



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Review

A Survey of Noncooperative Game Theory with Reference to Agricultural Markets: Part 1. Theoretical Concepts

Richard J. Sexton*

This paper is the first of a two-part survey on noncooperative game theory relevant to agricultural markets. Part 1 discusses types of noncooperative games and reviews important developments in noncooperative game theory solution concepts, including Nash equilibrium, subgame perfect equilibrium, and perfect Bayesian equilibrium. Strengths and weaknesses of game theory as a modelling tool are also assessed. Part 2 of the survey will discuss specific applications to agricultural markets.

1. Introduction

The advent of game theory is considered to be the publication of von Neumann and Morgenstern's book, *The Theory of Games and Economic Behavior* in 1944. In the immediately succeeding years important advances in game theoretic analysis were made by game theory's other pioneers including Nash (1950, 1951) and Shapley (1953). The state of the art during this era was summarized in Luce and Raiffa's classic book, *Games and Decisions: Introduction and Critical Survey*, published in 1957. However, few results useful to economics were developed over the next twenty years, and Luce and Raiffa's book remained a definitive source on basic game theory.

An upsurge of interest in pure and applied game theory in economics began in the mid 1970s as research began to emphasize decision makers who were rational but had limited information and who interacted with others in explicitly dynamic settings. Game theory texts published today bear little resemblance to Luce and Raiffa's book. With the publication in 1990 of David Kreps text, *A Course in Microeconomic Theory*, game theory is now being integrated into the training of most new Ph.Ds in economics and agricultural economics.

This paper is part 1 of a two-part survey on noncooperative game theory for agricultural economists. In part 1, I review recent conceptual advances in game theoretic analysis relevant to economics and assess its successes and failures. In part 2 to appear in the next issue of this *Review*, I consider applications of noncooperative game theory to agricultural markets. To date, the methodology has been little used by agricultural economists. Agricultural economics is an applied field and game theory is a tool of economic theory, so perhaps the infrequency of usage is not surprising. Another factor may be that agricultural markets are often regarded as prototype competitive markets, and game theory is a tool of imperfect competition.

I reject this latter argument, but I do agree that agricultural economics is and should remain an applied field. Most, however, would accept theory's role in guiding application, and agriculture as an industry is sufficiently unique that we cannot necessarily rely upon theory developed without reference to these distinctive features of agricultural markets. For example, concerns about monopsony or oligopsony power are relatively unique to agriculture, given the typical immobility of the raw product and fewness of processors. The fact that the marketing process for agricultural products is initiated by the production and sale of a particular raw product that is relatively nonsubstitutable for

* Associate professor, Department of Agricultural Economics, University of California-Davis. The author thanks Julian Alston, Garth Holloway, and the *Review* referees for helpful comments.

Review coordinated by John Zeitsch.

other inputs is also unique. Third, at the retail level, the emerging power of large food chains is important and relatively distinctive. Given that food manufacturers are also often powerful, this consideration raises important bilateral monopoly/ oligopoly and principal-agent issues. Fourth, agriculture is quite unique among industries in that producers are allowed, even encouraged or forced, to form coalitions for the purposes of procuring inputs and marketing production.

For these reasons, I argue that the potential for application of game theory to agricultural markets is quite high. Perhaps then a survey emphasizing applications in agriculture can stimulate interest in the topic among agricultural economists. The goal here is not to provide a comprehensive introduction to noncooperative game theory. Rather, I hope to describe and illustrate in part 1 some of the key concepts in use today, and then in part 2 to demonstrate their relevance to analysis of agricultural markets. A number of book-length treatments of the subject have appeared in recent years for those interested in detailed study.¹

2. Some Basic Classifications and Concepts

Games are partitioned into two broad classes: *cooperative* and *noncooperative*. Players in cooperative games can make binding commitments, whereas in noncooperative games they cannot. This distinction must be interpreted narrowly. For example, communication among players can be modelled under either game structure. And players in a noncooperative game setting can agree to cooperate and sign contracts if the game structure allows it. However, if it is individually desirable for a player to defect from an agreement or breach his/her contract, he/she will do so in a noncooperative game setting. Cooperative game theory is most useful in settings where players can form groups or coalitions. The analysis then focuses on what these coalitions can accomplish with little or no emphasis on the processes whereby these outcomes are achieved within the coalition. Most of the recent progress and interest in game theory has been in the area of noncooperative games, and, hence, those games are the focus of this review.²

Noncooperative games are analyzed in either their *normal* or *extensive form*. The extensive form is manifest as the familiar game tree. It specifies the order of play, information, and actions available to each player and the ensuing payoffs that are contingent upon the players' actions. A player's *strategy* specifies his/her action at each point (node) in the game tree where the player has to move. The normal or *strategic form* is a summarized description of the extensive form. It usually is depicted as a matrix associating payoffs with each possible combination of (pure) strategy choices by the players.

Every extensive form has a corresponding normal or strategic form, but different extensive forms may be represented by the same normal form. A main reason is that the normal form necessarily abstracts from the dynamic aspects of most interesting games. Kreps (1990a) argues that the "great successes of game theory in economics" have arisen primarily due to the opportunity to think about the dynamic character of competitive interactions afforded by the extensive form. Constructing the extensive form is the very essence of the art of game theoretic modelling.

¹ These include Kreps' microeconomic theory text (1990a) and a second book by Kreps (1990b) that is not concept oriented, but, rather, is a thoughtful discussion of noncooperative game theory's successes, failures, and future prospects. Rasmusen (1989) is an excellent, modern introduction to noncooperative game theory. Tirole's text (1988) in industrial organization is a masterful presentation of noncooperative game theory applications. The *Handbook of Industrial Organization* (Schmalensee and Willig 1989) focuses heavily on noncooperative game theory applications and includes a chapter on noncooperative game theory methods by Fudenberg and Tirole, who have also recently published a book on the subject (Fudenberg and Tirole 1991).

Books that treat both cooperative and noncooperative games include Friedman (1986) and the two volume treatise by Shubik (1982, 1984). For readers primarily interested in cooperative game theory, Luce and Raiffa remains an excellent reference.

² This focus is for brevity and is not to suggest that cooperative games do not provide a useful tool for analysis of agricultural markets. Indeed institutions such as agricultural cooperatives, marketing orders, and marketing boards enable coalitions of farmers to organize and make the type of binding agreements that are fundamental to cooperative game theory.

Because the discussion here will focus on games in extensive form, it is useful to review terminology relating to the extensive form using an example. Figure 1 is a simple model of *moral hazard*. There are two players, a farmer (the principal) and a marketer (the agent). If the farm product is marketed effectively (e.g., no spoilage), it is worth 3.0 at retail. A marketing agent can provide these services at a cost of 0.5, or the farmer, who is less efficient at marketing, can provide them at a cost of 1.0. The farm product net of marketing costs is worth 2.5 if the agent expends a high effort in marketing it. I assume that there are many competing agents, so that agents' services are priced at cost. The product is worth 2.0 if the farmer vertically integrates and markets the product him/herself. The product is only worth 1.5 if the agent shirks and expends low effort.

