

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale University
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 983

NOTE: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than acknowledgment that a writer had access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

REPEATED GAMES: COOPERATION AND RATIONALITY

by

David G. Pearce

June 1991

REPEATED GAMES: COOPERATION AND RATIONALITY

by

David G. Pearce^{*}
Cowles Foundation for Research in Economics
Yale University

ABSTRACT

The paper is a survey written for the Sixth World Congress of the Econometric Society. It is devoted largely to a discussion of the progress made in the last decade in understanding the structure of self-enforcing agreements in discounted supergames of complete information. Perfect and imperfect monitoring models are considered in turn, with attention given to the case of substantial impatience as well as to the various "folk theorems." The emphasis is on the features of constrained-optimal perfect equilibria, causes of inefficiency, and some relationships among different strands of the literature. The remainder of the paper is a critical and comparative consideration of recent work on renegotiation in repeated games.

^{*}I am grateful to Dilip Abreu, Al Klevorick and Ennio Stacchetti for their expert and patient help. Many thanks to the Alfred P. Sloan Foundation and the National Science Foundation for their financial support.

I. INTRODUCTION

In economic, political and personal life, the terms on which individuals or institutions interact are rarely determined fully by explicit, enforceable contracts. Within the bounds of the law, there is enormous scope for variation in the way in which commercial rivalries, international relations and social affairs are conducted. Often, the same parties interact repeatedly. As a consequence, there is a large role for implicit, self-enforcing contracts to play: agents have an incentive to conform to an implicit agreement today because they believe that this will influence the nature of subsequent interactions. Repeated games provide perhaps the simplest model in which self-enforcing cooperative arrangements can be studied formally. It is this aspect of repeated game theory that I attempt to survey here. The paper focuses on the structural and conceptual issues that have arisen in recent years in the study of repeated discounted games of complete information.

This choice of subject matter embraces a large literature, but excludes some important topics in repeated games. There is a substantial and challenging body of work on repeated games of incomplete information, much of which is surveyed by Mertens (1987). Following Kreps and Wilson (1982) and Milgrom and Roberts (1982), many papers have explored the effects of reputation formation in finitely repeated games with (initially) small amounts of incomplete information. These are covered by Fudenberg (1992) in a companion paper to the present one. The latter survey also touches on the growing literature that investigates how play evolves as success is rewarded by survival.

The first part of the paper chronicles the progress that was made in the past decade in understanding super-game equilibria from a technical point of view.¹ Many problems that had been considered intractable yielded to systematic analysis. Whereas earlier work on discounted repeated games had to content itself with studying artificially restricted behavior, a number of papers revealed that it was possible to drop those restrictions and still obtain strong results. Theorists began to explore more complicated and satisfying models, suggested by features of various economic situations. Players may observe different parts of the history of play, and some of their information may be stochastic, for example. They could meet different partners or rivals over time, or have different time horizons.

Section 2 considers models in which players receive information without any stochastic disturbance, while Section 3 is devoted to games with imperfect monitoring. In each case I begin with the analysis for an arbitrary

¹Numerous references can be found in subsequent sections.

discount factor (or rate of interest) and later address the important case in which players are very patient, relative to the delays between successive plays of the game. Finally, representative applications to applied fields are discussed, as well as some recent attempts to compare the theory with data in various ways.

The second part of the paper is devoted to some conceptual issues associated with repeated game theory, especially the problem of renegotiation. In a supergame equilibrium involving short-run sacrifices by some player for the good of the group, for example, there is an implicit threat that if the player fails to cooperate, he will be punished in some way. But *ex post*, will the threat actually be carried out, or will the continuation equilibrium be "renegotiated"? At issue here is the fundamental question of which threats are credible. Game theorists maintained an uncomfortable silence on this point for many years. Recently there has been a small riot of proposals regarding the appropriate formulation of a "renegotiation-proof" solution concept. We are left with an embarrassment of riches, since many of the suggestions are at odds with one another and lead to entirely different predictions. Section 4 reviews some of the solution concepts. I argue that the diversity of ideas on renegotiation-proofness is natural, given the essentially psychological nature of the problem, and suggest that any solution concept in this area be interpreted cautiously. Section 5 concludes briefly.

2. PERFECT MONITORING

This section and the following one are addressed not to the specialist in repeated games, but to scholars who would like a reader's guide to the literature on discounted repeated games of complete information. The emphasis is on the overall picture and the connections among papers in the field. Those looking for a "nuts and bolts" treatment of the material should consult the expert and up-to-date coverage in Chapter 5 of Fudenberg and Tirole (1990). Also enthusiastically recommended are the concise, specialized piece on folk theorems by Krishna (1987) and the wide-ranging survey of complete information supergames by Sabourian (1989).

In most of the models of this section, all players learn at the end of each period the actions taken in that period by other players. Usually, we are studying situations in which some simultaneous game G is played by the same set $N = \{1, \dots, n\}$ of players. The *stage game* (or component game) $G = (A_1, \dots, A_n; \Pi_1, \dots, \Pi_n)$ is described by the nonempty action sets (or pure strategy sets) $A_i, i \in N$, and the payoff functions $\Pi_i : A \rightarrow \mathbb{R}$, where $A = A_1 \times \dots \times A_n$. If each A_i is finite, G is called a finite game. Extend the functions Π_i in the usual way to the product $M = M_1 \times \dots \times M_n$ of the sets of mixed strategies. Let $F = \text{co } \Pi(A)$ be the convex hull of the set of payoff vectors from action profiles in A . Elements in F are called *feasible values*.

Playing G repeatedly produces a stream of payoffs for each player, which in most cases will be discounted by the factor δ , assumed for simplicity to be the same for each player. Payoffs are received at the end of each

period and discounted to the beginning of the first period, period 1. The finitely repeated game consisting of T plays of G , with discount factor $\delta \in (0,1]$, is denoted $G^T(\delta)$. When G is repeated indefinitely, and $\delta \in (0,1)$, we have the infinite horizon game $G^\infty(\delta)$. A pure strategy for player i in a repeated game specifies an action for i in each period t as a function of the actions chosen by all players in all preceding periods.² A mixed strategy (more properly, a behavior strategy) in the supergame allows the contingent choices to be stochastic. In this case it must be specified whether other players observe only the outcome of the randomization (this is the standard assumption, and usually the only plausible one) or also the random device used.

Often it is convenient to normalize supergame payoffs so that they are directly comparable to payoffs of a stage game: the average (discounted) value of a stream of payoffs is that number which, if received in every period, would have the same present discounted value as that of the original stream. For any strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ of the supergame let $\Psi(\sigma)$ denote the associated vector of (total) present discounted payoffs, and let $v(\sigma) = \frac{1-\delta}{\delta} \Psi(\sigma)$ denote the vector of average values.

Self-enforcing agreements in repeated games are the subjects of inquiry here, but it must be admitted that there is a great deal of controversy about how "self-enforcing" should be *defined* (see Section 4). For the moment, let us simply require that the agreements (whether spoken or unspoken) be *subgame perfect equilibria* (Selten, 1965, 1975) of the supergame. This means that following any $t-1$ period history of play, the agreed-upon strategy profile gives players instructions that constitute a Nash equilibrium (Nash, 1950) of the subgame beginning in period t . In other words, after no history should a player have an incentive to deviate unilaterally from his part of the strategy profile. When a particular player has no incentive to deviate following any history, we say that his strategy is a *perfect best response* to the other players' strategies. The set of average values of subgame perfect equilibria of $G^T(\delta)$ and $G^\infty(\delta)$ are written $V^T(\delta)$ and $V^\infty(\delta)$, respectively. When there is no danger of ambiguity, I simply write V .

Unimprovability

Consider the following requirement that, at first glance, looks much weaker than the perfect best response condition. A strategy for i is *unimprovable* against a vector of strategies of his opponents if there is no $t-1$ period history (for any t) such that i could profit by deviating from his strategy in period t only (conforming thereafter). To verify the unimprovability of a strategy, then, one checks only "one-shot" deviations from the

²It is often remarked that it is unnecessary to allow i to condition on his own past actions. It is easy to show by example that this is false.

strategy, rather than arbitrarily complex deviations (such as defecting in every period t such that t is prime). The following result simplifies the analysis of subgame perfect equilibria immensely. It is the exact counterpart of a well-known result from dynamic programming due to Howard (1960), and was first emphasized in the context of self-enforcing cooperation by Abreu (1988).

PROPOSITION. *Let the payoffs of G be bounded. In the repeated game $G^T(\delta)$ or $G^\infty(\delta)$ with $\delta \in (0,1)$, a strategy σ_i is a perfect best response to a profile γ of strategies if and only if σ_i is unimprovable against that profile.*

The proof is simple and generalizes easily to a wide variety of dynamic and stochastic games with discounting and bounded payoffs. If σ_i is not a perfect best response, there must be a history after which it is profitable to deviate to some other strategy σ'_i . If the deviation involves defection at infinitely many nodes, then for sufficiently large T , the strategy $\hat{\sigma}_i$ that agrees with σ'_i until time T and conforms to σ_i thereafter, is also a profitable deviation (because of discounting and boundedness, anything that happens in the distant future has almost no impact on payoffs today). Consider a profitable deviation $\hat{\sigma}_i$ involving defection at the smallest possible number of nodes, and let x be a node at which $\hat{\sigma}_i$ disagrees with σ_i for the last time. Not conforming to σ_i at x must be profitable, or else one would have had a profitable deviation with fewer defection nodes than $\hat{\sigma}_i$, a contradiction. The profitability of deviating from σ_i at x (and never again) means that σ_i is not unimprovable. Thus, if σ_i is unimprovable, σ_i is a perfect best response. The converse is trivial, since the requirements for unimprovability are a subset of those for a perfect best response.

The above equivalence will be exploited frequently in both this and the succeeding section; it does not depend on the structure of players' information.

