

Noncooperative Game Theory
For Industrial Organization:
An Introduction and Overview

by
Drew Fudenberg
and
Jean Tirole

Number 445

April 1987



Digitized by the Internet Archive
in 2011 with funding from
MIT Libraries

<http://www.archive.org/details/noncooperativega445fude>

NONCOOPERATIVE GAME THEORY
FOR INDUSTRIAL ORGANIZATION:
AN INTRODUCTION AND OVERVIEW*

Drew Fudenberg
University of California, Berkeley

Jean Tirole
Massachusetts Institute of Technology

*Forthcoming in Handbook of Industrial Organization, R. Schmalensee and R. Willig, editors. The authors would like to thank Freddie Kofman, Jennifer Reinganum, John Roberts, and the editors for helpful comments on a previous draft. Financial support from NSF Grants SES 85-09697 and SES 97619 is gratefully acknowledged.

Introduction

Non-cooperative game theory is a way of modelling and analyzing situations in which each player's optimal decisions depend on his beliefs or expectations about the play of his opponents. The distinguishing aspect of the theory is its insistence that players should not hold arbitrary beliefs about the play of their opponents. Instead, each player should try to predict his opponents' play, using his knowledge of the rules of the game and the assumption that his opponents are themselves rational, and are thus trying to make their own predictions and to maximize their own payoffs. Game-theoretic methodology has caused deep and wide-reaching changes in the way that practitioners think about key issues in oligopoly theory, much as the idea of rational expectations has revolutionized the study of macroeconomics. This essay tries to provide an overview of those aspects of the theory which are most commonly used by industrial organization economists, and to sketch a few of the most important or illuminating applications. We have omitted many interesting game-theoretic topics which have not yet been widely applied.

1. Games, Strategies, and Equilibria

This section introduces the two formalisms used to represent noncooperative games, and then discusses what we might mean by a "reasonable prediction" for how a game will be played. This will lead us to the ideas of Nash and subgame-perfect equilibria.

The Extensive and Normal Forms

There are two (almost) equivalent ways of formulating a game. The first is the extensive form.¹ An extensive form specifies: (1) the order of play; (2) the choices available to a player whenever it is his turn to move; (3) the information a player has at each of these turns; (4) the payoffs to each player as a function of the moves selected; and (5) the probability distributions for moves by "nature."

The extensive form is depicted by a "game tree," such as those in Figures 1 and 2. Game trees are the multi-player generalization of the decision trees used in decision theory. The open circle is the first or initial node. The tree's structure says which nodes follow which, and the numbers at each node indicate which player has the move there. (Part of what is meant by "tree" is that this structure is an ordering--two distinct nodes cannot have the same successor. Thus for example in chess, two different sequences of moves which lead to the same position on the board are assigned different nodes in the tree. See Kreps-Wilson (1982b) for a more formal discussion of this and other details of extensive games. See also the classic book by Luce and Raiffa (1957) which addresses most of the topics of this section.) The dotted line connecting two of player two's nodes indicate that these two nodes are in the same "information set," meaning that player two cannot tell which of the two actions has occurred when it is his turn to move. Players must know when it is their turn to move, so different players' information sets cannot intersect, and players must know which choices are feasible, so all nodes in the same information set must allow the same choices. We will restrict attention throughout to games of perfect recall, in which each player always knows what he knew previously, including

¹The following description is freely adapted from Kreps-Wilson (1982b).

his own previous actions. This implies an additional restriction on the information sets.

Players are assumed to maximize their expected utility, given their beliefs about the actions of their opponents and of "Nature." The payoffs corresponding to each sequence of actions are depicted at the terminal nodes or "outcomes" of the tree; (x,y) at a terminal node means that player one gets x and player two gets y . The different initial nodes in Figure 3 correspond to different moves by Nature, i.e. different "states of the world." (Note that this is a one-player game.) There is no loss in generality in placing all of Nature's moves at the start, because players need not receive information about these moves until later on. The initial assessment p is a probability measure over the initial nodes. The formal models we will discuss will always assume that this assessment, the terminal payoffs, and the entire structure of the tree is "common knowledge," meaning that all players know it, and they know that their opponents know it, and so on. This does not mean that all players are perfectly informed, but rather that we have explicitly depicted all the differences in information in our tree.² The extensive form will be taken to fully describe the real situation--all possible moves and observations will be explicitly specified. For example, if the "same game" is played three times, the "real game" to be analyzed is the three-fold replication. The idealized situation we have in mind is that, possibly after some "pre-play communication," players are in separate rooms. They are informed of the course of play only by signals corresponding to the information structure of the tree, and push various buttons corresponding to the feasible actions at their various information sets. Once

² See Aumann (1976) and Bradenburger-Dekel (1985) for a formal treatment of common knowledge, and also the Mertens-Zamir (1985) paper we mention in Section 3.

play begins, players cannot explicitly communicate, except as provided by the rules of the game. (In many situations, it is difficult to explicitly model all the possible means of communication. This has spurred interest in shorthand descriptions of the effects of communication. See our discussion of correlated equilibrium.) A behavioral strategy for player i is a map that specifies for each of his information sets, a probability distribution over the actions that are feasible at that set. A behavioral pure strategy specifies a single action at each information set, as opposed to a probability mixture. (Later we will discuss whether it might be reasonable for a player to randomize.) A given specification of behavioral strategies and an initial assessment generates a probability distribution over terminal nodes, and thus over payoffs, in the obvious way.

The distinguishing feature of game theory is that each player's beliefs about his opponents actions are not arbitrarily specified. Instead, each player is assumed to believe that his opponents are "rational", and to use that information in formulating his predictions of their play. Any predictions that are inconsistent with this presumed, but vaguely specified, rationality are rejected.

To help clarify what we mean, let us return to the game depicted in Figure 1. Is there a reasonable prediction for how this game should/will be played? One way to look for a prediction is to apply backwards induction. If player two's information set is reached, and the payoffs are as specified, then two should play L . Then if player one knows that player two knows the payoff, player one should play U . Is this a good prediction? If all is as in Figure 1, player two should not expect player one to play D . What should two tell himself if D is nevertheless observed? If the payoffs are guaranteed to be as specified, the only possible explanation is that player one made a

"mistake"--he meant to play U but somehow he failed to do so. This analysis falls apart if we take Figure 1 as a shorthand description for a game which is probably as depicted, but might not be, so that playing D could convey information to player two. We'll say more about this in Section 3. The key for now is that the game must be taken as an exact description of reality for our arguments to be sound.

In Figure 1, all (both) the information sets are singletons, so that each player knows all previous actions at each of his turns to move. Games like this are called "games of perfect information." The backwards induction argument used above is called "Kuhn's algorithm" (1953). It always "works" (yields a conclusion) in finite games of perfect information, and yields a unique conclusion as long as no two terminal nodes give any player exactly the same payoff. Backwards induction will not yield a conclusion in games of imperfect information, such as that in Figure 2. Player two's optimal choice at his information set depends on player one's previous move, which player two has not observed. To help find a reasonable prediction for this game we introduce the idea of the normal form.

The normal form representation of an extensive game condenses the details of the tree structure into three elements: The set of players, I ; each player's strategy space, which is simply the set of his behavioral pure strategies; and a payoff function mapping strategy selections for all of the players to their payoffs. We will use S_i to denote player i 's strategy space, S to be the product of the S_i , and $\pi^i: S \rightarrow R$ to be player i 's payoff function. A triple (I, S, π) completely describes a normal form.

Normal forms for two-player games are often depicted as matrices, as in Figure 4. The left-hand matrix is the normal form for Figure 1, while the right-hand one corresponds to Figure 2. Note that different extensive forms

can have the same normal form. For example, Figure 2 is a "simultaneous-move" game, in which neither player observes his opponent's action before choosing his own. We could represent this game equally well with an extensive form in which player two moved first.

A mixed strategy is a probability distribution over the normal-form strategies. Payoffs to mixed strategies are simply the expected value of the corresponding pure-strategy payoffs. We will denote mixed strategies by Σ , and the space of player i 's mixed strategies by Σ_i . Although different mixed strategies can give rise to the same behavior strategies, Kuhn showed that the two concepts are equivalent in games of perfect recall--any probability distribution over outcomes that can be generated using one kind of randomization can be duplicated by using the other.³

In the normal form corresponding to Figure 1, choosing L gives player two at least as high a payoff as choosing R regardless of player one's choice, and gives strictly more if player one plays D. In such a case we say that D is a (weakly) dominant strategy for player two. (Strict dominance means that the strategy is strictly better for all choices by opponents.) It seems reasonable that no player should expect an opponent to play a dominated strategy, which means that one should expect that two will play L. This is just rephrasing our backwards induction argument. The analog of rolling backwards through the tree is the iterated elimination of dominated strategies: making optimal choices at the last nodes is simple dominance, folding back one step is first-order iterated dominance, and so on. (Actually iterated dominance

³ Two strategies for a player which differ only at information sets which follow a deviation by that player yield the same probability distribution over outcomes for any strategy selections of the other players. Some authors define the normal form as identifying such equivalent strategies.

is a more general technique, as it can be applied to games of imperfect information.)

Nash Equilibrium

The normal form for Figure 2 does not have dominant strategies. Here to make predictions we will have to accept a weaker notion of "reasonableness," that embodied in the concept of a Nash equilibrium. A Nash equilibrium is a strategy selection such that no player can gain by playing differently, given the strategy of his opponent. This condition is stated formally as

Definition: Strategy selection s^* is a pure-strategy Nash equilibrium of the game (I, S, π) if for all players i in I and all s_i in S_i ,

$$(1) \quad \pi^i(s^*) \geq \pi^i(s_i, s^*_{-i}) .$$

Here, the notation (s^*_{-i}, s_i) represents the strategy selection in which all players but i play according to s^* , while i plays s_i . Note that s^* can be an equilibrium if there is some player i who is indifferent between s^*_i and an alternative, s_i . We view Nash equilibrium as a minimal requirement that a proposed solution must satisfy to be "reasonable." If a strategy selection is not a Nash equilibrium, then all players know that some player would do better not to play as the selection specifies. If "reasonable" is to mean anything, it should rule out such inconsistent predictions. Not all Nash equilibria are reasonable, as is revealed by examining the extensive and normal forms of Figure 5. The backwards-induction equilibrium (D, L) is a Nash equilibrium, but so is (U, R) . We will soon discuss the idea of a "perfect equilibrium," which is designed to formalize the idea that (U, R) is not

reasonable. The perfection notion and other refinements of Nash equilibrium do not help with the following problem. Consider a game like that in Figure 6. The only Nash equilibrium is (U,L), yet is this a reasonable prediction? It depends on whether the players are sure that the payoffs are exactly as we've specified, and that their opponents are "rational." If one plays U against L, his payoff is 5, which is better than the 4.9 that one gets from D. However, playing D guarantees that 1 gets 4.9, while if the outcome is (U,R) then one gets 0. And similarly, player two can guarantee 4.9 by playing R. Yet if player one's not sure that player two might not prefer R to L, then D could be attractive. And even if player one is sure of player two's payoffs, if player one's not sure that player two knows player one's payoffs, then player one might still fear that player two will play R. The point is that the logic of Nash equilibrium relies on every player knowing that every player knows that ... the payoffs are as specified. Technically, the payoffs should be "common knowledge," (as should the Nash concept itself.) The closer the payoffs guaranteed by D and R come to the equilibrium payoffs, the more we need to insist on the common knowledge. Ideally, equilibria should be subjected to this sort of informal check or "sensitivity analysis."

Returning to Figure 2, the game there has two pure strategy equilibria, (U,L) and (D,R). If there is a reasonable outcome in this game, both players must be able to predict it, and predict that their opponents will predict it, and so on. If players cannot so coordinate their expectations, there is no reason to expect observed play to correspond to either equilibrium--for example, we might see the outcome (U,R). Not all games have reasonable solutions, and on the data given so far this could be one. However, Schelling's (1960) theory of "focal points" suggests that in some "real life" situations players may be able to coordinate on a particular equilibrium by using information that is

abstracted away in the standard game formulation. For example, the names of the strategies may have some commonly-understood "focal" power. An example is two players who are asked to name an exact time, with the promise of a reward if their choices match. Here "12 noon" is focal, while "1:43" is not. The payoffs may also help coordinate expectations. If both players did better with (U,L) than (D,R), then (U,L) seems a natural outcome to expect one's opponent to expect that.... Some authors (including us!) have argued that if there is a unique Pareto optimum among the set of equilibria, it should be a focal point. While this intuition seems sound for two-player games, a recent example of Bernheim, Peleg and Whinston shows that with more than two players the intuition is suspect. In response, they have introduced the concept of "coalition-proofness", which we discuss at the end of this section.

The idea of a Nash equilibrium is implicit in two of the first games to have been formally studied, namely the Cournot and Bertrand models of oligopoly. Let us emphasize that despite the common practice of speaking of Cournot and Bertrand equilibrium, the models are best thought of as studying the Nash equilibria of two different simultaneous move games. In the Cournot model, firms simultaneously choose quantities, and the price is set at the market-clearing level by a fictitious auctioneer. In the Bertrand model, firms simultaneously choose prices, and then must produce to meet demand after the price choices become known. In each model, firms choose best responses to the anticipated play of their opponents.

For concreteness, we remind the reader of the Cournot model of a duopoly producing a homogeneous good. Firm 1 and Firm 2 simultaneously choose their respective output levels, q_1 and q_2 from feasible sets F_i . They sell their output at the market-clearing price $p(Q)$, where $Q = q_1 + q_2$. Firm i 's cost of production is $c_i(q_i)$, and firm i 's total profit is then $\pi^i(q_1, q_2)$.

$q_2) = q_i p(Q) - c_i(q_i)$. The feasible sets F_i and the payoff functions π^i determine the normal form of the game; the reader should check that he/she knows how to construct an equivalent extensive form. The "Cournot reaction functions" $R^1(q_2)$ and $R^2(q_1)$ specify each firm's optimal output for each fixed output level of its opponent. If the π^i are differentiable and strictly concave, and the appropriate boundary conditions are satisfied, we can solve for these reaction functions using the first-order conditions. The intersections of the two reaction functions (if any exist) are the Nash equilibria of the Cournot game: neither player can gain by a change in output, given the output level of its opponent.

The Cournot game is often contrasted to the situation in which one firm, say firm one, is a "Stackelberg leader" and the other firm is the "Stackelberg follower." The Stackelberg leader moves first, and chooses an output which is observed by the follower before the follower makes its own choice. Thus the Stackelberg game is one of perfect information. In the backwards induction (i.e. "perfect" - see page 17) equilibrium to this game, firm two's output is along its reaction curve. Knowing this, firm one chooses its own output to maximize its payoff along the graph of R^2 . The first-order condition for this choice is that

$$\partial \pi^1(R^1(q_2), q_2) / \partial q_1 + (\partial \pi^1(R^1(q_2), q_2) / \partial q_2) dR^2(q_1) / dq_1 = 0 .$$

The backwards-induction equilibrium to the Stackelberg game is called the "Stackelberg equilibrium." This terminology can be confusing to the beginner. The Stackelberg equilibrium is not an alternative equilibrium for the Cournot game, but rather a shorthand way of describing an equilibrium of an alternative

extensive form. While the prevailing terminology is too well established to be changed, the student will do well to keep this distinction in mind.

The Cournot and Bertrand models are all static games, in which firms make their choices once and for all. Section 2A discusses a dynamic version of these games. Also, even as static games the Cournot and Stackelberg models must be thought of as reduced forms, unless one literally believes in the existence of the price-setting auctioneer. Kreps-Scheinkman (1983) have shown that the auctioneer in the Cournot model can be replaced by a second period in which firms choose prices, taking their production as fixed (at least if the rationing scheme is "efficient" and the demand function is concave). Thus in both models firms choose both prices and outputs; the difference is in the timing of these two decisions. (See Gertner (1985a) for simultaneous choices.)

Existence of Nash Equilibria

We will now take up the question of the existence of Nash equilibria. Not all games have pure-strategy Nash equilibria. A simple example is "matching pennies": players one and two simultaneously announce either "heads" or "tails." If the announcements match, then player one gains a util, and player two loses one. If the announcements differ, it is player two who wins the util, and player one who loses. If the predicted outcome is that the announcements will match, then player two has an incentive to deviate, while player one would prefer to deviate from any prediction in which announcements do not match. The only "stable" situation is one in which each player randomizes between his two strategies, assigning equal probability to each. In this case each player is completely indifferent between his possible choices. A mixed-strategy Nash equilibrium is simply a selection of mixed strategies such that no player prefers to deviate, i.e. the strategies must satisfy equation (1). Since

expected utilities are "linear in the probabilities," if a player uses a non-degenerate mixed strategy (one that puts positive weight on more than one pure strategy) then that player cannot strictly prefer not to deviate--the inequality in (1) must be weak. (For the same reason, it suffices to check that no player has a profitable pure-strategy deviation.) This raises the question of why a player should bother to play a mixed strategy, when he knows that any of the pure strategies in its support would do equally well. In matching pennies, if player one knows that player two will randomize, player one has a zero expected value from all possible choices. As far as his payoff goes, he could just as well play "heads" with certainty, but if this is anticipated by player two the equilibrium disintegrates. Some authors have suggested that for this reason there is no "reasonable" prediction for matching pennies, or, equivalently, that all possible probability mixtures over outcomes are equally reasonable. (See e.g. Bernheim (1984) and Pearce (1984).) Harsanyi (1973) followed by Aumann et al. (1981) and Milgrom-Weber (1986) have offered the defense that the "mixing" should be interpreted as the result of small, unobservable variations in the player's payoffs. Thus in our example, sometimes player one might prefer matching on T to matching on H, and conversely. Then for each value of his payoff player one would play a pure strategy. This "purification" of mixed-strategy equilibria is discussed in Section 3C. Despite some controversy, mixed strategies have been widely used both in "pure" game theory and in its applications to industrial organization.

One reason is that, as shown by Nash (1950), mixed-strategy equilibria always exist in finite games (games with a finite number of nodes, or, equivalently, a finite number of normal-form pure strategies per player and a finite number of players.)

Theorem (Nash): Every finite n-player normal form game has a mixed-strategy equilibrium.

This can be shown by applying the Kakutani fixed-point theorem to the player's reaction correspondences, as we now explain. A good reference for some of the technical details involved is Green-Heller (1981).

Define player i 's reaction correspondence, $r^i(\sigma)$, to be the correspondence which gives the set of (mixed) strategies which maximize player i 's payoff when his opponents play σ_{-i} . This is just the natural generalization of the Cournot reaction functions we introduced above. Since payoffs are linear functions of the mixing probabilities, they are in particular both continuous and quasiconcave. This implies that each player's reaction correspondence is non-empty valued and convex-valued. Moreover, we can show that the reaction correspondences are "upper hemi-continuous": if $\sigma^n \rightarrow \sigma$ and $\sigma_i^n \in r^i(\sigma^n)$, then there is a subsequence of the σ_i^n which converges to a $\sigma_i \in r^i(\sigma)$. Now define the correspondence r to be the Cartesian product of the r_i . This correspondence satisfies the requirements of the Kakutani fixed-point theorem: it maps a compact convex subset of Euclidean space (the relevant probability simplex) into its subsets, and it is non-empty valued, convex-valued, and upper hemi-continuous. Hence r has a fixed point, and by construction the fixed points of r are Nash equilibria.

Economists often use models of games with an uncountable number of actions. Some might argue that prices or quantities are "really" infinitely divisible, while others that "reality" is discrete, and the continuum is a mathematical abstraction, but it is often easier to work with a continuum of actions rather than a large finite grid. Moreover, as Dasgupta-Maskin (1986) argue, when the continuum game does not have an equilibrium, the equilibria corresponding to

fine, discrete grids could be very sensitive to exactly which finite grid is specified. These fluctuations can be ruled out if the continuum game has an equilibrium. The existence of equilibria for infinite games is more involved than for finite ones. If payoffs are discontinuous there may be no equilibria at all. If the payoffs are continuous, then the Fan (1952) fixed-point theorem can be used to show that a mixed-strategy equilibrium exists. If payoffs are quasiconcave as well as continuous, then there exist equilibria in pure strategies, as shown by Debreu (1952) and Glicksberg (1952).