The points in Figure 1 at which either player takes an action are referred to as *nodes*. A *successor* to a node is any node that may occur later in the game if the given node has been reached. An *end node* is a node with no successors. A *branch* is one action from among a player's set of potential ac-

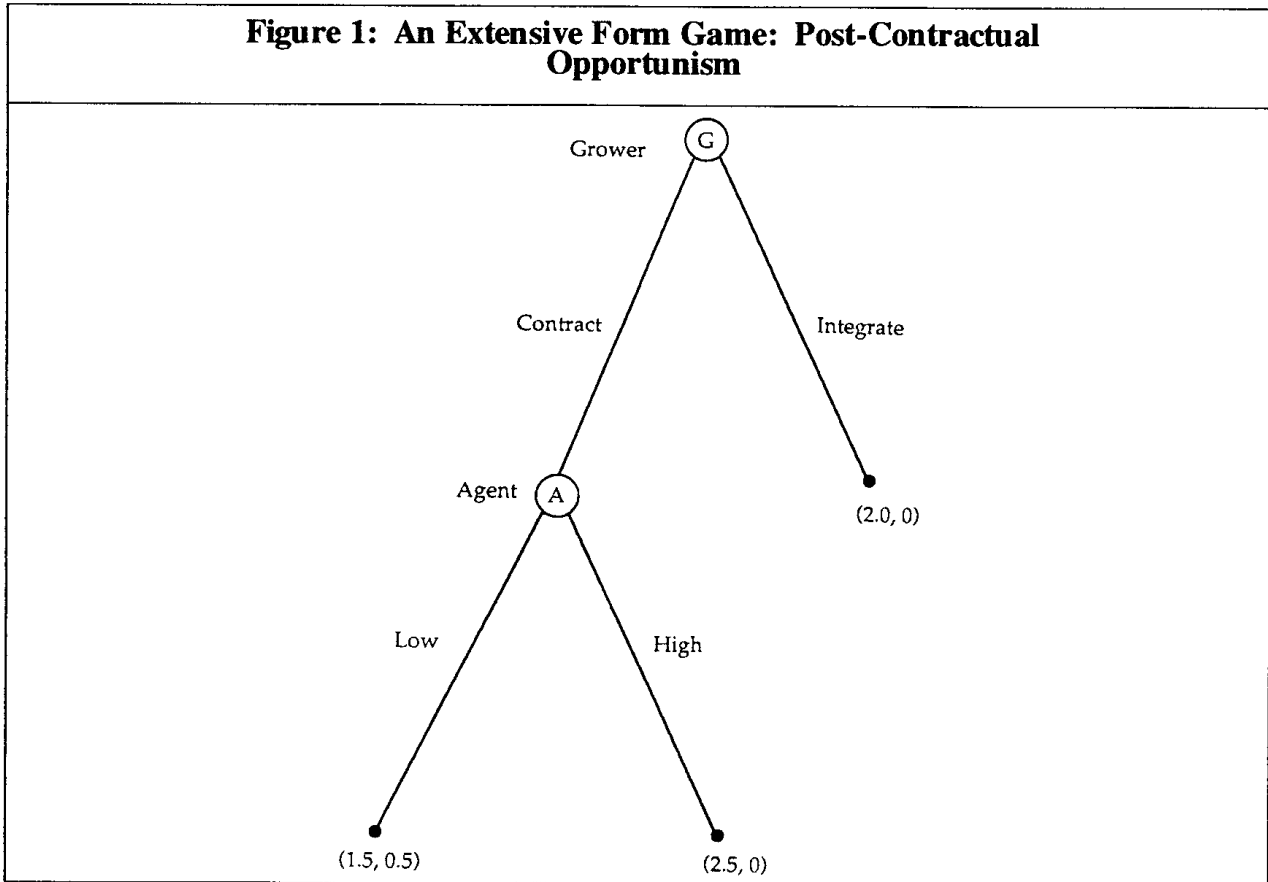
tions at a particular node. A *path* is a sequence of nodes and branches from the starting node to an end node. *Payoffs* for (grower, agent) are denoted at each terminal node.

The cornerstone solution concept for noncooperative games is the *Nash equilibrium*. A strategy combination s_1, \dots, s_n is a Nash equilibrium if no player $i = 1, \dots, n$ would wish to deviate from his/her strategy, given that no other player(s) deviate. In other words, taking his/her opponents' actions as given, if no player would wish to change his/her own action, the resulting strategy combination is a Nash equilibrium. To state the concept formally, define strategy sets S_i and payoff functions $\pi_i(s_1, \dots, s_n)$ for each player $i = 1, \dots, n$. The strategy combination $s^* = \{s_1^*, \dots, s_n^*\}$ is a Nash equilibrium if

$$\pi_i(s_1^*, \dots, s_n^*) \geq \pi_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*),$$

for all $s_i \in S_i$, and for all $i = 1, \dots, n$.

Many well-known results in economics are Nash equilibria of their associated games. The most



famous is mutual defection or "finking" in the various incarnations of the prisoners' dilemma game.³ The Cournot equilibrium is the Nash equilibrium to the static game where oligopolists choose quantities, and the Bertrand equilibrium is the Nash equilibrium to the static game where they set prices. Von Stackelberg's leader-follower equilibrium is a Nash equilibrium to a dynamic game where the leader moves first and then the follower moves. The Nash equilibrium to the moral hazard game in Figure 1 is for the agent to expend low effort (if given an opportunity to play) and for the grower to vertically integrate.

A number of existence results for Nash equilibria have been proven, many of which are summarized by Friedman (1986). A fundamental result due to Nash (1951) is that every game with a finite number of pure strategies has at least one Nash equilibrium, possibly in mixed strategies. Mixed strategies involve a player randomizing among his/her pure strategies.⁴ Similar existence results can be proven for games with a continuum of actions (such as the choice of a price or quantity), but complications enter when payoff functions are discontinuous or nonquasi-concave in the strategy choices (Dasgupta and Maskin 1986).

The process of finding pure strategy Nash equilibria is usually quite straightforward. The analyst merely proposes a candidate equilibrium strategy combination and then checks for each player if his/her strategy is optimal given the candidate strategies for all other players. If so, the candidate strategy combination is a Nash equilibrium.

It is worth commenting upon the Nash equilibrium as a solution concept because its problems have inspired refinements of the equilibrium concept that have comprised much of the recent progress in pure noncooperative game theory. The mutual best reply property of a Nash equilibrium is indeed an appealing property. However, two important criticisms of the Nash equilibrium as a solution concept can be raised:

1. Many games have multiple Nash equilibria, raising the question of how to choose among them.

2. Nash equilibria are very "noncooperative" in that the solutions they characterize often involve players doing distinctly worse than if they were somehow able to coordinate their actions.

I will consider each argument in turn. The games mentioned above in introducing the Nash equilibrium concept generally have a unique equilibrium, but many games have a multiplicity of equilibria in pure and/or mixed strategies. Consider, for example, the simple game of entry and entry deterrence illustrated in extensive form in Figure 2. In this game the entrant moves first and chooses to be IN the market or OUT. The incumbent then responds by choosing either PREDATE or ACCOMMODATE, where the former implies a price war and the latter might imply either Cournot or collusive behavior. Denote the entrant and incumbent by subscript E and I, respectively and monopoly, predation, and accommodation by superscripts M, P and A respectively. Then

$$\pi_I M > \pi_I A > \pi_I P, \text{ and}$$

$$\pi_{EA} > 0 > \pi_{EP}.$$

The Nash equilibria for this game are (IN, ACCOMMODATE) and (OUT, PREDATE).

³ Everyone is familiar with the two prisoners whose finking on each other produces long prison terms for each. However, the term "prisoners' dilemma" is applied broadly to contexts where cooperation is in players' mutual interests, but individually each has incentive to behave noncooperatively. Examples are duopolists setting prices or output levels, nations choosing trade policies, or communities competing for industry through tax incentives. A stimulating book by Axelrod (1984) is devoted to the study of prisoners' dilemma situations.