Cooperation Enforced by Nash Threats

During the 1950's and 1960's there developed a verbal tradition amongst game theorists to the effect that if players in an infinitely repeated game considered the future sufficiently important compared to the current period, an extremely wide variety of behavior could be supported in equilibrium. Friedman (1971) formalized this in the context of Cournot oligopoly. How can firms be persuaded to overcome the free-rider problem that drives them away from the joint monopoly output to the less lucrative Cournot-Nash solution? Consider a strategy profile specifying that (i) firms produce some vector of quantities each period that is more profitable for each firm than the Cournot-Nash equilibrium, as long as there has been no deviation from the vector in past periods, and (ii) following any deviation, they revert to playing the static Cournot-Nash solution forever. If players are very patient, a firm's temptation to increase profits today by cheating are outweighed by the permanent loss in profits in succeeding periods. Hence, it is not profitable to deviate once if no one has done so

before. If a deviation has occurred earlier, again there is no profitable one-shot deviation, because players are already playing a myopic best response to one another's actions. By the unimprovability criterion described above, we may conclude that the strategy profile is a subgame perfect equilibrium. Thus, although Friedman was concerned only with Nash equilibrium, he actually exhibited "cooperative" strategies that satisfied additional requirements of credibility.

PROPOSITION (Friedman, 1971): *Let $G = (A_1, \dots, A_n; \Pi_1, \dots, \Pi_n)$ have a Nash equilibrium $e = (e_1, \dots, e_n) \in A$, and let $q = (q_1, \dots, q_n) \in A$ satisfy $\Pi_i(q) > \Pi_i(e)$ for each $i \in N$. Then for δ sufficiently close to 1, there is a subgame perfect equilibrium of $G^\delta(\delta)$ in which q is played in every period on the equilibrium path.*

It is convenient (and frequently realistic) to convexify the set of equilibrium values by enriching the structure of the supergame as follows: at the beginning of each period, the realization of some continuous random variable is commonly observed, so players can make their choices conditional on the outcome. Modifying the game in this way is usually called "allowing for public randomization." If we do so, Friedman's argument implies that any element of F that strictly Pareto-dominates some Nash equilibrium of G is the average payoff of some subgame perfect equilibrium of $G^\delta(\delta)$, for sufficiently high δ . Indeed, arguments of Sorin (1986) and Fudenberg and Maskin (1986) can be used to show that this limit result holds even without public randomization: for values of δ near 1, convexification can be accomplished by varying play appropriately over time.

The Folk Theorem

A still stronger result was suggested by the verbal tradition alluded to earlier, one eventually proved by Aumann and Shapley (1976) and Rubinstein (1977). Their celebrated "folk theorem" for infinitely repeated games confirms the most optimistic conjecture one could reasonably make regarding which values are average payoffs of (subgame perfect) equilibria when players are ideally patient. Clearly, an average value must be feasible in the physical sense, that is, it must lie in F . The fact that a player always has the option of playing a myopic best response to other players' strategies in each period gives him a *security level* that must also be respected. Formally, let $\underline{v}_i = \min_{\alpha_{-i} \in M_{-i}} \max_{a_i \in A_i} \Pi_i(a_i, \alpha_{-i})$, where $\Pi_i(a_i, \alpha_{-i})$ means $\Pi_i(\alpha_1, \dots, \alpha_{i-1}, a_i, \alpha_{i+1}, \dots, \alpha_n)$. Any vector giving each player i at least his security level \underline{v}_i is called *individually rational*. Let F^+ denote the set of feasible and individually rational vectors. Remarkably, if there is "no discounting" in the sense that players

care only about their long-run average³ payoffs, the set of feasible, individually rational payoffs coincides with the set of long-run average payoffs of equilibria of the infinitely repeated game.

PROPOSITION. The Perfect Folk Theorem of Repeated Games (Aumann and Shapley (1976), Rubinstein (1977)):
Let G^∞ be the supergame in which G is repeated indefinitely and payoffs are evaluated according to the limit of means criterion. Then v is the average payoff of some subgame perfect equilibrium of G^∞ if and only if it is feasible and individually rational.

The essence of the proof can be conveyed by looking at the simplest case, namely a feasible and individually rational value v for which there exists a pure action profile $c \in A$ with $\Pi(c) = v$. Consider a supergame profile that instructs players to begin by playing c in each period, and to respond to any deviation by forcing the deviant player down to his security level ("minimaxing him") for t^2 periods, where t is the date of his deviation, and then returning to playing c (unless there is some further deviant at time t' , who will be minimaxed for $(t')^2$ periods, and so on).⁴ If a player deviates an infinite number of times, his long run average will be at best v_i , so no such defection would be profitable. If he deviates only a finite number of times, play eventually returns to c forever, and again he has not profited (the limit of means is insensitive to payoff changes in any finite set of periods). To summarize, after no history can anyone gain by a unilateral deviation, so the profile is a subgame perfect equilibrium of G^∞ .

Notice that with no discounting, the criteria "perfect best response" and "unimprovable" are *not* equivalent. The strategy profile according to which each firm in a symmetric Cournot duopoly produces half of the joint monopoly output, regardless of the history of play, is certainly not subgame perfect, and yet the strategies are unimprovable according to the limit of means criterion.

Rubinstein (1979a) also proved a perfect folk theorem akin to the one just discussed, for games with payoffs evaluated according to the overtaking criterion.⁵ While also capturing the idea of extreme patience, this criterion seems closer than the limit of means to the case of very little discounting (δ near 1) because it makes

³Given a stream of payoffs $\{x_t\}$, one can define the sequence of average payoffs $\{y_t\}$ by $y_t = \frac{1}{t} \sum_{k=1}^t x_k$. The sequence $\{y_t\}$ may not have a limit, but the limit inferior is always defined, and this is what is meant here by the "long run average" or "limit of means" associated with the original stream $\{x_t\}$.

⁴Simultaneous deviations are ignored, since they are irrelevant for checking subgame perfection.

⁵A payoff stream $\{w_t\}$ is strictly preferred to $\{x_t\}$ under the overtaking criterion if $\liminf_{T \rightarrow \infty} \sum_{t=1}^T (w_t - x_t) > 0$.

players sensitive to what happens in any single period. With this increased realism comes additional complication: the set of subgame perfect equilibrium values is not closed, and the statement of the theorem must be weakened slightly. When one uses the overtaking criterion, unimprovability is again not useful for checking subgame perfection.

The perfect folk theorems provided an important impetus for further research on *discounted repeated games*, because they suggested vividly that punishments more severe (and hence more effective as deterrents) than permanent reversion to static Nash equilibrium could be credible. Ironically, the proofs of the same theorems probably also threw researchers off track, because the line of attack that eventually proved successful in the discounted case was rather different from methods in the absence of discounting.

Simple Strategy Profiles

Abreu's work in the early 1980's (ultimately published as Abreu (1986, 1988)) marked a breakthrough in the study of the pure strategy perfect equilibria of discounted supergames. It reduced an ostensible tangle of intertemporal incentive constraints and punishment hierarchies to a comparatively orderly, manageable problem. The first step was to formalize an alternative to viewing a supergame strategy profile as a vector of infinite sequences of functions from histories into action sets. Notice that a strategy profile implicitly specifies what path⁶ should be followed, what new path should be followed if someone deviates from the original path, and so on. Indeed, the profile can be thought of as a *collection of paths* and a *rule* governing how to switch amongst them in the event of deviations. On the face of it, this perspective does not look promising: the collection of paths could be infinite and the rule arbitrarily complex. Abreu (1988) justified the reformulation, however, by showing that for any pure strategy subgame perfect equilibrium of G^δ , there is another perfect equilibrium that has the same value, and can be described by $n+1$ paths and an extremely elementary rule. For any $n+1$ paths Q_0, Q_1, \dots, Q_n , define the associated *simple strategy profile* $\sigma(Q_0, Q_1, \dots, Q_n)$ by the rule

- (i) Q_0 is the initial path
- (ii) after a deviation by a single player i from any ongoing path, play switches to following the path Q_i from the beginning (so if i deviates part way through path Q_i , for example, the path Q_i is restarted).

Working in the space of paths rather than supergame strategies affords a nice proof of the compactness of the equilibrium value set (henceforth, except in the statements of proofs, I often omit the qualifier "subgame perfect"). One implication is that *severest credible punishments for each player exist*. Let Q_1, \dots, Q_n be the

⁶A *path* is a sequence of action profiles, one for each period.

respective paths of some severest equilibria for each player. A central result of Abreu (1988) is that if Q_0 is the equilibrium path of *any* (perfect) equilibrium γ (simple or not), the simple profile $\sigma(Q_0, Q_1, \dots, Q_n)$ is also a perfect equilibrium (clearly with the same equilibrium path). Let us check that the profile satisfies the criterion of unimprovability, that is, from no point of any of the $n+1$ paths would a player i wish to deviate, given that path Q_i will subsequently be followed. Each path Q in the set $\{Q_0, Q_1, \dots, Q_n\}$ is the path of some perfect equilibrium, and hence each player i was deterred from cheating at any point on Q by the threat that play would switch to the path of some perfect continuation equilibrium. But at the same point on Q in $\sigma(Q_0, Q_1, \dots, Q_n)$, player i is faced with a threat at least as severe (because Q_i is by construction the worst perfect path for i). Thus player i cannot gain by deviating in any contingency.

A proposition summarizes our discussion.

PROPOSITION (Abreu, 1988): *Let $G = (A_1, \dots, A_n, \Pi_1, \dots, \Pi_n)$ have at least one equilibrium in pure strategies, and for each i , suppose A_i is compact and Π_i is continuous. Then*

- (a) *the pure strategy subgame perfect equilibrium value set $V^s(\delta)$ is nonempty and compact, and*
- (b) *for any equilibrium γ , there is a simple strategy profile that is a perfect equilibrium with the same path (and hence the same value).*

Why are strategy profiles of the form $\sigma(Q_0, Q_1, \dots, Q_n)$ called simple? In general the $n+1$ paths might themselves be highly nonstationary and complex. But the way in which deviations are responded to, that is, the implicit punishment hierarchy, is simple in the extreme. A deviation by player i is always treated the same way, regardless of the nature of the deviation, the period in which it occurred, the particular path in progress, or the point on the path at which the defection occurred. There is no need to "tailor the punishment to fit the crime."

