poker will ordinarily include many hands. Behavior in an evening of (say) 120 hands of poker may be different from behavior for a single hand of poker. In a game that consists of many single hands of poker, a player might bluff in early hands to establish a reputation that could be exploited in later hands.

For such an expanded game, we would have to include in the rules of the game those rules that determine how many hands are played. The description of information also changes for the expanded game. In an evening of poker, the information available to each player is greater than for a single hand – in an evening of poker, each player knows what all players have done in all preceding hands, as well as what all players have done in the current hand. This expanded information set creates the possibility for bluffing mentioned above.

We may need to consider a still more expanded game. The Marx brothers may meet once a week – every Saturday night – to play poker. In this case, the game is really a large series of hands of poker, spread over several weeks. We would want to expand the list of players to include potential players – Zeppo, for example – who might show up on any particular Saturday night. Once we admit the possibilities of potential players and successive evenings or rounds of play, we have to consider the possibility that the game might go on forever, even though the identities of the players might change.³

Poker is too difficult a game to analyze for present purposes.⁴ But it ought to be recognized that poker, as described above, has much in common – at least in a formal sense – with oligopoly. In oligopoly, there are a limited number of actual and potential players – firms. They know the rules. Each firm ordinarily knows its own “hand” – its costs, its market – and has only guesses or inferences about the “hands” – costs and markets – of other firms.

³Recall the musical *Guys and Dolls*, which revolves around “the longest-running floating crap game in the City of New York.”

⁴For analyses of versions of poker, see von Neumann and Morgenstern (1944); Nash (1951); Luce and Raiffa (1957, p. 456).

Each firm must make decisions about current actions – the current “hand” – but knows that the game – the search for profit – will go on. And if the players are firms, which have lives of indeterminate length, it is reasonable to consider games that go on forever.

3 Formal representations of a game

The elements of the poker game described above include

1. a list of players;
2. a specification of actions available to each player at each point in the game;
3. the relation between actions and payoffs;
4. the information structure of the game – what each player knows, when he knows it, and what each player knows about what each other player knows.
5. other rules of the game – concerning shuffling, cutting the deck, determining the dealer, and so on.

Depending on the purpose for which a game is being analyzed, some of these elements will be more important than others. Corresponding to these differences in emphasis, game theorists have developed three ways of describing a game – the extensive form, the strategic form, and the coalitional form.

Each highlights different aspects of a game, and each may be most useful depending on the problem at hand. We confine our attention here to the extensive and strategic forms, which are most widely used in industrial economics.⁵

3.1 The extensive form⁶

The *extensive form* provides the most comprehensive description of a game. It provides a complete specification of the interactive process of a game (Shubik, 1985, p. 91).

⁵Shubik (1981, 1987); Friedman (1986, pp. 14–8).

⁶For discussions of the extensive form, see von Neumann (1928), Kuhn (1953), Luce and Raiffa (1957, pp. 44–9), Friedman (1977, pp. 146–8; 1986, p. 11); Shubik (1981, 287–9; 1982, ch. 3). For formal definitions in a research context, see Selten (1974, pp. 25–7) or Kreps and Wilson (1982a, pp. 865–9).

As an example, consider the extensive form of a game that is rich in lessons for industrial economics – the Prisoners’ Dilemma.⁷ Scotland Yard has called in Professor Moriarty and Colonel Moran for questioning in connection with a number of heinous crimes. As the investigation has proceeded under the skilled supervision of Inspector Lestrade, however, evidence is a little weak. The two suspects are questioned separately, and are not permitted to communicate.

In this game, there are two players: Professor Moriarty and Colonel Moran. Each player has one turn, and the turns are simultaneous. Simultaneous play need not mean that plays are made at the same instant of time, although it may mean this for some games (see Problem A.1 and the description of scissors–paper–stone). The essential element of simultaneous play is that each player moves without knowing the move of the other.

Each player has two possible strategies – to maintain an air of outraged innocence (“I demand to see my solicitor at once!”) or to cravenly confess (“e’ardly screamed when I ’it ’im with my ’atchet, guv’nor”).

If both players maintain innocence, then Scotland Yard has only enough evidence to put them on probation for a year. This will interfere somewhat with their professional activities but is much preferred to going to prison.

If only one player confesses, he will be placed on probation for six months, while the other will be sent to prison for five years. If both players confess, they will both be sent to prison for two years.

Most of the essential elements of the Prisoners’ Dilemma game can be illustrated with the game tree shown in Figure 1.

The game tree begins at decision node D_1 , where it is the move of Professor Moriarty. Two branches, each representing a possible action, lead from D_1 to Colonel Moran’s *information set* – the cartouche containing decision nodes D_2 and D_3 . The branch on the left, labeled I , shows the course of the game if Professor Moriarty chooses to maintain his innocence. The branch on the right, labeled C , shows the course of the game if Professor Moriarty confesses.

When the game reaches Colonel Moran’s decision set, he knows that it is his turn, but he does not know Professor Moriarty’s move – he does not know whether he is at D_2 or D_3 . He only knows he is in his information set (which, as noted by Shubik (1981, p. 288), it would be more appropriate to call a “lack-of-information set”).

As an analytical device, the information set permits us to model simul-

⁷Axelrod (1984, p. 216) attributes the Prisoner’s Dilemma game to Merrill Flood and Melvin Dresher. For discussions, see Luce and Raiffa (1957), Axelrod (1984), Anthony (1986, pp. 266–9), Rapoport (1987), and Rasmusen (1989, p. 38).

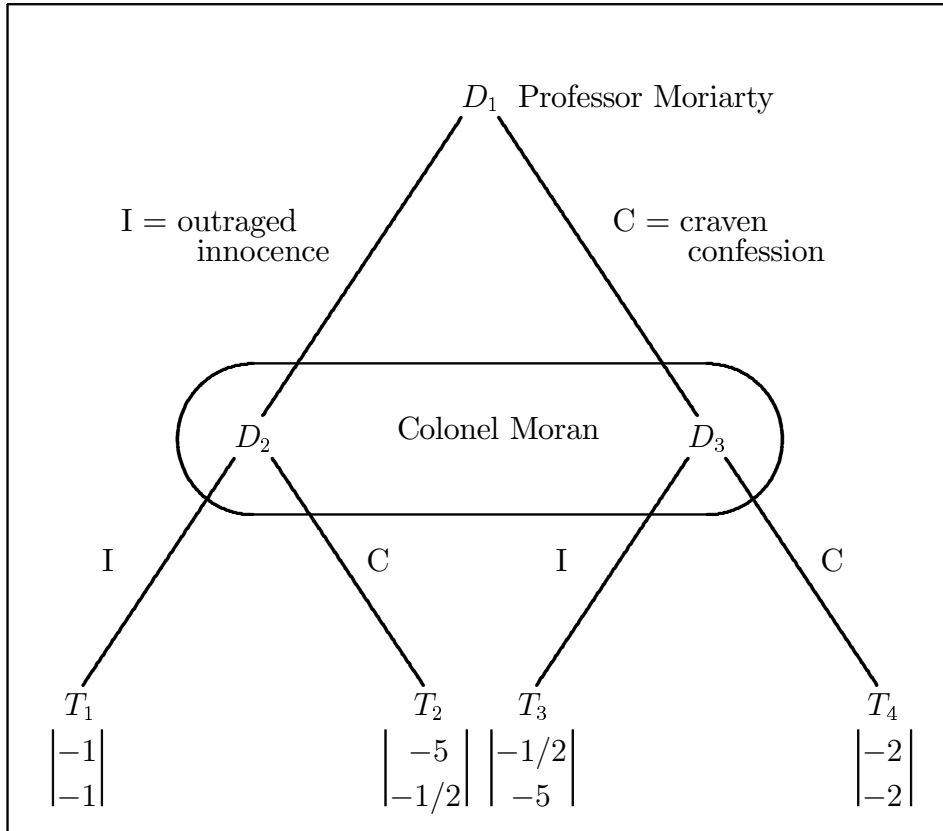


Figure 1: Game tree for the Prisoners' Dilemma: D_1 , D_2 , and D_3 , decision nodes; T_1 , T_2 , T_3 , and T_4 , terminal nodes; D_1 (Professor Moriarty), D_2 and D_3 (Colonel Moran), information sets; strategies, confess (C) or maintain innocence (I); payoff vectors, first element shows probation or prison time for Professor Moriarty, second element probation or prison time for Colonel Moran.

taneous moves. We could relabel Figure 1, placing Colonel Moran at D_1 and Professor Moriarty at the information set containing D_2 and D_3 , and the game tree would represent the same extensive game. The essential element is that neither player knows the other's move when he makes his own decision.

This example illustrates a more general point: game trees are not unique. It is often the case that different game trees can be drawn to represent the same game in extensive form.

There are four branches from Colonel Moran's information set – two from D_2 and two from D_3 . Since a player does not know exactly where he is in his information set, his choices must be specific to the information set rather than to decision nodes within the information set. Thus the same number of branches, indicating the same choices, must descend from each decision node in an information set. The branches going to the left indicate the course of the game if Colonel Moran maintains an air of outraged innocence, while the branches going to the right show the course of the game if Colonel Moran cravenly confesses.

The four branches coming from Colonel Moran's information set lead to four terminal nodes – T_1 , T_2 , T_3 , and T_4 . The game must end at one of these nodes. The column payoff vector below each node shows the payoffs to the two players if the game ends at that node. Players are modeled as seeking to maximize their payoffs. Payoffs are thus expressed as negative numbers, since prison or probation time is a bad, not a good. The upper element in the vector is Professor Moriarty's payoff, and the lower element is Colonel Moran's payoff. For example, at T_2 Professor Moriarty has maintained his innocence but Colonel Moran has confessed. If the game ends at T_2 , Professor Moriarty goes to jail for five years, while Colonel Moran receives six months (1/2 year) probation.

3.1.1 Information concepts

The extensive form includes a complete description of the information characteristics of the game. Some of the information characteristics can be illustrated on the game tree; others cannot.

Perfect recall If players always remember their own moves and other information of which they were once aware, then the game is one of *perfect recall*. Otherwise, it is a game of *imperfect recall*.⁸

⁸Shubik (1982, pp. 37–8); Friedman (1986, pp. 9–10).




Figure 2: Game tree for one play of Odds & Evens.

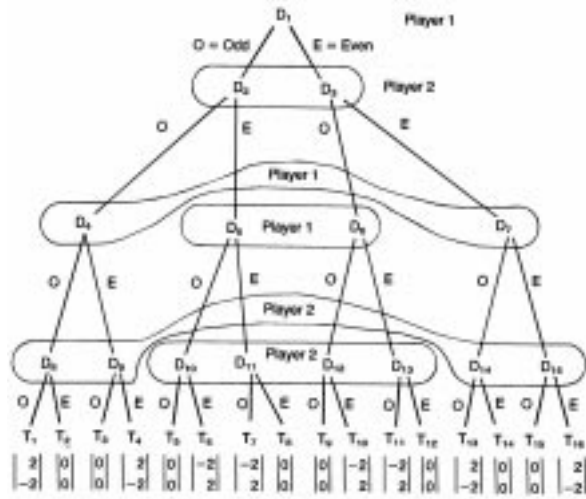


Figure 3: Game tree for two plays of Odds & Evens with imperfect recall

Figure 2 shows a game tree for one round of Odds & Evens. Two players simultaneously hold out any number of fingers on their right hands. The possible strategies are to hold out an odd number or an even number of fingers. If the total number of fingers is even, player 1 wins and player 2 pays player 1 a euro. If the total number of fingers is odd, player 2 wins and player 1 pays player 2 a euro. As with the information set for Colonel Moran in Figure A.1, the information set for player 2 in Figure 2 indicates that when player 2 makes his move he does not know the move of player 1.

Figure 3 shows a game tree for a repeated game that consists of two rounds of Odds & Evens with imperfect recall.⁹ The first two layers of Figure 3 are essentially the same as Figure 2. The terminal nodes T_1 , T_2 , T_3 , and T_4 in Figure 2 have been relabeled as decision nodes D_4 , D_5 , D_6 , and D_7 in Figure 3. At any of these decision nodes, it is player 1's move.