⁴ Most often economists are interested in pure strategy equilibria because mixed strategies are often difficult to interpret from an economic perspective. Many games may have both pure and mixed strategy equilibria, and the modeller will emphasize the pure strategy equilibria. See Fudenberg and Tirole (1989) and Sutton (1990) for discussion of alternative interpretations of mixed strategy equilibria. Rubinstein (1991) expresses the view that nonexistence of equilibrium in pure strategies should not necessarily cause the modeller to turn to analysis of mixed strategies. Rather, nonexistence of a solution should alert the modeller to possible deficiencies in the game description or assumptions underlying the solution concept.

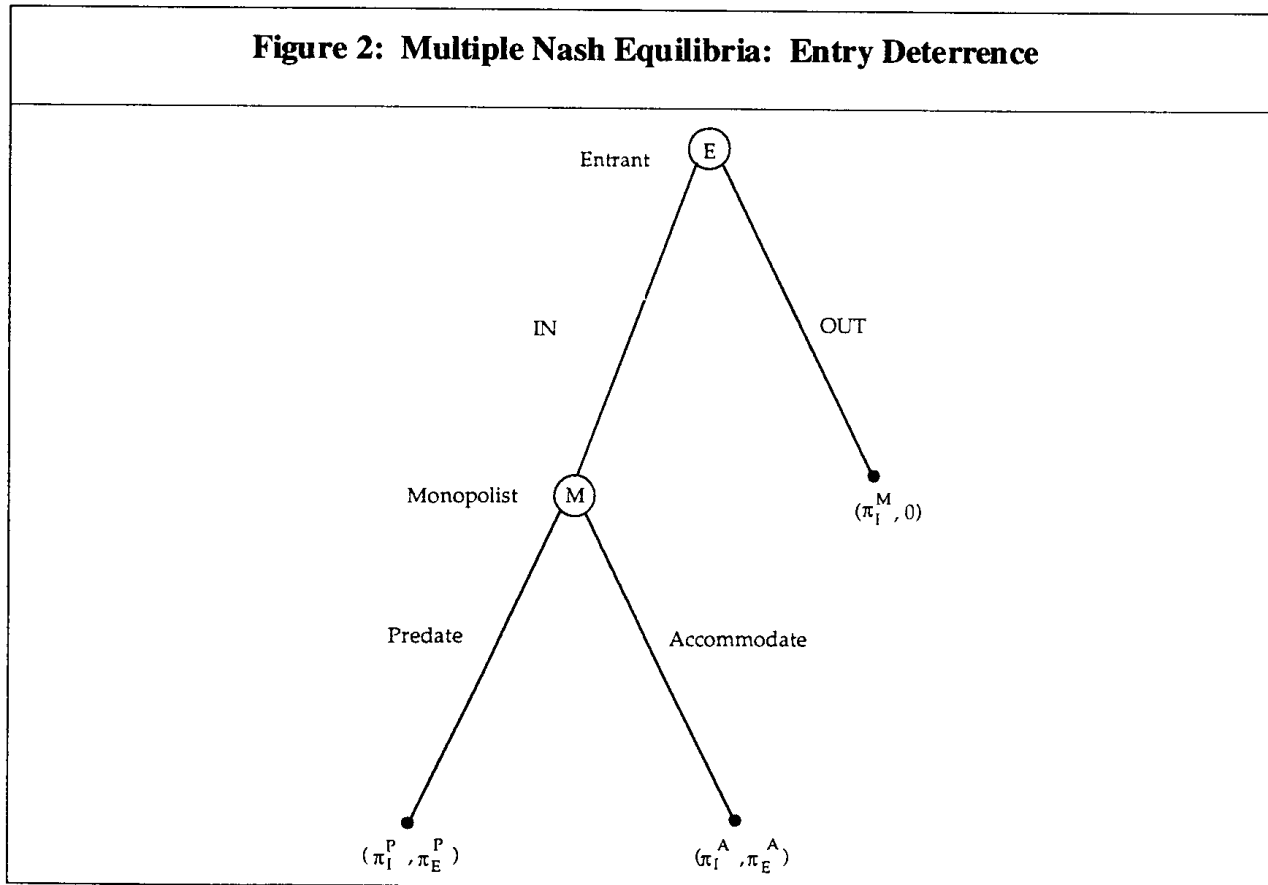
A multiplicity of Nash equilibria might signal either that the formal game specification fails to capture real-world elements that might suggest an obvious way to play the game or that the Nash equilibrium concept is ill-suited to analyze the game at hand. This is the case in the entry-deterrence game, where the equilibrium (OUT, PREDATE) involves a noncredible threat by the incumbent, i.e., if actually called upon to choose between PREDATE and ACCOMMODATE by the entrant's choice of IN, the incumbent rationally chooses ACCOMMODATE. Situations such as this have inspired refinements of Nash Equilibrium that we will examine shortly.

The notion of an obvious way to play a game is based on the pioneering work by Schelling (1960). The idea is that in many games that have multiple Nash equilibria, players may still know what to do. These equilibria are called *focal points*. They are Nash equilibria that are compelling for psychological reasons not easily incorporated in the formal game specification. Focal points may be based on past experience or a general sense of how people will behave.

The concern about the extreme "noncooperativeness" of Nash equilibria is that they often predict a distinctly suboptimal outcome from the perspective of the collective welfare of the players. All of the games mentioned at the outset are this way. The "prisoners" in the prisoners' dilemma game both get long jail sentences from finking on each other, the Bertrand and Cournot equilibria both earn the oligopolists less than the joint profit maximum output. And in the moral hazard game in Figure 1, the Nash equilibrium outcome with vertical integration is Pareto dominated by contracting with an agent who expends high effort.

Two comments are in order. First, in these games' static contexts, the noncooperative outcomes are probably realistic. Although superior outcomes to the Nash equilibrium are available in each instance, players have unilateral incentives to defect from these solutions. People can be their own worst enemies. Second, the divergence between equilibrium and optimum (in the sense of maximizing total payoffs) behavior may signal that the model is a poor representation of real-world behavior. For example, in single play games, reputation is not an

Figure 2: Multiple Nash Equilibria: Entry Deterrence



issue, nor are players able to make precommitments that might subsequently bind them to a more advantageous course of action. These considerations suggest the importance of including dynamics and information in game specifications, which, in fact, have been important dimensions of recent game theory research.

3. Information and Extensive Form Games

Having established the Nash equilibrium as a foundation, we now consider the advancements that have led to the recent years' explosion of interest in game theory modelling. A player's *information set* at any point in the game consists of the different nodes in the game tree that he/she knows might be the actual node but cannot distinguish among by direct observation. Consider the simple coordination problem among farmers illustrated in Figure 3. There are two market periods, early and late, and either farmer can plant a perishable crop for harvest during one but not both periods. The early harvest period is more lucrative due to greater demand, and Farmer A, who runs a larger scale operation is better able to take advantage of the early market than is Farmer B. However, if the farmers can coordinate their plantings to smoothen supply across market periods, they will each do better than if they harvest for the same period and create a glut. A similar coordination story might involve scheduling harvests to best utilize fixed processing capacity. The payoffs under the alternative outcomes are listed at the end nodes in Figure 3.

Panels (a) and (b) in Figure 3 illustrate two alternative ways this game might be played. In panel (a) the players commit to planting decisions simultaneously. Thus, although Farmer A is depicted first on the game tree, Farmer B does not know A's choice when it is time to make his/her own choice, i.e., he/she does not know whether B_1 or B_2 is the actual node. His/her information set consists of $\{B_1, B_2\}$. Information sets are depicted on game trees by either encircling nodes that comprise an information set as in panel (a) or connecting the nodes with a dashed line.

Panel (b) depicts a case where Farmer A is able to move first. How he/she achieves this position might be an interesting strategic question. For example, he/she could sign a labour contract specifying an early planting cycle and containing a large penalty for breach. In this case Farmer B knows what action farmer A has taken when it is time to make his/her decision. Every information set in panel (b) consists of a single node or in game theory parlance is a *singleton*.