The preceding analysis does not apply to mixed strategy equilibria, because it is not possible to tell from observing the actions played, whether or not the correct mixed strategies were employed. In other respects, however, the theory is comprehensive in its scope, covering for example all finite games G and all discount factors. A demonstration of its practical value in applied fields was given by Abreu (1986) in a study of optimal collusion among Cournot oligopolists. He considered n identical firms with positive, constant marginal costs c , no fixed costs, and strategy spaces $[0, \infty)$. They face a smooth inverse market demand function P satisfying $\lim_{q \rightarrow \infty} P(q) = 0$. The conditions of the last proposition above are assumed to hold, except for boundedness of the strategy space; the set of quantities q that a firm could conceivably play in equilibrium is bounded, and this is all that is needed.

To derive strong results about the shape of the paths of the worst punishments, it is necessary to restrict attention to *strongly symmetric equilibria*, that is, equilibria which after no history give any player different

instructions from any other player. It is easy to check that without loss of generality, the equilibrium paths of constrained optimal strongly symmetric equilibria may be taken to be *stationary*. But this is emphatically not true for severest punishment paths. The latter have a "stick and carrot" structure that is quite striking: the payoff in the first period is dismal, but the path starting in period 2 is constrained Pareto efficient. In other words, the misery is front-loaded to the maximum extent possible.

PROPOSITION (Abreu, 1986): *In a symmetric Cournot oligopoly satisfying the conditions described above, there is a most severe strongly symmetric equilibrium whose continuation value following the first period is constrained Pareto efficient. There is a critical value of δ above which the present value of profits in the severest symmetric equilibrium is zero.*

The idea here is that starting from any equilibrium profile not of the stick and carrot form, one can replace the continuation equilibrium by the Pareto efficient equilibrium, and restore the entire path to its original value by increasing first period production sufficiently. A firm's supergame payoff from conformity is the same as before, but the payoff to cheating in period 1 is generally lower (and never higher), because the firm's residual demand curve is lower. Hence, under the new arrangement incentives to conform are at least as strong as in the original equilibrium.

It is sometimes the case that in a stick and carrot regime, the deviant firm is cooperating in its own punishment, that is, in period 1 it is not playing a best response to other firms' production. One might have thought that this was impossible. After all, if you punish someone as severely as possible, how can you expect him to cooperate when there is nothing worse left to threaten him with? The answer is that you can threaten to restart the punishment, a sobering prospect in the case of stick and carrot punishments. One might also have guessed that it is impossible to have an equilibrium with value zero in a Cournot model, since a firm can always choose to play a best response in period 1, pocket the profits and produce nothing thereafter. The strategy can be foiled only by having first period production so high that the output of $n-1$ firms is enough to reduce price to marginal cost (or below), so that no firm can make money in period 1 by any choice of output level. Firms are in effect "mutually minimaxing" one another, a phenomenon that is impossible in some other models, as we shall see later in an example due to Fudenberg and Maskin (1986).

An attractive paper by Lambson (1987) uses simple strategy profiles to characterize optimal collusion in price-setting supergames with capacity constraints and alternative rationing rules. For some, but not all, of the rules considered it turns out that stick and carrot punishments are optimally severe, and using mixed strategies

would not expand the set of equilibrium values. In all cases the restriction to strongly symmetric equilibrium is without cost, in contrast to Cournot oligopoly.

There are numerous other applications of repeated games in particular areas, some using simple strategy profiles and others employing trigger strategies. A few examples are Barro and Gordon (1983), which stimulated much interest in strategic monetary theory (see the excellent survey by Rogoff (1989)); Weinberger (1990) on bargaining and delay to agreement, and Rotemberg and Saloner (1989) and Syropoulos (1989) on the relative merits of tariffs and quotas in dynamic trade policy.

Discount Factors Close to 1

Understanding behavior in repeated games with discount factors significantly different from 1 is important for several reasons, including the fact that the discount factor may represent both impatience in the usual sense and the chance that the strategic interaction may be interrupted by external factors (new laws, product innovations, and so on). But there are many examples in which the period length is sufficiently short that the players' primary concern is for the future. Thus the perfect discounted folk theorems of Fudenberg and Maskin (1986) occupy a special place in the literature. They demonstrate that, with two qualifications, the classical results of Aumann and Shapley and Rubinstein survive the introduction of a small amount of impatience. An example of Forges, Mertens and Neyman (1986) showed that values in which some players receive exactly their security levels may not be the payoff of any perfect equilibrium with discounting. For 2-person games, this is the only qualification that need be made to the earlier folk theorems.

PROPOSITION. Perfect Folk Theorem in Discounted 2-Person Games (Fudenberg and Maskin, 1986): *Let g be a finite 2-person game, and v be feasible and strictly individually rational (for each i , $v_i > v_i^s$). There exists $\underline{\delta}$ such that for all $\delta \in [\underline{\delta}, 1)$, v is the average discounted value of some subgame perfect equilibrium of $G^\delta(\delta)$.*

The proof depends critically on the possibility of players' simultaneously minimaxing one another. This cannot always be done in n -person games. For some such games, the folk theorem fails, as a neat example of Fudenberg and Maskin shows.

In the simultaneous 3-person game of Figure 1, player 1 chooses the row, 2 the column, and 3 determines which of the two matrices applies. Note that the three players' payoffs are always identical. Each person's security level is 0, but one can verify that for any strategy profile (pure or mixed), there is some player whose best response payoff is at least 1/4. Choose any $\delta \in (0, 1)$, and any subgame perfect equilibrium, and let i be a player whose myopic best response in the first period gives him at least 1/4. Since he will get no less than

		2	
	1,1,1	0,0,0	
1	0,0,0	0,0,0	

		2	
	0,0,0	0,0,0	
1	0,0,0	1,1,1	

Figure 1

the minimum equilibrium value, call it ξ , from the second period onward, regardless of what he does in the first period, we see that

$$\xi \geq (1-\delta)\frac{1}{4} + \delta\xi, \text{ that is, } \xi \geq \frac{1}{4}.$$

In other words, equilibrium payoffs are bounded away from the security level, uniformly in δ .

A sufficient condition for obtaining a full folk theorem in n -person games is that the set of feasible and individually rational payoffs of G be full-dimensional.

PROPOSITION. Perfect Folk Theorem with Discounting (Fudenberg and Maskin, 1986): *Let $G = (A_1, \dots, A_n; \Pi_1, \dots, \Pi_n)$ be a finite game such that the set F^+ of feasible, individually rational payoffs is of dimension n . Then for any feasible, strictly individually rational value v , there exists $\underline{\delta}$ such that for all $\delta \in [\underline{\delta}, 1)$, v is the average discounted value of some subgame perfect equilibrium of $G^\delta(\delta)$.*

To understand the idea of the proof, take the tidiest case, where the value v is in the interior of F^+ (this set would be empty if the dimensionality condition were violated) and there is some $c \in A$ with $\Pi(c) = v$. Choose $\epsilon > 0$ and n vectors

$$v(j) = (v_1 + \epsilon, \dots, v_{j-1} + \epsilon, v_j, v_{j+1} + \epsilon, \dots, v_n + \epsilon)$$

such that $v(j) \in F^+$, $j = 1, \dots, n$. For simplicity assume there are action profiles b^j and m^j , $j = 1, \dots, n$, where $\Pi(b^j) = v(j)$ and m^j minimaxes j (and has j playing a best response). Let $\Delta_j = \max_{a \in A} \Pi_j(a) - v_j$, $j = 1, \dots, n$. The simple strategy profile with paths as described below is a subgame perfect equilibrium of $G^\delta(\delta)$ for sufficiently high δ : on the equilibrium path, c is played indefinitely; j 's punishment path begins with k periods of m^j , and b^j thereafter, with k chosen large enough so that for each j , $k(v_j - v_j) > \Delta_j$. Checking for

unimprovability, we note first that no player j wants to deviate from the original path, because he then gets minimaxed for k periods. While being minimaxed, j cannot profitably cheat, because he is already playing a one-shot best response. In the second phase of the punishment path, j has no incentive to cheat because again, the result would be to be minimaxed for k periods instead of receiving v_j . At no time would a player other than j wish to deviate from the j^{th} punishment path, because for high δ , any short-run gains would be overwhelmed by the loss of an infinite stream of "bonuses" ϵ . (Note that the second phase of the punishment path is designed to reward players $i \neq j$ for minimaxing j , without also treating j favorably.)

Self-Generation

The sweeping characterization of equilibrium values when δ is near 1 has no analogue for arbitrary discount factors. There is, however, a useful sufficient condition for sets of values to be subsets of the supergame value set. The result, called "self-generation," was developed by Abreu, Pearce and Stacchetti (1986, 1990) for games with imperfect monitoring, but the principle behind it is quite general, and applies in the simple case of perfect monitoring (explicit treatments of self-generation in this setting can be found in Sabourian (1989) and in more detail in Cronshaw and Luenberger (1990)). Self-generation is in the spirit of dynamic programming, in the sense that it depends on the decomposition of a supergame profile into the induced behavior today and the *value* of behavior in the future, as a function of all possible actions today. The following discussion tries to motivate the result. The analysis can be done for mixed strategies (this is not the case in games with imperfect monitoring) but for ease of exposition, I consider pure strategies only.

What makes playing the first period of $G^\delta(\delta)$ different from playing G in isolation? In the former case, each player is interested in maximizing a weighted sum of his immediate payoff in G and his continuation payoff in the remainder of the game. In equilibrium, the vector of continuation payoffs after a particular first-period history is drawn from the (subgame perfect) equilibrium value set V of $G^\delta(\delta)$. Thus, $v \in V$ if and only if for some $a \in A$ (representing first-period actions) and $u : A \rightarrow V$ (contingent continuation payoffs),

$$(1) \quad v = (1 - \delta)\Pi_i(a) + \delta u_i(a)$$

and

$$(2) \quad (1 - \delta)\Pi_i(a) + \delta u_i(a) \geq (1 - \delta)\Pi_i(a'_i, a_{-i}) + \delta u_i(a'_i, a_{-i})$$

for all $a'_i \in A_i$, $i = 1, \dots, n$.