Player 1's situation, halfway through the game, is the same at decision nodes D_4 and D_7 – he is ahead €1. In Figure 2, these decision nodes are in the same information set – player 1 does not distinguish between D_4 and D_7 when he makes his choice in the final round. He does not remember the sequence of moves that put him ahead in the game – he only knows that he is ahead. Similarly, D_5 and D_6 are in a distinct information set – one in which player 1 is behind €1.

⁹Strategies here are ordered triples. (O, E, O) means player 1 plays odd on his first move, even on his second move if he is ahead €1, and odd on his second move if he is behind €1.

Since the players move simultaneously in the second round of the game, player 2 could have up to four information sets. One information set would contain D_8 and D_9 , one would contain D_{10} and D_{12} , one would contain D_{12} and D_{13} , and one would contain D_{14} and D_{15} . For such a game, player 2 would have perfect recall – he would know whether he was ahead or behind, and would know whether his previous move was odd or even. For the game tree in Figure 3, however, player 2 has only two information sets.

One – containing decision nodes D_8 , D_9 , D_{14} , and D_{15} – shows nodes where player 2 is behind at the start of the second round. This corresponds to the fact that player 1 is ahead at nodes D_4 and D_7 . Player 2's other information set – containing decision nodes D_{10} , D_{11} , D_{12} , and D_{13} – shows nodes at which player 2 is ahead at the start of the second round. For the game shown in Figure 3, player 2 does not remember the precise sequence of moves when he takes his second turn – he only knows whether he is ahead or behind at that point in the game.

Perfect Information If every information set contains exactly one decision node, then information is said to be *perfect*. If information is perfect, a player knows exactly where he is in the game tree whenever it is his turn. If some decision sets contain more than one decision node, as in Figure 2, then information is *imperfect*.

Imperfect information commonly arises in oligopoly. For example, Stigler (1964) identifies a particular kind of as a problem for collusion under oligopoly when colluding firms are unable to monitor each others' outputs. If a colluding firm notices a decline in its sales, that decline may reflect no more than the continuous random movement of customers from one supplier to another. Alternatively, the decline may result from secret price cuts offered by some other firm that is party to the collusive agreement. If we were to describe this in terms of a game tree, we would picture the typical firm as moving at an information set containing two decision nodes. At one decision node, all firms have maintained the collusive agreement. At the other decision node, at least one firm has broken the agreement. Information is imperfect because at its information set the firm does not know where it is in the information set – it does not know whether rivals are cheating or not.

Complete information Another essential distinction is that between *complete* and *incomplete* information. If information is complete, each player knows everything there is to know about the structure of the game – not only his own choices at his own information sets, but also the choices available to other players at their information sets. In a game of complete informa-

tion, the player knows the game tree and everything else that is part of the extensive form.

Harsanyi (1967–8) shows that it is possible to replace a game of incomplete information with a corresponding game of complete but imperfect information. For example, consider Stigler’s model of collusion in oligopoly. Players move simultaneously. Each player decides to adhere to the collusive agreement or to defect. Each player knows that demand is high or low, but does not know which. This is a game of incomplete information, since players do not know the state of demand when they make their moves.

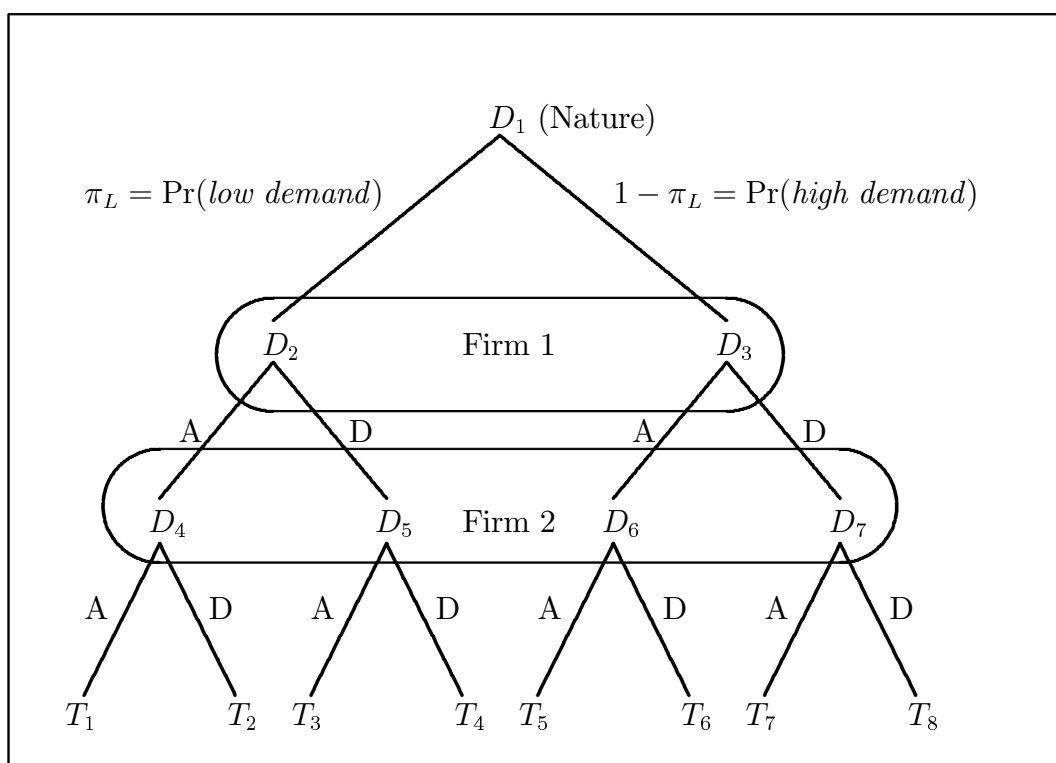


Figure 4: A move by nature in Stigler’s collusion game; A indicates player adheres to collusive agreement, D indicates player defects from collusive agreement.

This game of incomplete information is equivalent to the complete but imperfect information game shown in Figure 4. Here the first move belongs to an abstract player, nature, who (from the point of view of the players) sets demand at either a high level or a low level. In this reformulated game, it is natural to think of players independently maximizing their expected payoffs, using the probabilities they assign to each of nature’s moves. (Payoffs, which should follow the terminal nodes, have been left unspecified.) The advantage

of this formulation is that it allows analysis of games where players do not know everything there is to know about the structure of the game.¹⁰

Common knowledge¹¹ Closely related to the distinction between complete and incomplete information is the specification of *common knowledge*. Common knowledge includes all facts that are known to all players, that are known to be known to all players, that are known to be known to be known to all players, and so on. This is an infinite regress. The structure of the game itself, at least up to the kind of Harsanyi transformation illustrated in Figure 4, is taken to be common knowledge. Rationality of players is also taken to be common knowledge.

Summary The extensive form of a game consists of

1. the initial conditions of the game, including a specification of common knowledge;
2. the game tree and division of the game tree into information sets (which includes rules for determining whose turn it is and the moves available to a player whenever it is his turn);
3. the information a player has whenever it is his turn;
4. the payoffs to each player as a function of the moves of all players.

In essence, the extensive form is an exhaustive description of all elements of the game. A game in extensive form is said to be *finite* if it has a finite number of players, each of which moves at a finite number of points and has a finite number of moves at each point.

3.2 The strategic form¹²

The strategic form of a game emphasizes the relationship between strategies and payoffs, and leaves the detail of the extensive form aside. The strategic form of a game has three elements:

¹⁰See Kreps and Spence (1985) for a discussion of the use of Harsanyi's method in industrial economics.

¹¹See Aumann (1976), Milgrom (1981), Brandenburger and Dekel (1989), Rubinstein (1989a), and Binmore and Brandenburger (1990).

¹²For discussions of the strategic or normal form, see von Neumann and Morgenstern (1944), Luce and Raiffa (1957, Chs 3 and 7), Friedman (1983b, pp. 211–2; 1986, Ch. 2), Shubik (1981, pp. 289–91; 1982, Ch. 3), Aumann (1987, p. 460), Harrington (1987, p. 661).

1. the set of players;
2. the set of pure strategies available to each player;
3. a function for each player showing his payoff as a function of his own strategy and the strategies of all other players.

It is possible, and often convenient for finite games, to illustrate the strategic form by a payoff matrix. Figure 5 shows the payoff matrix for the Prisoners' Dilemma game illustrated in extensive form in Figure 1. The sides of the matrix are labeled with the names of the players. The rows of the matrix give the possible pure strategies of Professor Moriarty; the columns of the matrix give the possible pure strategies of Colonel Moran. The elements of the matrix – pairs of numbers – show the payoffs that result from different combinations of strategies. The payoff matrix compactly displays the elements of the of the game.

		Colonel Moran	
		Outraged Innocence	Craven Confession
Professor Moriarty	Outraged Innocence	-1, -1	-5, -1/2
	Craven Confession	-1/2, -5	-2, -2

Figure 5: Payoff matrix for the Prisoners' Dilemma game. Rows show strategies of Professor Moriarty, columns show strategies of Colonel Moran. The first element in each entry shows the payoff of Professor Moriarty, the second element in each entry shows the payoff of Colonel Moran.

A *pure strategy* is a plan for a player's actions throughout the course of the game. When each player has a single information set, as in a single round of the Prisoners' Dilemma game, each player's pure strategy is the same as its move at that information set. When a player has more than one turn in a game, then a pure strategy specifies what the player will do at each

information set. A pure strategy is a combination of moves, one move for each information set.

Mixed strategies are derived from pure strategies. Let S_i be the set of all possible pure strategies for player i . A mixed strategy π_i is a probability distribution over S_i . If there are three elements in S_i – s_1 , s_2 , and s_3 – then a probability distribution over S_i is a vector of three numbers – say, (π_1, π_2, π_3) – each nonnegative, that add up to unity. The first element of the vector gives the probability that player i will play s_1 , the second element gives the probability that player i will play s_2 , and the third element gives the probability that player i will play s_3 .

Let Π_i be the set of all such mixed strategies. One of the elements of Π_i is a “mixed strategy” $(1, 0, 0)$, under which the probability that player i plays s_1 is 1 and the probability that player i plays either of the other two strategies in S_i is 0. Thus Π_i includes pure strategies as special cases of mixed strategies. But Π_i includes truly random strategies as well. $(1/3, 2/3, 0)$ is a mixed strategy which has player i play s_1 with probability $1/3$, s_2 with probability $2/3$, and s_3 with probability 0.

Mixed strategies arise naturally as descriptive devices in games with players of limited rationality and games with rational players where the fact that players are rational is not common knowledge. If players follow pure strategies but there is some chance that each player will make a mistake, then each player can model the others as if they follow mixed strategies. Mixed strategies also arise naturally as a device for ensuring the existence of equilibrium in games for which there are no equilibria in pure strategies.

Figure 6 illustrates the strategic form for a repeated game or supergame made up of two rounds of Odds & Evens in which each player must decide moves for both rounds at the start of the game. Each player has two turns, and each player’s pure strategy specifies two moves – one for the player’s first turn and one for the player’s second turn. Thus $s_1 = (O, O)$, $s_2 = (O, E)$ is a combination of strategies according to which player 1 plays an odd number of fingers in each round while player 2 plays an odd number of fingers in the first round and an even number of fingers in the second round. As shown in the payoff matrix, both players break even if these strategies are played.

The strategic form masks much of the detail of the move and information structure of the game that is present in the extensive form. It highlights the relationship between what one player does, what other players do, and what each player receives. The cost of this emphasis is a certain loss of information about the underlying structure of the game.

		Player 2			
		(O,O)	(O,E)	(E,O)	(E,E)
P l a y e r 1	(O,O)	2, -2	0, 0	0, 0	-2, 2
	(O,E)	0, 0	2, -2	-2, 2	0, 0
	(E,O)	0, 0	-2, 2	2, -2	0, 0
	(E,E)	-2, 2	0, 0	0, 0	2, -2

Figure 6: Payoff matrix for two rounds of Odds & Evens. Rows show strategies of player 1, columns show strategies of player 2. The first element in each entry shows the payoff of player 1, the second element the payoff of player 2.

4 Solution concepts

We have outlined the description of games in extensive and strategic form. Now we turn our attention to what might reasonably be expected to happen when a game is played – to the analysis of solutions of a game.