Theorem (Debreu, Glicksberg, Fan): Consider an n -player normal form game whose strategy spaces S_i are compact convex subsets of an Euclidean space. If the payoff functions $\pi^i(s)$ are continuous in s , and quasiconcave in s_i , there exists a pure-strategy Nash equilibrium.

The proof here is very similar to that of Nash's theorem: we verify that continuous payoffs imply non-empty, upper hemi-continuous reactions, and that quasiconcavity in own actions implies that reactions are convex-valued.

Theorem (Glicksberg): Consider an n -player normal form game (I, S, π) . If for each i , S_i is a compact convex subset of a metric space, and π is continuous, then there exists a Nash equilibrium in mixed strategies.

Here the mixed strategies are the (Borel) probability measures over the pure strategies, which we endow with the topology of weak convergence.⁴ Once more,

⁴Fix a compact metric space A . A sequence of measures μ^n on A converges "weakly" to a limit μ if $\int f d\mu_n \rightarrow \int f d\mu$ for every real-valued continuous function on A .

the proof applies a fixed-point theorem to the reaction correspondences. One point to emphasize is that the mixed-strategy payoffs will be quasiconcave in own actions even if the pure-strategy payoffs are not. With infinitely many pure strategies, the space of mixed strategies is infinite-dimensional, so a more powerful fixed-point theorem is required. Alternatively, one can approximate the strategy spaces by a sequence of finite grids. From Nash's theorem, each grid has a mixed-strategy equilibrium. One then argues that since the space of probability measures is weakly compact, we can find a limit point of the sequence of these discrete equilibria. Since the payoffs are continuous, it is easy to verify that the limit point is an equilibrium.

There are many examples to show that if payoffs are discontinuous equilibria need not exist. Dasgupta-Maskin argue that this lack of existence is sometimes due to payoffs failing to be quasiconcave, rather than failing to be continuous. They show if payoffs are quasiconcave, then a pure strategy equilibrium will exist under a very weak condition they call "graph continuity." They also provide conditions for the existence of mixed-strategy equilibria in games without quasiconcave payoffs. The idea of their result is to provide conditions ensuring that the limits of the discrete-grid equilibria do not have "atoms" (non-negligible probability) on any of the discontinuity points of the payoff functions. Simon (1985) relaxes their condition by requiring only that at least one limit has this no-atoms property, instead of all of them.

A sizable literature has considered the existence of pure strategy equilibrium when payoffs are not quasiconcave, particularly in the Cournot model. Without quasiconcave payoffs, the reaction functions can have "jumps." To prove existence of equilibrium in this setting one must show that the jumps "do not matter." Roberts-Sonnenschein (1977) showed that "nice" preferences and technologies need not lead to quasiconcave Cournot payoffs, and provided

examples of the non-existence of pure-strategy Cournot equilibrium. McManus (1962) and Roberts-Sonnenschein (1976) show that pure strategy equilibria exist in symmetric games with real-valued actions if costs are convex. The key is that the convex-cost assumption can be shown to imply that all the jumps in the reaction functions are jumps up. Novshek (1985) has shown that pure-strategy equilibria exist in markets for a homogeneous good where each firm's marginal revenue is decreasing in the aggregate output of its opponents, for any specification of the cost functions. Topkis (1970) and Vives (1985) use a fixed-point theorem for non-decreasing functions due to Tarski (1955) to prove the existence of pure-strategy equilibria in games where the reactions are increasing. Tarski also proved that a function from $(0,1)$ to $(0,1)$ which has no downward jumps has a fixed point, even if the function is not everywhere non-decreasing. Vives uses this result to give a simple proof of the McManus/Roberts-Sonnenschein result. (In symmetric equilibria each firm's reaction function depends only on the sum of its opponents actions, and all firms have the the same reaction function. Thus if the actions are real-valued the second of the Tarski results can be applied.)

The converse of the existence question is that of the characterization of the equilibrium set. Ideally one would prefer there to be a unique equilibrium, but this is only true under very strong conditions. When several equilibria exist, one must see which, if any, seem to be reasonable predictions, but this requires examination of the entire Nash set. The reasonableness of one equilibrium may depend on whether there are others with competing claims. Unfortunately, in many interesting games the set of equilibria is difficult to characterize.

Correlated Equilibria

The Nash equilibrium concept is intended to be a minimal necessary condition for "reasonable" predictions in situations where the players must choose their actions "independently." Let us return to our story of players who may have pre-play discussion, but then must go off to isolated rooms to choose their strategies. In some situations, both players could gain if they could build a "signalling device" that sent signals to the separate rooms. Aumann's (1974) notion of a correlated equilibrium captures what could be achieved with any such signals. (See Myerson (1983) for a fuller introduction to this concept, and for a discussion of its relationship to the theory of mechanism design.)

To motivate this concept, consider Aumann's example, presented in Figure 7. This game has three equilibria: (U,L), (D,R), and a mixed-strategy equilibrium that gives each player 2.5. If they can jointly observe a "coin flip" (or sunspots, or any other publicly observable random variable) before play, they can achieve payoffs (3,3) by a joint randomization between the two pure-strategy equilibria. However, they can do even better (still without binding contracts) if they can build a device that sends different, but correlated, signals to each of them. This device will have three equally likely states, A, B, and C. Player one's information partition is (A,(B,C)). This means that if A occurs, player one is perfectly informed, but if the state is B or C, player one does not know which of the two prevails. Player two's information partition is ((A,B),C). In this transformed game, the following is a Nash equilibrium: player one plays U when told A, and D when told (B,C); player two plays R when told C, and L when told (A,B). Let's check that player one does not want to deviate. When he observes A, he knows that two observes (A,B), and thus that two will play L; in this case U is player

one's best response. If player one observes (B,C), then conditional on his information he expects player two to play L and R with equal probability. In this case player one will average 2.5 from either of his choices, so he is willing to choose D. So player one is choosing a best response; the same is easily seen to be true for player two. Thus we have constructed an equilibrium in which the players' choices are correlated: the outcomes (U,L), (D,L), and (D,R) are chosen with probability one-third each, while the "bad" outcome (U,R) never occurs. In this new equilibrium the expected payoffs are $3 \frac{1}{3}$ each, which is better than in any of the equilibria of the game without the signalling device. (Note that adding the signalling device does not remove the "old" equilibria: since the signals do not influence payoffs, if player one ignores his signal, player two may as well ignore hers.)

If we had to analyze each possible signalling device one at a time, we would never be done. Fortunately, if we want to know what could be done with all possible devices, we can dispense with the signals, and work directly with probability distributions over strategies. In our example, players need not be told about the states A, B, and C. They could simply be given recommended strategies, as long as the joint distribution over recommendations corresponds to the joint distribution over outcomes that we derived. Player one could be told "play D" instead of (B,C), as long as this means there's a 50-50 chance of player two playing L.

Definition: A correlated equilibrium is any probability distribution $p(s)$ over the pure strategies $S_1 \times \dots \times S_n$ such that, for every player i and every function $d_i(s_i)$ that maps S_i to S_i ,

$$\pi^i(p) \geq \sum p(s) \pi^i(d_i(s_i), s_{-i}) .$$

That is, player i should not be able to gain by disobeying the recommendation to play s_i if every other player obeys the recommendations.

A pure-strategy Nash equilibrium is a correlated equilibrium in which the distribution $p(s)$ is degenerate. Mixed-strategy Nash equilibria are also correlated equilibria: just take $p(s)$ to be the joint distribution over actions implied by the equilibrium strategies, so that the recommendations made to each player convey no information about the play of his opponents.

Inspection of the definition shows that the set of correlated equilibria is convex, so the set of correlated equilibria is at least as large as the convex hull of the Nash equilibria. Since Nash equilibria exist in finite games, correlated equilibria do too. Actually, the existence of correlated equilibria would seem to be a simpler problem than the existence of Nash equilibria, because the set of correlated equilibria is defined by a system of linear inequalities, and is therefore convex. Recently, Hart and Schmeidler (1986) have provided an existence proof that uses only linear methods (as opposed to fixed-point theorems.) One might also like to know when the set of correlated equilibria differs "greatly" from the convex hull of the Nash equilibria, but this question has not yet been answered.

We take the view that the correlation in correlated equilibria should be thought of as the result of the players receiving correlated signals, so that the notion of correlated equilibrium is particularly appropriate in situations with pre-play communication, for then the players might be able to design and implement a procedure for obtaining correlated, private signals. However, we should point out that Aumann (1986) and Brandenburger-Dekel (1985b) argue that

the correlated equilibrium notion is more "natural" than the Nash one from the point of view of subjective probability theory.

Coalition-Proof Equilibria and Strong Equilibria

While no single player can profitably deviate from a Nash equilibrium, it may be that some coalition could arrange a mutually beneficial deviation. If players can engage in pre-play communication, then some coalitions of players might hope to arrange for joint deviations from the specified play. The notion of a "strong equilibrium" (Aumann (1959)) requires that no subset of players, taking the actions of the others as given, could jointly deviate in a way that benefits all of its members. As this requirement applies to the grand coalition of all players, strong equilibria are Pareto-efficient. Because no restrictions are placed on the play of a deviating coalition, the conditions for a strong equilibrium are quite stringent, and these equilibria fail to exist in many games of interest for industrial organization, such as, for example, Cournot oligopoly. Recently, Bernheim, Peleg, and Whinston (1986) (B-P-W) have proposed the idea of a "coalition-proof" equilibrium, which, they argue, is a more natural way to take account of coalitional deviations.

The best way to explain their concept is to use their example, which also serves the important function of showing why the criterion of Pareto-dominance may not be a good way to select between equilibria when there are more than two players. In Figure 10, player one chooses rows, player two chooses columns, and player three chooses matrices. This game has two pure-strategy Nash equilibria, (U,L,A) and (D,R,B), and an equilibrium in mixed strategies. B-P-W do not consider mixed strategies, so we will temporarily restrict attention to pure ones. The equilibrium (U,L,A) Pareto-dominates (D,R,B). Is (U,L,A) then the obvious focal point? Imagine that this was the expected solution, and hold

player three's choice fixed. This induces a two-player game between players one and two. In this two-player game, (D,R) is the Pareto-dominant equilibrium! Thus, if players one and two expect that player three will play A, and if they can coordinate their play on their Pareto-preferred equilibrium in matrix A, they should do so, which would upset the "good" equilibrium (U,L,A).

The definition of a coalition-proof equilibrium proceeds by induction on the coalition size. First one requires that no one-player coalition can deviate, i.e. that the given strategies are a Nash equilibrium. Then one requires that no two-player deviation can deviate, given that once such a deviation has "occurred", either of the deviating players (but none of the others) is free to deviate again. That is, the two-player deviations must be Nash equilibria of the two-player game induced by holding the strategies of the others fixed. And one proceeds in this way up to the coalition of all players. Clearly (U,L,A) in Figure 10 is not coalition-proof; brief inspection shows that (D,R,B) is. However, (D,R,B) is not Pareto-optimal, and thus is not a strong equilibrium; no strong equilibrium exists in this game.

The idea of coalition-proofness is an interesting way to try to model the possibility of coalitional deviations. However, the assumption that only subsets of the deviating coalitions can be involved in further deviations can be questioned, and the general properties of the concept are unknown. For these reasons, and because coalition-proof equilibria need not exist (even with mixed strategies), we feel that at this time the B-P-W paper is more important for the issues it raises than for its solution concept. We should mention here that Bernheim-Whinston (1986) apply coalition-proofness to several well-known games with interesting results.

2. Dynamic Games of Complete Information

Most of the examples in the last section were static games: each player's choice of actions was independent of the choices of his opponents. Many of the interesting strategic aspects of the behavior of firms are best modelled with dynamic games, in which players can observe and respond to their opponents' actions. This is true not only of inherently dynamic phenomena such as investment, entry deterrence, and exit, but also of the determination of price and output in a mature market. Section 2 discusses a few special kinds of dynamic games that have been frequently used in the study of oligopoly theory. These are all games of complete information, i.e. the payoff functions are common knowledge. Section 3 discusses games of incomplete information, which have become increasingly common in the literature.

Subgame Perfection

In dynamic games a question arises that is not present in static ones: What beliefs should players have about the way that their current play will affect their opponents' future decisions? Recall that the game in Figure 1 had two Nash equilibria, (D,R) and (U,L). We argued that (U,L) was unreasonable, because L was dominated by R for player two. Alternatively, we arrived at (D,R) as our prediction by working backwards through the tree. Another way of putting this is that player one should not be deterred from playing D by the "threat" of player two playing L, because if player two's information set was actually reached, two would back off from his "bluff" and play D. This approach is useful for thinking about situations in which backwards induction and/or weak dominance arguments do not give sharp conclusions. Selten's (1965) notion of a subgame-perfect equilibrium generalizes the backwards-induction idea to rule out empty threats in more general situations.

Subgame-perfect equilibrium strategies must yield a Nash equilibrium, not just in the original game, but in every one of its "proper subgames." We'll define this more formally in Section 3, but for now think of a proper subgame as a subset of the initial game tree which: 1) is closed under succession--if a node is in the subgame, so are all of its successors; 2) "respects information sets" which means roughly that all of the information sets of the subgame are information sets of the initial game; and 3) begins with an information set that contains only one node. This last requirement is in a general sense very restrictive, which is one of the motivations for the various refinements of the perfection concept. However, most of the games we discuss in this section are "deterministic multi-period games," which have a very simple structure that makes subgame-perfection a useful tool. These games have extensive forms that can be divided into periods so that: (1) at the start of the k th period all play in periods 1 through $(k-1)$ is common knowledge, (the initial information sets in each period are all singletons); and (2) no information set contained in the k th period provides any knowledge of play within that period. Any game of perfect information is a multi-period game: just take all the successors of the initial nodes to belong to period 1, their successors to period 2, and so on. The Cournot and Bertrand models are 1-period games. If the same players play a Cournot game twice in a row, and all players observe the "first-period" quantities before making their second choice, we have a two-period game.

In a multi-period game, the beginning of each period marks the beginning of a new subgame. Thus for these games we can rephrase subgame-perfection as simply the requirement that the strategies yield a Nash equilibrium from the start of each period.

Figure 5 is actually the game Selten used to introduce subgame perfection. Here there are two proper subgames: the whole game, and the game beginning

in the second "period" if one played D . In this subgame, the only Nash equilibrium is for player two to choose L , so that any subgame perfect equilibrium must prescribe this choice, and only (D,L) is subgame-perfect. More generally, in any game of perfect information subgame perfection yields the same answer as backwards induction. In finite-period simultaneous move games, subgame-perfection does "backwards induction" period by period: at the last period, the strategies must yield a Nash equilibrium, given the history. Then we replace the last period with the possible last-period equilibria, and work backwards. For example, a subgame-perfect equilibrium of a two-period Cournot model must yield Cournot equilibrium outputs in the second period, regardless of first-period play. Caution: if there are several Cournot equilibria, then which of them prevails in the second period can depend on first-period play. We will say more about this when we discuss Benoit-Krishna (1985).

2A. Repeated Games and "Implicit Collusion"

Infinitely Repeated Games

Chamberlin (1956) criticized the Cournot and Bertrand models of oligopoly for assuming that firms were myopic. He argued that in an industry with few, long-lived firms, firms would realize their mutual interdependence and thus play more "cooperatively" than the Cournot and Bertrand models suggested. The theory of repeated games provides the simplest way of thinking about the effects of long-term competition.

This theory shows that, under the proper circumstances, Chamberlin's intuition can be partially formalized. Repetition can allow "cooperation" to be an equilibrium, but it does not eliminate the "uncooperative" static equilibria, and indeed can create new equilibria which are worse for all players

than if the game had been played only once. Thus to complete the Chamberlin argument, one must argue that the "cooperative" equilibria are "reasonable."

In an infinitely repeated game, players face the same constituent game in each of infinitely many periods. There is no direct physical link between the periods; each period's feasible actions and per-period payoffs are exactly as in the constituent game. This rules out important phenomena such as investment in productive machinery, so few interesting industries can be modelled exactly as repeated games. Nevertheless, if the history-dependent aspects of the industry are not too important, the repeated game model may be a reasonable approximation. Also, many of the qualitative predictions about the importance of repeated play and the nature of equilibria are useful in thinking about more general dynamic games, as we discuss in Section 2B. Of course, the main reason that repeated games have received so much attention is their simplicity.

The Constituent Game g is a finite n -player game in normal form, (I, Σ, π) where Σ_i is the probability distributions over a finite set S_i of pure strategies. In the repeated version of g , each player i 's strategy is a sequence of maps $(\sigma_i(t))$ mapping the previous actions of all players to a $\sigma_i \in \Sigma_i$. Let us stress that it is the past actions that are observable, and not past choices of mixed strategies.

Players maximize the average discounted sum of their per-period payoffs with common discount factor δ . (We use the average discounted sum rather than simply the sum so that payoffs in the one-shot and repeated games are comparable--if a player receives payoff 5 every period his average discounted payoff is 5, while the discounted sum is, of course, $5/(1-\delta)$).

Player i 's reservation utility is

$$v_i^* \equiv \min_{\sigma_{-i}} \max_{\sigma_i} \pi^i(\sigma_i, \sigma_{-i}) .$$

In any equilibrium of the repeated game, player i 's strategy must be a best response to the strategies of his opponents. One option player i has is to play myopically in each period, that is to play to maximize that period's payoff, ignoring the way this influences his opponents' future play. This static maximization will give player i at least v_i^* in each period, so that in any equilibrium, player i 's expected average payoff must be at least v_i^* . A payoff vector v is individually rational if for all players $v_i > v_i^*$.

Notice that the equilibria of the constituent game (the "static equilibria") remain equilibria if the game is repeated: If each player's play is independent of the past history, then no player can do better than to play a static best response. Notice also that if the discount factor is very low, we'd expect that the static equilibria are the only equilibria--if the future is unimportant, then once again players will choose static best responses. (This relies on g being finite.)

The best-known result about repeated games is the celebrated "folk theorem." This theorem asserts that if the game is repeated infinitely often and players are sufficiently patient, then "virtually anything" is an equilibrium outcome. By treating the polar case of extreme patience, the folk theorem provides an upper bound for the effects of repeated play, and thus a benchmark for thinking about the intermediate case of mild impatience.

The oldest version of the folk theorem asserts that if players are sufficiently patient (the discount factors are near enough to one) then any feasible individually rational payoffs are supportable by a Nash equilibrium.

The idea of the proof is simple: any deviation from the prescribed path by player i leads the other players to play to "minmax" him (i.e., using the strategies that attain the minimum in the definition of v_i^*) for the rest of the game. In a repeated Cournot game, this would correspond to all players choosing the largest possible output forever. Given this threat, players will indeed choose not to deviate as long as

- (1) never deviating yields more than v_i^* , and
- (2) the discount factor is large enough that the gains to any one-period deviation are outweighed by the never ending ("grim") punishment.

The strategies sketched above clearly need not be subgame perfect--no firm would choose to produce a huge amount if the market price were zero! However, the "perfect folk theorem" shows that the same outcome can be enforced by a perfect equilibrium, so that restricting attention to perfect equilibria does not reduce the limit set of equilibrium payoffs. (It does, of course, rule out some Nash equilibria.)