Notice that when one is allowed to affect first-period behavior using continuation values from V , one "generates" exactly the elements of V as values of equilibria created in the augmented static games. More generally, think of augmenting payoffs by values drawn from an arbitrary set $W \subseteq \mathbb{R}^n$, and call the values generated $B(W)$:

$$B(W) = \{(1-\delta)\Pi(a) + \delta u(a) \mid u : A \rightarrow W, \text{ and } (a,u) \text{ satisfies (2)}\}.$$

We see immediately that V is a fixed point of the map $B : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$. Let $B^t(W)$ denote the t^{th} iteration of B on W . For example, $B^2(W) = B(B(W))$.

A nonempty bounded set $W \subseteq \mathbb{R}^n$ is called *self-generating* if $W \subseteq B(W)$. If W is self-generating, there is enough variety in the payoffs in W to create incentives for different equilibria in the corresponding augmented games, indeed enough to generate any value of W . This leads to the conjecture that the values in W are actually equilibrium values, because they are able to generate themselves, just the way supergame equilibria generate equilibrium values by using supergame equilibrium values as continuation payoffs.

PROPOSITION. Self-Generation (Abreu, Pearce and Stacchetti, 1990). *Let G be a finite game and $\delta \in (0,1)$, and let $B : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$ be as defined above. Then if $W \subseteq \mathbb{R}^n$ is self-generating, $B(W) \subseteq V$ (indeed, for $t = 1, 2, \dots$, $B^t(W) \subseteq V$).*

Self-generation has many applications, both theoretical and practical, and will be encountered again in subsequent sections. Here I record one implication that will be helpful later in unifying results from different papers.

PROPOSITION. Algorithm: *Let G be a finite game, $\delta \in (0,1)$, and B be the associated generation map. For any bounded $W \subseteq \mathbb{R}^n$ with $V \subseteq W$,*

$$\bigcap_{t=1}^{\infty} B^t(W) = V.$$

The Proposition gives an algorithm for computing the equilibrium value set: choose any set that is "large enough" (F will do, for example), and apply the map B repeatedly. The limit of this process is V . Recently Cronshaw and Luenberger (1990) have given conditions under which the strongly symmetric equilibrium value set of symmetric repeated games may be computed with a non-iterative procedure. Their technique involves finding the largest solution of a scalar equation and uses the dynamic programming approach.

Relationships to Finitely Repeated Games

I turn now to finitely repeated games and their relationship to infinitely repeated games. It was long thought that finite horizon repeated games were of little theoretical interest because backward induction arguments could be used to show that subgame perfect equilibrium behavior in $G^T(\delta)$ could involve only a string of one-shot equilibria of G . While this is true if G has a unique equilibrium, Benoit and Krishna (1985) and Friedman (1985) showed resoundingly that more generally, the presumption was false. Benoit and Krishna showed that if for each player, not all of the equilibria of G have the same value, then folk theorems similar to those of Fudenberg and Maskin (1986) hold for $G^T(1)$ as T becomes large. (Unlike Fudenberg and Maskin (1986), Benoit and Krishna restrict attention to pure strategies.)

PROPOSITION. Folk Theorem for Finitely Repeated Games (Benoit and Krishna, 1985): *Suppose that*

- (i) *for each player i , there are two equilibria of G with different payoffs for i , and*
- (ii) *$n = 2$, or $\dim F^+ = n$.*

Then for any value v that is feasible and individually rational (relative to pure strategies) and any $\epsilon > 0$, there exists T_0 such that for each $T > T_0$, there exists a subgame perfect equilibrium of $G^T(1)$ with average value within ϵ of v .

This striking theorem is not only a result about finite horizon games, but also a testament to the intimate connection between infinitely and finitely repeated games. Both the statement of the theorem and the line of proof resemble closely those of the perfect folk theorem for discounted infinitely repeated games. In fact, using arguments mimicking those of Benoit and Krishna, one can strengthen their statement as follows: if for each player i there are two equilibria of G with different payoffs for i , then $\lim_{T \rightarrow \infty} V^T(1) = \lim_{\delta \rightarrow 1} V^\infty(\delta)$. The equivalence holds regardless of the number of players or the dimension of F^+ .

In games G having only one equilibrium, $G^T(\delta)$ has a unique subgame perfect equilibrium. But Radner (1980) pointed out that even in this case, cooperation is possible if the solution concept is (perfect) ϵ -equilibrium, that is, if after each history, any player's strategy is within $\epsilon > 0$, in average value terms, of the best strategy available from then on. In a game $G^T(\delta)$ with T large, anything that happens in the last few periods matters little in average terms, so that there are ϵ -equilibria in which cooperation is induced by the threat that endgame behavior will depend on play earlier in the game. Radner proves that for arbitrarily small $\epsilon > 0$, asymptotically efficient average payoffs can be obtained in ϵ -equilibria with patient players as T grows

large. This is a valuable technical result, but I find the interpretations in Radner (1980), Section 8 in terms of search costs and bounded rationality to be forced and unconvincing.

Fudenberg and Levine (1983) elaborated on Radner's idea to produce a powerful equivalence result for finite and infinite horizon games. For present purposes it is specialized to strictly repeated games and stated in terms of equilibrium values.⁷

PROPOSITION (Fudenberg and Levine, 1983): *For any finite game G and $\delta \in (0,1)$, $v \in V^\infty(\delta)$ if and only if there is a sequence $(\epsilon_T, v_T)_{T=1}^\infty$ such that*

- (i) $\epsilon_T > 0$ and v_T is the value of some ϵ_T -equilibrium of $G^T(\delta)$ for each T , and
- (ii) $\epsilon_T \rightarrow 0$ and $v_T \rightarrow v$.

Thus, supgame equilibrium values are exactly the limits of ϵ -equilibria of $G^T(\delta)$ as $\epsilon \rightarrow 0$ and T grows large.

To see why relaxation of the incentive constraints by an arbitrarily small $\epsilon > 0$ suffices to admit in $G^T(\delta)$ behavior associated with equilibria of $G^\infty(\delta)$, choose any subgame perfect equilibrium σ of $G^\infty(\delta)$, and T large enough so that the average value of any strategy differing from σ only after T is within ϵ of $v(\sigma)$. Let $\sigma(T)$ be the profile induced on $G^T(\delta)$ by σ . Because σ_i is a perfect best response to σ_{-i} in $G^\infty(\delta)$, our choice of T ensures that σ_i is an ϵ -perfect best response to $\sigma(T)_{-i}$ in $G^T(\delta)$, that is, $\sigma(T)$ is ϵ -perfect.

This paragraph is quite difficult, and can be skipped without loss of continuity. A couple of years ago David Kreps suggested to me that the Fudenberg-Levine limit result and the algorithm discussed earlier are related. There are a number of ways of explaining the connection; here is one. Recall that $B(W; \delta)$ is the set of equilibrium values generated by creating new games from G by modifying the payoffs with values from W . Now $B^2(W; \delta)$ is the set of values obtained by augmenting G with the continuation value set $B(W; \delta)$. Equivalently, $B^2(W; \delta)$ is the set of values of perfect equilibria of the *two-period games* obtained by modifying the terminal payoffs of $G^2(\delta)$ by values from W . The same is true for $B^T(W; \delta)$ and the W -augmentation of $G^T(\delta)$. For any positive integer T , and $\epsilon > 0$, let $W_{T, \epsilon} = \{w \in \mathbb{R}^n \mid 0 \leq w_i \leq \delta^{-T} \epsilon, i = 1, \dots, n\}$, which is a "cube" of size $\delta^{-T} \epsilon$. After some reflection, one sees that the ϵ -equilibria of $G^T(\delta)$ are exactly the equilibria of the $W_{T, \epsilon}$ -augmentation of $G^T(\delta)$. This prompts two observations. First, for a self-generating set $W \subset B(W)$, since the $B^T(W)$ converge to the limits of ϵ_T -equilibria with $\epsilon_T \rightarrow 0$ and $T \rightarrow \infty$, the ϵ -equilibrium limit result says that $B^\infty(W) \subset V$, thereby providing another perspective on the self-generation theorem. Conversely, the Fudenberg-Levine limit theorem can also be appreciated from the point of view of value iteration.

⁷Further work on equilibria of convergent sequences of games include Harris (1985) and Fudenberg and Levine (1986).

I would like to make some informal remarks aimed at creating an intuition that applies to several of the papers discussed so far. In period t of an infinite horizon game $G^\delta(\delta)$ the one-period incentive constraints of G are loosened to an extent that depends on the size of the continuation value set. This loosening of incentive constraints is mimicked in Fudenberg and Levine (1983) by the use of ϵ -equilibria, and in Abreu, Pearce and Stacchetti (1986, 1990) by the extra payoffs drawn from the set W . If the value of ϵ , or the set W , is too small, some incentive constraints may be violated; the algorithm does not necessarily work "from below" (contrast this with value iteration (Howard, 1960) in dynamic programming, where any initial values are acceptable). How, then, do Benoit and Krishna (1985) guarantee that there is enough "punishment power" to support cooperative behavior in $G^T(1)$ for large T , even if the multiple equilibria of G differ in value only minutely for some or all players? With $T = 10,000$ and $n = 5$, for example, play in the last 500 periods could consist of 100 periods of each player's respective favorite equilibrium. In the period preceding this 500-period endgame, each player has a lot to lose (remember that $\delta = 1$). Thus, values of $G^{10,000}(1)$ will be a lot like values of $B^{9,500}(W,1)$ where W is a large set. Chou and Geanakoplos (1987) show that in a generic class of games with continuous action spaces, allowing arbitrary behavior *in the last period only* of $G^T(1)$ is enough to generate a folk theorem, even though G may have a unique equilibrium. The subtlety here is that the leeway created by the arbitrary end-period behavior can, by the envelope theorem, be used to disturb behavior slightly in many preceding periods; when summed (without discounting), these changes have large value consequences and hence can create substantial incentives.

Cooperation amongst Mortals

Although the abstraction of a world that continues indefinitely is a useful device, modelling individuals as infinitely lived is less attractive. It seems important to inquire, then, into the possibilities of self-enforcing agreements amongst finitely lived agents in an infinite horizon world. One could avoid the question by arguing that people, no matter how old, have a good chance of living in period $t+1$ given that they are alive in t , especially if the period is a day or a week, for example. Another escape route is to note that reputation can be vested in an institution such as a firm, whose mortal owners behave in accordance with implicit understandings (even just before selling the firm) to protect the firm's market value; this is highly plausible and is one of the ideas explored by Kreps (1987) in his paper on corporate culture.