As far as game theory is concerned, the term “solution” gives a misleading impression of automaticity or exactitude, particularly for those with a certain background in mathematics. If we are asked for solutions to the equation

$$x^2 - 1 = 0, \tag{1}$$

there is really not much leeway in formulating a response. The solutions are +1 and –1. These are the values of x for which equation (1) is correct, and that is that.¹³

When we talk about solutions to noncooperative games, things are by no means so automatic. For a strategic game, a solution is a collection of strategies – one for each player – that it is reasonable or plausible to think would be played. What is plausible will depend critically on the amount of information and the degree of reasoning power or computing ability that players are assumed to have. Solution concepts that appear reasonable under conditions of complete and perfect information and unlimited reasoning power sometimes yield unattractive results when these assumptions are relaxed. Much interesting work is involved in refining solution concepts to deal with games played under conditions of limited information and less than complete reasoning ability.

4.1 Noncooperative equilibrium – pure strategies

We begin by assuming complete and unlimited reasoning ability. The defining characteristic of a noncooperative game is that players are unable to make binding commitments about the strategies they will play. If binding commitments are impossible, then for a collection of strategies to be a solution of the game, every player must be willing to play his or her part of the collection. Each player must prefer to play his part of the proposed equilibrium strategy, given that all other players play their parts of the proposed equilibrium strategy. Otherwise, some player would change his strategy and the game would follow a course different from that specified by the proposed equilibrium strategy.

¹³Of course, it is only training that makes these solutions seem automatic. Contemplation of equations like (1) led to the analysis of negative numbers; contemplation of equations like $x^2 + 1 = 0$ led to the analysis of imaginary numbers. So it goes.

The formalization of this idea leads to the concept of an equilibrium. A vector of pure strategies $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is said to be a *noncooperative equilibrium* for an n -person game in strategic form if for every player i ($i = 1, 2, \dots, n$), s_i^* maximizes player i 's payoff, given that all other players play the pure strategies specified by s^* .

Formally, let (s^*/s_i) be the vector s^* with s_i substituted for s_i^* :

$$(s^*/s_i) = (s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*). \quad (2)$$

Then s^* is a noncooperative for an n -player game if

$$p_i(s^*) \geq p_i(s^*/s_i) \quad \forall s_i \in S_i \text{ and } \forall i \in N, \quad (3)$$

where S_i is the set of pure strategies open to player i , $p_i(\cdot)$ is player i 's payoff function, and N is the set of players.

The assumption of complete information enters in that each player is assumed to know the equilibrium strategies of all other players. For relatively simple games, like a single play of the Prisoners' Dilemma, this is perhaps conceivable. For complex games in which a strategy specifies a player's move whenever it is his turn, the assumption is unrealistic. The assumption of unlimited reasoning ability enters in that each player is assumed to be able to calculate an individual strategy that maximizes his return, given that other players follow their own noncooperative equilibrium strategies. This assumption too is unlikely to be met unless the game is relatively simple.¹⁴

Also present is the idea that each player seeks to maximize his payoff. If s_i^* maximizes player i 's payoff, then player i is assumed to be willing to play s_i^* .

The modern analysis of noncooperative equilibrium begins with Nash (1950, 1951). The Nash solution concept was in a certain sense anticipated by that of Cournot (1838). In the Cournot model of quantity-setting duopoly, each firm picks the output that maximizes its own profit, given the output of its rival. The equilibrium output of each firm is a best reply to the equilibrium output of the other. In game theoretic terms, the firm's strategy is its output level, the firm's payoff is its profit. Cournot equilibrium is a noncooperative equilibrium as defined above – each firm maximizes its own profit, given

¹⁴See Nash (1953, p. 130):

Each player is assumed fully informed on the structure of the game and on the utility function of his co-player . . . These information assumptions should be noted, for they are not generally perfectly fulfilled in actual situations. The same goes for the further assumption we need that the players are intelligent, rational individuals.

the outputs of other firms. For this reason, noncooperative equilibrium is sometimes called Nash-Cournot equilibrium, particularly in the context of models of quantity-setting oligopoly.

To illustrate the notion of a noncooperative equilibrium, consider the payoff matrix for a single round of the Prisoners' Dilemma game (Figure 5). If Professor Moriarty chooses to maintain his innocence, then Colonel Moran maximizes his payoff (minimizes his time on probation or in jail) by cravenly confessing. Similarly, the strategy pair in which both prisoners maintain innocence is not a noncooperative equilibrium.

If Colonel Moran chooses to cravenly confess, then Professor Moriarty maximizes his payoff by cravenly confessing as well (this makes his sentence two years rather than five years). Similarly, neither strategy pair in which one prisoner maintains innocence while the other confesses is a noncooperative equilibrium. The noncooperative equilibrium of the Prisoners' Dilemma is the strategy pair in which both players confess, and both players go to jail for two years, even though by cooperating they could get off with a year's probation each.

		Player 2	
		Odd	Even
P l a y e r 1	Odd	1, -1	-1, -1
	Even	-1, 1	1, -1

Figure 7: Payoff matrix for one round of Odds & Evens. Rows show strategies of player 1, columns show strategies of player 2. The first element in each entry shows the payoff of player 1, the second element in each entry shows the payoff of player 2.

It is easy to verify that there are games that do not have an equilibrium in pure strategies. Figure 7 shows the payoff matrix for a single round of Odds & Evens. The strategy pair (*Odd, Odd*) is not a noncooperative equilibrium: if player 1 chooses *Odd*, player 2 will maximize his own payoff by playing

Even. In the same way, $(\textit{Even}, \textit{Even})$ is not a noncooperative equilibrium: if player 1 plays *Even*, player 2 will prefer to play *Odd*. But $(\textit{Odd}, \textit{Even})$ is not a noncooperative equilibrium either: if player 2 plays *Even*, player 1 will prefer to play *Even* as well.

4.2 Noncooperative equilibrium – mixed strategies

The problem of nonexistence of equilibrium can be handled by extending what we mean by a strategy to allow for random moves. This implies a corresponding alteration in the definition of a noncooperative equilibrium.

We have defined mixed strategies π_i as probability distributions over the set of a player's pure strategies, S_i . We can extend the definition of noncooperative equilibrium from pure strategies to mixed strategies. A vector of mixed strategies $\pi^* = (\pi_1^*, \pi_2^*, \dots, \pi_n^*)$ is said to be a noncooperative equilibrium for an n -person noncooperative game in strategic form if for every player i ($i = 1, 2, \dots, n$), π_i^* maximizes player i 's expected payoff, given that all other players play the mixed strategies in π^* .

For example, consider a single round of Odds & Evens. Player 1's expected payoff is the sum of player 1's payoffs for each strategy pair multiplied by the probabilities that the strategies in the pair are chosen:

$$E(P_1) = \begin{aligned} & 1\pi_1(\textit{Odd})\pi_2(\textit{Odd}) - 1\pi_1(\textit{Odd})\pi_2(\textit{Even}) \\ & -1\pi_1(\textit{Even})\pi_2(\textit{Odd}) + 1\pi_1(\textit{Even})\pi_2(\textit{Even}). \end{aligned} \quad (4)$$

Suppose $\pi_2(\textit{Odd}) = \pi_2(\textit{Even}) = 1/2$, that is, suppose player 2 plays *Odd* or *Even* randomly but with equal probability. Then player 1's expected return is

$$E(P_1) = 0, \quad (5)$$

which does not depend on $\pi_1(\textit{Odd})$ and $\pi_1(\textit{Even})$. Thus any values of $\pi_1(\textit{Odd})$ and $\pi_1(\textit{Even})$ will maximize player 1's expected return (probabilities must be nonnegative and add up to unity). The maximum expected return is zero. In particular, $\pi_1(\textit{Odd}) = \pi_1(\textit{Even}) = 1/2$ will maximize player 1's expected return.

By exactly the same sort of argument, if $\pi_1(\textit{Odd}) = \pi_1(\textit{Even}) = 1/2$, then $\pi_2(\textit{Odd}) = \pi_2(\textit{Even}) = 1/2$ will maximize player 2's expected return (and player 2's maximum expected return will be zero). Thus $\pi_1(\textit{Odd}) = \pi_1(\textit{Even}) = \pi_2(\textit{Odd}) = \pi_2(\textit{Even}) = 1/2$ is a noncooperative equilibrium in mixed strategies for a single round of Odds & Evens.

In this equilibrium, it is as if each player tosses a fair coin before the game is played. If the coin comes up heads, the player plays an odd number of fingers. If the coin comes up tails, the player plays an even number of fingers. Each player expects to break even.

Nash (1951) shows that all noncooperative games with a finite number of players and a finite number of strategies have a mixed-strategy noncooperative equilibrium.¹⁵ Thus the problem of nonexistence of equilibrium does not arise if players are permitted to use mixed strategies.

But equilibria in mixed strategies are more fragile – less plausible – than equilibria in pure strategies. For example, consider the strategy pair (*Confess*, *Confess*), the unique equilibrium in pure strategies of the Prisoners' Dilemma game. It is plausible that if this equilibria is reached it will be self-enforcing – it is in the interest of each player to follow his part of this strategy pair, and only his part of this strategy pair, given that the other prisoner does the same. But for a single round of Odds & Evens, if player 2 plays the mixed strategy $\pi_2(Odd) = \pi_2(Even) = 1/2$, then any random strategy will maximize player 1's expected return, not just $\pi_1(Odd) = \pi_1(Even) = 1/2$ (Aumann, 1985a, pp. 43–4). In particular, player 1 could just as well play a pure strategy (play *Odd* with certainty, for example) and still expect to break even.

From a less technical point of view, and with specific reference to applications of game theory to industrial economics, models of firm decision making in which firms are supposed to decide what to do randomly – in effect, by tossing coins – are not intuitively appealing. Models in which firms pick output or location randomly do not sit well. We shall see, however, that mixed strategies arise naturally in game theoretic models under conditions of incomplete and imperfect information.

4.3 Subgame-perfect equilibrium

Figure 8 shows a simple game in extensive and strategic form. Each player has one information set, and each information set consists of a single decision node. Thus the game is one of perfect information. The game begins at decision node D_1 , where player 1 may choose left (L) or right (R). If player 1 plays right, each player receives €2. If player 1 plays left, the game moves to decision node D_2 , where player 2 may choose left or right, with the payoffs indicated.

This game has two noncooperative equilibria – (L, L) and (R, R). We

¹⁵For a formal analysis of conditions for the existence of equilibrium in noncooperative games, see Friedman (1986, Ch. 2).

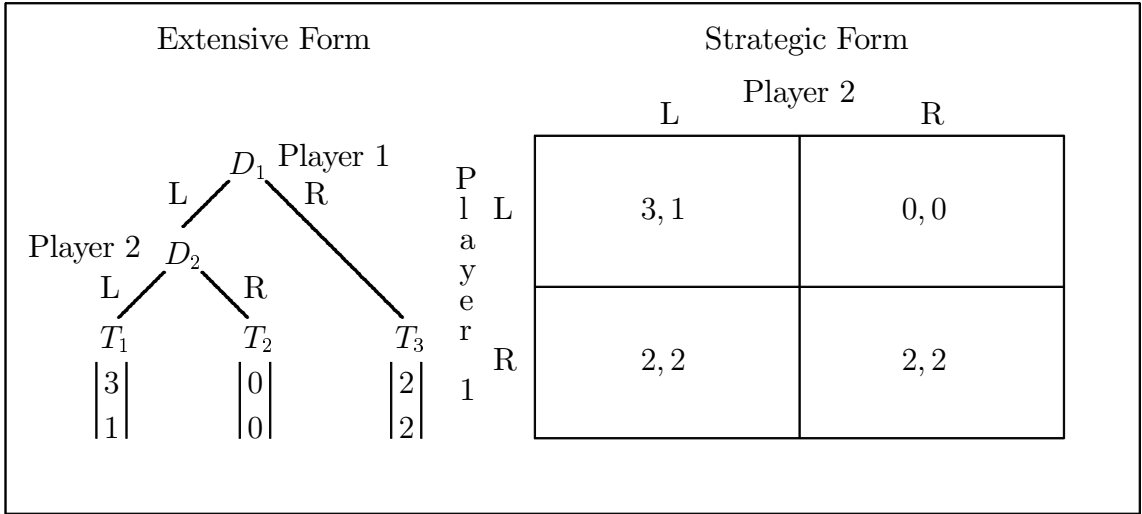


Figure 8: A game to illustrate subgame-perfect equilibrium; L indicates left, R indicates right. Source: Selten [1965, Figure I].

can show this formally by examining the players' expected payoffs for mixed strategies. Write π_1^L for the probability that player 1 will pick left, π_1^R for the probability that player 1 will pick right, and so on.