Friedman (1971) proved a weaker version of this theorem which showed that any payoffs better for all players than a Nash equilibrium of the constituent game are the outcome of a perfect equilibrium of the repeated game, if players are sufficiently patient. The desired play is enforced by the "threat" that any deviation will trigger a permanent switch to the static equilibrium. Because this "punishment" is itself a perfect equilibrium, so are the overall strategies. This result shows, for example, that patient, identical, Cournot duopolists can "implicitly collude" by each producing one-half the monopoly output, with any deviation triggering a switch to the Cournot outcome. This would be "collusive" in yielding the monopoly price. The collusion is "implicit" (or "tacit") in that the firms would not need to enter into binding

contracts to enforce their cooperation. Instead, each firm is deterred from breaking the agreement by the (credible) fear of provoking Cournot competition. If this equilibrium is suitably "focal," as it might be with two identical firms, then the firms might be able to collude without even communicating! This possibility has grave implications for anti-collusion laws based on observed conduct. How could two non-communicating firms be charged with conspiracy?

Whether collusion can be enforced in a particular oligopoly then depends on whether the "relevant" discount factor is sufficiently large. This discount factor measures the length of the observation lag between periods, as well as the player's impatience "per unit time." In a market where orders are large but infrequent, a single order might represent several years of full-time production. Here the short-run gains to cheating might well outweigh the costs of (greatly delayed) punishments. In the other extreme, with frequent, small orders, implicit collusion is more likely to be effective.

The Friedman result is weaker than the folk theorem because of its requirement that both players do better than in a static equilibrium. As a Stackelberg follower's payoffs are worse than a Cournot duopolist's, Friedman's result does not show that the Stackelberg outcome can be enforced in a repeated Cournot game. That this is however true is shown in the "perfect folk theorems" of Aumann-Shapley (1976), Rubinstein (1979), and Fudenberg-Maskin (1986a). Aumann-Shapley and Rubinstein consider the no-discounting models in which players are "completely" patient. Fudenberg-Maskin show that, under a mild "full-dimensionality" condition, the result continues to hold if the discount factors are sufficiently close to one. They also strengthen earlier results by allowing players to use mixed strategies as punishments. Aumann-Shapley and Rubinstein had restricted attention to pure strategies, which leads to higher individually-rational payoff levels, and thus a weaker theorem. (Their work

can also be interpreted as allowing mixed strategies as long as the mixing probabilities themselves, and not just the actions actually chosen, are observable at the end of each period.)

One might wish to characterize the set of perfect equilibria when there is "substantial" impatience. A fascinating paper by Abreu (1984) provides a tool for this purpose. (See also Harris (1986), who gives a clearer exposition and simpler proofs of Abreu's results.) Call strategies "simple" if they have the following form: there is an "equilibrium path" and n "punishment paths," one for each player. Play follows the equilibrium path as long as no one has deviated. If player i was the most recent player to deviate, and did so at period t , then play at period $(t+k)$ is given by the k^{th} element of the "punishment path" corresponding to player i . (What happens if two or more players deviate simultaneously is irrelevant.)⁵ The force in the restriction to simple strategies is that player i 's punishment path is independent of the history before i 's deviation and also of the nature of the deviation itself. Simple strategies are optimal if each player's average discounted utility at the beginning of his punishment phase is the lowest payoff he receives in any perfect equilibrium.

As the set of equilibria is closed (Fudenberg-Levine [1983]), there is a worst perfect equilibrium $w(i)$ for each player i . Any equilibrium path that can be enforced by the threat that player i 's deviations will be punished by switching to some equilibrium can clearly be enforced by the threat player i 's deviations will be punished by switching to $w(i)$. Therefore, as Abreu shows, optimal simple strategies exist, and any perfect equilibrium outcome can be

⁵ Simultaneous deviations can be ignored, because in testing for Nash or subgame-perfect equilibria, we ask only if a player can gain by deviating when his opponents play as originally specified.

enforced by such strategies. Thus, to characterize the set of equilibria in any game it suffices to find the worst possible perfect equilibrium payoffs for each player. In general this may be a difficult problem, but the set of symmetric equilibria of symmetric games is more easily characterized.

[Caution -- symmetry here requires not only that the payoffs along the equilibrium path be identical, but that the payoffs be identical in the punishment phases as well.] Abreu's thesis (1983) uses the idea of optimal simple strategies to characterize the symmetric equilibria of repeated Cournot games. Shapiro's essay in this Handbook explains this characterization in detail.

Another case in which the lowest perfect equilibrium payoffs can be pinned down is when equilibria can be constructed that hold players to their reservation values. Fudenberg-Maskin (1987) provide conditions for this to be true for a range of discount factors between some δ and 1. Because the reservation values are of course the worst possible punishments, any equilibrium outcome (Nash or perfect) can be enforced with the threat that deviations will switch play to an equilibrium in which the deviator is held to his reservation value.

Repeated Games with Imperfect Monitoring

One drawback of repeated games as a model of collusion is that they do not explain price wars: In equilibrium, no firm ever deviates. This lack motivated the Green-Porter (1984) model of "Noncooperative Collusion under Imperfect Price Information." The Green-Porter model is an infinitely repeated quantity-setting game in which firms do not observe the outputs of their opponents. Instead, firms only observe the market price, $p(Q, \theta)$, which is determined by aggregate output Q and a stochastic disturbance, θ . The θ 's in the different periods are identically and independently distributed according to a density $f(\theta)$, which is such that the set of possible prices (the support of $p(Q, \theta)$) is independent of Q . All firms are identical, and there is a symmetric ("Cournot") equilibrium of the constituent game in which each firm produces output q^C . As with ordinary repeated games, one equilibrium of the repeated game is for all firms to produce q^C each period. Could the firms hope to improve on this outcome if they are patient?

Green-Porter show that they can, by constructing a family of "trigger-price" equilibria of the following form: Play begins in the "cooperative" phase, with each firm producing some output q^* . Play remains in the cooperative phase as long as last period's price exceeded a trigger level p^* . If the price falls below p^* , firms switch to a "punishment phase" in which each firm produces output q^C . Punishment lasts for T periods, after which play returns to a cooperative phase. For a triple (q^*, p^*, T) to generate a Nash equilibrium, each firm must prefer not to cheat in either phase. Since q^C is a static equilibrium, no firm will cheat in the punishment phases, so we need only check the cooperative phase. Setting $q^* = q^C$ results in a trivial trigger-price equilibrium. If the firms are somewhat patient they can do better by setting $q^* < q^C$. In such an equilibrium, p must be high enough that punishment

occurs with positive probability. Otherwise, a firm could increase its output slightly in the cooperative phase without penalty. Thus punishment will occur even if no firm has deviated. On seeing a low price, all firms expand their output not out of concern that an opponent has cheated, but rather in the knowledge that if low prices did not sometimes trigger punishment, then their collusive scheme would not be self-enforcing. (See Rotemberg-Saloner (1986) for a repeated game model with perfect monitoring in which price wars are voluntary.)

The trigger-price equilibria constructed by Green-Porter have an appealing simplicity, but they need not be optimal--other equilibria may yield higher expected payoffs (for the firms): Abreu-Pearce-Stacchetti (1986) investigated the structure of the optimal symmetric equilibria in the Green-Porter model. In the process, they develop a tool which is useful for analyzing all repeated games with imperfect monitoring. This tool, which they call "self-generation," is extended in their 1986 paper.

Self-generation is a sufficient condition for a set of payoffs to be supportable by equilibria. It is the multi-player generalization of dynamic programming's principle of optimality, which provides a sufficient condition for a set of payoffs, one for each state, to be the maximal net present values obtainable in the corresponding states. Abreu-Pearce-Stacchetti's insight is that the "states" need not directly influence the player's payoffs, but can instead reflect (in the usual self-confirming way) changes in the play of opponents. Imagine for example that, in the Green-Porter model, there are only three possible values of the market price -- $p_1 > p_2 > p_3$. Price p_i occurs with probability $m_i(Q)$, where Q is total industry output. Note that past prices do not directly influence current payoffs or transition probabilities.

Nevertheless, we can construct equilibrium strategies that use the realized prices to determine the transitions between "fictitious" states.

For example, imagine that we are told that there are two fictitious states a and b , with associated payoffs for both firms of u_a and u_b . (We will look at symmetric equilibria; otherwise we would need to specify each firm's payoffs.) We are also given the following transition rule: the state switches from a to b if p_3 occurs, remaining at a if $p = p_1$ or p_2 . State b is absorbing: once it is reached, it prevails from then on. As we will see, state b corresponds to an infinite "punishment phase" in Green-Porter. The values u are self-generating if, in each state $i=a,b$, when players believe that their future payoffs are given by u , there is an equilibrium s_i in current actions with average (over current and future payoffs) payoff u_i . In the language of dynamic programming, this says that for each player the payoff u_i is unimprovable, given the specified continuation payoffs and his opponents' current actions. To show that self-generating payoffs are sustainable by Nash equilibria, we first must define strategies for the players. To do this, trace out the succession of single-period equilibria, i.e. if play begins in state a , and p_1 occurs in the first period, the state is still a , so the second-period outputs are again given by the s_a . By construction, no player can gain by deviating from strategy s_i in state i for one period and then reverting to them thereafter. The standard dynamic programming argument then shows that unimprovability implies optimality: By induction, no player can improve on u_a or u_b by any finite sequence of deviations, and the payoff to an infinite sequence of deviations can be approximated by finitely many of them. In our example, since state b is absorbing, for (u_a, u_b) to be self-generating, u_b must be self-generating as a singleton set. This means that u_b must be the payoffs in a static equilibrium, as in Green-Porter's

punishment phase. In state a , today's outcome influences the future state, so that players have to trade off their short run incentive to deviate against the risk of switching to state b . Thus state a corresponds to a "cooperative" phase, where players restrict output to decrease the probability of switching to the punishment state.

The self-generation criterion not only provides a way of testing for equilibria, it also suggests a way of constructing them: one can construct state spaces and transition rules instead of working directly with the strategy spaces. Fudenberg-Maskin (1986b) use this technique to investigate when "folk theorems" obtain for repeated games with imperfect monitoring.

Returning to the topic of implicit collusion in oligopolies, what lessons do we learn from the study of repeated games? First, repetition matters more, and (privately) efficient outcomes are more likely to be equilibria, when the periods are short. Second, more precise information makes collusion easier to sustain, and lowers the costs of the occasional "punishments" which must occur to sustain it. Third, firms will prefer "bright-line" rules which make "cheating" easy to identify. For example, firms would like to be able to respond to changes in market conditions without triggering "punishment." Scherer (1980) suggests that the institutions of price leadership and mark-up pricing may be responses to this problem. (See also Rotemberg-Saloner (1985), who explain how price leadership can be a collusive equilibrium with asymmetrically-informed firms.)

While most applications of repeated games have been concerned with games with infinitely lived players, "implicitly collusive equilibria" can arise even if all the players have finite lives, as long as the model itself has an infinite horizon. Let us give two examples. First, a finitely lived manager of a firm becomes the equivalent of an infinitely lived player if he owns the firm,

because the latter's value depends on the infinite streams of profits (as in Kreps (1985)). Second, overlapping generations of finite lived players can yield some cooperation between the players. A player who cheats early in his life will be punished by the next generation, which in turn will be punished by the following generation if it does not punish the first player, etc. (Cremer (1983)).

We conclude this section with three warnings on the limitations of the repeated game model. First, by focusing on stationary environments, the model sidesteps the questions of entry and entry deterrence. These questions can in principle be studied in games whose only time-varying aspect is the number of entrants, but serious treatments of the entry process more naturally allow for factors such as investment. Second, because repetition enlarges the set of equilibria, selecting an equilibrium becomes difficult. If firms are identical, an equal division of the monopoly profits seems an obvious solution; however if one complicates the model by, for instance, introducing a prior choice of investment, most subgames are asymmetric, and the quest for a focal equilibrium becomes harder. However, the selection criterion of picking a date-zero Pareto optimal equilibrium outcome is not "meta-perfect": date-zero Pareto optimal outcomes are typically enforced by the threat of switching to a non-Pareto optimal outcome if some player deviates. Just after the deviation, the game is formally identical to the period-zero game, yet it is assumed that players will not again coordinate on the focal Pareto-optimal outcome. Third, implicit collusion may not be enforceable if the game is repeated only finitely many times. What then should we expect to occur in finite-lived markets?

Finite-Horizon Games

Infinite-horizon repeated games are used as an idealization of repeated play in long-lived markets. Since actual markets are finite-lived, one should ask whether the infinite-horizon idealization is sensible. One response is that we can incorporate a constant probability μ of continuing to the next period directly into the utility functions: the expected present value of ten utils tomorrow, if tomorrow's utils are discounted by δ , and tomorrow arrives with probability μ , is simply $\delta\mu$.

Then if both δ and μ are near to one the folk theorem applies. This specification implies that the game ends in finite time with probability one, but there is still a positive probability that the game exceeds any fixed finite length. Thus one may ask what the theory predicts if the game is certain to end by some very far-distant date. It is well known that in some games the switch from an infinite horizon to a long finite one yields dramatically different conclusions--the set of equilibrium payoffs can expand discontinuously at the infinite-horizon limit. This is true for example in the celebrated game of the "prisoner's dilemma," which is depicted in Figure 8. When played only once, the game has a unique Nash equilibrium, as it is a dominant strategy for each player to "fink." "Never fink" is a perfect equilibrium outcome in the infinitely-repeated game if players are sufficiently patient.

With a finite horizon, cooperation is ruled out by an iterated dominance argument: Finking in the last period dominates cooperating there; iterating once, both players fink in the second period, etc. The infinite-horizon game lacks a last period and so the dominance argument cannot get started. Should we then reject the cooperative equilibria as technical artifacts, and conclude that the "reasonable solution" of the finitely-repeated prisoner's dilemma is "always fink"? Considerable experimental evidence shows that subjects do tend

to cooperate in many if not most periods. Thus, rather than reject the cooperative equilibria, we should change the model to provide an explanation of cooperation. Perhaps players derive an extra satisfaction from "cooperating" beyond the rewards specified by the experimenters. While this explanation does not seem implausible, it seems a bit too convenient. Other explanations do not add a payoff for cooperation per se, but instead change the model to break the backwards-induction argument, which is argued to be unreasonable. One way of doing this is developed in the "reputation effects" models of Kreps, Milgrom, Roberts, and Wilson, which we discuss in Section 3. These models assume, not that all players prefer cooperation, but that each player attaches a very small prior probability to the event that his opponent does.

Radner (1980) provides another way of derailing the backwards induction in the finitely-repeated game. He observes that the best response against an opponent who will not fink until you do, but will fink thereafter (the "grim" strategy) is to cooperate until the last period, and then fink. Moreover, as the horizon T grows, the average gain (the gain divided by T) to playing this way instead of always cooperating goes to zero. Formally, in an ϵ -equilibrium, player's strategy gives him within ϵ of his best attainable payoff (over the whole horizon); in a subgame-perfect ϵ -equilibrium this is true in every subgame. Radner shows that cooperation is the outcome of a perfect ϵ -equilibrium for any $\epsilon > 0$ if players maximize their average payoff and the horizon is sufficiently long. Radner's result relies on "rescaling" the player's utility functions by dividing by the length of the game. Thus one-period gains become relatively unimportant (compared to the fixed ϵ) as the horizon grows.

Fudenberg-Levine (1983) show that if players discount the future then the ϵ -equilibrium, finite horizon approach gives "exactly" the same conclusions as

the infinite-horizon one: the set of infinite-horizon (perfect) equilibria coincides with the set of limit points of finite horizon (perfect) ε -equilibria, where ε goes to zero as the horizon T goes to infinity. That is, every such limit point is on infinite-horizon equilibria, and every infinite-horizon equilibrium can be approximated by a convergent sequence of finite horizon ε -equilibria. Fudenberg-Levine defined the "limits" in the above with respect to a topology that requires the action played to be uniformly close in every subgame. In finite-action games (games with a finite number of actions per period) this reduces to the condition that $(s_n) \rightarrow s$ if s_n and s exactly agree in the first k_n periods for all initial histories, where $k_n \rightarrow \infty$ as $n \rightarrow \infty$. Harris (1985a) shows that this simpler convergence condition can be used in most games, and dispenses with a superfluous requirement that payoffs be continuous.

With either of the Fudenberg-Levine or Harris topology, the strategy spaces are compact in finite-action games, so that the limit result can be restated as follows: Let $\Gamma(\varepsilon, T)$ be the correspondence yielding the set of ε -equilibria of the T -period game. Then Γ is continuous at $(0, \infty)$. This continuity allows one to characterize infinite-horizon equilibria by working with finite-horizon ones. Backwards-induction can be applied to the latter, albeit tediously, but not to the former, so that working with the finite horizon ε -equilibria is more straightforward. The continuity result holds for discounted repeated games, and for any other game in which players are not too concerned about actions to be taken in the far-distant future. (It does not hold in general for the time-average payoffs considered by Radner.) Specifically, preferences over outcome paths need not be additively separable over time, and there can be links between past play and future opportunities. In particular the result covers the non-repeated games discussed later in this

section. The intuition is simply that if players are not too concerned about the future, the equilibria of the infinite-horizon game should be similar to the equilibria of the "truncated" game in which no choices are allowed after some terminal time T . So for any equilibrium s of the infinite horizon game and $\varepsilon > 0$, by taking T long enough, the difference in each player's payoff between the play prescribed by s and that obtained by truncating s at time t will be of order ε .

We should point out that the "epsilons" are not always needed to ensure continuity at the infinite-horizon limit. One example is Rubinstein's (1982) bargaining game, which even with an infinite horizon has a unique perfect equilibrium. (Rubinstein allows players to choose from a continuum of sharing rules between 0 and 1. With a finite grid of shares, the uniqueness result requires that each player prefers the second-largest partition today to the largest one tomorrow, so that the grid must be very fine if the discount factors are near to one.) Benoit-Krishna (1985) provide conditions for continuity to obtain in the "opposite" way, with the set of finite-horizon equilibria expanding as the horizon grows, and approaching the limit set given by the folk theorem. (Friedman (1984) and Fraysse-Moreaux (1985) give independent but less complete analyses.) For Nash equilibria this is true as long as the static equilibria give all players more than their minmax values. Then any individually-rational payoffs can be enforced in all periods sufficiently distant from the terminal date by the threat that any deviations result in the deviator being minmaxed for the rest of the game. Such threats are not generally credible, so proving the analogous result for perfect equilibria is more difficult. Benoit-Krishna show that the result does hold for perfect equilibria if each player has a strict preference for one static equilibrium as opposed to another (in particular there must be at least two static

equilibria) and the Fudenberg-Maskin full dimensionality condition is satisfied. The construction that Benoit-Krishna use to prove this is too intricate to explain here, but it is easy to see that there can be perfect equilibria of a finitely-repeated game which are not simply a succession of static equilibria. Consider the game in Figure 9.

There are two pure-strategy static equilibria, (U,L) and (M,M). In the twice-repeated game (without discounting, for simplicity) there is an equilibrium with total payoffs (-1,-1). These payoffs result from the strategies "play (D,R) in the first period; play (U,L) in the second iff (D,R) was played in the first, otherwise, play (M,M)."

2B. Continuous-time Games

Frequently, continuous-time models seem simpler and more natural than models with a fixed, non-negligible period length. For example, differential equations can be easier to work with than difference equations. As in games with a continuum of actions, continuous-time games may fail to have equilibria in the absence of continuity conditions. More troublesome, there are deep mathematical problems in formulating general continuous-time games.

As Anderson (1985) observes, "general" continuous-time strategies need not lead to a well-defined outcome path for the game, even if the strategies and the outcome path are restricted to be continuous functions of time. He offers the example of a two-player game where players simultaneously choose actions on the unit interval. Consider the continuous-time strategy "play at each time t the limit as $r \rightarrow t$ of what the opponent has played at times r previous to t ." This limit is the natural analog of the discrete-time strategies "match the opponent's last action." If at all times before t the players have chosen matching actions, and the history is continuous, there is no problem in

computing what should be played at t . However, there is not a unique way of extending the outcome path beyond time t . Knowing play before t determines the outcome at t , but is not sufficient to extend the outcome path to any open interval beyond t . As a result of this problem, Anderson opts to study the limits of discrete-time equilibria instead of working with continuous time.