Figure 3 illustrates the distinction in game theory between *perfect information* and *imperfect information*. In a game of perfect information each information set is a singleton; otherwise it is a game of imperfect information.

What are the pure strategy Nash equilibria to the coordination games in Figure 3? The game in panel (a) has two equilibria for (A,B): (EARLY, LATE) and (LATE, EARLY). The total payoff from (EARLY, LATE), exceeds that from (LATE, EARLY), but there is no way in this noncooperative game structure for Farmer A to necessarily persuade Farmer B to undertake that option.


Farmer B's strategy choices are complicated somewhat in the game depicted in panel (b). They must specify his/her move in response to either of A's possible actions. Three Nash equilibrium strategy combinations emerge:

1. (EARLY, if EARLY then LATE; if LATE then EARLY) with outcome that A plays EARLY and B plays LATE.
2. (LATE, if EARLY then EARLY; if LATE then EARLY) with outcome that A plays LATE and B plays EARLY.
3. (EARLY, if EARLY then LATE; if LATE then LATE) with outcome that A plays EARLY and B plays LATE.

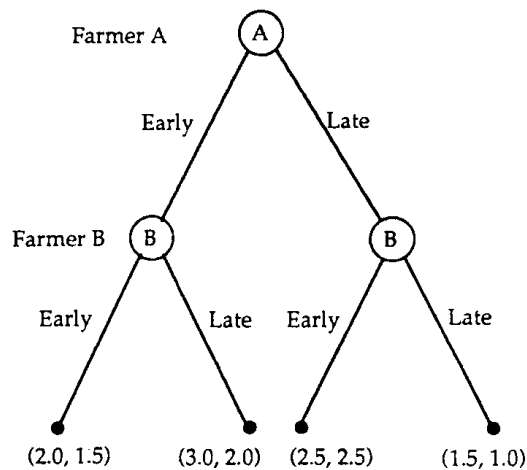
An important refinement of Nash equilibrium is the concept of *subgame perfect equilibrium (SPE)* due to Selton (1975). The game depicted in Figure 3(b) is dynamic in that A moves first and B observes his/her move. Yet the construct of Nash equilibrium requires A to take B's strategy as given in

Figure 3: Coordination Games Between Farmers

(a) Simultaneous Choices



(b) Sequential Choices



choosing his/her own move. This fact tends to produce Nash equilibria in dynamic games that involve noncredible threats on the part of some player(s). Both the second and third equilibrium to the game in panel (b) involve such threats. Equilibrium 2 involves a threat by B to play EARLY regardless of A's action. Taking this strategy as given, A's best reply is LATE. However, if A chose EARLY so that it was *fait accompli*, B's optimal response is to choose LATE, not EARLY. Similarly, the threat to play LATE if LATE in equilibrium 3 makes no sense, yet because B is never called upon to make that move in equilibrium, the strategy combination is a Nash equilibrium.

Subgame perfection works to eliminate noncredible threats. To understand the concept it is necessary to define a *subgame*. A subgame is a game consisting of a node that is a singleton for all players, that node's successors and the payoffs at the associated end nodes. The game in Figure 3(b) has three subgames: the complete game itself and the games beginning at nodes B₁ and B₂. Conversely in panel (a) the only subgame is the game itself. The game of entry and entry deterrence in Figure 2 has two subgames: the game itself and the game beginning at the node following the entrant's choice of IN. The moral hazard game in Figure 1 also has two subgames.

A SPE is a set of strategies for each player such that the strategies comprise a Nash equilibrium for the entire game and also for every subgame. Subgame perfection requires strategies to be in equilibrium everywhere along the game tree, not only along the equilibrium path.

The concept is exceedingly useful for analyzing dynamic games of perfect information such as those depicted in Figures 1,2 and 3(b) and also games of 'almost perfect' information. These are dynamic games where at a given date t players choose actions simultaneously knowing all actions taken during the preceding periods $1, \dots, t-1$. The within-periods simultaneity is a deviation from perfect information. The most common example of these games are *repeated games* where players repeatedly play a simultaneous single period game, such as a prisoners' dilemma or choices of price or quantity by oligopolists in a static market environment.

The virtues of the SPE concept are twofold: SPE are usually straightforward to derive using *backward induction*, and requiring subgame perfection is often very effective at eliminating nonplausible Nash equilibria in dynamic games. Solution by backwards induction involves proceeding to the final play (a node whose successors are all end nodes) and deriving the optimal behavior for the player who has the move at that node. The solution at this point will be simple common sense; the player will choose whatever option maximizes his/her payoff among the alternatives. That portion of the game tree can then be replaced with the optimal action to take place there and the associated payoffs, and the analyst can proceed up the game tree to the next node or set of nodes. Optimal play can be derived here given that it is now known what will transpire subsequently. In this manner the game can continue to be folded back and solved. The manner in which the solution is derived insures that the properties of a SPE are satisfied, i.e., optimal behavior was derived at each node.⁵

The backwards induction algorithm can be used to solve the dynamic games posited thus far in this paper. In Figure 1's moral hazard game, if the agent gets the move, his/her best response is to exert LOW effort. Given the Nash equilibrium to

this subgame, the grower's best response at his/her move is to vertically integrate. Thus (INTEGRATE, LOW) is the unique SPE.

Subgame perfection eliminates one of the equilibria in the entry-deterrence game. Given a choice of IN by the entrant, the monopolist's best response in the ensuing subgame is to ACCOMMODATE. Given accommodation, the entrant's best move at his/her play is to choose IN. Thus, (IN, ACCOMMODATE) is the unique SPE, and the Nash equilibrium (OUT, PREDATE) is eliminated as a noncredible threat.

Finally, the coordination game in Figure 3b had three Nash equilibria. Two of them involve noncredible threats by B, and, thus, do not satisfy the requirements of subgame perfection. These are the threat to play EARLY in response to EARLY by A in the second equilibrium, and the threat to play LATE in response to LATE by A in the third. The unique SPE then involves A playing EARLY and B playing LATE.

Consider now dynamic games with "almost perfect" information. Two classic examples are the iterated prisoners' dilemma and the *chainstore game* due to Selton (1978). They are useful to consider because they suggest the failure of subgame perfection in certain contexts which has led to the search for further refinements of equilibrium.

Consider playing a prisoners' dilemma game over some large but finite number of periods. Whereas

⁵ This solution algorithm is effective so long as the game tree isn't too big or complicated. Circumstances where players are indifferent among alternatives can also create problems because the manner in which ties are resolved likely will effect play of the game. Usually the analyst has leeway to resolve ties, and some justification from theory can often be given for a particular resolution. Figure 1 illustrates this point. In many games one type of player will be assumed to behave competitively and earn just some reservation level of payoff, usually normalized to zero. The agent in Figure 1 earns zero both from accepting a contract and expending high effort and from staying out of the market under grower integration. Any payoff to the agent strictly above his/her reservation payoff cannot be an equilibrium because another payoff that paid him/her slightly less could be proposed and would be accepted.

the Nash equilibrium of mutual finking and joint punishment is intuitive in any single play of the game, it seems sensible that as the players repeated the game several times they would learn eventually to cooperate and, thus, each achieve a better payoff. Such is not the case. Solving the game via backward induction, it is clear that mutual finking is the unique Nash equilibrium in the final period, because there can be no gain from playing a cooperative strategy. Since the final period's play is now determinate, there is no gain from cooperating in the penultimate period, so mutual finking ensues there also. And so the game unravels to produce a unique SPE wherein each player finks at any and every opportunity.