The expected payoffs of the two players are

$$H_1 = 3\pi_1^L\pi_2^L + 0\pi_1^L\pi_2^R + 2\pi_1^R = 2 + 3\pi_1^L\pi_2^L - 2\pi_1^L \quad (6)$$

$$H_2 = 1\pi_1^L\pi_2^L + 0\pi_1^L\pi_2^R + 2\pi_1^R = 2 + \pi_1^L\pi_2^L - 2\pi_1^L \quad (7)$$

(using the identity $\pi_1^L + \pi_1^R \equiv 1$ to eliminate π_1^R from the expressions for the payoffs).

Suppose player 1 chooses left. Then $\pi_1^L = 1$ and player 2's expected payoff is $H_2 = \pi_2^L$. To maximize his expected payoff, player 2 should make π_2^L as large as possible. Hence player 2 should set $\pi_2^L = 1$ and play left as well.

Now suppose player 2 plays left, so that $\pi_2^L = 1$. Player 1's expected return is $H_1 = 3\pi_1^L$. To maximize his expected payoff, player 1 should make π_1^L as large as possible, set $\pi_1^L = 1$, and play left. This shows that (L, L) is a noncooperative equilibrium.

Now suppose player 2 picks right, so that $\pi_2^L = 0$. Then player 1's expected payoff is $H_1 = 2 - 2\pi_1^L$. To maximize this payoff, player 1 should make π_1^L as small as possible, that is, player 1 should set $\pi_1^L = 0$ and play right.

What if $\pi_1^L = 0$, so that player 1 plays right with certainty? Then player 2's expected payoff is $H_2 = 2$, independent of player 2's strategy. Any mixed

strategy of player 2 will maximize 2's expected payoff, if 1 plays right. In particular, playing right will maximize 2's expected payoff. This shows that (R, R) is also a noncooperative equilibrium for the game shown in Figure 8.

Examining the payoff matrix in Figure 8, it is evident that player 2 will prefer the equilibrium (R, R) to (L, L) . In the former, 2's payoff is €2, in the latter only €1. But what would player 1 have to believe about the way player 2 would behave in order for (R, R) to be a credible solution? For (R, R) to be a plausible solution, 1 would have to believe that player 2, if at decision node 2, would play R . But if the game reaches D_2 , player 2 has no hope of getting to terminal node 3. If the game should reach D_2 , the best player 2 can do for himself is play left, take the €1 payoff, and walk away grumbling. For (R, R) to be a credible solution, player 1 would have to believe that player 2 would prefer a payoff of zero to a payoff of €1, if the game should reach D_2 . It is unreasonable to expect player 1 to believe this, which makes it unreasonable to regard (R, R) as a solution to the game in Figure 8, even though (R, R) is a noncooperative equilibrium in pure strategies.

Selten (1965) makes the following observation. If both players are at D_2 (never mind how they got there), then it is possible to think of the part of the game tree that branches out from D_2 as representing a game in its own right, a game that is a component of the original game. From this point of view, the problem with (R, R) as a solution to the original game is that (R, R) commits player 2 to a strategy that player 2 would prefer to avoid in the game beginning at D_2 .

Define a subgame of a game in extensive form as any set of nodes and branches which, starting from some decision node, is itself a game in extensive form. The initial node of the subgame must be the only decision node in its information set. If any information set is reached from the initial node, all nodes in the information set must be reachable from the initial node.¹⁶

We can then define a subgame perfect noncooperative equilibrium in pure strategies as a pure strategy combination that is a noncooperative equilibrium for every subgame of the original game. Subgame perfection requires strategies to be privately optimal even from decision nodes in the game that will not be reached if the strategies are played. For the game illustrated in Figure 8, only (L, L) is a subgame perfect noncooperative equilibrium.

¹⁶For formal definitions of a subgame, see Selten (1975, pp. 32–3) or Kreps and Wilson (1982a, pp. 868–9).

4.4 Trembling-hand perfect equilibrium

Imposing subgame perfection as an equilibrium condition may eliminate a number of implausible noncooperative equilibria – those that commit a player to moves he would be unwilling to make from some points in the game tree. There are other ways of approaching the same problem.

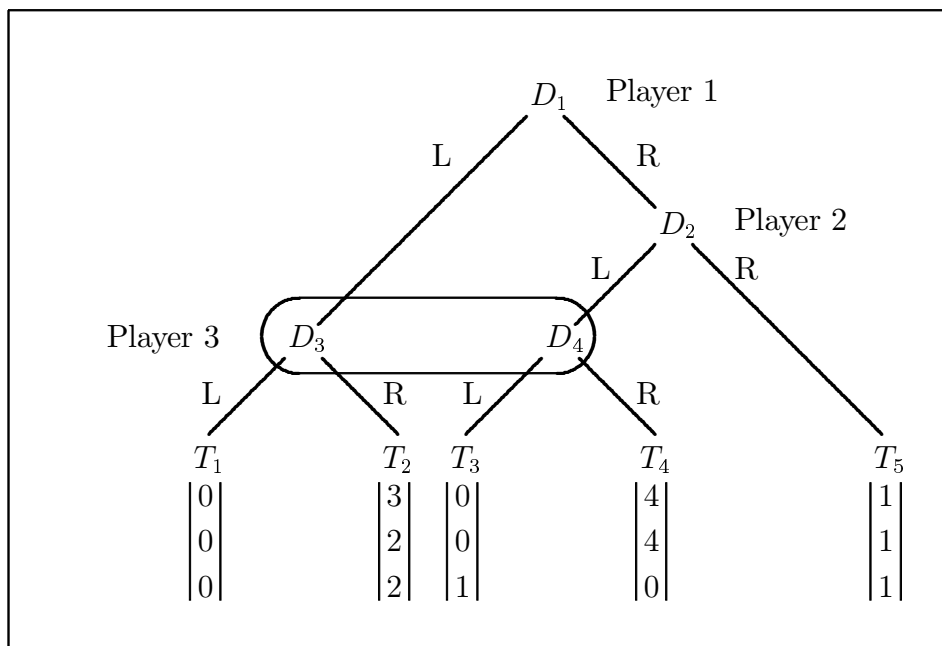


Figure 9: A game tree to illustrate trembling-hand perfect equilibrium. L indicates left, R indicates right.

Figure 9 illustrates a game in extensive form (due to Selten, 1975). There are three players in the game. The game is one of perfect recall but imperfect information. Players 1 and 2 each have an information set that consists of just one decision node, but player 3's information set has two decision nodes. If player 3's information set is reached, he does not know whether player 1 has played L (so that player 3 is at D_3) or player 1 has played R and player 2 has played L (so that player 3 is at D_4). The game in Figure 9 has only one subgame – itself.¹⁷

¹⁷ D_2 cannot be the initial node of a subgame – if we start at D_2 and play left, we reach player 3's information set. But player 3's information set contains D_3 , which cannot be reached starting from D_2 .

4.4.1 Noncooperative equilibria for the unperturbed game

There are two types of noncooperative equilibria in mixed strategies for the game shown in Figure 9.

The players' expected payoffs are

$$\begin{aligned} H_1 &= 0\pi_1^L\pi_3^L + 3\pi_1^L\pi_3^R + 0\pi_1^R\pi_2^L\pi_3^L + 4\pi_1^R\pi_2^L\pi_3^R + 1\pi_1^R\pi_2^R \\ &= 3(1 - \pi_1^R)\pi_3^R + 4\pi_1^R(1 - \pi_2^R) + \pi_1^R\pi_2^R \end{aligned} \quad (8)$$

$$\begin{aligned} H_2 &= 0\pi_1^L\pi_3^L + 2\pi_1^L\pi_3^R + 0\pi_1^R\pi_2^L\pi_3^L + 4\pi_1^R\pi_2^L\pi_3^R + 1\pi_1^R\pi_2^R \\ &= 2(1 - \pi_1^R)\pi_3^R + 4\pi_1^R(1 - \pi_2^R) + \pi_1^R\pi_2^R \end{aligned} \quad (9)$$

$$\begin{aligned} H_3 &= 0\pi_1^L\pi_3^L + 2\pi_1^L\pi_3^R + 1\pi_1^R\pi_2^L\pi_3^L + 0\pi_1^R\pi_2^L\pi_3^R + 1\pi_1^R\pi_2^R \\ &= 2(1 - \pi_1^R)\pi_3^R + \pi_1^R(1 - \pi_2^R)(1 - \pi_3^R) + \pi_1^R\pi_2^R \end{aligned} \quad (10)$$

(using that fact that a single player's probabilities of moving left and of moving right must add up to one to express expected payoffs solely in terms of the probabilities of moving right).

Now consider the derivative of each player's expected payoff with respect to his own probability of moving right:

$$\frac{\partial H_1}{\partial \pi_1^R} = \pi_2^R + \pi_3^R - 4\pi_2^R\pi_3^R \quad (11)$$

$$\frac{\partial H_2}{\partial \pi_2^R} = \pi_1^R(1 - 4\pi_3^R) \quad (12)$$

$$\frac{\partial H_3}{\partial \pi_3^R} = 2 - \pi_1^R(3 - \pi_2^R). \quad (13)$$

The right-hand side of (11) is either zero, positive or negative. If the right-hand side of (11) is zero, $\partial H_1/\partial \pi_1^R$ is zero, player 1's expected payoff is not affected by its strategy, and we can costlessly assume that player 1 will be willing to play whatever strategy is indicated in cases when $\partial H_1/\partial \pi_1^R$ is either strictly positive or strictly negative.

If the right-hand side of (11) is positive, $\partial H_1/\partial \pi_1^R$ is positive, and to maximize his expected payoff, player 1 should make π_1^R as large as possible – player 1 should play right. If the right-hand side of (11) is negative, $\partial H_1/\partial \pi_1^R$ is negative, and to maximize his expected payoff, player 1 should make π_1^R as small as possible – player 1 should play left. We consider the two cases in turn.

Case 1: $\frac{\partial H_1}{\partial \pi_1^R} \geq 0$, $\pi_1^R = 1$ If $\pi_1^R = 1$, then (12) becomes

$$\frac{\partial H_2}{\partial \pi_2^R} = 1 - 4\pi_3^R. \quad (14)$$

If $\pi_3^R < 0$, then $\partial H_2/\partial \pi_2^R$ is positive, and player 2 should make π_2^R as large as possible, consistent with staying in Case 1, that is, consistent with keeping $\partial H_1/\partial \pi_1^R \geq 0$.

If $\pi_1^R = 1$ and $0 \leq \pi_3^R \leq 1/4$, then (11) becomes

$$\frac{\partial H_1}{\partial \pi_1^R} = 1 - 3\pi_3^R \geq \frac{1}{4}, \quad (15)$$

which is positive. Thus $\pi_1^R = 1$, $\pi_2^R = 1$, $0 \leq \pi_3^R \leq \frac{1}{4}$ is consistent with being in Case 1, and $\pi_1^R = 1$, $\pi_2^R = 1$, maximize the expected payoffs of players 1 and 2, given that player 3's strategy falls in the indicated range.

Now turn to player 3's strategy. If $\pi_1^R = \pi_2^R = 1$, (13) becomes

$$\frac{\partial H_3}{\partial \pi_3^R} = 0. \quad (16)$$

Thus if players 1 and 2 play right, any value of π_3^R will maximize player 3's expected payoff. The reason is evident from Figure 9: if players 1 and 2 play right, then player 3's decision set is never reached.