Continuous time formulations are fairly tractable when strategies depend on a "small" set of histories. This is the case in stopping-time games, open-loop games, and in situations where players use "state-space" strategies. These games or strategies are not restricted to continuous time, and discrete-time versions of all of them have been used in the industrial organization literature.

2C. State-Space or Markov Equilibria

Consider games in which players maximize the present value of instantaneous flow payoffs, which may depend on state variables as well as current actions. (The feasible actions may also depend on the state.) For example, current actions could be investment decisions, and the state could be the stocks of machinery. Or current actions could be expenditures on R&D, with the state variables representing accumulated knowledge. The strategy spaces are simplified by restricting attention to "state-space" (or "Markov") strategies that depend not on the complete specification of past play, but only on the state (and, perhaps, on calendar time.) A state-space or Markov equilibrium is an equilibrium in state-space strategies, and a perfect state-space equilibrium must yield a state-space equilibrium for every initial state. Since the past's influence on current and future payoffs and opportunities is summarized in the state, if one's opponents use state-space strategies, one could not gain by conditioning one's play on other aspects of the history. Thus

a state-space equilibrium is an equilibrium in a game with less restricted strategies. The state-space restriction can however rule out equilibria, as shown by the infinitely-repeated prisoner's dilemma. Since past play has no effect on current payoffs or opportunities, the state-space is null, and all state-space strategies must be constants. Thus the only state-space equilibrium is for both sides to always fink. (Caution: this conclusion may be due to a poor model, and not to the wrong equilibrium concept. Section 4E shows how the conclusion is reversed in a slightly different model.)

Maskin-Tirole (1985), using the Markov restriction, obtain collusion in a repeated price game in which prices are locked in for two periods. They argue that what is meant by "reaction" is often an attempt by firms to react to a state that affects their current profits; for instance, when facing a low price by their opponents, they may want to regain market share. In the classic repeated game model, firms move simultaneously, and there is no physical state to react to. If, however, one allows firms to alternate moves, they can react to their opponent's price. (Maskin-Tirole derive asynchronicity as the (equilibrium) result of the two-period commitments.) The possibility of reaction leads to interesting Markov equilibria. However, although equilibrium payoffs are bounded away from the competitive levels (in contrast to the folk theorem approach), they are still many equilibria (Maskin-Tirole use renegotiation-proofness to select one which exhibits the classic "kinked demand curve.") Gertner (1985b) formalizes collusion with Markov strategies when commitment (inertia) takes the form of a fixed cost of changing prices.

The literal definition of a state says that strategies can depend "a lot" on variables with very little influence on payoffs, but they cannot depend at all on strategies that have no influence. This can generate rather silly discontinuities. For example, we can restore the cooperative equilibria in the

repeated prisoner's dilemma by adding variables that keep track of the number of times each player has finked. If these variables have an infinitesimal effect on the flow payoffs, the cooperative equilibria can be restored.

The state-space restriction does not always rule out "supergame"-type equilibria, as shown in Fudenberg-Tirole (1983a). They reconsidered a model of continuous-time investment that had been introduced by Spence (1979). Firms choose rates of investment in productive capacity. The cost of investment is linear in the rate up to some upper bound, with units chosen so that one unit of capital costs one dollar. If firms did not observe the investment of their rivals, each firm would invest up to the point where its marginal productivity of capital equalled the interest rate. The capital levels at this "Cournot" point exceed the levels the firms would choose if they were acting collusively, because each firm has ignored the fact that its investment lowers its rivals' payoffs. Now, if firms observe their rivals' investment (in either discrete or continuous time) they could play the strategy of stopping investment once the collusive levels are reached. This "early stopping" is enforced by the (credible) threat that if any firm invests past the collusive level, all firms will continue to invest up to the "Cournot" levels. The state-space restriction seems to have little force in this game. There are no general results on when the restriction is likely to have a significant impact.

State-space games closely resemble control problems, so it is not surprising that they have been studied by control theorists. Indeed, the idea of perfection is just the many-player version of dynamic programming, and it was independently formulated by Starr-Ho (1967) in the context of nonzero-sum differential games. The differential games literature restricts attention to state-space equilibria in which the equilibrium payoffs are continuous and almost-everywhere differentiable functions of the state. These conditions

obtain naturally for control problems in smooth environments, but they impose significant restrictions in games: It might be that each player's strategy, and thus each player's payoff, change discontinuously with the state due to the self-fulfilling expectation that the other players use discontinuous strategies. This was the case in the "early-stopping" equilibria of the last paragraph, so those equilibria would not be admissible in the differential games setting. Perhaps the continuity restriction can be justified by the claim that the "endogenous discontinuities" that they prohibit require excessive coordination, or are not robust to the addition of a small amount of noise in the players' observations. We are unaware of formal arguments along these lines.

The technical advantage of restricting attention to smooth equilibria is that necessary conditions can then be derived using the variational methods of optimal control theory. Assume that player i wishes to choose a_i to maximize the integral of his flow payoff π^i , subject to the state evolution equation

$$(2) \quad \dot{k}(t) = f(k(t)) , k(0) = k_0 .$$

Introducing costate variables λ_i , we define H_i , the Hamiltonian for player i , as

$$(3) \quad H_i = \pi^i(k, a, t) + \lambda_i f(k(t)) .$$

A state-space equilibrium $a(t)$ must satisfy

$$(4) \quad a_i = a_i(k, t) \text{ maximizes } H_i(k, t, a, \lambda_i) .$$

and

(5) $\dot{\lambda}_i - \partial H_i / \partial k_i - \sum_{j \neq i} \partial H_i / \partial a_j \partial a_j / \partial k_i$, along with the appropriate transversality condition.

Notice that for a one-player game the second term in (5) vanishes, and the conditions reduce to the familiar ones. In the n-player case, this second term captures the fact that player i cares about how his opponents will react to changes in the state. Because of the cross-influence term, the evolution of λ is determined by a system of partial differential equations, instead of by ordinary differential equations as in the one-player case. As a result, very few differential games can be solved in closed form. An exception is the linear-quadratic case, which has been studied by Starr-Ho among others. Hanig (1985) and Reynolds (1985) consider a linear-quadratic version of the continuous-time investment game. (Their model is that of Spence-Fudenberg-Tirole, except that the cost of investment increases quadratically in the rate.) They show that the "smooth" equilibrium for the game has higher steady-state capital stocks and so lower profits, than the static "Cournot" levels. Is this a better prediction than the collusive levels? We do not know.

Judd (1985) offers an alternative to the strong functional form assumptions typically invoked to obtain closed form solution to differential games. His method is to analyze the game in the neighborhood of a parameter value that leads to a unique and easily computed equilibrium. In his examples of patent races, he looks at patents with almost zero value. Obviously if the patent has exactly zero value, in the unique equilibrium players do no R&D and have zero values. Judd proceeds to expand the system about this point, neglecting all

terms over third order in the value of the patent. Judd's method gives only local results, but it solves an "open set" in the space of games, as opposed to conventional techniques that can be thought of as solving a lower-dimensional subset of them. We encourage the interested reader to consult Judd's paper for the technical details.

2D. Games of Timing

In a game of timing, each player's only choice is when and whether to take a single pre-specified action. Few situations can be exactly be described this way, because players typically have a wider range of choices. For example, firms typically do not simply choose a time to enter a market, but also decide on the scale of entry, the type of product to produce, etc. This detail can prove unmanageable, which is why industrial organization economists have frequently abstracted it away to focus on the timing question in isolation.

We will not even try to discuss all games of timing, but only two-player games which "end" once at least one player has moved. Payoffs in such games can be completely described by six functions $L_i(t)$, $F_i(t)$, and $B_i(t)$, $i=1,2$. Here L_i is player i 's payoff if player i is the first to move (the "leader"), F_i is i 's payoff if j is the first to move (the "follower"), and B_i is i 's payoff if both players move simultaneously. This framework is slightly less restrictive than it appears, in that it can incorporate games which continue until both players have moved. In such games, once one player has moved, the other one faces a simple maximization problem, which can be solved and "folded back" to yield the payoffs as a function of the time of the first move alone. A classic example of such a game is the "war of attrition," first analyzed by Maynard Smith (1974): Two animals are fighting for a prize of value v ; fighting costs one util per unit time. Once one animal

quits, his opponent wins the prize. Here $L(t)$ is $-t$, and $F(t)$ is $v-t$. With short time periods $B(t)$ will turn out to not matter much; let's set it equal to $v/q-t$, $q \geq 2$. If $q=2$ then each player has probability $1/2$ of winning the prize if both quit at once; if $q=\infty$ this probability is zero. Let us solve for a symmetric equilibrium of the discrete time version with period length Δ . Let p be the probability that either player moves at t when both are still fighting. For players to use stationary mixed strategies, the payoff to dropping out, pv/q , must equal that to fighting one more period and then dropping out, $pv+(1-p)pv/q - \Delta$. Equating these terms yields

$$p = (1 - (1 - 4\Delta/qv)^{1/2})q/2.$$

[Dropping out with probability p per period is a "behavioral strategy;" the corresponding mixed strategy is an exponential distribution over stopping times.]

Let us note that as $\Delta \rightarrow 0$, $p \rightarrow \Delta/v$, independent of q . More generally, a war of attrition is a game of "chicken", in which each player prefers his opponent to move ($F(t) > L(t)$), and wishes that he would do so quickly (F and L decrease over time.) Weiss-Wilson (1984) characterize the equilibria of a large family of discrete-time wars of attrition; Hendricks-Wilson do the same for the continuous-time version. Section 3 describes some of the many incomplete-information wars of attrition that have been applied to oligopoly theory.

Preemption games are the opposite case, with $L(t) > F(t)$, at least over some set of times. Here the specification of $B(t)$ is more important, as if B exceeds F we might expect both players to move simultaneously. One example

of a preemption game is the decision of when and whether to build a new plant or adopt a new innovation, when the market is only big enough to support one such addition. (If each firm will eventually build a plant, but the second mover would optimally choose to wait until long after the first one, we can "fold back" the second mover's choice to get the payoffs as a function of the time of the first move alone.)

The relationship between L and F can change over time, and the two players may have different "types" of preferences, as in Katz-Shapiro (1984). No one has yet attempted a general classification of all games of timing. Because the possible actions and histories are so limited, it is easy to formulate continuous-time strategies for these games, in a way that permits a well-defined map from strategies to outcomes. We develop these strategies below. However, we will see that the simplicity of this formulation is not without cost, as it is not rich enough to represent some limits of discrete-time strategies. That is, there are distributions over outcomes (who moves and when) that are the limits of distributions induced by discrete-time strategies, but which cannot be generated by the "obvious" continuous time strategies.

The usual and simple continuous-time formulation is that each player's strategy is a function $G_i(t)$ which is non-decreasing, right-continuous, and has range in $(0,1)$. This formulation was developed in the 1950's for the study of zero-sum "duels", and was used by Pitchik (1982), who provides several existence theorems. The interpretation is that G_i is a distribution function, representing the cumulative probability that player i has moved by time t conditional on the other player not having moved previously. These distribution functions need not be continuous; a discontinuity at time t implies that the player moves with non-zero probability at exactly time t . Where G is differentiable, its derivative dG is the density which gives the probability

of a move over a short time interval. With this notation, player one's payoff to the strategies G_1, G_2 is

$$V^i(G_1, G_2) = \int_0^\infty [L(s)(1-G_j(s))dG_i(s) + F(s)(1-G_i(s))dG_j(s)] + \sum \alpha_i(s)\alpha_j(s)B(s),$$

where $\alpha_i(s)$ is the size of the jump in G_i at s .

This formulation is very convenient for wars of attrition. In these games there are "nice" discrete-time equilibria in which the probability of moving in each period is proportional to the period length. In the example computed above, the equilibrium strategies converged to the continuous-time limit $G(t) = 1 - \exp(-t/v)$. (For the case $q=2$, the sum of the two players' payoffs is upper hemi-continuous, and the fact that the equilibria converge is a consequence of Theorem 5 of Dasgupta-Maskin (1986).) More complex wars of attrition can have continuous-time equilibria with "atoms," ($\alpha_i(t) > 0$ for some t), but as the periods shrink these atoms become isolated, and again admit a nice continuous-time representation.

Preemption games are markedly different in this respect, as shown in Fudenberg-Tirole (1985). Consider the discrete-time "grab-the dollar" game: $L(t)=1$, $F(t)=0$, and $B(t) = -1$. The interpretation is that "moving" here is grabbing a dollar which lies between the two players. If either grabs alone, he obtains the dollar, but simultaneous grabbing costs each player one. There is a symmetric equilibrium in which each player moves with (conditional) probability $1/2$ in each period. Note well that the intensity of the randomization is independent of the period length. The corresponding payoffs are $(0,0)$, and the distribution over outcomes is that with identical probability $(1/4)^{(t+1)}$ either player one wins (moves alone) in period t , or player two wins in period t , or both move at once at t . As the length of the period converges

to zero, this distribution converges to one in which the game ends with probability one at the start, with equal probabilities of $1/3$ that player one wins, that player two does, or that they both move at once. This distribution cannot be implemented with the continuous-time strategies described above, for it would require a correlating device a la Aumann. Otherwise, at least one player would move with probability one at the start, which would make it impossible for his opponent to have a $1/3$ probability of winning. The problem is that a great many discrete-time strategies converge to a continuous-time limit in which both players move with probability one at time zero, including "move with probability $1/2$ each period," and "move with probability one at the start." The usual continuous-time strategies implicitly associate an atom of size one with an atom of that size in discrete time, and thus they cannot represent the limit of the discrete-time strategies. Fudenberg-Tirole offered an expanded notion of continuous time strategies that "works" for the grab-the-dollar game, but they did not attempt a general treatment of what the strategy space would need to be to handle all games of timing.

The moral of this story is that while continuous time is often a convenient idealization of very short time periods, one should keep in mind that a given formulation of continuous time may not be adequate for all possible applications. When confronted with, for example, the non-existence of equilibria in a seemingly "nice" continuous-time game, it can be useful to think about discrete-time approximations. Simon and Stinchcombe (1985) provide a general analysis of when the usual continuous-time strategies are in fact appropriate.

2E. Discrete vs. Continuous Time, and the Role of Period Length

The discussion above stressed that one must be careful that a given continuous time model is rich enough to serve as the "appropriate" idealization of very short time periods. Now we'd like to point out that new equilibria can arise in passing to continuous time, and that these should not be discarded as pathological. The simplest, and oldest, example of this fact is the Kreps-Wilson (1982a) stopping-time version of the prisoner's dilemma. In this version, players begin by cooperating, and once either finks they both must fink forever afterwards. Thus the only choice players have is when to fink if their opponent has not yet done so. In discrete time with a finite horizon, the familiar backwards induction argument shows that the only equilibrium is to both fink at once. However, the gain to finking one period ahead of one's opponent is proportional to the period length, and in the continuous-time limit, there is no gain to finking. Thus cooperation is an equilibrium in the continuous-time game.

The analogy with the finite-to-infinite horizon limit is more than suggestive. In a generalization of their earlier work, Fudenberg-Levine (1986) showed that, in cases such as stopping-time games and state-space games where the continuous time formulation is not too problematic, any continuous-time equilibrium is a limit of discrete-time epsilon-equilibria, where the epsilon converges to zero with the length of the period.

Little is known in general about the effect of period length on equilibrium play, but several examples have been intensively studied. The best known is the work of Coase (1972), Bulow (1982), Stokey (1981), and Gul-Sonnenschein-Wilson (1986) who argue with varying degrees of formality that the monopolistic producer of a durable good loses the power to extract rents as the time period shrinks, thus verifying the "Coase conjecture." (See also Sobel-Takahashi (1983) and Fudenberg-Levine-Tirole (1985).)

2F. Open-Loop Equilibria

The terms "open-loop" and "closed-loop" refer to two different information structures for multi-stage dynamic games. In an open-loop model, players cannot observe the play of their opponents; in a closed-loop model all past play is common knowledge at the beginning of each stage. Like "Cournot" and "Bertrand" equilibria, open- and closed-loop equilibria are shorthand ways of referring to the perfect equilibria of the associated model. (Caution: This terminology is widespread but not universal. Some authors use "closed-loop equilibrium" to refer to all the Nash equilibria of the closed-loop model. We prefer to ignore the imperfect equilibria.) Open and closed loop models embody different assumptions about the information lags with which players observe and respond to each other's actions, and thus about the length of time to which players can "commit" themselves not to respond to their opponents. In an open-loop model, these lags are infinite, while in a closed-loop model, a player can respond to his opponents. Because dynamic interactions are limited in open-loop equilibria, they are more tractable than closed-loop ones. For this reason, economists have sometimes analyzed the open-loop equilibria of situations which seem more naturally to allow players to respond to their opponents. One possible justification for this is that, if there are many "small" players, so that no one player can greatly affect the others, then optimal reactions should be negligible. When this is true, the open-loop equilibria will be a good approximation of the closed-loop ones. Fudenberg-Levine (1987) explore this argument, and find that its validity in a T-period game requires strong conditions on the first through T-th derivatives of payoffs.

Section 3: Static Games of Incomplete and Imperfect Information

3A. Bayesian Games and Bayesian Equilibrium

Players in a game are said to have incomplete information if they do not know some of their opponent's characteristics (objective functions); they have imperfect information if they do not observe some of their opponent's actions. Actually, the distinction between incomplete and imperfect information is convenient, but artificial. As Harsanyi [1967] has shown, at a formal level, one can always transform a game of incomplete information into a game of imperfect information. The idea is the following: let the original game be an n -player game with incomplete information. Assume that each player's characteristic is known by the player, but, from the point of view of the $(n-1)$ other players, is drawn according to some known probability distribution. (See below for a discussion of this representation).

Harsanyi's construction of a transformed game introduces nature as a $(n+1)$ st player, whose strategy consists in choosing characteristics for each of the n original players at the start of the game, say. Each player observes his own characteristics, but not the other players'. Thus, he has imperfect information about nature's choice of their characteristics. (One can endow nature with an objective function in order for it to become a player. One way of doing so is to assume that nature is indifferent between all its moves. To recover the equilibria of the original game (i.e., for given initial probability distributions), one takes the projection of the equilibrium correspondence for these probability distributions).

The notion of "type": The "characteristic" or "type" of a player embodies everything which is relevant to this player's decision making. This includes the description of his objective function (fundamentals), his beliefs about the other player's objective functions (beliefs about fundamen-

tals), his beliefs about what the other players believe his objective function is (beliefs about beliefs about fundamentals), etc. As this is a bit abstract, it is helpful to begin with Harsanyi's simple representation (this representation is used in virtually all applications). Suppose that in an oligopoly context, each firm's marginal cost c_i is drawn from an "objective" distribution $p_i(c_i)$ (n.b.: we will write probability distributions as if the number of potential types were finite. Continuous type spaces are also allowed; summation signs should then also be replaced by integral signs); c_i is observed by firm i , but not by the other firms; p_i is common knowledge; everybody knows that c_i is drawn from this distribution; that everybody knows that c_i is drawn from this distribution, etc. ...⁵ In this case firm i 's type is fully summarized by c_i : because the probability distributions are common knowledge, knowing c_i amounts to knowing everything known by firm i . By abuse of terminology, one can identify firm i 's type with the realization of c_i .

More generally, Harsanyi assumed that the player's types $\{t_i\}_{i=1}^n$ are drawn from some objective distribution $p(t_1, \dots, t_n)$, where t_i belongs to some space T_i . For simplicity, let us assume that T_i has a finite number $|T_i|$ of elements. t_i is observed by player i only. $p_i(t_{-i} | t_i)$ denotes player i 's conditional probability about his opponent's types $t_{-i} = (t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_n)$ given his type t_i .