The chainstore game is essentially a many period replication of the entry-deterrence game of Figure 2. Whereas accommodation of a single entrant makes sense, the intuition is that a firm facing entry in different markets in successive periods ought to respond aggressively early in the game (choose FIGHT) in hopes of deterring subsequent entrants. Such is not the case, however, as the SPE calls for accommodation and entry in every period, a solution verified easily by backward induction.

3.1 Infinitely Repeated Games

If the game is repeated infinitely, the backward induction algorithm that generated the SPE described above breaks down; there is no final period to solve to begin folding the game back. The fundamental result for infinitely repeated games is the *folk theorem* which asserts that almost any outcome can be a Nash equilibrium provided players are sufficiently patient (don't discount the future too heavily). The idea is that any feasible, individually rational payoffs can be supported as a Nash equilibrium by the players espousing strategies to punish anyone who deviates from the prescribed equilibrium path. These strategies will satisfy the properties of a Nash equilibrium if the one period gain from cheating does not exceed the subsequent discounted losses from punishment.

Such strategies need not be subgame perfect, i.e., players may not have incentive to play their threat strategies if actually called upon to do so. How-

ever, restricting attention to SPE is not helpful in infinitely repeated games as another version of the folk theorem shows that this refinement does not reduce the limit set of equilibrium payoffs.

What are the implications of repeated games and the folk theorems for applied researchers who may wish to use game theory? Most fundamentally, considerable suspicion is called for if anyone puts much emphasis on a particular equilibrium for an infinitely repeated game. A second point is that infinitely repeated games are not very reflective of real-world contexts. Most decision makers do not have infinite horizons, but it is notable that this feature does not undermine the message of the folk theorems because the theorems also hold for games with a finite probability of ending in any period, provided this probability is sufficiently low.⁶

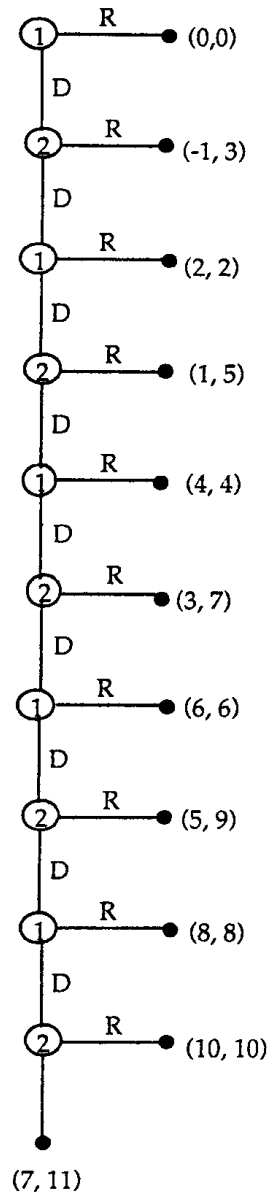
A more significant indictment of repeated games (whether finite or infinite) is that life does not usually involve repeated play of the same game. Consider, for example, repeated play of Figure 1's moral hazard game. LOW effort by an agent may be interpreted to mean letting product quality deteriorate. Consequentially, consumers may be alienated from the product in subsequent periods, and, hence, the structure of those games is altered. In other words, what happens today usually affects the games to be played in the future.

The main virtue of repeated games lies not in their value as realistic modelling paradigms, but, rather, in suggesting through the stark results they generate that richer and more realistic specifications of the game environment are called for. Providing richer game structures has also inspired further refinements in equilibrium that we now examine.

3.2 Games of Incomplete or Imperfect Information

An element missing from either the iterated prisoners' dilemma or chainstore games is reputation. It

⁶ If $\gamma \leq 1$ is the discount parameter and $\theta \leq 1$ is the probability that play continues at each period, then players should merely use the factor $\gamma\theta$ to discount the future.

Figure 4: The Centipede Game

would seem that the "prisoners" have a great interest in acquiring a cooperative reputation. Similarly the chainstore should value a reputation as one who responds aggressively to entry. These elements have no way of emerging in the prototype finite-horizon versions of these games. Another important game that illustrates a shortcoming of finite-period, perfect-information games is Rosenthal's *centipede game* (1981) illustrated in Figure 4. By playing their cards right (i.e., choosing DOWN), players A and B can each secure payoffs of 10 in

this game. Yet the unique SPE results in A playing RIGHT at his/her first opportunity, leading to payoffs of (0,0).


The intuition in the iterated prisoners' dilemma or centipede games is that a player might "take a chance" on playing cooperatively at the outset just to see what might happen. The backward induction algorithm of subgame perfection does not permit this intuition to emerge. The environment where it can emerge is in games of *incomplete information*. Analysis of these games was facilitated by Harsanyi's observation (1967) that a game with incomplete information could be transformed into a game with imperfect information by introducing *Nature* as a player who moves first at the outset of a game. The choices made by Nature define a player's *type*, including possibly his/her strategy set, payoff functions, and knowledge concerning locations on the game tree—*information partition* in game theory parlance. When nature moves in these environments, this is said to establish a *state of the world*.

I will now illustrate the modelling procedure for games with incomplete information and describe the refinements in equilibrium they have inspired. We can then demonstrate how incomplete information can be used to unravel the logic that produces the paradoxical equilibria in the games just discussed.

Figure 5 illustrates the modelling process for the sequential-choice version of the coordination game among farmers. The incomplete information concerns player B's type. He/she might be either a "profit maximizer" or "mean spirited." A profit-maximizing B has the same payoffs as in Figure 3. A mean-spirited B, however, obtains utility from inflicting pain upon his/her neighbour, and, hence, will always time his/her planting to diminish A's payoff. The way to model this uncertainty is to let Nature choose between (maximizer, mean) with probabilities (P, 1-P).

Moves by Nature at the outset of a game convert the game to one of incomplete information whenever at least one of the players is uninformed of Nature's choice. If some players observe nature's choice and others do not, then the game involves

Figure 5: Coordination Between Farmers Under Incomplete Information



asymmetric information, and some players have valuable private information.⁷

In Figure 5 the more sensible alternative is that A is uninformed, which produces the extensive form in Figure 5(a). The less realistic alternative in this particular example but the alternative with more important consequences for game theoretic modeling is that B is uninformed as illustrated in Figure

⁷ In technical terms private information means that some player's information partition is *finer* than some other player's partition. Games of asymmetric information are necessarily games of imperfect information because if the players' information partitions differ, the information sets cannot all be singletons. Games can have asymmetric information without having incomplete information. For example, players may undertake moves at the outset of a game that are not revealed to other players but which influence the way they play subsequently in the game.

5(b). The dotted lines depict information sets which are not singletons. In Figure 5(a) Farmer A does not know Nature's choice and, hence, whether the actual node is A₁ or A₂. Player B's information sets are all singletons because he/she observes both Nature's and A's move.

In Figure 5(b) B cannot distinguish between B₁ and B₃ or between B₂ and B₄. The introduction of incomplete information in the manner depicted in Figure 5(a) does not complicate solving the game in any meaningful way. A knows that Maximizer B will choose the opposite of A's choice of EARLY or LATE, and Mean B will choose the same as A. To solve this type of game, A is assumed to have a von Neumann-Morgenstern utility function and choose between {EARLY, LATE} to maximize his/her expected payoff. In this case EARLY is a dominant choice for A regardless of the value of P, so equilibrium involves A choosing EARLY and B choosing EARLY (LATE) if he/she is mean spirited (a profit maximizer).

The type of game depicted in Figure 5(b) is interesting because it possibly allows the uninformed player to update his/her information based upon the informed player's move.⁸ This type of scenario has prompted further important refinements of Nash equilibrium.