It follows that one type of noncooperative equilibrium for the game in Figure 9 has players 1 and 2 choosing right and player 3 choosing right with probability not greater than one-quarter:

$$\pi_1^R = 1, \pi_2^R = 1, 0 \leq \pi_3^R \leq \frac{1}{4}. \quad (17)$$

Case 2: $\frac{\partial H_1}{\partial \pi_1^R} \leq 0$, $\pi_1^R = 0$

If $\pi_1^R = 0$, then (12) and (13) become

$$\frac{\partial H_2}{\partial \pi_2^R} = 0 \quad \frac{\partial H_3}{\partial \pi_3^R} = 2 > 0. \quad (18)$$

Since $\partial H_3/\partial \pi_3^R$ is positive, player 3 should make π_3^R as large as possible, consistent with being in Case 2. If $\pi_3^R = 1$, then

$$\frac{\partial H_1}{\partial \pi_1^R} = 1 - 3\pi_2^R, \quad (19)$$

and this will be negative, as required in Case 2, provided $\pi_2^R \geq 1/3$.

From (18), if player 1 plays left any value of π_2^R will maximize player 2's expected profit. Again, the reason is evident from Figure 9. If player 1 plays left, player 2's information set (decision node D_2) is never reached.

Thus a second type of noncooperative equilibrium for the game shown in Figure 9 has player 1 choosing left, player 3 choosing right, and player 2 choosing right with a probability of at least one-third:

$$\pi_1^R = 0, \frac{1}{3} \leq \pi_2^R \leq 1, \pi_3^R = 1. \quad (20)$$

Reasonableness of the equilibria Pick a specific Case 2 equilibrium: $\pi_1^R = 0, \pi_2^R = 1, \pi_3^R = 1$. Player 1 plays L , while players 2 and 3 play R . *Given* that player 1 moves left, there is no difficulty with player 2's part of this equilibrium strategy: player 2's information set is never reached. But if, somehow, the game were at D_2 and player 2 really believed that player 3 would play R , is it reasonable to expect player 2 to play R ? By playing left from D_2 , player 2 would move the game to D_4 . If player 3 complies with his part of the equilibrium strategy, and plays R from D_4 , then player 2 ends up with €4, while player 2 gets only €1 by playing R from D_2 .

The difficulty with Case 2 equilibria is that they prescribe unreasonable behavior for player 2 *if his information set should be reached*. This does not prevent them from being noncooperative equilibria, because if they are followed then player 2's information set is not reached. It does prevent them from being *reasonable* noncooperative equilibria: we ought to expect that player 2 will act in his own self-interest, given that he has a chance to act at all. Case 2 equilibria suppose that player 2 will act against his own self-interest, if he has a chance to act.

4.4.2 The perturbed game

The reason that case 2 equilibria show up in the example above is because players are assumed to be absolutely rational, and each player is assumed to make decisions in the knowledge that other players are absolutely rational. But no matter how confident one is in one's own ability to assess facts, calculate the odds, and make decisions, there are many situations – precisely those of complete and imperfect information – when it seems imprudent to assume that other players are also completely rational (Selten, 1975; 1985; Rosenthal, 1981).

Selten (1975) models the possibility of irrational behavior by adding to the game a small probability that players will make a mistake. In such a game, no move has zero probability. Even if a fully rational player would prefer not to make a move, in the perturbed game the move might occur by

mistake. In such a perturbed game there is always some possibility, however slight, that every information set will be reached. This rules out equilibria similar to the case 2 type for the game shown in Figure 9, equilibria that have a player making a move he would not care to make at his own information set because he is absolutely certain that his information set will not be reached.

After calculating the noncooperative equilibria in the perturbed game, one can let the probability that errors occur go to zero. The resulting equilibria for the original game – called *trembling hand perfect noncooperative equilibria* – do not allow irrational behavior at unreached information sets, because in the game with trembles there is some probability that every information set is reached.

In the perturbed version of the game shown in Figure 9, the probabilities that players will move left or right differ from the probabilities in the underlying (unperturbed) game by a term that reflects the chance that a player will make a mistake. Write the probabilities in the perturbed game as

$$\widehat{\pi}_1^L = (1 - \varepsilon)\pi_1^L + \varepsilon q^L \quad \widehat{\pi}_1^R = (1 - \varepsilon)\pi_1^R + \varepsilon q^R \quad (21)$$

$$\widehat{\pi}_2^L = (1 - \varepsilon)\pi_2^L + \varepsilon q^L \quad \widehat{\pi}_2^R = (1 - \varepsilon)\pi_2^R + \varepsilon q^R \quad (22)$$

$$\widehat{\pi}_3^L = (1 - \varepsilon)\pi_3^L + \varepsilon q^L \quad \widehat{\pi}_3^R = (1 - \varepsilon)\pi_3^R + \varepsilon q^R. \quad (23)$$

Here a “hat” denotes the probability of a move in the perturbed game. q^L is the probability of a move to the left if a player makes a mistake; q^R is the probability of a move to the right if the player makes a mistake. Naturally, $q^L + q^R \equiv 1$. Finally, ε is the probability of making a mistake.

In the perturbed game, the probability that a player moves to the left is a weighted average of the probability that the player intends to move left and the probability that the player moves left by mistake. We assume ε to be small. As ε goes to zero, the perturbed game reduces to the underlying, unperturbed, game.

The probabilities of moves in the underlying game must lie between zero and one. From (21)–(23), this implies the following bounds on probabilities of moves in the perturbed game:

$$\varepsilon q^R \leq \widehat{\pi}_1^R \leq 1 - \varepsilon q^L \quad (24)$$

$$\varepsilon q^R \leq \widehat{\pi}_2^R \leq 1 - \varepsilon q^L \quad (25)$$

$$\varepsilon q^R \leq \widehat{\pi}_3^R \leq 1 - \varepsilon q^L. \quad (26)$$

There are corresponding bounds implied for the probabilities of moves to the left in the perturbed game. All moves in the perturbed game have positive probability.

As we remarked above, it is not at first encounter intuitively plausible to model players as making random moves. But suppose all players follow pure strategies in the underlying game. If each player thinks there is a chance that other players will make a mistake, then he will treat the other players as if they are following mixed strategies. The randomness ascribed to the choices of others comes from the belief that they might err, not from the belief that their actions are inherently random. So long as players might make mistakes, one can justify an interest in mixed strategies.

Player 3's strategy Returning to the analysis of the perturbed game, expected payoffs have the same form as (8)–(10), substituting probabilities in the perturbed game (hatted probabilities) for probabilities in the underlying game. Similarly, the derivatives of each player's expected payoff with respect to his own strategy probability in the perturbed game have the same form as in (11)–(13):

$$\frac{\partial H_1}{\partial \hat{\pi}_1^R} = \hat{\pi}_2^R + \hat{\pi}_3^R - 4\hat{\pi}_2^R \hat{\pi}_3^R \quad (27)$$

$$\frac{\partial H_2}{\partial \hat{\pi}_2^R} = \hat{\pi}_1^R (1 - 4\hat{\pi}_3^R) \quad (28)$$

$$\frac{\partial H_3}{\partial \hat{\pi}_3^R} = 2 - \hat{\pi}_1^R (3 - \hat{\pi}_2^R). \quad (29)$$

It is evident from (28) that $\hat{\pi}_3^R = 1/4$ is a critical value, since it makes $\partial H_2 / \partial \hat{\pi}_2^R = 0$. Further, by analogy with case 1 solutions to the underlying game, we are led to suspect that $\hat{\pi}_3^R \leq 1/4$ will be part of a solution. Let us begin by assuming $\hat{\pi}_3^R > 1/4$ and seeking a contradiction.

If $\hat{\pi}_3^R > 1/4$, then (28) implies that $\partial H_2 / \partial \hat{\pi}_2^R < 0$. It follows that player 2 should make $\hat{\pi}_2^R$ as small as possible. In the perturbed game, this means he should set $\hat{\pi}_2^R = \varepsilon q^R$.

Substituting this value in (27), we obtain

$$\frac{\partial H_1}{\partial \hat{\pi}_1^R} = \varepsilon q^R + \hat{\pi}_3^R (1 - 4\varepsilon q^R). \quad (30)$$

On the right-hand side of (30), εq^R and $\widehat{\pi}_3^R$ are positive. The final expression, $1 - 4\varepsilon q^R$, will be positive for ε sufficiently small (for example, for $\varepsilon < 1/4$; recall that $0 \leq q^L \leq 1$).

Since we assume ε to be small, we can take the right-hand side of (30) to be positive. To maximize his expected payoff in the perturbed game, player 1 should make $\widehat{\pi}_1^R$ as large as possible, which means $\widehat{\pi}_1^R = 1 - \varepsilon q^L$.

Substituting $\widehat{\pi}_1^R = 1 - \varepsilon q^L$ and $\widehat{\pi}_2^R = \varepsilon q^R$ into (29), we obtain

$$\frac{\partial H_3}{\partial \widehat{\pi}_3^R} = -1 + \varepsilon (3q^L + q^R - \varepsilon q^L q^R). \quad (31)$$

The right-hand side is negative for ε sufficiently small. To maximize his expected return, therefore, player 3 should make $\widehat{\pi}_3^R$ as small as possible. But this contradicts our initial assumption that $\widehat{\pi}_3^R > 1/4$. We have therefore established that $\widehat{\pi}_3^R \leq 1/4$.

Now suppose that $\widehat{\pi}_3^R < 1/4$. From (28), $\partial H_2 / \partial \widehat{\pi}_2^R > 0$, so player 2 should make $\widehat{\pi}_2^R = 1 - \varepsilon q^L$ to maximize his expected profit.

Write (27) as

$$\frac{\partial H_1}{\partial \widehat{\pi}_1^R} = \widehat{\pi}_2^R (1 - 4\widehat{\pi}_3^R) + \widehat{\pi}_3^R, \quad (32)$$

which is positive for $\widehat{\pi}_3^R < 1/4$. Thus player 1 should set $\widehat{\pi}_1^R = 1 - \varepsilon q^L$ to maximize his expected return.

Substitute $\widehat{\pi}_1^R = 1 - \varepsilon q^L$ into (29) to obtain

$$\frac{\partial H_3}{\partial \widehat{\pi}_3^R} = \varepsilon q^L (1 + \varepsilon q^L) > 0. \quad (33)$$

To maximize his payoff, player 3 should make $\widehat{\pi}_3^R$ as large as possible. But this contradicts our initial assumption that $\widehat{\pi}_3^R < 1/4$.

We have therefore established that the only possible noncooperative equilibrium value for $\widehat{\pi}_3^R$ in the perturbed game is $1/4$. Let us examine the equilibrium strategies of the other players, given $\widehat{\pi}_3^R = 1/4$.

Strategies of players 1 and 2 Substitute $\widehat{\pi}_3^R = 1/4$ into (27) and (28) to obtain

$$\frac{\partial H_1}{\partial \widehat{\pi}_1^R} = \frac{1}{4} > 0 \quad (34)$$

$$\frac{\partial H_2}{\partial \widehat{\pi}_2^R} = 0. \quad (35)$$

From (34), player 1 should make $\widehat{\pi}_1^R$ as large as possible to maximize his expected return. Substitute $\widehat{\pi}_1^R = 1 - \varepsilon q^L$ into (29) to obtain

$$\frac{\partial H_3}{\partial \widehat{\pi}_3^R} = 2\varepsilon q^L - (1 - \varepsilon q^L) \widehat{\pi}_2^L. \quad (36)$$

In order for player 3 to be willing to play $\widehat{\pi}_3^R = 1/4$ (which we have already determined to be the only possible equilibrium value), the right-hand side of (36) must equal zero, in which case player 3's return is independent of his strategy. Setting (36) equal to zero, we obtain

$$\widehat{\pi}_2^R = 1 - \frac{2\varepsilon q^L}{1 - \varepsilon q^L} \quad (37)$$

as the equilibrium value for $\widehat{\pi}_2^R$ in the perturbed game. By (35), player 2 will be willing to play this strategy if the other two players make the indicated choices.

In summary, the noncooperative equilibrium strategies in the perturbed game are

$$\widehat{\pi}_1^R = 1 - \varepsilon q^L \quad \widehat{\pi}_2^R = 1 - \frac{2\varepsilon q^L}{1 - \varepsilon q^L} \quad \widehat{\pi}_3^R = \frac{1}{4}. \quad (38)$$

If any two of the players follow the indicated strategies, then the remaining player will maximize his expected returned by following his part of the strategy given in (38).