But what of situations where the participants' limited horizons are known precisely and reputations reside exclusively with the individual? If equilibrium in the component game is unique, as in many free-rider problems, things look discouraging at first glance. Since it is impossible to induce a person to cooperate in the final period

of his life, misbehavior having no future repercussions for him, presumably incentives unravel by the backward induction argument familiar from finite horizon games. Crémer (1986) showed that this need not happen, and his work has been generalized by Cooper and Daughety (1988), Salant (1988), Kandori (1989a), and Smith (1990). In an overlapping generations model, suppose that society acknowledges that in the final 3 periods of his life, say, no individual will act cooperatively. Hence, selfish behavior by the aged is part of the implicit agreement. But if any young person fails to cooperate, the accord is broken and everyone subsequently optimizes myopically. Young persons will choose not to defect, because they would lose the benefits of social cooperation for the rest of their lives. Folk theorems similar to those discussed earlier hold here, and Kandori (1989a) shows that if successive individuals are born far enough apart in time, there is no need to invoke any full dimension restriction. Recall that this assumption is usually made to ensure that punishers can be rewarded for minimaxing a defector, without incidentally also rewarding the defector himself. In Kandori's construction, the punishers wait until the defector dies, and then celebrate their earlier self-discipline.

Cooperation in Matching Models

When large numbers of players are partitioned into pairs who interact strategically perhaps for only one period before the pairings are rearranged, a particular player i may observe, at the end of period t , exactly what his partner j did in that period, while others may be uninformed or only partly informed about j 's action. This makes it harder to sustain cooperation, because the group as a whole does not have the information needed to respond immediately and concertedly to a transgression by one individual. Nonetheless, self-enforcing agreements are sometimes possible even under such poor informational conditions. Kandori (1989b) studies trigger strategies in a repeated prisoners' dilemma matching game (always cooperate until someone you meet plays tough, and then play tough against everyone you subsequently meet). He shows that these strategies are a perfect equilibrium when δ is near 1. (The delicate constraint to check here is that when a person is cheated for the first time, he is willing to accelerate the decay of goodwill in the community by treating his next partner ungenerously.)

Okuno-Fujiwara and Postlewaite (1989) focussed attention on environments with somewhat better information flows, which they call "local information processing." Each person has a "label" observable by his partner. The label in period $t+1$ depends only on the labels and actions of the individual and his t -partners in period t . (Examples of labels include membership in an organization and possession of a license or credit card.) Folk theorems hold for communities with local information processing and infinite populations (Okuno-Fujiwara and Postlewaite, 1989) or, under additional assumptions, finite populations (Kandori, 1989b.) Community enforce-

ment of social norms for bilateral strategic behavior has become the subject of much interest amongst economic historians. Recently a number of papers have traced the development of institutions that promoted community enforcement of fair trade practices in the absence of adequate legal sanctions (see especially Greif (1989), Milgrom, North and Weingast (1990) and Greif, Milgrom and Weingast (1990).

3. IMPERFECT MONITORING

Even during the early development of the theory of cooperation in games with perfect monitoring, researchers became dissatisfied with its scope. In many economic examples of practical interest, the assumption that players observe one another's past actions is inappropriate. Instead, player i observes the outcome of some random variable (team output, number of product failures or consumer complaints, market price, and so on) whose distribution is affected by the private actions of some or all of the players. Positive results for models of this kind again appeared first for games without discounting. The pioneering papers by Rubinstein (1979b), Radner (1981) and Rubinstein and Yaari (1983) proved that in infinitely repeated principal-agent games of various kinds, it is possible to overcome the inefficiency associated with the moral hazard problem in the static model. Rubinstein and Yaari also remarked that their arguments could be extended to yield a perfect folk theorem for agency games with imperfect monitoring.

The Green-Porter Model

Green and Porter (1984) and Porter (1983a) were the first papers to study discounted repeated games in which players receive information related only stochastically to others' actions. Whereas the work without discounting had concentrated on one-sided imperfect monitoring (the principal's actions were not private), Green and Porter were interested in seeing whether n players all of whose actions are taken privately, could sustain cooperative (non-myopic) behavior by making their actions conditional on a relevant, commonly observed random variable. They answered the question in the affirmative in a Cournot oligopoly with random shocks to market price. By producing less following some observed prices than others, firms can create an implicit reward function (mapping observed prices into supergame continuation payoffs). For economists, one of the most attractive features of the model is that it escapes the prediction of dynamically uniform behavior on the most collusive equilibrium path, thereby offering a possible interpretation of observed phenomena such as price wars.

Porter investigated symmetric equilibria that are optimally collusive among a restricted set of 'trigger strategy' profiles. A trigger strategy is described by a quantity q , a trigger price p , and a positive integer T . Firms begin by each producing q , and do so in every period until the price falls below p . A price realization of

less than p triggers a T -period phase of Cournot-Nash behavior, after which cooperation resumes (until the Cournot phase gets triggered again). Should incentives to produce q in the cooperative phase be provided by punishing frequently, or infrequently but with greater severity (larger T)? Porter found that the answer varies with the family of distributions used for stochastic demand, but often it is optimal to set $T = \infty$, that is, revert permanently to the stage-game Cournot-Nash equilibrium.

Constrained Optimal Solutions

Abreu, Pearce and Stacchetti (1986) dropped the restriction to trigger-strategy profiles and characterized optimal pure strategy symmetric equilibria of a class of games that generalize the Green-Porter model. They found that a constrained efficient solution is described by two "acceptance regions" Ω_1 and Ω_2 in the signal space (price space, in the oligopoly example) and two actions q_1 and q_2 . In the efficient equilibrium, players choose q_1 as long as the value of the signal falls in Ω_1 . Otherwise, they switch to q_2 and keep playing q_2 as long as the signal falls in Ω_2 (when it falls outside Ω_2 they switch back to q_1 , and so on). Thus, behavior on the optimal equilibrium path is a simple first-order Markov process with two states, indexed by the current "target action" q_1 or q_2 . Why should the efficient solution take this form?

The value of an equilibrium of the supergame is the weighted sum of the first-period expected payoff and the expectation of the continuation values from period 2 onward. The latter values are drawn from the *symmetric* subgame perfect equilibrium value set $V \subset \mathbb{R}$ (elsewhere V was used for equilibrium values in \mathbb{R}^n). Thus, in the oligopoly example, with expected payoff function Π and price density function f , we want to choose a first-period quantity q_1 and a continuation reward function $u(p)$ with values in V , to maximize $(1-\delta)\Pi(q, \dots, q) + \delta \int u(p)f(p, nq)dp$, subject to the incentive constraint that (given the immediate and future rewards) there is no alternative quantity that a firm would prefer to q_1 . If it weren't for the need to provide incentives, one would choose $u(p) = \max V$ everywhere. A subset of price space is a good place to assign a lower reward value if its occurrence is much more likely when myopically tempting deviations take place than when q_1 is produced. For example, if there is only one incentive constraint (that is, only one tempting alternative q'), the best places to punish are where the likelihood ratio $\frac{f(p, (n-1)q_1 + q')}{f(p, nq_1)}$ is high. In this case, assign the value $\min V$ to prices with very high likelihood ratios, and keep adding regions of price space (in decreasing order of likelihood ratio) to the punishment region until the incentive constraint is satisfied. This procedure concentrates the "punishment" into a region Ω_1^c that is as informationally efficient as possible. Using a larger region and a less severe punish-

ment will generally result in a loss of efficiency because of the region's poorer ability to discriminate between good and bad behavior.⁸

Thus, after one period of the best equilibrium, players will be instructed either to begin the worst equilibrium (if price fell in the punishment region) or to restart the best equilibrium (play q_1). Now the worst equilibrium corresponds to the problem of choosing an action q_2 and a reward function $w(p)$ from V , to *minimize* the sum of the current and continuation payoffs, while providing for adequate incentives. Again, we would like to give the minimum reward everywhere, but to create incentives efficiently we give rewards $\max V$ in a region Ω_2^c chosen for its discriminatory power. At the end of period 1, players are told to restart the worst equilibrium if price fell in Ω_2 , and to start the best equilibrium otherwise. Notice that in every contingency, players are duplicating the behavior of the first period of one of two equilibria (the best or the worst), so that only two quantities are ever produced. Switches between regimes are governed by the regions Ω_1 and Ω_2 , as specified earlier.

The requirements this solution places on players' memories is unexpectedly modest. They need only remember which of two quantities they were supposed to produce last period and what price arose.

Self-Generation under Imperfect Monitoring

Self-generation and related techniques were first developed in the context of unrestricted symmetric equilibria of the Green-Porter model, and then presented in greater generality in Abreu, Pearce and Stacchetti (1990). Suppose that players take private actions $a_i \in A_i$ (finite), $i \in N$, that determine the density $f(p; a)$ of a commonly observed random variable p with constant support Ω . Player i 's realized payoff depends on his action a_i and on the realization p . Let i 's *expected* payoff be $\Pi_i(a)$. In pure strategy equilibria of the repeated game, one-shot incentives are supplemented by continuation values drawn from the equilibrium value set $V \subset \mathbb{R}^n$. The continuation equilibria in effect create a (measurable) reward function mapping Ω into V . Hence, the natural value generation function to look at in this case is $B : 2^{\mathbb{R}^n} \rightarrow 2^{\mathbb{R}^n}$ defined by:

$$B(W) = \left\{ w \in \mathbb{R}^n \mid \exists (a, \mu) \in A \times L^\infty(\Omega, W) \text{ s.t. } w = (1-\delta)\Pi_i(a) + \delta \int u(p)f(p; a)dp \text{ and} \right. \\ \left. (1-\delta)\Pi_i(a) + \delta \int u(p)f(p; a)dp \geq (1-\delta)\Pi_i(b_i, a_{-i}) + \delta \int u(p)f(p; b_i, a_{-i})dp \quad \forall b_i \in A_i, i \in N \right\}.$$

Again, if W is nonempty and bounded and $W \subset B(W)$, W is called *self-generating*. This value-generation approach led to a number of results summarized below.