Figure 5(b) illustrates the problem that arises for subgame perfection as a solution concept for these types of games. Because of the imperfect information, the nodes where B moves are no longer subgames; none of nodes B₁ - B₄ are singletons. Thus, the only subgame is the entire game itself, and requiring subgame perfection does not eliminate either of the Nash equilibria that involve noncredible threats.

It is natural that a refinement of Nash equilibrium to accommodate games of incomplete and asymmetric information should consider both players' strategies and their beliefs and the manner in which those beliefs are updated as the game is played. A refinement that accomplishes this objective is *perfect Bayesian equilibrium* (PBE). In a PBE, players' strategies are optimal given their beliefs and beliefs are obtained from strategies and observed actions using Bayes' rule whenever possible.⁹

The following is a formal definition of a PBE based on Rasmusen (1989): A PBE consists of a strategy combination and a set of beliefs such that at each node of the game: (1) the strategies are Nash for the remainder of the game, given the beliefs and strategies of the other players, and (2) the beliefs at each information set are rational given the evidence, if any, from previous play in the game. Condition (1) is a perfectness condition, and condition (2) says that beliefs should be formed using Bayesian updating whenever possible.¹⁰

⁸ Notice that this happens not to be the case in the Figure 5(b) game because A has the dominant strategy of EARLY regardless of B's type.

⁹ Credit for the development of perfect Bayesian equilibrium is somewhat hard to pinpoint. The concept is aligned with Selton's work (1975) on perfection and Kreps and Wilson's work (1982a) on *sequential equilibrium*. Early signalling models such as Akerlof (1970) and Spence (1973) implicitly use the concept. The first explicit application is Milgrom and Roberts (1982a). Kreps (1990b) credits Fudenberg and Tirole (1988) with formalizing the concept.

¹⁰ The following example illustrates using Bayes rule to calculate posterior probabilities. It is bad form and perhaps illegal to inquire about the marital status of an applicant for a faculty position. Still, however, inquiring minds want to know. Suppose an interviewer's prior probability that an applicant is married (M) is:

$$P(M) = 0.4.$$

The data observed by the interviewer is that the applicant is a homeowner, a fact revealed in casual conversation. The interviewer knows the conditional probabilities of observing this information for a married or unmarried (UM) person of the applicant's age:

$$\begin{aligned} P(H/M) &= 0.6 \\ P(H/UM) &= 0.2. \end{aligned}$$

The marginal probability of observing home ownership among this applicant's age cohort is

$$\begin{aligned} P(H) &= [P(H/M) * P(M)] + [P(H/UM) * P(UM)] \\ 0.36 &= (0.6 * 0.4) + (0.2 * 0.6). \end{aligned}$$

In other words, homeowners are twice as likely to be married as not. Thus, the posterior probability that the applicant is married is

$$P(M/H) = P(H/M) * P(M) / P(H) = 2/3.$$

Because the interviewer observed data more consistent with M than UM, it is intuitive that the prior on M should be revised upward. The above equations can be converted to general formulae by replacing H with "data," M with "event," and UM with "not the event."

There is no general solution method to calculate PBE comparable to the backward induction algorithm for SPE. Rather, solution is a thought process that involves proposing plausible strategy combinations and testing to see if they are best responses (i.e., Nash). Then each player's strategy is tested at each node to see if it is a best response given the player's beliefs at each node. Out-of-equilibrium beliefs and strategies are an important part of constructing a PBE. In particular, the analyst must check whether any player would like to take an out-of-equilibrium action in order to influence other players' beliefs. I illustrate application of the PBE concept to an important class of incomplete information games known as *signalling games*.

3.3. Signalling Games

The basic signalling game is a two-period dynamic game. The player who moves first (the leader) has private information about his/her type that affects the player who moves last (the follower). Signalling's origin is Spence's model (1973) of education. The model has proven to be rich in application in the succeeding years.

The following definition of PBE for a signalling game is from Fudenberg and Tirole (1989). Player 1 (the leader) observes private information as to his/her type t_1 and chooses action a_1 . Player 2 observes a_1 and chooses action a_2 . Payoffs for each player are $\pi_i(a_1, a_2, t_1)$. Prior to play, player 2 has beliefs $P_1(t_1)$ concerning player 1's type. Player 2 can update his/her belief about t_1 based upon his/her observation of 1's action, a_1 . Denote this posterior probability as $P_1^*(t_1/a_1)$. However, player 1 anticipates that his/her action will influence player 2's posterior beliefs and, hence, his/her action. A PBE is a set of strategies $a_1^*(t_1)$ and $a_2^*(a_1)$ and posterior beliefs $P_1^*(t_1/a_1)$ that satisfy the following conditions:

1. $a_1^*(t_1)$ maximizes $\pi_1(a_1, a_2(a_1), t_1)$,
2. $a_2^*(a_1)$ maximizes $\sum_{t_1} P_1^*(t_1/a_1) \pi_2(a_1, a_2, t_1)$
3. $P_1^*(t_1/a_1)$ is derived from the prior P_1 , a_1 , and Bayes rule whenever possible.

Conditions 1 and 2 are perfectness conditions, and condition 3 is the Bayesian updating requirement. Notice that condition 1 requires player 1 to take into account his/her role in influencing player 2's action. The qualifier on condition 3 is important because Bayes rule is not applicable for events that occur off the equilibrium path. These events occur with zero probability, which implies a division by zero in Bayes formula (see footnote 10), making the posterior undefined. Any posterior beliefs are compatible with Bayes rule in these cases. This result, in turn, admits many perfect Bayesian equilibria for some games and has inspired a search in recent years for further refinements to eliminate some of the equilibria.

To illustrate the application of PBE, consider the entry deterrence model of Milgrom and Roberts (1982a). Milgrom and Roberts wished to show that limit pricing might emerge as a rational strategy under incomplete information. The asymmetric information concerns the incumbent firm's unit costs, which may be either HIGH or LOW and denoted respectively as c_H and c_L . If the entrant enters, he/she incurs a sunk cost $K > 0$, and post-entry play is assumed to be Cournot. Let the entrant's profits net of K be denoted by π_E and assume that

$$\pi_E(c_H) > 0 > \pi_E(c_L),$$

i.e., entry is profitable if the incumbent is high cost but not if he/she is low cost.¹¹

Signalling enters the Milgrom-Roberts model because a low-cost incumbent produces more and charges less than a high-cost counterpart under normal conditions. For example, denote the static profit-maximizing monopoly outputs for high- and low-cost incumbents as $q^M(c_H)$ and $q^M(c_L)$, respectively. However, producing $q^M(c_L)$ may not be sufficient for a low-cost incumbent to signal its type because a high-cost incumbent may be willing to produce this output, thereby reducing its period

¹¹ A low-cost incumbent will produce more in a Cournot equilibrium than will a high-cost version, and, thus, post-entry profits will be lower if the incumbent is low cost.

1 profit in order to masquerade as low cost in hopes of deterring entry.

Whether or not players succeed in signalling their types is an important dimension of signalling models. A PBE where signalling does distinguish among types is known as a *separating equilibrium*. A PBE where the types remain undistinguished is known as a *pooling equilibrium*. Many signalling games have both types of equilibria.¹²

In general three types of constraints must be satisfied to establish a separating equilibrium. In the context of the limit pricing application, they are: (1) Participation--the payoffs available in equilibrium must be financially viable for the uninformed player; (2) Incentive compatibility-- a high-cost firm must not have incentive to choose the low-cost firm's output;¹³ and (3) Nonpooling--the low cost firm must earn higher profits through signalling its type than through pooling.