As ε goes to zero, the probability of an error in the perturbed game becomes smaller and smaller. In the limit, setting $\varepsilon = 0$ in (38), we obtain the unique trembling-hand perfect equilibrium strategy for the game shown in Figure 9:

$$\widehat{\pi}_1^R = 1 \quad \widehat{\pi}_2^R = 1 \quad \widehat{\pi}_3^R = \frac{1}{4}. \quad (39)$$

The trembling-hand equilibrium strategy for the underlying game is a Case 1 noncooperative equilibrium. Case 2 noncooperative equilibria are entirely ruled out.

4.5 Sequential equilibrium

The unsatisfactory nature of some noncooperative equilibria for the game in Figure 9 involves equilibrium strategies for players whose information sets are never reached. Selten (1975) deals with this by adding small errors to

moves. Kreps and Wilson (1982a) take a different approach to the same problem. They explicitly model the beliefs players hold at different points in the game tree. They require strategies to maximize expected payoffs from any information set in the tree, given the player's beliefs about where he is in the information set if the information set is reached. At the same time, they require a player's beliefs about where he is in an information set to be consistent with strategies being played.

The resulting *sequential equilibrium* includes both strategies and beliefs (Kreps and Wilson, 1982b, p. 257):¹⁸

There are three basic parts to the definition of a sequential equilibrium:

(a) Whenever a player must choose an action, that player has some probability assessment over the nodes in its information set, reflecting what the player believes has happened so far.

(b) These assessments are consistent with the hypothesized equilibrium strategy. For example, they satisfy Bayes' rule whenever it applies.

(c) Starting from every information set, the player whose turn it is to move is using a strategy that is optimal for the remainder of the game against the hypothesized future moves of its opponent (given by the strategies) and the assessment of past moves by other players . . .

4.5.1 The consistency of beliefs in sequential equilibrium

Consider the game shown in Figure 10. Player 1 has three possible moves – left, center, or right. If player 1 chooses left or center, then the game moves to player 2's information set. But the game is one of imperfect information – once at his information set, player 2 does not know whether he is at decision

¹⁸See also Kreps and Wilson (1982a, p. 871)

This is the substance of sequential rationality: The strategy of each player starting from there according to some assessment over the nodes in the information set and the strategies of everyone else.

and Kreps (1987, p. 586)

The basic idea is that behaviour in all parts of a game tree should be rationalized by some beliefs as to the play of the game that are not contradicted by what the player knows for sure.

For a formal definition, see Kreps and Wilson (1982a, p. 872).

node D_2 or decision node D_3 . A sequential equilibrium specifies the beliefs that player 2 entertains about his location if his information set is reached.




Figure 10: A game to illustrate sequential equilibrium. L indicates left, C indicates center, R indicates right. Source: Kreps & Wilson (1982a, Figure 1).

Bayes' rule – which relates conditional and unconditional probabilities – imposes certain constraints on these beliefs, if it can be applied. Thus in Figure 10, the unconditional probability of being at D_2 is the product of the probability of being at D_2 given that the game is in player 2's information set and the probability of being in the player 2's information set:

$$\Pr(\text{at } D_2) = \Pr(\text{at } D_2 | \text{at } D_2 \text{ or } D_3) \Pr(\text{at } D_2 \text{ or } D_3). \quad (40)$$

But we can pin down two out of three of the terms in (40), which will allow us to get at the third – $\Pr(\text{at } D_2 | \text{at } D_2 \text{ or } D_3)$, the probability implied by the strategies that player 2 is at D_2 if the game reaches player 2's information set.

If the game is at D_2 , then the game must end at T_1 or T_2 . Thus the unconditional probability of being at D_2 is the same as the probability that the game ends at T_1 or T_2 . In the same way, if the game is in player 2's information set, then the game must end at one of the terminal nodes T_1 , T_2 , T_3 , or T_4 . Thus the unconditional probability that the game reaches player 2's information set is the same as the probability that the game ends at one of T_1 , T_2 , T_3 , or T_4 .

Thus we can rewrite (40) as

$$\begin{aligned} \mu(D_2) &= \Pr(\text{at } D_2 | \text{at } D_2 \text{ or } D_3) \\ &= \frac{\text{Probability of reaching terminal nodes that follow } D_2}{\text{Probability of reaching terminal nodes that follow player 2's information set}} \\ &= \frac{\Pr(T_1) + \Pr(T_2)}{\Pr(T_1) + \Pr(T_2) + \Pr(T_3) + \Pr(T_4)} = \frac{\Pr(T_1) + \Pr(T_2)}{1 - \Pr(T_5)} \end{aligned} \quad (41)$$

The probabilities of reaching the five terminal nodes are implied by the players' strategies:

$$\begin{aligned} \Pr(T_1) &= \pi_1^L \pi_2^L & \Pr(T_2) &= \pi_1^L \pi_2^R \\ \Pr(T_3) &= \pi_1^C \pi_2^L & \Pr(T_4) &= \pi_1^C \pi_2^R . \\ \Pr(T_5) &= \pi_1^R \end{aligned} \quad (42)$$

Substituting (42) into (41), we obtain the conditional probabilities of being at D_2 and D_3 , given that the game is in player 2's information set, as implied by the player's strategies:

$$\mu(D_2) = \frac{\pi_1^L}{\pi_1^L + \pi_1^C} \quad \mu(D_3) = \frac{\pi_1^C}{\pi_1^L + \pi_1^C}, \quad (43)$$

where $\mu(D_2)$ is the conditional probability of being at D_2 and $\mu(D_3)$ is the conditional probability of being at D_3 . Naturally, $\mu(D_2) + \mu(D_3) \equiv 1$. If player 2 is in his information set, he must be at D_2 or D_3 .

From either (41) or (43), we see a limitation to the scope of Bayes' rule. If Bayes' rule is to be used to define player 2's beliefs, π_1^R must be less than 1. Otherwise, it will be impossible to carry out the division required by (43). If $\pi_1^R = 1$, then strategies imply that the game will not reach player 2's information set. If strategies imply that the game will not reach player 2's information set, then strategies cannot be used to define conditional probabilities within player 2's information set. This is a specific example of a general characteristic of Bayesian expectations (Binmore, 1990, p. 88) :¹⁹

¹⁹See also Binmore (1991) and Fudenberg and Tirole (1991).

Bayesian principles do not in themselves provide an adequate foundation for game theory. The essential reason is that “Bayesian rationality” is concerned only with the consistency of the beliefs held by players. It has nothing to say about the origins of such systems of beliefs.

Kreps and Wilson (1982a, p. 873) argue

What happens when a player reaches an information set . . . with [the probability of all terminal successor nodes] = 0? It is plausible to suppose that the player will construct some hypothesis as to how the game has been played, in the form of a strategy . . . that satisfies [the probability of all successor nodes] > 0 and then use [this new strategy] and Bayes’ rule to compute [beliefs] for [decision nodes in the information set].

and continue (1982a, pp. 873–4):

Fix a player i . His “primary hypothesis” as to how the game will be played is π , and . . . he applies π to compute [beliefs] whenever possible. We might assume that when π does not apply – when he comes to an information set . . . with [the probability of all terminal successor nodes] = 0 – then he has a “second most likely hypothesis” $\pi(2)$ that he attempts to apply. If that fails, he tries his “third most likely hypothesis” $\pi(3)$, and so on. . . the sequence of “alternative hypotheses: is independent of [the particular information set] $\pi(2)$ is the player’s second most likely hypothesis for all [his information sets]. A further strengthening of this requires that all players use the same finite sequence [of alternative hypotheses]. This requirement is in the spirit of the “common knowledge” hypothesis of Nash equilibrium — if there are rational secondary hypotheses, they should be unanimously held, just as is the primary hypothesis π .

Thus the definition of sequential equilibrium imposes strong assumptions about consistency of strategies and beliefs. Kreps and Wilson (1982a, p. 876) acknowledge that these requirements might be overly strong.²⁰ But these assumptions lead to a way of modeling strategic behavior that fits into the framework of decision theory (Kreps and Wilson, 1982a, p. 864).

²⁰For further discussion, see van Damme (1987, pp. 108–9).

4.5.2 Sequential equilibrium for the game of Figure 10

For the game shown in Figure 40, the expected payoffs of the two players are

$$H_1 = \pi_1^L(3\pi_2^L - 2\pi_2^R) - \pi_1^C \quad (44)$$

$$H_2 = -(\pi_2^L + 2\pi_2^R)\mu(D_2) + 2\pi_2^R(D_3). \quad (45)$$

If $\pi_2^L = 1$, then player 1's expected payoff becomes

$$H_1 = 3\pi_1^L - \pi_1^C. \quad (46)$$

If player 1 knows 2 will play left, then 1's choices are to end up at T_1 and collect €3 or to end up at T_3 and lose €1. Clearly, 1's best strategy is to play left.²¹

On the other hand, if 1 plays left, then 2's beliefs are defined by Bayes' rule – 2 knows he is at D_2 (from (43), if $\pi_2^L = 1$, then $\mu(D_2) = 1$). His expected payoff is then

$$H_2 = -(\pi_2^L + 2\pi_2^R). \quad (47)$$

Player 2, who evidently will not enjoy this game very much, can play left and lose €1 or play right and lose €2. His best strategy is to play left.

Thus $\mu(D_2) = 1$, $\mu(D_3) = 0$, $\pi_1^L = \pi_2^L = 1$, is a sequential equilibrium for the game shown in Figure 10.

4.5.3 Sequential equilibrium for the game of Figure 9

The unique trembling-hand equilibrium for the game shown in Figure 9, $\pi_1^R = \pi_2^R = 1$, $\pi_3^R = 1/4$, is also a sequential equilibrium supported by the belief $\mu(D_3) = 1/3$. But $\mu(D_3) = 1$, $\mu(D_4) = 0$, $\pi_1^R = 0$, $1/3 \leq \pi_2^R \leq 1$ is also a sequential equilibrium for the game shown in Figure 9.

This is a specific example of a more general phenomenon. Kreps and Wilson (1982a, section 7) show that every trembling hand is a sequential equilibrium, but not vice versa. They define a *weak sequential equilibrium* as one in which there are some unused strategies that do as well, in the sense of maximizing some player's expected payoff, as strategies that are used. Sequential equilibria that are not weak are *strict*.

In the second sequential equilibrium for the game in Figure 9, any strategy of player 2 will maximize his expected payoff. He is limited to the range

²¹Formally, 1's decision should be modeled as a constrained optimization problem – to maximize the expected payoff, subject to the constraint that $\pi_1^L + \pi_1^C + \pi_1^R \equiv 1$.

$1/3 \leq \pi_2^R \leq 1$ because this restriction is necessary to induce player 1 to set $\pi_1^R = 0$. But as far as player 2 is concerned, he would be just as well off playing strategies in the range $0 \leq \pi_2^R < 1/3$ — these strategies also maximize his expected payoff.

Kreps and Wilson (1982a, theorem 3) show that strict sequential equilibria are also trembling hand perfect equilibria.

It is only weak sequential equilibria that may fail to be trembling hand equilibria.

5 Supergames²²

A *supergame* is a game that is built up out of repetitions of some simpler, component game. A game in which four people play 100 rounds of poker is a supergame; so is a game that consists of two rounds of Odds & Evens.

		Player 2				
		Cooperate	Defect			I take
P l a y e r 1	Cooperate	3, 3	0, 4	Generous	0	3
	Defect	4, 0	1, 1	Selfish	1	0

Figure 11: Payoff matrices for a Prisoners' Dilemma game and decomposed game with cash payoffs. Source: Selten (1978, Figure 4, Figure 10).

Figure 11 shows a version of the Prisoners' Dilemma game with payoffs expressed in cash rather than prison time. If both players cooperate, each receives €3. If 1 player defects while the other player cooperates, the defecting player receives €4 while the cooperating player receives nothing. If both players defect, each receives €1.

²²See Luce and Raiffa (1957, pp. 97–102); Selten (1978, pp. 42–53); Aumann (1985b); Friedman (1986, Chs 3 and 4); Harrington (1987a); Mertens (1987a, b); and Sabourian (1989).