⁸Often the optimal region will consist of prices below some critical value p . When a tail test is *not* optimal, but is imposed arbitrarily, raising the critical value may add some informationally more efficient points to the region. This explains Porter's finding that sometimes $T < \infty$ is optimal; a unique interior solution of this kind can occur only when imposing a tail test is inappropriate.

PROPOSITIONS (Abreu, Pearce and Stacchetti, 1990).

Self-Generation. If W self-generating, then $W \subseteq \bigcup_{t=1}^{\infty} B^t(W) \subseteq V$.

Bang-bang Rewards. V is compact, and for all $v \in V$ there exists an equilibrium whose implicit reward functions after each history take only values in the set of extreme points of V .

Algorithm. If W is bounded and $V \subseteq W$, then $\bigcap_{t=1}^{\infty} B^t(W) = V$.

Monotonicity. If $0 < \delta_1 < \delta_2 < 1$, then $V(\delta_1) \subseteq V(\delta_2)$.

Under certain conditions the "bang-bang sufficiency" result given above can be strengthened to a necessity result: an equilibrium that maximizes a linear combination of player payoffs (including negative combinations) *must* have implicit reward functions that use only extreme points of V . The rough intuition is the same as the one given earlier for the Green-Porter model: if you are creating incentives by moving rewards in a direction that reduces the objective function of the problem, do so aggressively (move until you can't go any further in V) but in as small and informative a region of signal space as possible. This advice cannot be applied literally in a model with a discrete signal space, so the bang-bang necessity result does not hold. The sufficiency result can be restored trivially in an essentially discrete model if the signal space is taken to include the outcome space of a public randomization device.

The scope of the preceding analysis is limited in three ways to preserve the "recursive structure" of the supergame equilibria: players receive no private signals, they use only pure strategies, and the commonly observed signal has constant support. When any of these restrictions is relaxed, some equilibria may, after certain histories, have continuation profiles that are not Nash equilibria of the supergame. This arises because imperfect correlation may develop in the actions of different players who are conditioning their behavior on private signals from earlier periods (including realizations of their own mixed strategy randomizing device). Fudenberg, Levine and Maskin (1989) impose none of the three restrictions, but they avoid the messy consequences in one superbly pragmatic stroke. They consider only the *public equilibria* of the supergame, that is, profiles of strategies that are perfect best responses to one another and use information from earlier periods only if it is publicly observed. The continuations of these equilibria are again public equilibria, and a straightforward dynamic programming approach can be used. I return to Fudenberg, Levine and Maskin's work in some detail later in this section.

Discontinuity at $\delta = 1$

Following the appearance of the efficiency results for undiscounted repeated agency problems mentioned at the beginning of this section, Radner (1985a) demonstrated the existence of fully efficient perfect equilibria in a class of partnership games with the limit of means criterion. Especially once Radner (1985b) had shown that asymptotic efficiency could be attained in repeated discounted agency problems as δ approaches 1, it seemed likely that the same could be proved for discounted partnerships. Thus, theorists were particularly intrigued when Radner, Myerson and Maskin (1986) produced an example of a two-person repeated partnership game whose equilibria are bounded away from the efficient frontier, uniformly in δ . Each player has two strategies: work or shirk (the latter is a dominant strategy in the component game). The commonly observed signal is the shared output, which may be either high or low; the probability of low output is f_w or f_s if both players work or only one works, respectively, where $0 < f_w < f_s < 1$. Restrict attention to symmetric equilibria (for expositional ease), and let \bar{v} be the value of the maximal equilibrium. The best way to get players to work is to give a continuation value of \bar{v} when high output is observed, and a lower value $\bar{v} - x$ when low output is observed. Choose x just large enough that the expected loss in continuation value equals the average value of the absolute myopic gain (say g) from shirking:

$$(f_s - f_w)x = g \frac{(1-\delta)}{\delta} .$$

Even when both players work, low output occurs with probability f_w , so that if Π is the expected payoff in G when both work, we have (if δ is not too low):

$$\begin{aligned} \bar{v} &= (1-\delta)\Pi + \delta(\bar{v} - f_w x) \\ &= (1-\delta)\Pi + \delta \left(\bar{v} - \frac{f_w}{f_s - f_w} g \frac{(1-\delta)}{\delta} \right) \\ \therefore \bar{v} &= \Pi - \frac{g}{\ell - 1} , \end{aligned}$$

where ℓ is the likelihood ratio f_s/f_w . Since δ does not appear in the expression for \bar{v} , we see that the average payoff does not approach the first best Π as δ approaches 1. The average efficiency loss $g/(\ell-1)$ is proportional to the one-shot gain from cheating, and inversely proportional to the (transformed) likelihood ratio of the punishment region. Abreu, Milgrom and Pearce (in press) show that this formula applies quite generally to symmetric equilibria of repeated partnership problems; I explain later how they use this to study the effects of changing information and timing in such games.

The limit of means criterion and discounting with δ near 1 are alternative ways of modelling very patient players. Together, the papers by Radner (1985a) and Radner, Myerson and Maskin (1986) show that they are by no means equivalent; this is sometimes called a "discontinuity at $\delta = 1$." In my opinion the repeated partnership (and most repeated games) are better modelled with discounting than with the limit of means criterion, and the example under discussion illustrates this well. If players are to have incentives to work in period t , they must be punished (sooner or later) if period t output is low. Since low output may occur even under good behavior, this imposes a real cost, one which must be borne every time players are supposed to work. The per-period nature of the problem is nicely reflected in the discounting case, where the loss gets capitalized in the value set. Without discounting, it is *not* necessary to deter shirking period by period: if a player cheats for k periods, it has no effect on his long-run average payoff. Only infinite strings of deviations are a problem, and these Radner detects using a "review strategy" that, according to the law of the iterated log, will yield a first-best equilibrium average payoff. I can think of few economic problems that are well modelled by the assumption that it is safe to ignore incentives in any particular 50,000 periods. For this reason I consider discounted folk theorems (and counterexamples) important advances, even in the presence of the comprehensive theory for $\delta = 1$.

The need to deter single deviations (with discounting) and its absence (with the limit of means) probably explains the difference in methodologies in the literatures with and without discounting. Statistical methods are ideally suited to guarding against long-run deviations, whereas dynamic programming methods are largely inapplicable at $\delta = 1$ (recall, for example, the failure of the equivalence of unimprovability and perfect best responses). With discounting, the problem of deterring current deviations leads naturally to the decomposition of a supergame profile into behavior today and continuation values for the future. The dynamic programming perspective has the benefit of unifying the treatment of patient and impatient players, infinite and finite horizon games, and implicit and explicit contracts (of which, more later). This is not to say that the statistical approach cannot be used to advantage when payoffs are discounted; see, for example, the work of Fudenberg and Levine (in press) on folk theorems for approximate equilibria.

Information and Timing

In models with perfect monitoring, fixing the players' rate of time preference and shortening the length of the period (of fixed actions) is equivalent to letting the discount factor approach 1: in either case, today's payoff becomes a small part of total payoffs, so that the folk theorems have two interpretations. With imperfect monitoring, shrinking the period length still implies less discounting from one period to the next, but it also leaves less time for players to observe signals relevant to behavior. For example, if signals arrive according to a Poisson

process in continuous time, with the arrival rate determined by players' current behavior, the quality of information (in a sense relevant for incentives, as explained below) available over a period of time of length s deteriorates as s decreases. Hence, there are two effects of reducing the period length: an effective increase in patience, which we know from the monotonicity result stated earlier tends to increase the average value set, and a worsening of information, which Kandori (1988) has elegantly shown to decrease the set of equilibrium values. Either of these two effects can dominate in a particular case.

The upper bound for \bar{v} developed above for a simple partnership problem holds as stated for symmetric equilibria of repeated partnerships with arbitrary signal spaces, as shown by Abreu, Milgrom and Pearce (in press), and with slight modification for more than two actions. Attaching different punishments (continuation values) to different signal values is equivalent to using the severest punishment with different probabilities, which in turn simply amounts to choosing a region (say Ω_0) of extended signal space (the product of the natural signal space and the range of a public randomizing device) in which to punish uniformly. Once Ω_0 has been chosen optimally, which can be accomplished by solving a linear program, the efficiency loss from providing incentives for cooperation is $g/(t-1)$, as before (if cooperation is possible at all).

Now think of the Poisson example with arrivals interpreted as "good news," such as research breakthroughs or the winning of major contracts. If all members of the team are working, perhaps the arrival rate over a year's time is 10, whereas it drops dramatically to 1 if anyone shirks. If the period of fixed action is a year, under plausible parameter values cooperation could be sustained very profitably. But suppose that instead actions can be changed daily. The only way to encourage cooperation is to punish the event that there is no good news (zero arrivals), which has probability near 1 whether anyone shirks or not. Thus, the likelihood ratio is little more than 1, so that the efficiency loss is enormous. More precisely, as the period length shortens, the value falls until cooperation is no longer possible, and the formula ceases to apply. Ironically, in this case the players' ability to respond quickly to information destroys all possibility of cooperation. This suggests that *delaying* the release of information might actually be valuable in partnerships; Abreu, Milgrom and Pearce show that for high δ , information delays can virtually eliminate the inefficiency that Radner, Myerson and Maskin identified.

Folk Theorems with Imperfect Monitoring

The prospects for a general folk theorem for discounted repeated games with multi-sided moral hazard seemed dim, in the face of the Radner, Myerson and Maskin counterexample. But gradually a number of papers challenged the presumption that the troublesome example was representative. First, Williams and Radner (1987) showed that efficiency could be approached in generic static partnerships with enforceable contracts. Matsushima

(1989) subsequently used a first-order approach and some fairly palatable assumptions on information and the value set to generate asymptotic efficiency in equilibria of infinitely repeated partnerships. Next, Fudenberg, Levine and Maskin (1989) independently demonstrated that under remarkably weak conditions on primitives, a folk theorem holds for a wide class of games including those with moral hazard on all sides. Demougin and Fishman (1988) also showed that under reasonable conditions, oligopolies with imperfect monitoring could enjoy efficient collusion.