To complete the description of a Bayesian game, we must specify an action set A_i (with elements a_i) and an objective function $\Pi_i(a_1, \dots,$

⁵Aumann [1976] formalizes common knowledge in the following way. Let (Ω, p) be a finite probability space and let P and Q denote two partitions of Ω representing the informations of two players. Let R denote the meet of P and Q (i.e., the finest common coarsening of P and Q). An event E is common knowledge between the two players at $\omega \in \Omega$ if the event in R that includes ω is itself included in E .

a_n, t_1, \dots, t_n) for each player i . The action spaces A_i , the objective functions Π_i and the probability distribution p are common knowledge (every player knows them, knows that everybody knows them, ...). In other words, everything which is not commonly known is subsumed in the type.

The Harsanyi formulation looks a priori restrictive because it presumes a large common knowledge base. However, as Mertens and Zamir [1983] have shown (see also Brandenburger-Deckel [1985]), one can always define type spaces that are large enough to describe every element of player i 's private information. Coming back to our original idea, player i 's type then includes his beliefs about the other players' beliefs about payoff relevant information, his beliefs about the other players' beliefs about payoff relevant information etc. Mertens and Zamir essentially show that this infinite regression is well defined (under some weak assumptions, the space of types is compact for the product topology).

In this section we consider only one-shot simultaneous move games of incomplete information. The n players first learn their types and then simultaneously choose their actions (note that the game is also a game of imperfect information). The game is static in that the players are unable to react to their opponent's actions. The inference process as to the other players' types is irrelevant because the game is over at the time each player learns some signal related to his opponents' moves. Section 4 considers dynamic games and the associated updating process.

Each player's decision naturally depends on his information, i.e., his type. For instance, a high cost firm chooses a high price. Let $a_i(t_i)$ denote the action chosen by player i when his type is t_i (this could also denote a mixed strategy, i.e., a randomization over actions for a given type). If he knew the strategies adopted by the other players $\{a_j(t_j)\}_{j \neq i}$ as a func-

tion of their types, player i would be facing a simple decision problem; given his type t_i , he ought to maximize:

$$\sum_{t_{-i}} p_i(t_{-i} | t_i) U_i(a_1(t_1), \dots, a_i, \dots, a_n(t_n), t_1, \dots, t_i, \dots, t_n).$$

Harsanyi extended the idea of a Nash equilibrium by assuming that each player correctly anticipates how each of his opponents behaves as a function of his type:

Definition: A Bayesian equilibrium is a set of (type contingent) strategies $\{a_i^*(t_i)\}_{i=1}^n$ such that $a_i^*(t_i)$ is player i 's best response to the other strategies when his type is t_i :

$$a_i^*(t_i) \in \arg \max_i \sum_{t_{-i}} p_i(t_{-i} | t_i) U_i(a_1(t_1), \dots, a_i, \dots, a_n(t_n), t_1, \dots, t_i, \dots, t_n).$$

Thus, the Bayesian equilibrium concept is a straightforward extension of the Nash equilibrium concept, in which each player recognizes that the other player's strategies depend on their types.

Proving existence of a Bayesian equilibrium turns out to involve a simple extension of the proof of a Nash equilibrium. The trick is the following: since player i 's optimal action depends on type t_i , everything is as if player i 's opponents were playing against $|T_i|$ different players, each of these players being drawn and affecting his opponent's payoffs with some probability. Thus, considering different types of the same player as different players leads to transform the original game into a game with $\{\sum_{i=1}^n |T_i|\}$ players. Each "player" is then defined by a name and a type. He does not care (directly) about the action of a player with the same name and a different type (another incarnation of himself), but he does care about the other players' actions. If a_{i,t_i} denotes the action chosen by player $\{i, t_i\}$, play-

er $\{i, t_i\}$'s objective is to maximize over a:

$$\sum_{t_{-i}} p_i(t_{-i} | t_i) U_i(a_1, t_1, \dots, a_n, t_n, t_1, \dots, t_i, \dots, t_n).$$

Thus, existence of a Bayesian equilibrium of a game with $|T_i|$ players stems directly from the existence of a Nash equilibrium for a game with $\{\Sigma_i | T_i | \}$ players, as long as the numbers of players and types are finite.

With a continuum of types, some technicalities appear about whether there exists a measurable structure over the set of random variables (Aumann [1964]). One is then led to define a mixed strategy as a measurable function from $[0, 1] \times T_i$ into A_i . Or, equivalently, one can define it, as Milgrom and Weber [1985] do, as a measure on the subjects of $T_i \times A_i$ for which the marginal distribution on T_i is p_i . Milgrom and Weber give sufficient conditions for the existence of an equilibrium in such settings.

Example 1: Consider a duopoly playing Cournot (quantity) competition. Let firm i 's profit be quadratic: $\Pi_i = q_i(t_i - q_i - q_j)$, where t_i is the difference between the intercept of the linear demand curve and firm i 's constant unit cost ($i = 1, 2$) and q_i is the quantity chosen by firm i ($a_i = q_i$). It is common knowledge that, for firm 1, $t_1 = 1$ ("firm 2 has complete information about firm 1", or "firm 1 has only one potential type"). Firm 2, however, has private information about its unit cost. Firm 1 only knows that $t_2 = 3/4$ with probability $1/2$. Thus, firm 2 has two potential types, which we will call the "low cost type" ($t_2 = 5/4$) and the "high cost type" ($t_2 = 3/4$). The two firms choose their outputs simultaneously. Let us look for a pure strategy equilibrium. Firm 1 plays q_1 , firm 2 plays q_2^L (if $t_2 = 5/4$) or q_2^H (if $t_2 = 3/4$). Let us start with firm 2:

$$q_2(t_2) \in \arg \max_{q_2} \{q_2(t_2 - q_2 - q_1)\} \Rightarrow q_2(t_2) = (t_2 - q_1)/2.$$

Let us now consider firm 1, which does not know which type it faces.

$$q_1 \in \arg \max_{q_1} \left\{ \frac{1}{2} q_1 (1 - q_1 - q_2^H) + \frac{1}{2} q_1 (1 - q_1 - q_2^L) \right\}$$

$$\Rightarrow q_1 = (1 - E q_2) / 2,$$

where $E(\cdot)$ denotes an expectation over firm 2's types. But $E q_2 = \frac{1}{2} q_2^H + \frac{1}{2} q_2^L = (E t_2 - q_1) / 2 = (1 - q_1) / 2$. One thus obtains $\{q_1 = 1/3, q_2^L = 11/24, q_2^H = 5/24\}$ as a Bayesian equilibrium (one can prove this equilibrium is unique). This simple example illustrates how one can compute the Bayesian equilibrium as a Nash equilibrium of a 3-player game ($|T_1| = 1, |T_2| = 2$).

Example 2: Consider an incomplete-information version of the war of attrition discussed in Section 2d. Firm i chooses a number a_i in $[0, +\infty)$. Both firms choose simultaneously. The payoffs are:

$$\Pi_i = \begin{cases} -a_i, & \text{if } a_j > a_i \\ t_i - a_j, & \text{if } a_j < a_i. \end{cases}$$

t_i , firm i 's type, is private information and take values in $[0, +\infty)$ with cumulative distribution function $P_i(t_i)$ and density $p_i(t_i)$. Types are, as in example 1, independent between the players. t_i is the price to the winner, i.e., the highest bidder. The game resembles a second-bid auction in that the winner pays the second bid. However, it differs from the second-bid auction in that the loser also pays the second bid.

Let us look for a Bayesian equilibrium of this game. Let $a_i(t_i)$ denote firm i 's strategy. Then, we require

$$a_i(t_i) \in \arg \max_{a_i} \{-a_i \text{Prob}(a_j(t_j) > a_i) + \int_{\{t_j | a_j(t_j) < a_i\}} (t_i - a_j)\}$$

A few tricks make the problem easy to solve. First, one can write the "self-selection constraints": by definition of equilibrium, type t_i prefers $a_i(t_i)$ to $a_i(t'_i)$, and type t'_i prefers $a_i(t'_i)$ to $a_i(t_i)$. Writing the two corresponding inequalities and adding them up shows that a_i must be a non-decreasing function of t_i . Second, it is easy to show that there can not be an atom at $a_i > 0$, i.e., $\text{Prob}(a_j(t_j) = a_i > 0) = 0$. To prove this, notice that if there were an atom of types of firm j playing a_i , firm i would never play in $[a_i - \epsilon, a_i)$ for ϵ small: it would be better off bidding just above a_i (the proof is a bit loose here, but can be made rigorous). Thus, the types of firms that play a_i would be better off playing $(a_i - \epsilon)$, because this would not reduce the probability of winning and would lead to reduced payments.

Let us look for a strictly monotonic, continuous function $a_i(t_i)$ with inverse $t_i = \Phi_i(a_i)$. Thus, $\Phi_i(a_i)$ is the type that bids a_i . We then obtain:

$$a_i(t_i) \in \arg \max_{a_i} \{-a_i(1 - P_j(\Phi_j(a_i))) + \int_0^{a_i} (t_i - a_j) p_j(\Phi_j(a_j)) \Phi'_j(a_j) da_j\}.$$

By differentiating, one obtains a system of two differential equations in $\Phi_1(\cdot)$ and $\Phi_2(\cdot)$ (or, equivalently, in $a_1(\cdot)$ and $a_2(\cdot)$). Rather than doing so, let us take the following intuitive approach: If firm i , with type t_i , bids $(a_i + da_i)$ instead of a_i , it loses da_i with probability 1 (since there is no atom), conditionally on firm j bidding at least a_i (otherwise this increase has no effect). It gains $t_i (= \Phi_i(a_i))$ with probability $\{p_j(\Phi_j(a_i)) \Phi'_j(a_i) / (1 - P_j(\Phi_j(a_i)))\} da_i$. Thus, in order for firm i to be indifferent:

$$\Phi_i(a_i)P_j(\Phi_j(a_i))\Phi_j'(a_i) = 1 - P_j(\Phi_j(a_i)) .$$

We leave it to the reader to check that, for a symmetric exponential distribution $P_i(t_i) = 1 - e^{-t_i}$, there exists a symmetric equilibrium: $\Phi_i(a_i) = \sqrt{2a_i}$, which corresponds to $a_i(t_i) = \frac{t_i^2}{2}$ (as Riley [1980] has shown, there also exists a continuum of asymmetric equilibria:

$$\Phi_1 = K\sqrt{a_1} \text{ and } \Phi_2 = \frac{2}{K} \sqrt{a_2} \text{ for } K > 0).$$

Let us now give an industrial organization interpretation of the game. Suppose that there are two firms in the market; they both lose 1 per unit of time when they compete; they make a monopoly profit when their opponent has left the market, the present discounted value of which is t_i (it would make sense to assume that the duopoly and monopoly profit are correlated, but such a modification would hardly change the results). The firms play a war of attrition. a_i is the time firm i intends to stay in the market, if firm j has not exited before. At this stage, the reader may wonder about our dynamic interpretation: if firms are free to leave when they want and are not committed to abide by their date 0 choice of a_i , is the Bayesian equilibrium "perfect"? It turns out that the answer is "yes"; the dynamic game is essentially a static game (which is the reason why we chose to present it in this section). At any time a_i , either firm j has dropped out (bid less than a_i) and the game is over, or firm j is still in the market and the conditional probability of exit is the one computed earlier. Thus the equilibrium is

perfect as well.⁶

3B. Using Bayesian Equilibria to Justify Mixed Equilibria

In Section 1, we saw that simultaneous move games of complete information often admit mixed strategy equilibria. Some researchers are unhappy with this notion because, they argue, "real world decision makers do not flip a coin." However, as Harsanyi [1973] has shown, mixed strategy equilibria of complete information games can often be vindicated as the limits of pure strategy equilibria of slightly perturbed games of incomplete information. Indeed, we have already noticed that in a Bayesian game, once the players' type-contingent strategies have been computed, each player behaves as if he were facing mixed strategies by his opponents (nature creates uncertainty through its choice of types rather than the choice of the side of the coin).

To illustrate the mechanics of this construction, let us consider the one-period version of the "grab-the-dollar" game introduced in Section 2. Each player has two possible actions: investment, no investment. In the complete information version of the game, a firm gains 1 if it is the only one to make the investment (wins), loses 1 if both invest, and breaks even if it does not invest. (We can view this game as an extremely crude representation of a natural monopoly market.) The only symmetric equilibrium involves mixed strategies: each firm invests with probability $1/2$. This clearly is an equilibrium: each firm makes 0 if it does not invest, and $\frac{1}{2}(1) + \frac{1}{2}(-1)$

⁶The war of attrition was introduced in the theoretical biology literature (e.g., Maynard Smith [1974], Riley [1980] and has known many applications since. It was introduced in industrial organization by Kreps-Wilson [1982a]. (See also Nalebuff [1982] and Ghemawat-Nalebuff [1985]). For a characterization of the set of equilibria and a uniqueness result with changing duopoly payoffs and/or large uncertainty over types, see Fudenberg-Tirole [1986]. See also Hendricks and Wilson [1985a, 1985b].

= 0 if it does not. Now consider the same game with the following type of incomplete information: Each firm has the same payoff structure except that, when it wins, it gets $(1+t)$ where t is uniformly distributed on $[-\epsilon, +\epsilon]$. Each firm knows its type t , but not that of the other firm. Now, it is easily seen that the symmetric pure strategies: "a($t < 0$) = do not invest, a($t > 0$) = invest" form a Bayesian equilibrium. From the point of view of each firm, the other firm invests with probability $1/2$. Thus, the firm should invest if and only if $\frac{1}{2}(1+t) + \frac{1}{2}(-1) > 0$, i.e., $t > 0$. Last, note that, when ϵ converges to zero, the pure strategy Bayesian equilibrium converges to the mixed strategy Nash equilibrium of the complete information game.

As another example, the reader may want to study the symmetric war of attrition. Under complete information and symmetric payoffs, it is easily shown that in a symmetric equilibrium, each player's strategy is a mixed strategy with exponential distribution over possible times.⁷ The symmetric incomplete information equilibrium (computed in Section 3A) converges to this mixed strategy equilibrium when the uncertainty converges to zero (see Milgrom-Weber for the case of a uniform distribution).

Milgrom-Weber [1985] offers sufficient (continuity) conditions on the objective functions and information structure so that the limit of Bayesian equilibrium strategies when the uncertainty becomes "negligible," forms a Nash equilibrium of the limit complete information game. (Note: the war of attrition does not satisfy their continuity conditions; but as Milgrom and

⁷Letting t denote the common payoff to winning, waiting da more yields $(x(a)t)da$ where $x(a)da$ is the probability that the opponent drops between a and $(a+da)$. This must equal the cost of waiting: da . Thus, $x(a) = 1/t$ is independent of time a .

Weber show, the result holds anyway). They also identify a class of (atomless) games for which there exists a pure strategy equilibrium.

We must realize that games of complete information are an idealization. In practice, everyone has at least a slight amount of incomplete information about the others' objectives; Harsanyi's argument shows that it is hard to make a strong case against mixed strategy equilibria on the grounds that they require a randomizing device.

4. Dynamic Games of Incomplete Information

We now study games in which, at some point of time, a player bases his decision on a signal that conveys information about another player. This type of game is dynamic in that a player reacts to another player's move. The tricky aspect of it is that, under incomplete information, the former must apply Bayes rule to update his beliefs about the latter's type. To do so, he uses the latter's choice of action (or a signal of it) and equilibrium strategy, as we shall see shortly. The equilibrium notion for dynamic games of incomplete information is naturally a combination of the subgame perfect equilibrium concept that we discussed earlier and Harsanyi [1967]'s concept of Bayesian equilibrium for games of incomplete information. In this section we consider the simplest such notion, that of perfect Bayesian equilibrium concept, as well as some easy-to-apply (and sometimes informal) refinements.

In the next section, we will discuss more formal refinements of the perfect Bayesian equilibrium concept for finite games.

The notion of a perfect Bayesian equilibrium was developed under various names and in various contexts in the late 'sixties and the 'seventies. In economics, Akerlof [1970] and Spence [1974]'s market games and Ortega-Reichert [1967]'s analysis of repeated first-bid auctions make implicit use

of the concept. In industrial organization the first and crucial application is Milgrom-Roberts [1982a]'s limit pricing paper, followed by the work of Kreps-Wilson [1982a] -Milgrom-Roberts [1982b] on reputation. In game theory, Selten [1975] introduced the idea of trembles to refine the concept of subgame perfect equilibria in games without (many) proper subgames. (If each player's type is private information, the only proper subgame is the whole game, so subgame perfection has no force). Kreps-Wilson [1982b]'s sequential equilibrium is similar, but, in the tradition of the economics literature, it emphasizes the formation of beliefs, which makes the introduction of refinements easier to motivate. We should also mention the work of Aumann-Machler [1967] on repeated games of incomplete information.

We start this section with the simplest example of a dynamic game of incomplete information, the signaling game. This is a two-period leader-follower game in which the leader is endowed with private information that affects the follower. We give some examples of such games and introduce some refinements of the equilibrium concept. As the principles enunciated here for signaling games generally carry over to general games, we do not treat the latter in order to save on notation and space.

4A. The Basic Signaling Game

As mentioned earlier, the simplest game in which the issues of updating and perfection arise simultaneously has the following structure: There are two players; player 1 is the leader (also called "sender", because he sends a signal) and player 2 the follower (or "receiver"). Player 1 has private information about his type t_1 in T_1 , and chooses action a_1 in A_1 . Player 2, whose type is common knowledge for simplicity, observes a_1 and chooses a_2 in A_2 . Payoffs are equal to $\Pi_i(a_1, a_2, t_1)$ ($i = 1, 2$). Before the game begins,

player 2 has prior beliefs $p_1(t_1)$ about player 1's type.

Player 2, who observes player 1's move before choosing his own action, should update his beliefs about t_1 and base his choice of a_2 on the posterior distribution $\tilde{p}_1(t_1 | a_1)$. How is this posterior formed? As in a Bayesian equilibrium, player 1's action ought to depend on his type; let $a_1^*(t_1)$ denote this strategy (as before, this notation allows a mixed strategy). Thus, figuring out $a_1^*(\cdot)$ and observing a_1 , player 2 can use Bayes' rule to update $p_1(\cdot)$ into $\tilde{p}_1(\cdot | a_1)$. And, in a rational expectations world, player 1 should anticipate that his action would affect player 2's also through the posterior beliefs. Thus, the natural extension of the Nash equilibrium concept to the signaling game is:

Definition: A perfect Bayesian equilibrium (PBE) of the signaling game is a set of strategies $a_1^*(t_1)$ and $a_2(a_1)$ and posterior beliefs $\tilde{p}_1(t_1 | a_1)$ such that:

$$(P_1) \quad a_2^*(a_1) \in \arg \max_{a_2} \sum_{t_1} \tilde{p}_1(t_1 | a_1) \Pi_2(a_1, a_2, t_1)$$

$$(P_2) \quad a_1^*(t_1) \in \arg \max_{a_1} \Gamma_1(a_1, a_2^*(a_1), t_1)$$

(B) $\tilde{p}_1(t_1 | a_1)$ is derived from the prior $p_1(\cdot)$, a_1 and $a_1^*(\cdot)$ using Bayes' rule (when applicable).

(P_1) and (P_2) are the perfectness conditions. (P_1) states that player 2 reacts optimally to player 1's action given his posterior beliefs about t_1 . (P_2) demonstrates the optimal Stackelberg behavior by player 1; note that he takes into account the effect of a_1 on player 2's action. (B) corresponds to the application of Bayes' rule. The quantifier "when applicable" stems from

the fact that, if a_1 is not part of player 1's optimal strategy for some type, observing a_1 is a zero-probability event and Bayes rule does not pin down posterior beliefs. Any posterior beliefs $\tilde{p}_1(\cdot|a_1)$ are then admissible. Indeed, the purpose of the refinements of the perfect Bayesian equilibrium concept is to put some restrictions on these posterior beliefs.