In a separating equilibrium, observing the equilibrium choices of the informed players allows a complete inference to be made as to their types. The limit pricing model tends to have both pooling and separating equilibria. A separating equilibrium involves a low-cost incumbent producing an output, $q^*(c_L)$ sufficiently in excess of $q^M(c_L)$ that a high-cost version would not be tempted to pool (constraint (2) above) and, rather, would choose $q^M(c_H)$. The entrant correctly infers this result and chooses not to enter if it observes $q^*(c_L)$. To complete specification of the PBE, posterior beliefs, $P^*(\cdot)$ on the part of the entrant for outputs other than $q^*(c_L)$ or $q^M(c_H)$ must be specified that support the proposed equilibrium. These beliefs are arbitrary, so $P^*(HIGH/q') = 1$ for all $q' \in \{q^M(c_H), q^*(c_L)\}$ is a valid choice to support the equilibrium.

If the cost of signalling is sufficiently great, a low-cost incumbent will instead choose $q^M(c_L)$ (constraint (3) above is violated) and a pooling equilibrium will ensue where both high- and low-cost incumbent types produce the same output. In this case the entrant enters if its expected profit is positive, given its priors on the incumbent's type. An important implication of this type of model is that the introduction of just a small probability in

the entrant's mind that the incumbent is high cost possibly causes the rational low-cost incumbent to *discretely* increase its period 1 output above its profit-maximizing monopoly level to signal its type.

3.4 Reconsidering the Paradoxical Equilibria in Finite-Horizon Games

The preceding observation is the key to unravelling the paradoxical equilibria in the iterated prisoners' dilemma, chainstore, and centipede games. The key references are Kreps and Wilson (1982b) and Milgrom and Roberts (1982b) on the chainstore game and Kreps, Wilson, Milgrom and Roberts (1982) on the prisoners' dilemma. The modelling approach is similar in each case. The game is converted to one of incomplete and asymmetric information by introducing the probability that a player's type is not as modelled in the original specifications of the game. For example, Kreps *et al.* consider the possibility that one of the "prisoners" can only play a "tit-for-tat" strategy that calls for him/her to cooperate at the outset of play and at any subsequent period t if his/her opponent cooperated at period $t-1$. Or in the chainstore game, the possibility of a "rapacious" incumbent who enjoys predation is introduced by Kreps and Wilson.

A key facet of these (and any other) games is that the game structure is *common knowledge*. This means that each player knows the configuration of the game tree and the other player(s) know that he/she knows and so on. This point is important because it means that an informed player has an opportunity to exploit an uninformed player's uncertainty. For example a rational (non tit-for-tat) prisoners' dilemma player can play cooperatively at the outset of the game to give the impression that he/she is tit for tat. The other player is not fooled by this behavior, but, nonetheless, as long as his/her partner is playing cooperatively, it may be in his/

¹² In addition, a third type of equilibrium may exist, where, in the context of the limit pricing model, the high-cost firm randomizes between masquerading and not masquerading as a low-cost firm.

¹³ The fact that the high-cost firm's profits are lower at $q^M(c_L)$ than at $q^M(c_H)$ is the key feature in meeting this constraint.

her interest to play along by choosing to cooperate also.

Analogously, in the chainstore game, a nonrapacious incumbent has incentive to predate during the early periods of play of this game to perpetuate the possibility in entrants' minds that he/she is rapacious. Potential entrants, being aware that even a nonrapacious incumbent may fight entry during early periods of play, elect rationally not to enter.

Introducing uncertainty into these models is, thus, seen to rather drastically alter the equilibria from the stark results obtained by applying subgame perfection to the perfect information versions of these games. The new equilibria call for players in the prisoners' dilemma to cooperate in early periods and only fink towards the end of play, or in the chainstore game for the incumbent to fight entry in early periods and accommodate only towards the end of play. These outcomes sit better with intuition and, moreover, with actual play of the games in experimental settings (see, for example, Axelrod 1984 and McKelvey and Palfrey 1992). A further key point is that these new equilibria are obtained even with very modest degrees of uncertainty, e.g., low probabilities that a prisoner is tit for tat or an incumbent is rapacious.

3.5 Further Refinements

I discuss briefly here other refinements to Nash equilibrium that have emerged in the literature in recent years. Two equilibrium concepts that were developed contemporaneously with PBE and have similar properties (and, hence, yield similar equilibria) to PBE are Selton's (1975) concept of *trembling-hand perfect equilibrium* and Kreps and Wilson's (1982a) *sequential equilibrium*. The idea behind trembling hand perfection is that players may make mistakes (their hands may tremble) during play of a game. A trembling-hand perfect equilibrium strategy continues to be optimal for a player even if there is a small chance that some other player will pick an out-of-equilibrium action.¹⁴

The concept of sequential equilibrium is also based upon the specification of strategy profiles that are

Nash for the remainder of the game, given the beliefs and strategies of the other players, and updating beliefs using Bayesian inference whenever possible. Kreps and Wilson add a further consistency requirement for sequential equilibrium which for some games limits the range of equilibria relative to perfect Bayesian equilibrium. The consistency requirement, for example, would require that two players observing another player's actions should form the same beliefs as to that player's type. It also imposes consistency of beliefs over time.¹⁵

The concepts of SPE, PBE, trembling-hand perfect equilibrium, and sequential equilibrium can be related as follows: Every sequential, perfect Bayesian, and trembling-hand perfect equilibrium is also subgame perfect. Every trembling-hand perfect equilibrium is a sequential equilibrium, and every sequential equilibrium is also a perfect Bayesian equilibrium but not vice-versa.

As noted, the problem of multiplicity of PBE due to the arbitrariness of out-of-equilibrium beliefs has stimulated the search for ways to restrict these beliefs and, hence, limit the admissible PBE. This has been an area of considerable on-going research and is beyond the scope of this survey. For interested readers, the book by Van Damme (1987) provides a comprehensive discussion, although some work has been accomplished since its publication.¹⁶

¹⁴ For an example of how trembling-hand perfection refines equilibrium consider the coordination game between farmers in Figure 3(b). One Nash equilibrium involves A, who moves first, playing EARLY and B playing (if EARLY then LATE; if LATE then LATE). As long as A plays EARLY, B's strategy is a best reply, but if there is a chance that A will tremble and play LATE, then it is certainly not optimal for B to respond with LATE, i.e., this Nash equilibrium is not trembling-hand perfect. The equilibrium where A plays LATE and B plays (if EARLY then EARLY, if LATE then EARLY) can be eliminated by the same argument.

¹⁵ The additional restrictions on equilibrium imposed by sequential equilibrium relative to PBE imply a mechanical check of the PBE to see whether they satisfy the consistency requirement.

¹⁶ Additional key references include McLennan (1985), Kohlberg and Mertens (1986), Grossman and Perry (1986), Banks and Sobel (1987), Cho and Kreps (1987), and Fudenberg, Kreps, and Levine (1988).

4. Problems in Noncooperative Game Theory

I end part 1 of this survey by summarizing what are considered to be some of modern game theory's major problems. Kreps (1990a) and Sutton (1990) provide a more complete discussion. A first observation is that game theory requires clear and precise specification of the rules of the game. This means that modes of "free-form" competition are not amenable to game theory analysis. More significant is the problem that the equilibria of games often shift dramatically due to seemingly minor modifications of the rules. This situation is observed most vividly in games of bargaining, a topic discussed in part 2 of this survey. Related to this point is Kreps' concern that the rules of the game are specified exogenously by the analyst and taken for granted. Where do the rules come from? Might they be endogenous? The response to these concerns is Rubinstein's (1991) point that careful specification of the rules of the game is the essence of game theoretic modelling and why indeed it is an "art."