I concentrate here on the paper by Fudenberg, Levine and Maskin (hereafter FLM) because it is by far the most general, and represents the state of the art in discounted folk theorems for a broad range of information structures. Anyone interested in repeated games should read it closely. To avoid the introduction of further notation, I shall simplify their model in a way that makes it easy to describe here. Start with the n -person repeated game with imperfect monitoring studied by Abreu, Pearce and Stacchetti, discussed earlier, and make two changes:

- (i) let the signal take on only k values, k finite, and
- (ii) drop the "constant support" assumption.

As FLM explain, this model embraces perfect monitoring games (where the signal is simply the vector of players' actions), oligopolies, partnerships, and principal-agent problems (where the agent's action in the component game is a plan contingent on the compensation function offered by the principal). What I omit here are adverse selection problems, discussed in the final section of FLM.

Logically prior to the possibility of *efficient* cooperation is the question of whether cooperation can be supported at all. If one is allowed to employ arbitrary continuation values in \mathbb{R}^n as threats and promises, can players necessarily be induced to take a particular desired vector of actions? With perfect monitoring, the answer is obviously yes: players will do anything to avoid sufficiently severe punishments. With imperfect monitoring, however, player 1 might have three actions a_1 , b_1 and c_1 such that given some profile of actions for other players, the distribution of the public signal is the same when 1 chooses c_1 as when he randomizes between a_1 and b_1 with probabilities .6 and .4, for example. If the component game payoff to 1 is higher for both a_1 and b_1 than for c_1 , then it is impossible to induce him to play c_1 . No matter what rewards are attached to signal realizations, switching from c_1 to the mixture raises player 1's immediate payoff and leaves the distribution of continuation rewards unchanged. Hence, FLM's first informational assumption is one ensuring that a player's different possible actions can be distinguished, and hence encouraged or discouraged. Specifically, they impose the *individual*

full rank condition: at each profile $a \in A$ and for each player i , the $k \times m_i$ matrix⁹ whose columns are the probability distributions induced by the respective action profiles (b_i, a_{-i}) , $b_i \in A_i$, has rank m_i , that is, the probability vectors corresponding to each pure action of i are linearly independent. It is easy to verify that this guarantees that any behavior can be *enforced* if arbitrary continuation payoffs can be used.

Enforceability of this kind is clearly not enough to yield a folk theorem, because the Radner, Myerson and Maskin counterexample satisfies individual full rank. The problem there was that the only way to enforce good behavior was to punish *both* players in the event that output is low. Efficient (or nearly efficient) cooperation in a model where no player's actions are observed, generally requires that when one player's continuation payoff is reduced, another's must be increased:¹⁰ surplus should be passed back and forth amongst players, not thrown away. For a transfer of surplus from i to j to be effective in creating incentives, it needs to be associated with information that discriminates statistically between deviations by i and deviations by j . The availability of such information is ensured by the *pairwise full rank* condition: for each pair of players i and j there is some profile $\alpha \in A$ such that the $k \times (m_i + m_j)$ -dimensional matrix whose first m_i columns are the respective public distributions induced by the vectors (b_i, α_{-i}) , $b_i \in A_i$ and whose final m_j columns are the distributions induced by the vectors (c_j, α_{-j}) , $c_j \in A_j$, is of rank $m_i + m_j - 1$. (There is inevitably one linear dependency among the columns, because both i and j can create the distribution associated with α by putting the appropriate weights on pure actions.)

It would have been reasonable to guess that to prove that a desired profile γ can be enforced (almost) efficiently, it would be necessary to impose pairwise full rank relative to deviations from γ . By contrast, all that is actually assumed is that for each i and j , there is some "distinguishing" α that allows i 's and j 's deviations to be distinguished, and not necessarily the same α for each pair of players! FLM demonstrates that a profile as close as desired to γ can be found that puts a little weight on the strategies used in the "distinguishing profiles" and discriminates as required between deviations of different players. They attribute the kernel of this idea to Legros¹¹ (1989).

⁹There are k possible signal values, and m_i pure strategies for each player i .

¹⁰This is true in the context of the present model, in which the range of the public signal is finite (assuming in addition that all realizations occur with positive probability in equilibrium). In games with richer signal spaces, it is sometimes possible as δ approaches 1 to construct a sequence of symmetric punishments that are asymptotically efficient, based on punishment regions whose likelihood ratios are exploding. Recall the famous example of Mirrlees (1976) in which the static agency problem is overcome in the same way.

¹¹Legros was concerned with a static incentive problem. Recently Legros and Matsushima (1990) have given a nice sufficient condition for the existence of efficient solutions of static partnership problems.

With the additional restrictions on the information structure guaranteed by the full rank conditions, FLM prove a folk theorem of virtually the same degree of generality as for perfect monitoring. There is no restriction to pure strategies.

PROPOSITION. Perfect Folk Theorem with Imperfect Monitoring (Fudenberg, Levine and Maskin, 1989). *For a finite game G satisfying individual full rank and pairwise full rank (see above) and $\dim F^+ = n$, for any closed set W in the relative interior of F^+ there exists $\underline{\delta}$ such that for all $\delta > \underline{\delta}$, $W \subseteq V(\delta)$.*

Fudenberg and Levine (1989) show how the folk theorem must be weakened if some of the participants in a supgame are "short run" players; they provide an upper bound for the payoffs attainable by the long run players, as a function of the information structure.

Agency and Repeated Contracting

The classic moral hazard problem with one principal and one agent is an important example of a game with imperfect monitoring on one side. Many principal-agent relationships are of an ongoing nature, and much effort has been devoted to understanding the implications of repetition for the shape and performance of optimal contracts. Some of this research was underway at the time of the Fifth World Congress of the Econometric Society, and was included in the authoritative survey by Hart and Holmström (1987). For a taste of what has been done since, see the relevant section of FLM, Malcomson and MacLeod (1988, 1989), Pearce and Stacchetti (1987), Phelan and Townsend (1989) and Rey and Salanie (1990). While not strictly repeated, some games of international debt repayment are closely related; see especially Atkeson (1988).

One of the benefits of the recent overlap of contract theory and repeated games has been a growing understanding of the relationship between what can be accomplished by implicit and explicit (legal) enforcement mechanisms, respectively. Naturally, implicit contracts are advantageous when the concerned parties share information that cannot (for legal reasons, or because verification costs are prohibitive, etc.) be used in an explicit contract. But when the contracts can be specified in terms of the same information, do self-enforcing agreements achieve what explicit contracts can? The answer depends on whether one can create the same variation in continuation payoffs in self-enforcing agreements as in explicit contracts (and hence provide incentives with the same degree of efficiency). This is not always possible: the equilibrium value set (in the implicit contract environment) might be of less than full dimension, or might be too small to allow efficient exploitation of the game's signal space (recall the earlier discussion of likelihood ratios). Suppose, however, that the conditions of the FLM folk theorem of this section are met. Then as δ approaches 1, self-enforcing agreements achieve almost any feasible, individually rational payoff, so that asymptotically, implicit contracts perform as well as their

explicit counterparts. Note that the preceding discussion has abstracted from transactions costs, which may differ for implicit and explicit contracts.

Confronting the Theoretical Predictions with Reality

This will be a brief subsection, because reality and I have been out of touch for a long time. A number of investigators have developed econometric tests of the Green-Porter model and applied them to data on the Joint Executive Committee railroad cartel (weekly aggregate time series for the period 1880 to 1886). Porter (1983b), Lee and Porter (1984) and Hajivassiliou (1989), respectively, use switching regression models of increasing sophistication to allow for collusion punctuated by price wars. Berry and Briggs (1988) use the same data to test the hypothesis that the alternation between regimes follows a Markov process, and Hajivassiliou (1989) compares the performance of the Abreu, Pearce and Stacchetti (1986) and Rotemberg and Saloner (1986) analyses of oligopolistic collusion. Slade (1986, 1987) tests a learning model (Slade (1989a)) using a daily time series on gasoline prices in Vancouver that she collected.

I value this body of work principally for its implementation of econometric methods appropriate to the study of collusive markets and for its organization of some facts about intertemporal strategic behavior in a few oligopolies. It seems entirely premature, however, to draw conclusions about the relevance of the particular models tested (and accordingly I do not summarize the results of the various tests here).¹² First, it is highly probable that none of the models comes close to capturing the strategic considerations at work in the oligopolies in question. The environments were far more complicated, in important ways, than any of the models tested, and I think Slade has the right attitude when she describes the process of relating supergame models to data as follows: "The object of the exercise is not to pick a winner. Instead, the role of industry characteristics in determining pricing dynamics is assessed, and the reasons why simple models may fail to explain complex pricing patterns are examined" (Slade, 1989b). A second reservation is that the collusive theories tested are quite naive from a conceptual point of view, ignoring renegotiation, coalition formation, and other considerations of equilibrium refinement. In my opinion, the pure theory of implicit collusion is at such a primitive stage that it is in no shape to be tested.

Still, I feel that there is a lot to be learned from studying collusion in specific industries while keeping in mind an assortment of questions provoked by modern theory. An exciting example of what can result is Levenstein's (1989) work on the bromine industry in the U.S.A. and Germany from 1880 until 1914. By analyzing the

¹²Useful summaries and discussion may be found in the surveys by Jacquemin and Slade (1989) and Bresnahan (1989).

internal documents of the Dow Chemical Company and its correspondence with other American and German producers, Levenstein gives us an extraordinary picture of the evolution of competition and collusion among the oligopolists as they gained experience, learned about their rivals, and faced changing market conditions. Other fascinating examples of self-enforcing contracts in the economic history literature include Greif (1989) on reputation among medieval Mediterranean traders, and Milgrom, North and Weingast (1990) on the role of the Law Merchant and the Champagne fairs in Europe in the Middle Ages.