Thus, a PBE is simply a set of strategies and beliefs such that, at any stage of the game, strategies are optimal given beliefs and beliefs are obtained from equilibrium strategies and observed actions using Bayes' rule. Two features of the concept developed thus far should be emphasized:

First, a PBE has a strong fixed point flavor. Beliefs are derived from strategies, which are optimal given beliefs. For this reason, there exists no handy algorithm to help us construct equilibria. Remember that for games of complete information, Kuhn's algorithm of backward induction gave us the set of perfect equilibria. Here we must also operate the Bayesian updating in a forward manner. This makes the search for equilibria rely on a few tricks (to be discussed latter) rather than on a general method.

Second, too little structure has been imposed on the type and action spaces and on the objective functions to prove existence of a PBE. Actually, existence theorems are available only for games with a finite number of types and actions (see subsection 4E). Most applications, however, involve either a continuum of types or/and a continuum of actions. Existence is then obtained by construction, on a case by case basis.

For more general games than the signaling game, the definition of a PBE is the same: At each information set posterior beliefs are computed using optimal strategies and the information at the information set. And, strategies are optimal given beliefs. We will not give the formal definition of

this because it involves nothing more than a (very heavy) extension of the notation.

Let us now give simple examples of PBE in signaling games. From now on, we delete the subscript on player 1's type, as there is no possible confusion.

4B. Examples

Example 1: a two-period reputation game. The following is a much simplified version of the Kreps-Wilson-Milgrom-Roberts reputation story. There are two firms ($i = 1, 2$). In period 1, they are both in the market. Only firm 1 (the "incumbent") takes an action a_1 . The action space has two elements: "prey" and "accomodate." Firm 2 (the "entrant")'s profit is D_2 if firm 1 accomodates and P_2 if firm 1 preys, such that $D_2 > 0 > P_2$. Firm 1 has one of two potential types t_1 : "sane" and "crazy." When sane, firm 1 makes D_1 when it accomodates and P_1 when it preys, where $D_1 > P_1$. Thus, a sane firm prefers to accomodate rather than preying. However, it would prefer to be a monopoly, in which case it would make M_1 per period. When crazy, firm 1 enjoys predation and thus preys (its utility function is such that it is always worth preying). Let p_1 (respectively, $(1-p_1)$) denote the prior probability that firm 1 is sane (respectively, crazy).

In period 2, only firm 2 chooses an action a_2 . This action can take two values: "stay" and "exit." If it stays, it obtains a payoff D_2 if firm 1 is actually sane, and P_2 if it is crazy (the idea is that unless it is crazy, firm 1 will not pursue any predatory strategy in the second period because there is no point building or keeping a reputation at the end. This assumption can be derived more formally from the description of the second-period competition). The sane firm gets D_1 if firm 2 stays and $M_1 > D_1$ if firm 2

exits. We let δ denote the discount factor between the two periods.

We presumed that the crazy type always preys. The interesting thing to study is thus the sane type's behavior. From a static point of view, it would want to accommodate in the first period; however, by preying it might convince firm 2 that it is of the crazy type, and thus induce exit (as $P_2 < 0$) and increase its second-period profit.

Let us first start with a taxonomy of potential perfect Bayesian equilibria. A separating equilibrium is an equilibrium in which firm 1's two types choose two different actions in the first period. Here, this means that the sane type chooses to accommodate. Note that in a separating equilibrium, firm two has complete information in the second period:

$$\tilde{p}_1(t = \text{sane} | a_1 = \text{accomodate}) = 1 \text{ and } \tilde{p}_1(t = \text{crazy} | a_1 = \text{prey}) = 1.$$

A pooling equilibrium is an equilibrium in which firm 1's two types choose the same action in the first period. Here, this means that the sane type preys. In a pooling equilibrium, firm 2 does not update its beliefs when observing the equilibrium action: $\tilde{p}_1(t = \text{sane} | a_1 = \text{prey}) = p_1$. Last, there can also exist hybrid or semi-separating equilibria. For instance, in the reputation game, the sane type may randomize between preying and accommodating, i.e., between pooling and separating. One then has

$$\tilde{p}_1(t = \text{sane} | a_1 = \text{prey}) \in (0, p_1) \text{ and } \tilde{p}_1(t = \text{sane} | a_1 = \text{accomodate}) = 1.$$

Let us first look for conditions of existence of a separating equilibrium. In such an equilibrium, the sane type accommodates and thus reveals its type and obtains $D_1(1+\delta)$ (firm 2 stays because it expects $D_2 > 0$ in the second period). If it decided to prey, it would convince firm 2 that it is crazy and would thus obtain $P_1 + \delta M_1$. Thus, a necessary condition for the

existence of a separating equilibrium is:

$$(6) \quad \delta(M_1 - D_1) < (D_1 - P_1).$$

Conversely, suppose that (6) is satisfied. Consider the following strategies and beliefs: the sane incumbent accommodates, and the entrant (correctly) anticipates that the incumbent is sane when observing accommodation; the crazy incumbent preys and the entrant (correctly) anticipates that the incumbent is crazy when observing predation. Clearly, these strategies and beliefs form a separating PBE.

Let us now look at the possibility of a pooling equilibrium. Both types prey; thus, as we saw, $\tilde{p}_1 = p_1$ when predation is observed. Now, the sane type, who loses $(D_1 - P_1)$ in the first period, must induce exit. Thus, it must be the case that:

$$(7) \quad p_1 D_2 + (1 - p_1) P_2 < 0.$$

Conversely, assume that (7) holds, and consider the following strategies and beliefs: both types prey; the entrant has posterior beliefs $\tilde{p}_1 = p_1$ when predation is observed and $\tilde{p}_1 = 1$ when accommodation is observed. The sane type's equilibrium profit is $P_1 + \delta M_1$ while it would become $D_1(1 + \delta)$ under accommodation. Thus, if (6) is violated, the proposed strategies and beliefs form a pooling PBE (note that if (7) is satisfied with equality, there exists not one, but a continuum of such equilibria). So the equilibrium is that the entrant never enters and the incumbent never preys.

We leave it to the reader to check that if both (6) and (7) are violated, the unique equilibrium is a hybrid PBE (with the entrant's randomizing when observing predation).

Remark: The (generic) uniqueness of the PBE in this model is due to the fact that the "strong" type (the crazy incumbent) is assumed to prey no matter what. Thus, predation is not a zero probability event and, furthermore, accomodation is automatically interpreted as coming from the sane type if it belongs to the equilibrium path. The next example illustrates a more complex and a more common structure, for which refinements of the PBE are required. Example 2, which, in many respects, can be regarded as a generalization of example 1 also involves several cases resembling those in example 1.

Example 2: The Limit-Pricing Game

As mentioned earlier, the paper which introduced signaling games into the industrial organization field is Milgrom-Roberts [1982a]'s article on limit pricing. Let us take the following simple version of their two-period model. Firm 1, the incumbent, has in the first period a monopoly power and chooses a first-period quantity $a_1 = q_1$. Firm 2, the entrant, then decides to enter or to stay out in the second period (thus, as in the previous game, $a_2 = 0$ or 1 or $\varepsilon[0,1]$ if we allow mixed strategies). If it enters, there is duopolistic competition in period two. Otherwise, firm 1 remains a monopoly.

Firm 1 can have one of two potential types: its constant unit production cost is "high" (H) with probability p_1 and "low" (L) with probability $(1-p_1)$. We will denote by q_m^t the monopoly quantities for the two types of incumbent ($t = H, L$). Naturally, $q_m^H < q_m^L$. We let $M_1^t(q_1)$ denote the monopoly profit of type t when producing q_1 ; in particular, let $M_1^t = M_1^t(q_m^t)$ denote type i 's monopoly profit when it maximizes its short-run profit. We assume that $M_1^t(q_1)$ is strictly concave in q_1 .

Firm 1 knows t_1 from the start; firm 2 does not. Let D_2^t denote firm 2's duopoly profit when firm 1 has type t (it possibly includes entry costs). To make things interesting, let us assume that firm 2's entry decision is influenced by its beliefs about firm 1's type: $D_2^H > 0 > D_2^L$. The discount factor is δ .

Let us look for separating equilibria. For this, we first obtain two necessary conditions: that each type does not want to pick the other type's equilibrium action ("incentive constraints"). We then complete the description of equilibrium by choosing beliefs off-the-equilibrium path that deter the two types from deviating from their equilibrium actions. Thus, our necessary conditions are also sufficient, in the sense that the corresponding quantities are equilibrium quantities. In a separating equilibrium, the high-cost type's quantity induces entry. He thus plays q_m^H (if it did not, he could increase his first-period profit without adverse effect on entry). Thus, he gets $\{M_1^H + \delta D_1^H\}$. Let q_1^L denote the output of the low-cost type. The high-cost type, by producing this output, deters entry and obtains $\{M_1^H(q_1^L) + \delta M_1^H\}$. Thus, a necessary condition for equilibrium is:

$$(8) \quad M_1^H - M_1^H(q_1^L) > \delta(M_1^H - D_1^H).$$

The similar condition for the low-cost type is:

$$(9) \quad M_1^L - M_1^L(q_1^L) < \delta(M_1^L - D_1^L).$$

To make things interesting, we will assume that there is no (separating) equilibrium in which each type behaves as in a full information context; i.e., the low-cost type would wish to pool:

$$(10) \quad M_1^H - M_1^H(q_m^L) < \delta(M_1^H - D_1^H).$$

To characterize the set of q_1^L satisfying (8) and (9), one must make more specific assumptions on the demand and cost functions. We will not do it here, and we refer to the literature for this. We just note that, under reasonable conditions, (8) and (9) define a region $[\tilde{q}_1, \bar{q}_1]$, where $\tilde{q}_1 > q_m^L$. Thus, to separate, the low-cost type must produce sufficiently above its monopoly quantity so as to make pooling very costly to the high-cost type. A crucial assumption in the derivation of such an interval is the Spence-Mirrlees (single-crossing) condition:

$$\frac{\partial}{\partial q_1} (M_1^L(q_1) - M_1^H(q_1)) > 0.$$

\tilde{q}_1 is such that (8) is satisfied with equality; it is called the "least-cost" separating quantity, because, of all potential separating equilibria, the low-cost type would prefer the one at \tilde{q}_1 .

Let us now show that these necessary conditions are also sufficient. Let the high cost type choose q_m^H and the low-cost type choose q_1^L in $[\tilde{q}_1, \bar{q}_1]$. When a quantity that differs from these two quantities is observed, beliefs are arbitrary. The easiest way to obtain equilibrium is to choose beliefs that induce entry; this way, the two types will be little tempted to deviate from their presumed equilibrium strategies; so let us specify that when q_1 does not belong to $\{q_m^H, q_1^L\}$, $\tilde{p}_1 = 1$ (firm 2 believes firm 1 has high-cost); whether these beliefs, which are consistent with Bayes' rule, are "reasonable," is discussed later on. Now, let us check that no type wants to deviate. The high-cost type obtains its monopoly profits in the first period

and, thus, is not willing to deviate to another quantity that induces entry. He does not deviate to q_1^L either from (8). And similarly for the low-cost type. Thus, we have obtained a continuum of separating equilibria.

Note that this continuum of separating equilibria exists for any $p_1 > 0$. By contrast, for $p_1 = 0$, the low-cost firm plays its monopoly quantity q_m^L . We thus observe that a tiny change in the information structure may make a huge difference. A very small probability that the firm has high cost may force the low cost firm to increase its production discontinuously to signal its type. Games of incomplete information (which include games of complete information!) are very sensitive to the specification of the information structure, a topic we will come back to later on.

Note also that Pareto dominance selects the least cost separating equilibrium among separating equilibria. [The entrant has the same utility in all separating equilibria (the informative content is the same); similarly, the high cost type is indifferent. The low cost type prefers lower outputs].

The existence of pooling equilibria hinges on whether the following condition is satisfied.

$$(11) \quad p_1 D_2^H + (1-p_1) D_2^L < 0.$$

Assume that condition (11) is violated (with a strict inequality -- we will not consider the equality case for simplicity). Then, at the pooling quantity, firm 2 makes a strictly positive profit if it enters (as $\tilde{p}_1 = p_1$). This means that entry is not deterred, so that the two types can not do better than choosing their (static) monopoly outputs. As these outputs differ, no pooling equilibrium can exist.

Assume, therefore, that (11) is satisfied so that a pooling quantity q_1

deters entry. A necessary condition for a quantity q_1 to be a pooling equilibrium quantity is that none of the types want to play his static optimum. If he were to do so, it would at worst deter entry. Therefore, q_1 must satisfy (9) and the analogous condition for the high-cost type:

$$(12) \quad M_1^H - M_1^H(q_1) < \delta(M_1^H - D_1^H).$$

Again, the set of outputs q_1 that satisfy both (9) and (12) depends on the cost and demand functions. Let us simply notice that, from (10), there exists an interval of outputs around q_m^L that satisfy these two inequalities.

Now it is easy to see that if q_1 satisfies (9) and (12), q_1 can be made part of a pooling equilibrium. Suppose that whenever firm 1 plays an output differing from q_1 (an off-the-equilibrium path action), firm 2 believes firm 1 has a high cost. Firm 2 then enters, and firm 1 might as well play its monopoly output. Thus, from (9) and (12), none of the types would want to deviate from q_1 .

We leave it to the reader to derive hybrid equilibria (the analysis is very similar to the previous ones). We now investigate the issue of refinements.

4C. Some Refinements

Games of incomplete information in general have many PBE. The reason why this is so is easy to grasp. Consider the basic signaling game and suppose that one wants to rule out some action a_1 by player 1 as an equilibrium action. If, indeed, a_1 is not played on the equilibrium path, player 2's beliefs following a_1 are arbitrary. In most games there exists some type t such that if player 2 puts all the weight on t , it takes an action that is detrimental for all types of player 1 (for instance, t is the high cost type

in the limit pricing game; it induces entry). As playing a_1 produces a bad outcome for player 1, not playing a_1 on the equilibrium path may be self-fulfilling. Some authors have noted that, while non credible actions were ruled out by the perfectness part of the PBE, players could still "threaten" each other through beliefs. This subsection and 4D discuss refinements that select subsets of PBEs.

Often, however, the very structure of the game tells us that some beliefs, while allowable because off-the-equilibrium path, "do not make sense." Over the years intuitive criteria for selection of beliefs have been developed for each particular game. We mention here only a few of these criteria. These criteria, which apply to all types of games (including games with a continuum of types or actions), are sometimes informal in that they have not been designed as part of a formal solution concept for which existence has been proved. But most of them are, for finite games, satisfied by the Kohlberg-Mertens [1986] concept of stable equilibria, which are known to exist (see subsection 4E below). Last, we should warn the reader that the presentation below resembles more a list of cookbook recipes than a unified methodological approach.

i) Elimination of Weakly Dominated Strategies

In the tradition of Luce-Raiffa [1957], Farquharson [1969], Moulin [1979], Bernheim [1984], and Pierce [1984], it seems natural to require that, when an action is dominated for some type, but not for some other, the posterior beliefs should not put any weight on the former type. This simple restriction may already cut on the number of PBE considerably. Consider the limit pricing game. Quantities above \tilde{q}_1 are dominated for the high cost type

(if this type chooses q_m^H , its intertemporal profit is $M_1^H + \delta D_1^H$; if it chooses q_1 , this profit does not exceed $M_1^H(q_1) + \delta M_1^H$; from the definition of \tilde{q}_1 , the second action is weakly dominated for $q_1 > \tilde{q}_1$). Thus, when q_1 belongs to $[\tilde{q}_1, \bar{q}_1]$, the entrant should believe that the incumbent's cost is low, and should not enter. Thus, the low cost incumbent need not produce above \tilde{q}_1 to deter entry. We thus see that we are left with a single separating PBE instead of a continuum (this reasoning is due to Milgrom-Roberts).

A small caveat here: playing a quantity above \tilde{q}_1 is dominated for the high cost type only once the second period has been folded back. Before that, one can think of (non-equilibrium) behavior which would not make such a quantity a dominated strategy. For instance, following q_m^H , the entrant might enter and charge a very low price. So, we are invoking a bit more than the elimination of dominated strategies. A quantity above \tilde{q}_1 is dominated conditional on subsequent equilibrium behavior -- a requirement in the spirit of perfectness. More generally, one will want to iterate the elimination of weakly dominated strategies.

Note that, in the limit pricing game, the elimination of weakly dominated strategies leaves us with the "least cost" separating equilibrium, but does not help us select among the pooling equilibria. This is because the equilibrium pooling quantities are not dominated for the high cost type.

ii) Elimination of Equilibrium Weakly Dominated Strategies (Intuitive Criterion)

The next criterion was proposed by Kreps [1984] to single out a property satisfied by the more stringent stability requirement of Kohlberg-Mertens [1986] and, thus, to simplify its use in applications of game theory. The

idea is roughly to extend the elimination of weakly dominated strategies to strategies which are dominated relative to equilibrium payoffs. So doing eliminates more strategies and thus refines the equilibrium concept further.

More precisely, consider the signaling game and a corresponding PBE and associated payoffs. Let a_1 denote an out of equilibrium action which yields for a subset J of types payoffs lower than their equilibrium payoffs whatever beliefs player 2 forms after observing a_1 .

More formally, let $\Pi_1^*(t)$ denote player 1's equilibrium payoff when he has type t . Let $BR(\tilde{p}_1, a_1) = \arg \max_{a_2 \in A_2} \{\Sigma \tilde{p}_1(t) \Pi_2(a_1, a_2, t)\}$ denote player 2's

best response(s) when he has posterior beliefs $\tilde{p}_1(\cdot)$; and let $BR(I, a_1) =$

$\bigcup_{\{\tilde{p}_1 : \tilde{p}_1(I)=1\}} BR(\tilde{p}_1, a_1)$ denote the set of player 2's best responses when his

posterior beliefs put all the weight in a subset I of types.

Suppose that there exists a subset J of T such that:

(1) For all t in J and for all a_2 in $BR(T, a_1)$, $\Pi_1(a_1, a_2, t) < \Pi_1^*(t)$.

(2) There exists a type t in $T-J$ such that for all a_2 in $BR(T-J, a_1)$,

$$\Pi_1(a_1, a_2, t) > \Pi_1^*(t).$$

From condition (1), we know that no type in J would want to deviate from his equilibrium path, whatever inference player 2 would make following the deviation. It thus seems logical that player 2 does not put any weight on types in J . But, one would object, no type outside J may gain from the deviation either. This is why condition (2) is imposed. There exists some type outside J , this type strictly gains from the deviation. The intuitive criterion rejects PBE that satisfy (1) and (2) for some action a_1 and some subset J .

One immediately sees that this criterion has most power when there are only two potential types (see below for an application to the limit pricing game). The subsets J and $T-J$ of the criterion are then necessarily composed of one type each. Thus, the requirement "for all a_2 in $BR(T-J, a_1)$..." in condition 2 is not too stringent, and the criterion has much cutting power. With more than two types, however, there may exist many a_2 in $BR(T-J, a_1)$ and, therefore, the requirement that some type prefers the deviation for all a_2 in $BR(T-J, a_1)$ becomes very strong. The refinement then loses some of its power.

Cho [1985] and Cho-Kreps [1987] invert the quantifiers in condition (2), which becomes:

(2') For all action a_2 in $BR(T-J, a_1)$, there exists t such that

$$\Pi_1(a_1, a_2, t) > \Pi_1^*(t).$$

In other words, whatever beliefs are formed by player 2 which do not put weight on J , there exists some type (in $T-J$) who would like to deviate. Condition (2') is somewhat more appealing than condition (2), as if (2') is satisfied, the players can not think of any continuation equilibrium which would satisfy (1) and deter any deviation. By contrast, condition (2), except in the two-type case, allows continuum equilibria that satisfy (1) and such that no player in $(T-J)$ wants to deviate from equilibrium behavior.

Cho and Cho-Kreps's "communicational equilibrium" is a PBE such that there does not exist an off-the-equilibrium action a_1 and a subset of types J that satisfy (1) and (2'). Banks and Sobel [1985] identify a condition that is equivalent to (2'); they require (among other things) that player 2's off-the-equilibrium path beliefs place positive probability only on player 1's types who might not lose from a defection. They go on to define the concept

of "divine equilibrium." A divine equilibrium thus satisfies the Cho-Kreps criterion and, for finite games, exists (because it is stable).