This game can be described in another way which yields insight into the incentives facing the players. Each player can be thought of as having two strategies. If a player chooses the generous strategy, then the other player receives €3. If a player chooses the selfish strategy, he himself receives €1.

It follows that if both players choose the generous strategy, then both receive €3. But if both players choose the selfish strategy, then both receive €1. If one player chooses the selfish strategy and one player chooses the generous strategy, then the selfish player receives €4 while the generous player receives nothing. Thus the payoffs in this decomposed game are identical to those shown in the standard Prisoners' Dilemma payoff matrix. The decomposed form of the game is also shown in Figure 11.

The situation facing the players is the same as in the earlier version of the Prisoners' Dilemma game (Figure 11). The noncooperative equilibrium is *(Defect, Defect)* (or *(Selfish, Selfish)*). Yet it is clear that each player would be individually better off if both players cooperate (are generous).

It is tempting to think that this outcome might change if the game is repeated. One might think that over time players could learn to trust each other, that they could learn from experience that cooperation was more rewarding than defection, could establish reputations for cooperative behavior (generosity), and so on. If the number of repetitions is finite and players are completely rational, however, this is not the case.

Suppose the game in Figure 11 is repeated ten times. Players seek to maximize their total payoff over the 10 rounds. Each player's strategy is a vector specifying his move in each of the 10 rounds.

Begin with the final round and work backward. If the players are at the start of the tenth round, then they are in exactly the same situation as they would be in a single round of the game. But we already know that the only noncooperative equilibrium for a single round of the Prisoners' Dilemma is *(Defect, Defect)*. Thus each player's equilibrium strategy must specify defection in the final round.

Now consider the ninth round. At this point, the tenth round is of no consequence – we already know that players must defect in the tenth round. But then at the start of the ninth round the players are in exactly the same situation as they would be in a single round of the game. The same logic leads to the same answer: players' equilibrium strategies must specify defection in the ninth round.

We can repeat this argument, moving backward, until we reach the first round. At this point, defection is specified for rounds 2 through 10. The only possible noncooperative equilibrium move for the first round is therefore *(Defect, Defect)*. By backward induction, we conclude that the noncooperative equilibrium for the 10-round supergame has each player defect in each

round. Each player collects €10, although with cooperation they could earn €30 each.

What drives this backward induction argument is the assumption that the number of repetitions is finite. Suppose the game shown in Figure 11 is repeated an infinite number of times, and that players maximize the discounted present value of their payoffs. In that case, $(Defect, Defect)$ in each period remains a noncooperative equilibrium. But in the infinitely-repeated game, there are other noncooperative equilibria as well.

Recall that a strategy specifies the player's move in each period. Define the following *trigger strategy* (Radner, 1980), recursively:

1. cooperate in period 1;
2. in later periods, cooperate if the other player cooperated in the previous period;
3. if the other player defects, defect in all subsequent periods.

The strategy pair in which both players follow this trigger strategy can be a noncooperative equilibrium for the infinitely repeated of the game in Figure 11.

To show this, suppose player 1 plays the trigger strategy. Is it in player 2's interest to follow the trigger strategy, or to defect in the first round? If player 2 follows the trigger strategy, his payoff (assuming payments are received at the end of the period) is

$$PDV_{trigger} = \frac{3}{1+r} + \frac{3}{(1+r)^2} + \frac{3}{(1+r)^3} + \dots = \frac{3}{r} \quad (48)$$

(where r is the interest rate used to discount payoffs). If player 2 defects, he receives more in the first period, but less each period thereafter. His payoff from defecting in the first round is

$$PDV_{defect} = \frac{4}{1+r} + \frac{1}{(1+r)^2} + \frac{1}{(1+r)^3} + \dots = \frac{1+4r}{r(1+r)}. \quad (49)$$

The difference between payoffs is

$$PDV_{trigger} - PDV_{defect} = \frac{2-r}{r(1+r)}. \quad (50)$$

If the interest rate is very high, the player puts great weight on the additional euro received in the defection period, and less – greatly discounted

– weight on income lost in later periods. If the interest rate is sufficiently low, cooperation beats defection.

But the logic that applies to player 2 also applies to player 1. If player 2 plays the trigger strategy, then player 1 will find it in his interest to play the trigger strategy as well. Thus in infinitely repeated noncooperative games it is possible to obtain the cooperative result, unless interest rates are too high.

This is a specific example of a general result known as the *Folk Theorem*, so called because its exact origins are unknown.²³ The Folk Theorem is that noncooperative behavior in a repeated game can sustain any strategy producing individual payoffs that exceed noncooperative payoffs in the stage game if the interest rate is sufficiently small.

There is a fundamental difference between finitely and infinitely repeated games. In an infinitely repeated game, there is always a tomorrow. There is always something that will be lost, tomorrow, by defecting today. There is always a marginal cost to defecting – future lost income – to be set against the current marginal benefit from defecting. In a finitely repeated game, there is always one period – the last one – in which there is no tomorrow. In the last period of a finitely repeated game, there is a marginal benefit to defecting but no marginal cost. Not surprisingly, it pays to defect in the final period. And this result, by itself, drives the analysis all the way back to the initial period.²⁴

Alternatively, cooperation may arise in the Prisoners' Dilemma if the players are rational but the fact that players are rational is not common knowledge. Kreps et al. (1982) consider finite repetitions of a modified Prisoners' Dilemma game in which one player believes there is a small probability that the other player, instead of maximizing his payoff, is playing a *Tit-for-Tat* strategy. Under a Tit-for-Tat strategy, a player begins by cooperating and thereafter plays whatever his rival played in the previous period. They show that in sequential equilibrium the number of rounds in which one or the other player defects is bounded above by a limit that is independent of the number of rounds played. This shows that, except for a finite number of periods, cooperative behavior can emerge if a player thinks his rival might behave cooperatively.

The underlying logic is much the same as in the infinite repetition case. If a player believes that his rival might behave cooperatively, then something

²³Aumann (1985a, p. 46; 1987, p. 468); Friedman (1986, pp. 103–4). See also Friedman (1990) and Fudenberg and Levine (1990).

²⁴It is possible to sustain cooperative outcomes in finite games if the underlying single-period game has more than one noncooperative equilibrium. In this case, what is lost by defecting is the difference between payoffs at the more remunerative and the less remunerative equilibria. See Benoit and Krishna (1985), Friedman (1985).

might be lost by defecting. But then defection will not occur if the expected marginal cost of defection – lost expected future income – is sufficiently great. At least in the early rounds of the game, it will pay to cooperate if you think your rival might respond to cooperation by cooperating. Once again, there are circumstances under which cooperative outcomes can be sustained through noncooperative strategies.

6 Experimental games²⁵

Simple games lend themselves to controlled experimentation in laboratory situations. Experimental subjects can be made to play against each other, or against computer programs designed to follow particular strategies. Payoffs can be varied, as can the way the game is presented to the subjects. The results of the play can be compared with theoretical predictions.

Of particular interest, in view of our discussion of repeated games, are the results of experiments that repeat the Prisoners' Dilemma game (Colman, 1982, pp. 115–6):²⁶

the most striking general finding is undoubtedly the [Defect, Defect] lock-in effect. . . . When the game is repeated many times there is a tendency for long runs of [Defect] choices by both players to occur.

. . . three phases typically occur in a long series of repetitions. On the first trial, the proportion of cooperative . . . choices is typically slightly greater than 1/2, but this is followed by a rapid decline in cooperation (a “sobering period”). After approximately 30 repetitions, cooperative choices begin to increase slowly in frequency (a “recovery period”), usually reaching a proportion in excess of 60 per cent by trial 300.

These results seem a clear refutation of the predictions of noncooperative game theory concerning cooperation in finite games. At least, they refute predictions of finite models that assume players are fully rational. As we have seen above, in a finite Prisoners' Dilemma supergame the only non-cooperative equilibrium has both players choosing (*Defect, Defect*) in all rounds. The fact that experiments typically show a long initial series of (*Defect, Defect*) plays does not vindicate the theoretical prediction if this

²⁵See Plott (1982, 1989); Colman (1982); Shubik (1981, pp. 316–7); Kreps (1987); Rapoport (1987); and Smith (1990).

²⁶See also Selten and Stoecker (1986).

initial series is followed by a period of learning and an eventual period of imperfect cooperation.

These results occur whether or not players are informed in advance of the number of plays (Rapoport, 1987, p. 975). If the number of plays is finite but unknown, the experiment takes on aspects of an infinite game (Selten, 1988b, p. 50) — there might always be another round, there might be a tomorrow. For this class of experiments, the results reported above can be taken as vindicating the predictions of noncooperative game theory as applied to infinite supergames: noncooperative behavior can sustain cooperative outcomes in such games.

Other aspects of experimental results raise questions about the predictions of models that rely on the assumption of super-rationality. In one way or another, these results all suggest that the outcomes of supergames are affected by things that would not matter if players were completely rational.

For example, consider the two ways of describing the game shown in Figure 11. They are absolutely equivalent; strategies and payoffs are effectively the same, although they are given different interpretations. The standard payoff matrix presents the strategies in terms of cooperation or noncooperation. The decomposed payoff matrix presents the strategies in terms of generous or selfish behavior. Experimental evidence confirms that cooperative behavior shows up more often if the game is presented in decomposed form (Colman, 1982, pp. 120–1). Perhaps players get an additional utility out of feeling generous, or perhaps they avoid a disutility by not appearing selfish. In any event, it seems clear that the way the game is presented affects the outcome of the game.

The possibility that players' payoffs might depend on how they feel about the game implies that the experimenter cannot really control the (Axelrod, 1984, pp. 117–8). The overall payoff is a combination of the monetary payoff and the way the player feels about the strategy pursued. Axelrod (1984, pp. 73–87) presents a fascinating case study of the emergence of cooperative behavior among opposing forces during trench warfare in the First World War. The example seems clearly consistent with the interpretation that after sustained periods of cooperation, soldiers came to place a value on cooperative behavior for its own sake, over and above the effect cooperation had on their chances of survival. Selten (1978, pp. 47–9) shows formally that secondary utilities of this kind can lead to the emergence of cooperative behavior in finite repetitions of the Prisoners' Dilemma game.

Axelrod's (1984) Prisoners' Dilemma tournaments are also worthy of mention. Axelrod invited game theorists to submit programs to be used to play repeated rounds of the Prisoners' Dilemma game. Each program was to play against every program submitted (including itself). The program with the

largest overall score would be declared the winner. In the first tournament, players were told that there would be 200 rounds. In the second tournament (held after releasing the results of the first tournament), the number of rounds played was random, with expected value 200. In both tournaments, a single player followed the Tit-for-Tat strategy, and in both tournaments, Tit-for-Tat had the largest score.

Tit-for-Tat is a semi-cooperative strategy. It begins by cooperating and responds to cooperation by cooperation, but it also responds to defection by defection. A defecting rival can gain for one period, by defecting, but after that a Tit-for-Tat opponent will defect as well. The only way the original defector can induce a Tit-for-Tat rival to return to cooperation is by making a cooperative move, knowing that for the round in which the cooperative move is made the Tit-for-Tat player will defect.

This will cancel the defector's initial gain, but changes the pattern of moves back to (*Cooperate, Cooperate*). The best a rival can do against Tit-for-Tat is to win by one move (for example, by cooperating on every round except the last and then defecting, or by defecting on the first round and defecting in every round thereafter).

Experimental tests of simple games suggest that models that assume fully rational behavior are not adequate descriptive devices. In experimental games, players do not mentally jump ahead 200 rounds, calculate optimal noncooperative behavior in that round, and work backward. Experimental players act in ways that suggest that they do not assume that their rivals are fully rational and that they maximize utility functions that go beyond the specified payoff matrix of the game.

7 Conclusion

Game theory offers economics, and industrial economics in particular, a framework for analyzing interactive behavior. Real-world oligopolistic markets can be thought of as noncooperative games – each player seeking to maximize a payoff that depends in part on the actions of other players. Models of oligopoly as games in extensive form and in strategic form are also ways of thinking about the reasonableness of the assumptions the modeler has made about the way players behave.