A problem discussed by both Kreps and Sutton is the multiplicity of equilibria that often emerge and the associated problems of choosing among them. As Sutton (p. 506) notes, "given any form of behaviour observed in the market, we are now quite likely to have on hand at least one model which. . . [derives] that form of behaviour as the outcome of individually rational decisions." This problem has led to the search for refinements as we have just seen, but Kreps and Sutton are also concerned with the method of most refinements. Most refinements focus upon out-of-equilibrium actions, but Kreps (p. 114) notes that most are "based on the assumption that observing a fact that runs counter to the theory doesn't invalidate the theory in anyone's mind for the rest of the game." This concern has led Kreps to focus on so-called *complete theories*, whereby no action is absolutely precluded, but out-of-equilibrium actions are held to be unlikely *a priori* (see Fudenberg, Kreps, and Levine 1988).

Kreps' final concern is with the mode of equilibrium analysis itself. Again, to quote (p. 139):

Equilibrium analysis is based formally on the presumptions that every player maximizes perfectly and completely against the strategies of his/her opponents, that the character of those opponents and their strategies are perfectly known (or any uncertainty on the part of one player about another player is fully appreciated by all the players and the strategy as a function of the other player's character is also known), and that players are able to evaluate all their options.

The point is that none of these conditions are met fully in reality, and the approximation may be appropriate in some cases but not others.

5. Conclusion

This paper, part 1 of a two-part survey, has reviewed noncooperative game theory concepts that might be used to analyze agricultural markets. In the next issue of this *Review*, part 2 will consider application of the concepts to agricultural markets. The review began by explicating the concept of Nash equilibrium, the cornerstone solution concept in noncooperative game theory. We then proceeded in section 3 to examine refinements of Nash equilibrium to handle dynamic games (subgame perfect equilibrium) and games of asymmetric information (perfect Bayesian equilibrium).

The key deficiencies of game theory, as judged by its leading practitioners, were discussed in section 4.

Despite its deficiencies, noncooperative game theory is certainly in vogue among economists and probably will become even more popular as it integrates fully into graduate curricula. It remains to be seen what role the subject will play in agricultural economics. Through the applications suggested in part 2 of the survey, I hope to show that it has potential to play an important role in research on agricultural markets.

References

- AKERLOF, G. (1970), "The market for lemons: quality uncertainty and the market mechanism," *Quarterly Journal of Economics* 84, 488-500.
- AXELROD, R. (1984), *The Evolution of Cooperation*, New York, Basic Books.
- BANKS, J. and SOBEL, J. (1987), "Equilibrium selection in signaling games," *Econometrica* 55, 647-61.
- CHO, I-K. and KREPS, D.M. (1987), "Signalling games and stable equilibria," *Quarterly Journal of Economics* 102, 179-221.
- DASGUPTA, P. and MASKIN, E. (1986), "The existence of equilibrium in discontinuous economic games, I: theory," *Review of Economic Studies* 53, 1-26.
- FRIEDMAN, J. (1986), *Game Theory with Applications to Economics*, New York, Oxford University Press.
- FUDENBERG, D. and TIROLE, J. (1988), "Perfect Bayesian and sequential equilibrium," Massachusetts Institute of Technology, mimeo.
- FUDENBERG, D. and TIROLE, J. (1989), "Noncooperative game theory for industrial organization: an introduction and overview," in *Handbook of Industrial Organization*, Vol 1, R. Schmalensee and R. Willig (eds.), Amsterdam, North Holland.
- FUDENBERG, D. and TIROLE, J. (1991), *Game Theory*, Cambridge, MA, MIT Press.
- FUDENBERG, D., KREPS, D.M. and LEVINE, D. (1988), "On the robustness of equilibrium refinements," *Journal of Economic Theory* 44, 354-80.
- GROSSMAN, S. and PERRY, M. (1986), "Perfect sequential equilibria," *Journal of Economic Theory* 39, 97-119.
- HARSANYI, J. (1967), "Games with incomplete information played by 'Bayesian' players, I: the basic model," *Management Science* 14, 159-82.
- KOHLBERG, E. and MERTENS, J. (1986), "On the strategic stability of equilibria," *Econometrica* 54, 1003-7.
- KREPS, D.M. (1990a), *Game Theory and Economic Modelling*, New York, Oxford University Press.
- KREPS, D.M. (1990b), *A Course in Microeconomic Theory*, Princeton, Princeton University Press.
- KREPS, D.M. and WILSON, R. (1982a), "Sequential equilibrium," *Econometrica* 50, 863-94.
- KREPS, D.M. and WILSON, R. (1982b), "Reputation and imperfect information," *Journal of Economic Theory* 27, 253-79.
- KREPS, D.M., WILSON, R., MILGROM, P. and ROBERTS, J. (1982), "Rational cooperation in the finitely repeated prisoners' dilemma," *Journal of Economic Theory* 27, 245-52.
- LUCE, R.D. and RAIFFA, H. (1957), *Games and Decisions: Introduction and Critical Survey*, New York, John Wiley.
- MCKELVEY, R.D. and PALFREY, T.R. (1992), "An experimental study of the centipede game," *Econometrica* 60, 803-36.
- MCLENNAN, A. (1985), "Justifiable beliefs in sequential equilibrium," *Econometrica* 53, 889-904.
- MILGROM, P. and ROBERTS, J. (1982a), "Limit pricing and entry under incomplete information: an equilibrium analysis," *Econometrica* 50, 443-59.
- MILGROM, P. and ROBERTS, J. (1982b), "Predation, reputation, and entry deterrence," *Journal of Economic Theory* 27, 280-312.
- NASH, J. (1950), "The bargaining problem," *Econometrica* 18, 155-62.
- NASH, J. (1951), "Non-cooperative games," *Annals of Mathematics* 54, 286-95.
- RASMUSEN, E. (1989), *Games and Information: An Introduction to Game Theory*, Oxford, Basil Blackwell.
- ROSENTHAL, R. (1981), "Games of perfect information, predatory pricing, and the chain-store paradox," *Journal of Economic Theory* 25, 92-100.
- RUBINSTEIN, A. (1991), "Comments on the interpretation of game theory," *Econometrica* 59, 909-24.
- SCHELLING, T. (1960), *The Strategy of Conflict*, Cambridge, MA, Harvard University Press.
- SCHMALENSEE, R. and WILLIG, R. (eds.) (1989), *Handbook of Industrial Organization*, Amsterdam, North Holland.
- SELTON, R. (1975), "Re-examination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory* 4, 25-55.
- SELTON, R. (1978), "The chain-store paradox," *Theory and Decision* 9, 127-59.
- SHAPLEY, L. (1953), "A value for n-person games," in *Contributions to the Theory of Games*, Vol. I, H. Kuhn and A. Tucker (eds.), Annals of Mathematics Studies, No. 28. Princeton, Princeton University Press.

- SHUBIK, M. (1982), *Game Theory in the Social Sciences: Concepts and Solutions*, Cambridge, MA, MIT Press.
- SHUBIK, M. (1984), *A Game-Theoretic Approach to Political Economy*, Cambridge, MA, MIT Press.
- SPENCE, A.M. (1973), "Job market signalling," *Quarterly Journal of Economics* 87, 355-74.
- SUTTON, J. (1990), "Explaining everything, explaining nothing? Game theoretic models in industrial economics," *European Economic Review* 34, 505-12.
- TIROLE, J. (1988), *The Theory of Industrial Organization*, Cambridge, MA, MIT Press.
- VAN DAMME, E. (1987), *Stability and Perfection of Nash Equilibrium*, Berlin, Springer-Verlag.
- VON NEUMANN, J. and MORGENSTERN, O. (1944), *The Theory of Games and Economic Behavior*, New York, John Wiley.