We should also mention the work by Farrell [1984] and Grossman and Perry [1986a,b] who offer a criterion similar to, but stronger than, the intuitive criterion. In a signaling game their criterion roughly says that, if there exists a deviation a_1 and a set of types J such that if the posterior beliefs are the same as the prior truncated to $(T-J)$, types in J (respectively, in $(T-J)$) lose (respectively, gain) relative to their equilibrium payoffs, the initial equilibrium is not acceptable. This requirement is stronger than the Cho-Kreps criterion because, in particular, it does not allow any leeway in specifying posterior beliefs within the support $(T-J)$. The refinement, however, is so strong that equilibrium may not exist; so it is restricted to a given (and yet unknown) class of games.

Let us now apply the intuitive criterion to the limit pricing game. As the intuitive criterion is stronger than iterated elimination of weakly dominated strategies, we get at most one separating equilibrium. The reader will check that this least-cost separating equilibrium indeed satisfies the intuitive criterion. Let us next discuss the pooling equilibria (when they exist, i.e., when pooling deters entry). Let us show that pooling at $q_1 < q_m^L$ does not satisfy the intuitive criterion: Consider the deviation to q_m^L . This deviation is dominated for the high-cost type ("J = H"), who makes a lower first-period profit and cannot increase his second-period profit. Thus, posterior beliefs after q_m^L should be $\tilde{p}_1 = 0$, and entry is deterred. But, then the low-cost type would want to produce q_m^L . This reasoning, however, does not apply to pooling equilibria with $q_1 > q_m^L$. Deviations to produce

less are not dominated for any type. Thus, one gets a (smaller) continuum of pooling equilibria (the intuitive criterion here has less cutting power than in the Spence signaling game -- see Kreps [1984]).

One can restrict the set of pooling equilibria that satisfy the intuitive criterion by invoking Pareto dominance: The pooling equilibrium at q_m^L Pareto dominates pooling equilibria with $q_1 > q_m^L$ (both types of player 1 are closer to their static optimum, and player 2 does not care). But, we are still left with a separating and a pooling equilibria, which cannot be ranked using Pareto dominance (player 2 prefers the separating equilibrium). For further refinements in the context of limit pricing, see Cho [1986].

(iii) Guessing Which Equilibrium One is in (McLennan [1985])

McLennan's idea is that a move is more likely if it can be explained by a confusion over which PBE is played. He calls an action "useless" if it is not part of some PBE path. Posterior beliefs at some unreached information set must assign positive probability only to nodes that are part of some PBE, if any (i.e., to actions which are not useless). One thus obtains a smaller set of PBE, and one can operate this selection recursively until one is left with "justifiable equilibria" (which, for finite games, are stable).

iv) Getting Rid of Out-of-Equilibrium Events

As we explained, the indeterminacy of beliefs for out-of-equilibrium events is often a factor of multiplicity. The previous criteria (as well as the one presented in the next section) try to figure out what posterior beliefs are reasonable in such events. An alternative approach, which was pioneered by Saloner [1981] and Matthews-Mirman [1983] consists in perturbing

the game slightly so that these zero-probability events do not occur. The basic idea of this technique is to let the action chosen by an informed player be (at least a bit) garbled before it is observed by his opponents. For instance, one could imagine that a firm's capacity choice is observed with an error or that a manufacturer's price is garbled at the retail level. By introducing noise, all (or most) potentially received signals are equilibrium ones and, thus, refinements are useless. Although the class of games to which this technique can be applied is limited (the noise must represent some reasonable economic phenomenon), this way of proceeding seems natural and is likely to select the "reasonable" equilibria of the corresponding ungarbled game in the limit (as Saloner, for instance, shows in the limit pricing game).

4D. Finite Games: Existence and Refinements in Finite Games

We now informally discuss refinements that are defined only for finite games. Some of these refinements (Selten, Myerson) rest on the idea of taking the limit of equilibria with "totally mixed strategies." One basically considers robustness of each PBE to slight perturbations of the following form: each agent in the game tree is forced to play all his potential actions with some (possibly small) probability, i.e., to "tremble." This way, Bayes' rule applies everywhere (there is no off-the-equilibrium-path outcome). To be a bit more formal, assume that an agent is forced to put weight (probability) $\sigma(a)$ on action a where $\sigma(a) > \varepsilon(a) > 0$ (for each action a). Then the agent can maximize his payoff given these constraints and pick a best perturbed strategy. A refined equilibrium is a PBE which is the limit of equilibria with totally mixed strategies, where the limit is taken for a given class of perturbations. The other two refinements we discuss (Kreps-

Wilson, Kohlberg-Mertens) employ somewhat similar ideas. We shall present the refinements in an increasing-strength order.

General existence results for equilibria of dynamic games with incomplete information have been provided only for games with a finite number of actions and types, starting with Selten. We sketch the proof of existence of a trembling hand equilibrium below. Proofs of existence for alternative refinements are similar.

i) Sequential Equilibrium (Kreps-Wilson [1982])

Kreps-Wilson look at PBE which satisfy a consistency requirement. The set of strategies and beliefs at each information set of the game must be the limit of a sequence of sets of strategies and beliefs for which strategies are always totally mixed (and beliefs are thus pinned down by Bayes' rule.) Moreover, the beliefs on all players are derived as the limit corresponding to a common sequence of strategies. The strategies and beliefs are not a priori required to form a PBE of a perturbed game. So, the check is purely mechanical; given a PBE, it suffices to show that it is or is not the limit of a sequence of totally mixed strategies and associated beliefs.

Let us now discuss the consistency requirement. In the simple signaling game considered above, it has no bite, and a PBE is also sequential, as is easily seen (by choosing adequately the trembles in player 1's strategy, one can generate any beliefs one wants). Sequential equilibrium has more cutting power in more complex games because it imposes consistent beliefs between the players (or agents) off the equilibrium path. For instance, if there are two receivers in the signaling game (players 2 and 3), these two players should form the same beliefs as to player 1's type when observing the latter's action. This property comes from the fact that at each stage of the converging

sequence, players 2 and 3's Bayesian updating uses the same trembles by player 1 and, thus, reach the same conclusion. Similarly, sequential equilibrium requires consistency of a player's beliefs over time. Kreps and Wilson have shown that for "almost all" games, the sequential equilibrium concept coincides with the perfect equilibrium concept (see below). For the other (non generic) games, it allows more equilibria. Selten requires the strategies in the perturbed game to be optimal given the perturbed strategies. But, unless the payoff structure exhibits ties, this condition has no more bite than the consistency requirement of Kreps-Wilson.

ii) Trembling-hand Perfect Equilibrium (Selten [1975])

In developing his notion of the "trembling hand" perfection Selten begins by working with the normal form. An equilibrium is "trembling-hand perfect in the normal form" if it is the limit of equilibria of " ϵ -perturbed" games in which all strategies have at least an ϵ probability of being played. That is, in an ϵ -perturbed game, players are forced to play action a with probability of at least $\epsilon(a)$, where the $\epsilon(a)$ are arbitrary as long as they all exceed ϵ . The $\epsilon(a)$ are called "trembles." The idea of introducing trembles is to give each node in the tree positive probability, so that the best responses at each node are well-defined. The interpretation of the trembles is that in the original game if a player unexpectedly observes a deviation from the equilibrium path he attributes this to an inadvertent "mistake" by one of his opponents.

To see how the trembles help refine the equilibrium set, let us once again consider the game in Figure 5 which Selten used to motivate subgame perfectness.

The Nash equilibrium $\{U,R\}$ is not the limit of equilibria with trembles: if player 1 plays D with some probability, player 2 puts as much weight as possible on L.

However, Selten notes that his refinement is not totally satisfactory. Consider Figure 11, which is a slight variation on the previous game. Player 1 moves at "dates" 1 and 3.

PUT FIGURE 11 HERE

The only subgame perfect equilibrium is $\{L_1, L_2, L'_1\}$. But the subgame-imperfect Nash equilibrium $\{R_1, R_2, R'_1\}$ is the limit of equilibria with trembles. To see why, let player 1 play (L_1, L'_1) with probability ϵ^2 and (L_1, R'_1) with probability ϵ . Then player 2 should put as much weight as possible on R_2 , because player 1's probability of "playing" R'_1 conditional on having "played" L_1 is $\epsilon/(\epsilon + \epsilon^2) \approx 1$ for ϵ small.

When perturbing the normal form, we are allowing for a type of correlation between a player's trembles at different information sets. In the above example, if a player "trembles" onto L_1 , he is very likely to tremble again. This correlation goes against the idea that players expect their opponents to play optimally at any point in the game tree, including those not on the equilibrium path.

To avoid this correlation, Selten introduces a second refinement, based on the "agent's normal form." The idea is to treat the two choices of player 1 in Figure 11 as made by two different players, each of whom trembles independently of the other. More precisely, the agent normal form for a given

game is constructed by distinguishing players not only by their names (i) and their types (t_i), but also by their location in the game tree. So, for instance, player 1 with type t_1 playing at date 1 is not the same agent as player 1 with type t_1 playing at date 3; or player 1 with type t_1 playing at date 3 should be considered as a different agent depending on his (her) information at that date. In the agent's normal form, each information set represents a different agent/player. However, different agents of a same player i with type t_i are endowed with the same objective function. A "trembling hand perfect" equilibrium is a limit of equilibria of ϵ -perturbed versions of the agent's normal form.

It is clear that a trembling-hand perfect equilibrium is sequential: We can construct consistent beliefs at each information set as the limit of the beliefs computed by Bayes rule in the perturbed games, and the equilibrium strategies are sequential given these beliefs. One might expect that the (constrained) optimality requirement along the converging sequence adds some cutting power. However, the arbitrariness of the $\epsilon(a)$ makes perfectness a weak refinement, as shown by Kreps-Wilson's result on that the sets of the sequential and perfect equilibria coincide for generic extensive-form payoffs.

Let us now sketch the proof of existence of a trembling-hand perfect equilibrium. Remember that the proof of existence of a Bayesian equilibrium consists of considering $\{\sum_i |T_i|\}$ players (i.e., in introducing one player per type), and applying standard existence theorems for Nash equilibrium. More generally, the proof for trembling-hand perfect equilibrium uses existence of a Nash equilibrium on the agents' normal form. Consider the perturbed game in which the agents are forced to play trembles (i.e., to put weight at least equal to $\epsilon(a)$ on action a). The strategy spaces are compact convex subsets

of a Euclidean space. Payoff functions are continuous in all variables and quasiconcave (actually, linear) in own strategy. So there exists a Nash equilibrium of the agents' normal form of the perturbed game. Now consider a sequence of equilibrium strategies when ϵ tends to zero. Because the strategy spaces are compact, there is a converging subsequence. The limit of such a subsequence is called a trembling-hand perfect equilibrium.⁸

We should also note that Selten works with the normal form or the agents' normal form; so do the next two refinements. Thus, beliefs are left implicit. Kreps and Wilson's paper is the first pure game theory article to put emphasis on the extensive form and on beliefs (although there is a current debate about whether defined on the normal or extensive form, the refinements that are currently easily applicable to industrial organizational models put constraints on beliefs -- see the previous section).

iii) Proper Equilibrium (Myerson [1978])

Myerson considers perturbed games in which, say, a player's second best action(s) get at most ϵ times the weight of the first best action(s), the third best action(s) get at most ϵ times the weight of the second best action(s), etc. The idea is that a player is "more likely to tremble" and put weight on an action which is not too detrimental to him; the probability of deviations from equilibrium behavior is inversely related to their costs. As the set of allowed trembles is smaller, a proper equilibrium is also perfect. [With such an ordering of trembles, there is no need to work on the agent's normal form. The normal form suffices.]

⁸Note that because payoffs are continuous, the limit is automatically a Nash equilibrium. But the converse, of course, is not true (for instance, for games of perfect information, a trembling-hand equilibrium is subgame perfect, as is easily seen).

To illustrate the notion of proper equilibrium, consider the following game (due to Myerson):

PUT FIGURE 12 HERE

This game has three pure strategy Nash equilibria: (U,L), (M,M) and (D,R). Only two of these are perfect equilibria: D and R are weakly dominated strategies and therefore cannot be optimal when the other player trembles. (M,M) is perfect: Suppose that each player plays M with probability $1-2\epsilon$ and each of the other two strategies with probability ϵ . Deviating to U for player one (or to L for player two) increases this player's payoff by $(\epsilon-9\epsilon)-(-7\epsilon) = -\epsilon < 0$. However, (M,M) is not a proper equilibrium. Each player should put much more weight (tremble more) on his first strategy than on his third, which yields a lower payoff. But if player one, say, puts weight ϵ on U and ϵ^2 on D, player two does better by playing L than by playing M, as $(\epsilon-9\epsilon^2) - (-7\epsilon^2) > 0$ for ϵ small. The only proper equilibrium in this game is (U,L).

iv) Stable Equilibrium (Kohlberg-Mertens [1986])

Ideally, one would wish a PBE to be the limit of some perturbed equilibrium for all perturbations when the size of these perturbations goes to zero. Such an equilibrium, if it exists, is labelled "truly perfect." Unfortunately, true perfection may be out of this world (truly perfect equilibria tend not to exist). Kohlberg and Mertens, to obtain existence, settled for "stability." Stability is a complex criterion, which encompasses the intuitive criterion mentioned in the previous section and other features as well. Let us give an example of the description of a stable equilibrium in the signaling game (this introduction follows Kreps [1984]). Consider two

totally mixed strategies $\tilde{\sigma}_1$ and $\tilde{\sigma}_2$ for players 1 and 2, and two strictly positive numbers ε_1 and ε_2 . A $\{\varepsilon_i, \tilde{\sigma}_i\}_{i=1}^2$ perturbation of the original game is such that, when player i chooses strategy σ_i , the strategy which is implemented for him is σ_i with probability $(1-\varepsilon_i)$ and $\tilde{\sigma}_i$ with probability ε_i . Let (σ_1, σ_2) be a PBE of the perturbed game. A subset E of PBE of the original game is stable if, for any $\eta > 0$, there exists an equilibria of the perturbed game that lies no more than ε from the set E . A stable component is then defined as a minimal connected stable set of equilibria. Kohlberg and Mertens have shown that every game has at least one stable component, and that, for almost every signaling game, all equilibria within a given connected component give rise to the same probability distribution on endpoints.

4E. Perturbed Games and Robust Equilibria

Our earlier discussion of the Saloner/Matthews-Mirman contribution emphasized the robustness of the solution to the introduction of noise. More generally, robustness to "reasonable" structural changes in the game seem desirable. This leads us to the discussion of the reputation-effects model of Kreps-Wilson-Milgrom-Roberts [1982], which is one of the most important applications of the theory of dynamic games of incomplete information.

This work actually started with a robustness issue: In the finite horizon repeated prisoners' dilemma the only equilibrium is "fink, fink" at each period. As we observed in Section 2, this conclusion seems extreme for long, finite games; in response, the four authors decided to perturb the prisoner's dilemma game slightly by introducing a small probability that each party is willing to play the suboptimal strategy tit-for-tat. Similarly, in the con-

text of example 1, one could introduce a probability that firm 1 enjoys preying (is crazy). Then, if the horizon is sufficiently long and the discount rate sufficiently small, it may be worthwhile for a sane type (one whose payoff is as originally specified) to pretend at the start that it is a crazy type. By cooperating in the repeated prisoners' dilemma game or preying in the predation game, the sane type invests in reputation that will induce the other player to take actions that are favorable to the former (cooperate; stay out). Thus, in games that are repeated for a long time, a small difference in information can make a big difference in terms of outcome.

Fudenberg-Maskin [1986] develop the reputation-effects model to its logical conclusion. They show that, for any ϵ , when the horizon goes to infinity, all individually rational payoffs of a finitely repeated, full-information game can arise as PBE of a slightly perturbed, incomplete information game, in which the objective function of each player is the one of the original game with probability $(1-\epsilon)$ and can be any "crazy" objective function with probability ϵ . In the Friedman tradition, the result that one can obtain any payoff Pareto superior to a Nash payoff is easy to derive: Consider a Nash equilibrium of the original game ("fink, fink" in the repeated prisoners' dilemma) and an allocation that dominates this Nash equilibrium, and the corresponding prescribed strategies ("cooperate, cooperate"). Suppose that with probability ϵ , each player has the following objective function: "I like to play the strategy corresponding to the superior allocation as long as the others have followed their corresponding strategies; if somebody has deviated in the past, my taste commands me to play my Nash equilibrium strategy forever." Now suppose that the horizon is long. Then by cooperating, each player loses some payoff at most over one period if the other player deviates. When deviating, he automatically loses the gain of

being able to cooperate with the crazy type until the end. So, as long as there remains enough time until the end of the horizon, ("enough" depends on ϵ) deviating is not optimal. The proof for points that do not dominate a Nash equilibrium is harder.

The reputation effects papers show that adding a small ϵ of incomplete information to a long but finitely repeated game could make virtually anything into a PBE. However, for any fixed horizon, a sufficiently small ϵ of the form they considered has no effect. If we admit the possibility that players have private information about their opponents' payoffs, then even in a fixed extensive form, the sequential rationality requirements of PBE completely lose their force. More precisely, any Nash equilibrium of an extensive form is a PBE (indeed, a stable PBE) of a perturbed game in which payoffs differ from the original ones with vanishingly small probability.

PUT FIGURE 13 HERE

Consider the game in Figure 13. Player 1 has two possible types t_1 and t_2 , with $\text{Prob}(t = t_1) = 1 - \epsilon$. When $t = t_1$, the game is just as in the game of Figure 5, where the backwards induction equilibrium was (D,C). When $t = t_2$, though, player 2 prefers R to L. The strategies (U_1, D_2, R) are a PBE for this game; if player 2 sees D, he infers that $t = t_2$. Thus, a "small" perturbation of the game causes a large change in play -- player 1 chooses U with probability $(1 - \epsilon)$. Moreover, this equilibrium satisfies all the currently known refinements.

Most of these refinements proceed by asking what sort of beliefs are "reasonable" -- what should players expect following unexpected events? If they have very small doubts about the structure of the game, the unexpected

may signal, as here, that things are indeed other than had previously seemed likely. This point is developed in Fudenberg-Kreps-Levine [1987].

Thus, small changes in information structure can always extend the set of predictions to include all of the Nash equilibria, and in long repeated games the "robustness" problem is even more severe. What then is the predictive content of game theory? In real world situations, it may be the case that only some types are unlikely (most types of "craziness" are not plausible). The players may then have a fairly good idea of what game is played. However, the economist, who is an outsider, may have a hard time knowing which information structure is the relevant one. Thus, one may think of a situation in which, at the same time, the players are following the Kreps-Wilson-Milgrom-Roberts strategies, and the reputation literature is of little help to the economist. If this is true, the economist should collect information about the way real world players play their games and which information structure they believe they face, and then try to explain why particular sorts of "craziness" prevail.

The above implies a fairly pessimistic view of the likelihood that game theory can hope to provide a purely formal way of choosing between PBEs. It would be rash for us to assert this position too strongly, for research on equilibrium refinements is proceeding quite rapidly, and our discussion here may well be outdated by the time it appears in print. However, at present we would not want to base important predictions solely on formal grounds. In evaluating antitrust policy, for example, practitioners will need to combine a knowledge of the technical niceties with a sound understanding of the workings of actual markets.

Concluding Remark

Our already incomplete discussion of equilibrium concepts for dynamic games of incomplete information is likely to be out of date very shortly, as the pace of activity in this field is very intense, and current refinements have not yet been tested for a wide class of models. Our purpose was only to provide an introduction, a survey and some cookbook recipes for readers who currently want to apply these techniques to specific games.