The basic solution concept for noncooperative games – Nash noncooperative equilibrium – emphasizes the requirement that equilibria be self-enforcing. Later refinements, such as subgame perfection, trembling-hand equilibrium, and sequential equilibrium, extend this idea to games of imperfect information. Extensions to repeated games show that cooperative

outcomes can emerge from noncooperative behavior.

Game theory adds powerful tools to the analytical methods available to industrial economists. But (Shubik, 1981, p. 312)

It is important to bear in mind that three different skills are called for in the investigation of oligopolistic markets. They are the skills of the economist at describing economic institutions and activities and selecting the relevant variables and relationships; the skills of the modeler in formulating a mathematical structure that reflects the pertinent aspects of the economic phenomenon; and the skills of the analyst in deducing the properties of the mathematical system that has been formulated.

Game theory offers rather more to the second and the third of these skills than to the first. For game theory to make the greatest possible contribution to industrial economics, it is essential that economic analysis not be overwhelmed by the technical skills required.

8 Problems

A1. These are the rules of scissors–paper–stone:

1. Two players simultaneously make shapes with their hands.
2. The shapes are
 - (a) scissors – two fingers spread apart;
 - (b) paper – palm facing down;
 - (c) stone – a fist.
3. If both players make the same sign, no one wins. Otherwise, the winner is determined according to the rules
 - (a) scissors cuts paper;
 - (b) paper covers stone;
 - (c) stone blunts scissors.

The winner of a round receives 1 euro.

(a) Suppose two players play a single round of scissors-paper-stone. Draw the game tree for the extensive form and the payoff matrix for the strategic form.

(b) Suppose the players play a supergame that consists of an infinite number of rounds of scissors-paper-stone. Is there an equilibrium strategy for the supergame?

See von Neumann, John and Morgenstern (1944, p. 94, p. 111).

A2 Consider a Prisoners' Dilemma game with general cash payoffs, as in Figure 12.

		Player 2	
		Cooperate	Defect
P l a y e r 1	Cooperate	R, R	S, T
	Defect	T, S	P, P

Figure 12: Payoff matrix for the Prisoners' Dilemma game with general cash payoffs.

Discuss the conditions $T > R > P > S$ and $R > (T + S)/2$. See Axelrod (1984, pp. 9–10).

A3 Figure 3 (Kreps and Wilson, 1982a, Figure 8) illustrates a game in extensive form.




Figure 13: Game tree for problem A3.

(a) Show that the following strategies are Nash equilibria and subgame perfect Nash equilibria:

- (1) $\pi^1(a) = 1, \pi^2(A) = 1, \pi^3(a') = 1$
- (2) $\pi^1(b) = 1, \pi^2(B) = 1, \pi^3(b') = 1$

(Hint: First calculate the realization probabilities – the probabilities of the different terminal outcomes. Then calculate the expected payoffs of the three players. Then show, for (1) and (2) above, that if any two players make the indicated choices, the remaining player maximizes his expected payoff by making the indicated choice.)

(b) Show that $\pi^1(b) = 1, \pi^2(B) = 1, \pi^3(b') = 1$ is a trembling-hand perfect equilibrium and a sequential equilibrium.

(Hint: for the latter, use the realization probabilities of the terminal nodes from the answer to (a) to compute the probabilities of the decision nodes according to Bayes' rule. Then compute the expected payoffs of the three players and show that each player maximizes his expected payoff by making the indicated choice.)

Remarks: The game may begin at Decision node D_1 or at decision node D_4 . The initial assessment of the probability that the game begins at D_1 is $\rho_1 = 1/2$. The initial assessment of the probability that the game begins at D_4 is $\rho_4 = 1 - \rho_1 = 1/2$. Row vectors give payoffs.

A4 Figure 14 (Kreps and Wilson, 1982a, Figure 14) illustrates a game in extensive form.




Figure 14: Game tree for problem A4.

(a) Show that $\pi_1(R), \pi_2(r) = 1$ is a sequential equilibrium.

(b) Show that $\pi_1(L), \pi_2(l) = 1$ is a sequential equilibrium.

Discuss.

9 References

- Anthony, Piers *Golum in the Gears*. New York: Ballantine Books, 1986.
- Aumann, Robert J. “Agreeing to disagree,” *Annals of Statistics* 4, 1976, pp. 1236–9.
- “What is game theory trying to accomplish?,” in Kenneth J. Arrow and Seppo Honkapohja, editors, *Frontiers of Economics*. Oxford: Basil Blackwell, 1985a, pp. 28–76.
- “Repeated games,” in George I. Feiwel, editor, *Issues in contemporary microeconomics and welfare*. London: Macmillan, 1985b, pp. 209–42.
- “Game theory,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987, Volume 2, pp. 460–83.
- Axelrod, R. *The Evolution of Cooperation*. New York: Basic Books, 1984.
- Benoit, Jean-Pierre and Krishna, Vijay “Finitely repeated games,” *Econometrica* 53(4), July 1985, pp. 905–22.
- Binmore, Ken *Essays on the Foundations of Game Theory*. Basil Blackwell, 1990.
- Binmore, Ken and Adam Brandenburger “Common knowledge and game theory,” in Ken Binmore *Essays on the Foundations of Game Theory*. Basil Blackwell, 1990.
- Brandenburger, Adam and Eddie Dekel “The role of common knowledge assumptions in game theory,” in Frank Hahn, editor, *The Economics of Missing Markets, Information, and Games*. Oxford: Clarendon Press, 1989, pp. 46–61.
- Colman, Andrew *Game Theory and Experimental Games*. Oxford: Pergamon Press, 1982.
- Cournot, Augustin *Researches into the Mathematical Principles of the Theory of Wealth*. original Paris: L. Hachette, 1838. English translation by Nathaniel T. Bacon. New York: The Macmillan Company, 1897; reprinted 1927 by The Macmillan Company, New York with notes by Irving Fisher; reprinted 1960, 1964, 1971 by Augustus M. Kelley, New York.

- van Damme, Eric *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag, 1987.
- Fisher, F. M. “Games economists play,” *Rand Journal of Economics* 20(1), Spring 1989, pp. 113–24.
- Friedman, James W. *Oligopoly and the Theory of Games*. Amsterdam: North Holland Publishing Company, 1977.
- *Oligopoly Theory*. Cambridge: Cambridge University Press, 1983b.
- *Game Theory with Applications to Economics*. Oxford: Oxford University Press, 1986.
- Fudenberg, Drew and David K. Levine “An approximate Folk Theorem with imperfect private information,” *Journal of Economic Theory* 54(1), June 1990, pp. 26–47.
- Harrington, Joseph E., Jr. “Non-cooperative games,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987, Volume 3, pp. 661–63.
- Harsanyi, John C. “Games with incomplete information played by ‘Bayesian’ players, I-III,” *Management Science* 14(3), November 1967, pp. 159–82, 14(5), January 1968, pp. 320–34, 14(7), March 1968, pp. 486–502 reprinted in John C. Harsanyi *Papers in Game Theory*. Dordrecht: D. Reidel Publishing Company, 1982, pp. 115–70.
- Kreps, David M. “Nash Equilibrium,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987, Volume 3, pp. 584–88.
- *Game theory and economic modelling*. Oxford: Clarendon Press, 1990a.
- *A Course in Microeconomic Theory*. New York: Harvester Wheatsheaf, 1990b.
- Kreps, David M. and A. Michael Spence, “Modeling the role of history in industrial organization and competition,” in George I. Feiwel, editor, *Issues in contemporary microeconomics and welfare*. London: Macmillan, 1985, pp. 340–78.
- Kreps, David M. and Robert Wilson “Sequential equilibria,” *Econometrica* Volume 50, Number 4, July 1982a, pp. 863–94.

- “Reputation and imperfect information,” *Journal of Economic Theory* 27 August 1982b, pp. 253–7.
- Kuhn, H. “Extensive games and the problem of information,” in H. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, Volume II. Princeton: Princeton University Press, 1953, pp. 193–216.
- Luce, R. D. and Raiffa, H. *Games and Decisions*. New York: John Wiley, 1957.
- Mertens, Jean-François “Repeated games,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987a, Volume 4, pp. 151–3.
- “Supergames,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987b, Volume 4, pp. 551–3.
- Milgrom, Paul “An axiomatic characterization of Common Knowledge,” *Econometrica* 49, 1981, pp. 219–22.
- Nash, John F. “Equilibrium points in n -person games,” *Proceedings of the National Academy of Sciences*, 36, 1950, pp. 48–9.
- “Non-cooperative games,” *Annals of Mathematics* 54(2), September 1951, pp. 286–95.
- von Neumann, John “Zur theorie der gesellschaftsspiele, ” *Mathematische Annalen* 100, 1928, pp. 295–320.
- von Neumann, John and Morgenstern, Oskar *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944.
- Plott, Charles R. “Industrial organization theory and experimental economics,” *Journal of Economic Literature* 20(4), December 1982, pp. 1485–527.
- “An updated review of industrial organization: applications of experimental methods,” in Richard C. Schmalensee and Robert D. Willig, editors, *Handbook of Industrial Organization* Amsterdam: North Holland, 1989, Volume II, Chapter 19, pp. 1109–76.
- Radner, Roy “Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives,” *Journal of Economic Theory* 22(2), 1980, pp. 136–54, reprinted in Andrew F. Daughety *Cournot*

- oligopoly: characterization and applications*. Cambridge: Cambridge University Press, 1988, pp. 122–41.
- Rapoport, Anatol “Prisoners’ dilemma,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987, Volume 3, pp. 973–6.
- Rasmusen, Eric *Games and Information*. Oxford: Basil Blackwell, 1989.
- Rosenthal, Robert W. “Games of perfect information, predatory pricing and the chain-store paradox,” *Journal of Economic Theory* 25, 1981, pp. 92–100.
- Roth, Alvin E. “Game theory as a part of empirical economics,” *Economic Journal* 101(404), January 1991, pp. 107–4.
- Rubinstein, Ariel “The electronic mail game: strategic behavior under ‘almost common knowledge’” *American Economic Review* 79(3), June 1989, pp. 385–91.
- Sabourian, H. “Repeated games: a survey,” in Frank Hahn, editor, *The Economics of Missing Markets, Information, and Games*. Oxford: Clarendon Press, 1989, pp. 62–105.
- Selten, Reinhard “Reexamination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory* 4(1), 1975, pp. 25–55, reprinted in Reinhard Selten *Models of Strategic Rationality*. Dordrecht: Kluwer Academic Publishers, 1988.
- Selten, Reinhard and Rolf Stoecker “End behavior in sequences of finite Prisoners’ dilemma supergames,” *Journal of Economic Behavior and Organization* 7, 1986, pp. 47–70.
- Shapiro, Carl “The theory of business strategy,” *Rand Journal of Economics*, 20(1), Spring 1989, pp. 125–37.
- Shubik, Martin “Game theory models and methods in political economy,” Chapter 7 in Kenneth J. Arrow and Michael D. Intriligator, editors, *Handbook of Mathematical Economics*, Volume 1. Amsterdam: North-Holland, 1981, pp. 285–330.
- “Cooperative games,” in John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave*. London: Macmillan Press Limited, 1987, Volume 1, pp. 662–4.

- Smith, Vernon L. “Experimental economics: behavioral lessons for microeconomic theory and policy,” 1990 Nancy L. Schwartz Memorial Lecture. J. L. Kellogg Graduate School of Management, Northwestern University, 1990.
- Stigler, George J. “A theory of oligopoly,” *Journal of Political Economy* 72(1) February 1964, pp. 44–61 reprinted in George J. Stigler *The Organization of Industry*. Homewood, Illinois: Richard D. Irwin, Inc., 1968, pp. 39–63.
- Sutton, John “Explaining everything, explaining nothing? Game theoretic models in industrial economics,” *European Economic Review* 34(2/3), May 1990, pp. 505–2.
- Tirole, Jean “Jeux dynamiques un guide de l'utilisateur,” *Revue d'Économie Politique* 93(4), 1983, pp. 551–75.
- “Noncooperative game theory: a user’s manual,” Chapter 11 in Jean Tirole, *The Theory of Industrial Organization*. Cambridge, Massachusetts: The MIT Press, 1988, pp. 423–59.