Figure 1




Figure 2


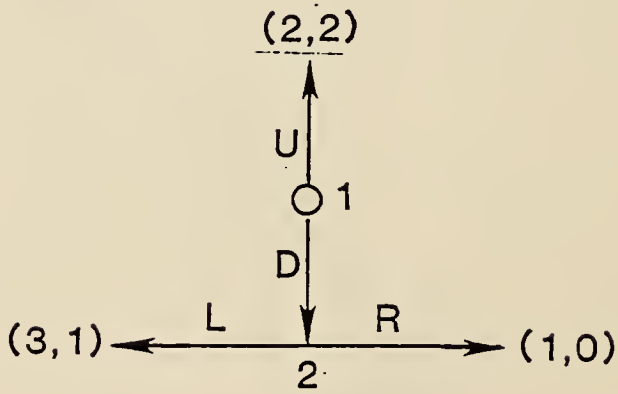


Figure 3

		Player 2	
		L	R
Player 1	U	2,1	2,1
	D	1,3	4,0

		Player 2	
		L	R
Player 1	U	2,1	0,0
	D	0,0	1,2

Figure 4



		Player 2	
		L	R
Player 1	U	2,2	2,2
	D	3,1	1,0

Figure 5

		<u>Player 2</u>		
		<u>L</u>	<u>M</u>	<u>R</u>
<u>Player 1</u>	<u>U</u>	<u>5,5</u>	<u>4,4</u>	<u>0,4.9</u>
	<u>M</u>	<u>3,8</u>	<u>1,1</u>	<u>6,5</u>
	<u>D</u>	<u>4.9,0</u>	<u>5,6</u>	<u>4.9,4.9</u>

Figure 6

		<u>Player 2</u>	
		<u>L</u>	<u>R</u>
<u>Player 1</u>	<u>U</u>	<u>5,1</u>	<u>0,0</u>
	<u>D</u>	<u>4,4</u>	<u>1,5</u>

Figure 7

		Player 2	
		don't fink	fink
Player 1	don't fink	2,2	-1,4
	fink	4,-1	0,0

Figure 8

		Player 2		
		L	M	R
Player 1	U	5,5	-1,-1	-2,-2
	M	-1,-1	0,0	-2,-2
	D	-2,-2	-2,-2	-6,-6

Figure 9

Player 1

		Player 2	
		L	R
U	0, 0, 10	-5, -5, 0	
D	-5, -5, 0	1, 1, -5	

A

		Player 2	
		L	R
U	-2, -2, 0	-5, -5, 0	
D	-5, -5, 0	-1, -1, 5	

B

Figure 10

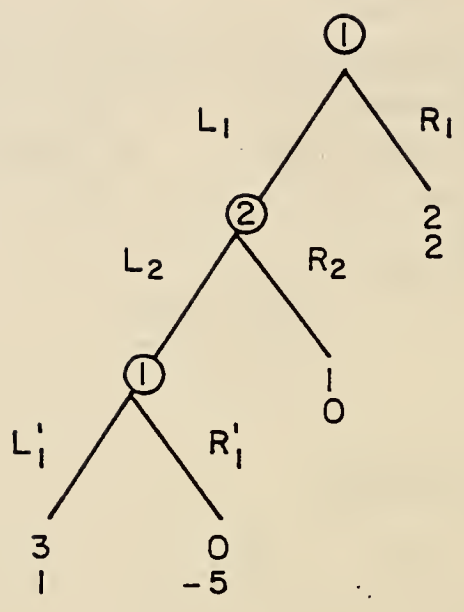


Figure II

		Player 2		
		L	M	R
Player 1	U	1, 1	0, 0	-9, -9
	M	0, 0	0, 0	-7, -7
	D	-9, -9	-7, -7	-7, -7

Figure 12




Figure 13

References

Abreu, D. (1983), Repeated Games with Discounting, Ph.D. dissertation, Department of Economics, Princeton University.

_____ (1986), "Extremal Equilibria of Oligopolistic Supergames," Journal of Economic Theory 39, 191-225.

Abreu, D., D. Pearce, and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring," Journal of Economic Theory 39, 251-269.

Akerlof, G. (1970), "The Market for 'Lemons'," Quarterly Journal of Economics 90, 629-650.

Anderson, R. (1985), "Quick-response Equilibria," mimeo.

Aumann, R. (1959), "Acceptable Points in General Cooperative n-Person Games," in Contributions to the Theory of Games IV, Princeton, NJ: Princeton University Press.

_____, (1964), "Mixed vs. Behavior Strategies in Infinite Extensive Games," Annals of Mathematics Studies No. 52, Princeton University Press, 627-630.

_____ (1974), "Subjectivity and Correlation in Randomized Strategies," Journal of Mathematical Economics 1, 67-96.

_____ (1976), "Agreeing to Disagree," Annals of Statistics 4, 1236-1239.

_____ (1986), "Correlated Equilibria as an Expression of Rationality," forthcoming in Econometrica.

Aumann, R. and L. Shapley (1976), "Long Term Competition: A Game Theoretic Analysis," mimeo.

Aumann, R., Y. Katznelson, R. Radner, R. W. Rosenthal, and B. Weiss (1981), "Approximate Purification of Mixed Strategies," mimeo, Bell Laboratories.

Banks, J., and J. Sobel (1985), "Equilibrium Selection in Signaling Games," DP 85-9, San Diego.

Baron, D., and D. Besanko (1984), "Regulation, Asymmetric Information, and Auditing," Rand Journal 15, 447-570.

Benoit, J. P., and V. Krishna (1985), "Finitely Repeated Games," Econometrica 53, 890-904.

Bernheim, D. (1984), "Rationalizable Strategic Behavior," Econometrica 52, 1007-1028.

Bernheim, D., B. Peleg, and M. Whinston (1985), "Coalition-Proof Nash

Equilibria, I: Concepts," mimeo, Stanford University.

Bernheim, D., and M. Whinston (1985), "Coalition-Proof Nash Equilibria, II: Applications," Stanford University.

Brandenburger, A., and E. Dekel (1985a), "Hierarchies of Beliefs and Common Knowledge," mimeo.

_____ (1985b), "Rationalizability and Correlated Equilibria," forthcoming in Econometrica.

Bulow, J. (1982), "Durable-goods Monopolists," Journal of Political Economy, 314-332.

Chamberlin, E. (1956), The Theory of Monopolistic Competition, Cambridge, Massachusetts: Harvard University Press.

Cho, I.-K. (1986), "Equilibrium Analysis of Entry Deterrence: Reexamination," mimeo, Princeton University.

Cho, I.-K., and D. Kreps (1985), "More Signalling Games and Stable Equilibrium," forthcoming, Quarterly Journal of Economics.

Coase, R. H. (1972), "Durability and Monopoly," Journal of Law and Economics 15, 143-149.

Crawford, V., and J. Sobel (1982), "Strategic Information Transmission," Econometrica 50, 1431-1452.

Cremer, J. (1983), "Cooperation in Ongoing Organizations," mimeo.

Dasgupta, P., and E. Maskin (1986), "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," Review of Economic Studies 53, 1-26.

Debreu, G. (1952), "A Social Equilibrium Existence Theorem," Proceedings of the National Academy of Sciences 38, 886-893.

Fan, K. (1952), "Fixed Point and Minimax Theorems in Locally Convex Topological Linear Spaces," Proceedings of the National Academy of Sciences 38, 121-126.

Farquharson, R. (1969), Theory of Voting, New Haven: Yale University Press.

Farrell, J. (1983), "Renegotiation-proof Equilibrium in Repeated Games," mimeo, MIT.

_____ (1984), "Credible Neologisms in Games of Communication," mimeo, MIT.

Frasse, J., and M. Moreaux (1985), "Collusive Equilibria in Oligopolies with Finite Lives," European Economic Review 45-55.

Friedman, J. (1971), "A Noncooperative Equilibrium for Supergames,"

Review of Economic Studies 38, 1-12.

_____ (1977), Oligopoly and the Theory of Games, Amsterdam: North-Holland.

_____ (1984), "Trigger Strategy Equilibria in Finite Horizon Supergames," mimeo.

Fudenberg, D., and D. Kreps (1986), "Reputation with Multiple Opponents," mimeo.

Fudenberg, D., D. Kreps, and D. Levine (1986), "On the Robustness of Equilibrium Refinements," forthcoming in the Journal of Economic Theory.

Fudenberg, D., and D. Levine (1983), "Subgame-Perfect Equilibria of Finite and Infinite Horizon Games," Journal of Economic Theory 31, 227-256.

_____ (1986), "Limit Games and Limit Equilibria," Journal of Economic Theory 38, 261-279.

_____ (1987), "Open-loop and Closed-loop Equilibria in Games with Many Players," forthcoming in the Journal of Economic Theory.

Fudenberg, D., D. Levine, and J. Tirole (1985), "Infinite Horizon Models of Bargaining with One-sided Incomplete Information," in Alvin Roth (ed.), Bargaining with Incomplete Information, Cambridge University Press.

Fudenberg, D., and E. Maskin (1986a), "Folk Theorems for Repeated Games with Discounting or with Incomplete Information," Econometrica 54, 533-554.

_____ (1986b), "Discounted Repeated Games with Unobservable Actions," mimeo.

_____ (1986c), "Nash and Perfect Equilibria of Discounted Repeated Games," mimeo.

Fudenberg, D., and J. Tirole (1983a), "Capital as a Commitment: Strategic Investment to Deter Mobility," Journal of Economic Theory 31, 227-250.

_____ (1983b), "Sequential Bargaining with Incomplete Information," Review of Economic Studies 50, 221-247.

_____ (1985), "Preemption and Rent Equalization in the Adoption of New Technology," Review of Economic Studies 52, 383-401.

_____ (1986), "A Theory of Exit in Oligopoly," IMSSS T. R. 429, Stanford University, Econometrica 54, 943-960.

Geanakoplos, J., and H. Polemarchakis (1978), "We Can't Disagree

Forever," IMSSS D. P. 277, Stanford University.

Gertner, R. (1985a), "Simultaneous Move Price-quantity Games and Non-market Clearing Equilibrium," mimeo, MIT.

_____ (1985b), "Dynamic Duopoly with Price Inertia," mimeo, MIT.

Ghemawat, P. and B. Nalebuff (1985), "Exit," Rand Journal of Economics, 16, 84-94.

Glicksberg, I. L. (1952), "A Further Generalization of the Kakutani Fixed Point Theorem with Application to Nash Equilibrium Points," Proceedings of the National Academy of Sciences 38, 170-174.

Green, E. (1984), "Continuum and Finite-players, Non-cooperative Models of Competition," Econometrica 52, 975-994.

Green, E. and R. Porter (1984), "Noncooperative Collusion under Imperfect Price Information," Econometrica 32, 87-100.

Green, J. and W. Heller (1981), "Mathematical Analysis and Convexity with Applications to Economics," pp. 15-52 in the Handbook of Mathematical Economics, K. Arrow and M. Intriligator, eds., Amsterdam: North-Holland Publishing Co.

Grossman, S., and O. Hart (1983), "An Analysis of the Principal-Agent Problem," Econometrica 51, 7-46.

Grossman, S., and M. Perry (1984), "Sequential Bargaining under Asymmetric Information," mimeo, Princeton University.

_____ (1986), "Perfect Sequential Equilibrium," Journal of Economic Theory 39, 97-119.

Gul, F., H. Sonnenschein, and R. Wilson (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture," Journal of Economic Theory 39, 155-190.

Hanig, M. (1985), "A Differentiable Game Model of Duopoly with Reversible Investment," mimeo, MIT.

Harris, C. (1985a), "A Characterization of the Perfect Equilibria of Infinite-horizon Games," Journal of Economic Theory 37, 99, 125.

_____ (1985b), "Existence and Characterisation of Perfect Equilibrium in Games of Perfect Information," Econometrica 53, 613-628.

_____ (1986), "A Note on the Existence of Optimal Simple Penal Codes," mimeo, Nuffield College.

Hart, S. (1979), "Lecture Notes: Special Topics in Game Theory," IMSSS Technical Report.

Harsanyi, J. (1964), "A General Solution for Finite Non-cooperative

Games, Based on Risk Dominance," in Advances in Game Theory (M. Dresher et al., Eds.), Annals of Mathematics, Study 52, 627-650.

_____ (1967-8), "Games with Incomplete Information Played by Bayesian Players," Management Science 14, 159-82, 320-34, 486-502.

_____ (1973), "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," International Journal of Game Theory, 2, 1-23.

Hendricks, K., and C. Wilson (1985a), "Discrete versus Continuous Time in Games of Timing," mimeo.

_____ (1985b), "The War of Attrition in Discrete Time," C. V. Starr Center Working Paper No. 85-32, New York University.

Holmstrom, B. (1982), "Moral Hazard in Teams," BJE 13, 324-340.

Judd, K. (1985), "Closed-loop Equilibrium in a Multi-stage Innovation Race," Discussion Paper No. 647, Kellogg Graduate School of Management, Northwestern University.

Katz, M., and Shapiro, C. (1984), "Perfect Equilibrium in a Development Game with Licensing or Imitation," Princeton DP #85, Woodrow Wilson School.

Kohlberg, E. (1981), "Some Problems with the Concept of Perfect Equilibrium," NBER Conference, U.C. Berkeley.

Kohlberg, E., and J. F. Mertens (1986), "On the Strategic Stability of Equilibria," Econometrica 54, 1003-1038.

Kreps, D. (1985), "Signalling Games and Stable Equilibrium," mimeo.

Kreps, D., and J. Scheinkman (1983), "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," Bell Journal of Economics 14, 326-337.

Kreps, D., and R. Wilson (1982a), "Sequential Equilibrium," Econometrica 50, 863-894.

_____ (1982b), "Reputation and Imperfect Information," Journal of Economic Theory 27, 253-279.

Kreps, D., et al. (1982), "Rational Cooperation in the Finitely-Repeated Prisoner's Dilemma," Journal of Economic Theory 27, 245-252.

Kuhn, H. (1953), "Extensive Games and the Problem of Information," Annals of Mathematics Studies, No. 28, Princeton University Press.

Lockwood, B. (1983), "Perfect Equilibria in Repeated Games with Discounting," mimeo.

Luce, R., and H. Raiffa (1957), Games and Decisions, Chapters 1, 3-5, New York: John Wiley and Sons.

Maskin, E., and J. Tirole (1985), "A Theory of Dynamic Oligopoly, II: Price Competition," forthcoming, Econometrica.

Maynard Smith, J. (1974), "The Theory of Games and the Evolution of Animal Conflicts," Journal of Theoretical Biology 47, 209, 274.

McManus, M. (1962), "Numbers and Size in Cournot Oligopoly," Yorkshire Bulletin of Social and Economic Research, 14.

Mertens, J. F., and S. Zamir (1985), "Formulation of Bayesian Analysis for Games with Incomplete Information," International Journal of Game Theory 14, 1-29.

Milgrom, P. (1981), "An Axiomatic Characterization of Common Knowledge," Econometrica 49, 2, 219-222.

Milgrom, P., and J. Roberts (1982a), "Limit Pricing and Entry under Incomplete Information," Econometrica 50, 443-460.

_____ (1982b), "Predation, Reputation, and Entry Deterrence," Journal of Economic Theory 27, 280-312.

Milgrom, P., and R. Weber (1986), "Topologies on Information and Strategies in Games with Incomplete Information," Northwestern MEDS D.P. 463.

_____ (1982), "A Theory of Auctions and Competitive Bidding," Econometrica 50, 1089-1122.

_____ (1986), "Distributional Strategies for Games with Incomplete Information," Mathematics of Operations Research 10, 619-631.

Moulin, H. (1979), "Dominance Solvable Voting Schemes," Econometrica 1337-1353.

Myerson, R. (1978), "Refinements of the Nash Equilibrium Concept," International Journal of Game Theory 7, 73-80.

_____ (1979), "Optimal Auction Design," Mathematics of Operations Research 6, 58-73.

_____ (1983), "Bayesian Equilibrium and Incentive Compatibility: An Introduction," Northwestern MEDS D.P. 548.

Myerson, R., and M. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading," Journal of Economic Theory 29, 265-281.

Nalebuff, B. (1982), "Brinkmanship," mimeo, Harvard University.

Nash, J. F. (1950), "Equilibrium Points in N-person Games," Proceedings

of the National Academy of Sciences, U.S.A. 36, 48-49.

- Novshek, W. (1985), "On the Existence of Cournot Equilibrium," Review of Economic Studies 52, 85-98.
- Ortega-Reichert, A. (1967), "Models for Competitive Bidding under Uncertainty," Ph.D. thesis, Stanford University.
- Pearce, D. (1984), "Rationalizable Strategic Behavior and the Problem of Perfection," Econometrica 52, 1029-1050.
- Pitchik, C. (1982), "Equilibria of a Two-person Non-Zerosum Noisy Game of Timing," International Journal of Game Theory 10, 207-221.
- Porter, R. (1983), "Optimal Cartel Trigger Price Strategies," Journal of Economic Theory 30, 313-338.
- Radner, R. (1980), "Collusive Behavior in Non-cooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives," Journal of Economic Theory 22, 136-154.
- Radner, R., R. Myerson, and E. Maskin (1986), "An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria," Review of Economic Studies 53, 59-70.
- Reynolds, S. (1985), "Capacity Investment, Preemption and Commitment in an Infinite Horizon Model," mimeo.
- Roberts, J., and H. Sonnenschein (1976), "On the Existence of Cournot Equilibrium without Concave Profit Functions," Journal of Economic Theory 13.
- _____ (1977), "On the Foundations of the Theory of Monopolistic Competition," Econometrica, 45, 101-113.
- Rotemberg, J. and G. Saloner (1986), "A Supergame-Theoretic Model of Business Cycles and Price Wars during Booms," American Economic Review.
- _____ (1985), "Price Leadership," mimeo.
- Rubinstein, A. (1977), "Equilibrium in Supergames," Center for Mathematical Economics and Game Theory, Hebrew University of Jerusalem.
- _____ (1979), "Equilibrium in Supergames with the Overtaking Criterion," Journal of Economic Theory 21, 1-9.
- _____ (1982), "Perfect Equilibrium in a Bargaining Model," Econometrica 50, 97-110.
- Schelling, T. (1960), The Strategy of Conflict, Cambridge, Massachusetts: Harvard University Press.
- Scherer, F. M. (1980), Industrial Market Structure and Economic

Performance, 2nd edition, Boston: Houghton-Mifflin.

Selten, R. (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," Zeitschrift für die gesamte Staatswissenschaft 12, 301-324.

_____ (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory 4, 25-55.

Simon, L. (1985), "Games with Discontinuous Payoffs I: Theory," mimeo.

Simon, L., and Maxwell Stinchcombe (1986), "Extensive Form Games in Continuous Time. Part I: Pure Strategies," mimeo.

Sobel, J., and I. Takahashi (1983), "A Multi-stage Model of Bargaining," Review of Economic Studies 50, 411-426.

Spence, A. M. (1974), Market Signalling, Cambridge: Harvard University Press.

_____ (1979), "Investment Strategy and Growth in a New Market," Bell Journal of Economics 10, 1-19.

Starr, R., and Y. C. Ho (1967), "Further Properties of Nonzero-sum Differential Games," Journal of Optimization Theory and Applications 3, 207-219.

Stokey, N. (1981), "Rational Expectations and Durable Goods Pricing," Bell Journal of Economics 12, 112-128.

Tarski, A. (1955), "A Lattice-theoretical Fixpoint Theorem and Its Applications," Pacific Journal of Mathematics 5, 285-308.

Topkis, D. (1970), "Equilibrium Points in Nonzero-sum N-person Subadditive Games," Technical Report ORC 70-38, Operations Research Center, U.C., Berkeley.

Vives, X. (1985), "Nash Equilibrium in Oligopoly Games with Monotone Best Responses," CARESS Working Paper #85-10, University of Pennsylvania.

Weiss, A., and C. A. Wilson (1984), "An Analysis of Games of Timing with Complete Information: An Application to Oil Exploration," mimeo.