4. RENEGOTIATION AND SELF-ENFORCING AGREEMENTS

If an agreement among players in a repeated game is truly self-enforcing, it must be able to withstand the possibility that the players could renegotiate the terms of the agreement after any history. This section principally concerns renegotiation involving all parties to the agreement, although the potential for defections by smaller coalitions is an important and difficult problem as well. An explosion of research in the last six or seven years has produced a baffling variety of criteria for "renegotiation-proofness." Rather than exhaust the space available here by reproducing the details of the many definitions, I will try to provide a conceptual overview of the literature, emphasizing the concerns that prompted the authors to formulate the new solution concepts. There will be no attempt to describe the technical characterizations of the solution sets; under moderate assumptions existence is not a problem, except where mentioned. Nonspecialists will find the discussion more meaningful if they first (or concurrently) refer to some of the original papers in the literature.

Although all the work on renegotiation is skeptical about the credibility of the kinds of equilibria described in Sections 2 and 3, there is an even more radical critique that deserves mention. In the spirit of Harsanyi and Selten (1988) one could say that the behavior of "ideally rational" players in a subgame depends only on the internal *structure* of the subgame and not on how it was reached. Since all subgames of an infinite horizon, strictly repeated game are identical, it follows that there is no scope for negotiation of any kind: the same (noncooperative) outcome will occur in each period. Güth, Leininger and Stephan (1988) provide a formal argument based on a generalization of subgame consistency (Selten, 1973). It is not clear why full rationality necessarily precludes the formation of agreements to vary behavior across isomorphic subgames if the result is beneficial to all concerned. But I have some sympathy with the radical critique insofar as I think (and argue later) that most other authors err in the opposite direction by overestimating the influence of verbal agreements on the course of play.

The first work on renegotiation in infinitely repeated games¹³ was done independently by Bernheim and Ray (1989) and Farrell and Maskin (1989). Their position is that the stationary structure of $G^\infty(\delta)$ implies that the set of credible (renegotiation-proof) equilibria is the same in every subgame. Moreover, they assume that after any history of play, an ongoing agreement would be renegotiated (abandoned) if and only if a Pareto superior credible equilibrium were available. That is, players will stick with the status quo unless everyone can credibly be made better off. In the terminology of Farrell and Maskin, a subgame perfect equilibrium is called *weakly renegotiation-proof* (WRP) if no two of its continuation values¹⁴ (after *any* history) are Pareto ranked. This can be translated into a criterion for sets of supgame (average discounted) values, to facilitate comparison with other solution concepts. A set W of values is WRP if it is self-generating (this imposes the discipline of subgame perfection) and if no two values in W are Pareto ranked. One might say that such a solution set is "Pareto-thin."

Unfortunately, a WRP set W may contain a point w that is Pareto-dominated by a point x in some other WRP set X . Why is the value w credible if players can propose the universally preferred continuation value x , which is itself credible according to the WRP criterion? A WRP set none of whose values is Pareto-dominated by an element of any other WRP set is called *strongly renegotiation-proof* (SRP). Such a set does not always exist: there may be no "greatest" WRP set, but rather two or more intersecting "maximal" WRP sets, such as W and W' in Figure 2. This possibility led to several interesting definitions intermediate between WRP and SRP, none of them completely satisfying. These include relative strongly renegotiation-proofness (Farrell and Maskin (1989)) and minimal consistency and simple consistency (Bernheim and Ray (1989)). For technical characterizations of WRP, SRP and their variants see the two papers just mentioned, and also van Damme (1989) and Evans and Maskin (1989).

¹³This work was being done around the same time Bernheim, Peleg and Whinston (1987) were developing their notion of coalition-proof equilibrium, whose extensive form expression can be specialized to a theory of renegotiation-proofness in finitely repeated games. An early note on renegotiation in infinitely repeated games was circulated by Farrell (1984). Cave (1987) deserves mention for studying the minimal punishments necessary to support a given degree of collusion in a dynamic fisheries model.

¹⁴For any profile σ , the set of continuation values of σ is $C(\sigma) = \{v(\sigma|_h) | h \text{ is some (possibly degenerate) history of play}\}$, where $\sigma|_h$ is the profile induced by σ on the subgame following h . Note that $v(\sigma)$ is included in $C(\sigma)$.




Figure 2

Bernheim and Ray have reservations about all of the foregoing definitions because in the following sense they require too little. If W is the solution set, it should not be possible to construct an equilibrium σ with value Pareto superior to some value in W that uses only continuation values from W after every nontrivial history. After all, how could one argue against renegotiating to σ (from an equilibrium with the Pareto dominated value), if in all future contingencies it specifies continuation values that are credible within the theory? Thus, one needs to require that the solution set W satisfy $W = \text{eff } B(W)$, where $\text{eff } X$ means "the (Pareto) efficiency frontier of X ." Such a set is called *internally renegotiation-proof* and is studied by Ray (1989). Existence is apparently problematic.

Pearce (1987) suggests a different approach to renegotiation, one that ties the influence of a negotiated agreement to the usefulness of such agreements in the future. For simplicity, think of a very special case, namely, an infinitely repeated game all of whose subgame perfect equilibrium average values are symmetric (if one player gets x in equilibrium, so does everyone else). Players at time t will not follow an equilibrium with continuation value 2, say, if the only reason for doing so is that the threat of triggering the value 2 was needed *earlier* to induce a certain pattern of cooperation. But the continuation equilibrium *will* stand if the value 2 is truly indispensable for the provision of incentives in the *future*, that is, if *every* subgame perfect equilibrium must use continuation values of 2 or less after some histories. In the special case under consideration, Pearce (1987)

calls an equilibrium σ *renegotiation-proof* if¹⁵ $\inf C(\sigma) \geq \inf C(\gamma)$ for every subgame perfect equilibrium γ . Implicitly some notion of precedent is being invoked: players should not believe that they can abandon a punishment of 2 now, but commit never to do so again in the future. I will omit the extension of the solution concept to general repeated games, because (like almost all of the literature) it uses the Pareto criterion in determining when a credible alternative will be adopted by the group of players. There is a need for more plausible criteria, a point discussed briefly below.

Greenberg (1990) presents the results of research over a number of years on the "theory of social situations," which expands on the "stable set" methodology of von Neumann and Morgenstern (1947) and applies it to a variety of strategic settings. This work has had an influence on the debate about renegotiation, particularly through the research of Asheim (1988, 1990), whose analysis of renegotiation relies explicitly on stable sets of equilibria. A solution in his sense is a set of equilibria for each subgame; the set associated with a given subgame is interpreted as those equilibria considered credible in that subgame. Asheim calls a solution "Pareto perfect" (following the terminology of Bernheim and Ray (1989)) if it is both internally and externally stable. Internal stability means that for any subgame g and element σ of the associated solution set, and any subsequent subgame h , no element of the solution set of h Pareto dominates the equilibrium induced on h by σ . (Thus, no credible equilibrium should be interrupted after some history by a Pareto superior equilibrium considered credible following that history.) External stability requires that for any equilibrium σ excluded from the solution set for subgame g , there must be some subsequent subgame h such that some element of the solution set of h Pareto dominates the continuation equilibrium induced by σ on h . An attractive feature of the theory, then, is that it explains why no further equilibria were included in the solution sets. Existence has not been established in general, and it seems unlikely that strong results are possible. Asheim (1990) shows that even in very simple examples, existence is incompatible with stationarity of the solution sets.

Instead of giving up stationarity, perhaps we can exploit it to arrive at an appropriate relaxation of external stability. If an equilibrium σ is not dominated after any history by an "included" equilibrium, σ might still be excluded on the grounds that it is internally inconsistent. Suppose that for some history h , the continuation equilibrium $\sigma|_h$ is dominated by σ itself. If players found σ credible, then after the history h they would unanimously agree to renegotiate away from $\sigma|_h$ to σ , which contradicts the hypothesis that such a σ could be found credible. One can find arguments for and against weakening external stability by allowing for exclusion on the basis of internal inconsistency. For those who would object that this is a departure from the original spirit of

¹⁵Here $C(\sigma)$ is treated as a set of scalars, since everyone's payoff is the same.

the von Neumann and Morgenstern definition, I should remark that they presumably were thinking principally of strict dominance relations that were irreflexive, so that the issue of self-dominating elements did not arise. It is only in the context of objects with a dynamic structure that the question becomes important.

Abreu and Pearce (in press) adopt the stable solution formulation with a number of changes. First, we weaken external stability in the way just described. Secondly, we argue that it is reasonable to look for stationary stable sets of *deviations*, rather than sets of *credible plans*. The interpretation is that since a deviation from an ongoing social agreement does not depend upon the old agreement for its legitimacy, it must stand on its own. Thus, if it is *credible* in one contingency, it is *credible* in any contingency. (This is not to say that it will be *adopted* independently of the context; the group may not *want* to adopt it if the status quo is sufficiently attractive.) Hence, it seems natural to impose stationarity on the set of deviations that are considered *credible*. Finally, we suggest that the Pareto criterion be replaced as dominance criterion by some ordering, preferably complete, reflecting considerations of bargaining power in the game.

The use of the Pareto criterion for determining when a *credible alternative* will be adopted imputes to each player veto power over changes from the old negotiated agreement to a suggested alternative. Why should a verbal agreement embody such commitment power? In my (minority) opinion, it is more plausible to posit that decisions to adopt *credible alternatives* are governed by some rule that takes into account the bargaining positions of the players, as determined by the structural features of the game. Abreu and Pearce (in press) have nothing constructive to say about precisely what the rule ought to be. The problem is a little easier in symmetric games, where Abreu, Pearce and Stacchetti (1989) modify the definition of renegotiation-proofness given in Pearce (1987) by replacing the Pareto criterion with a simple (many would say too simple) bargaining rule. DeMarzo (1990) raises another problem concerning veto power and the status quo. Suppose that the initial equilibrium σ is adhered to in the first period, and in the second players are pondering the credibility of an alternative σ' . If they adopt σ' , then in the third period can each player insist on the continuation of σ' (against a new proposal σ'' , say) or on the continuation of σ ? In other words, what should serve as the status quo? Departing from the earlier literature, DeMarzo says that in many circumstances σ is more appropriate, because the original equilibrium has the weight of tradition behind it (it is a "social norm").

Bergin and MacLeod (1989) present an axiomatic system within which they can generate a number of the alternative solution concepts in the renegotiation literature. Although in some cases (such as their axiomatization of my own definition) I find the particular decomposition of the definitions into principles and preferences unnatural or forced, in others it is very helpful, especially in understanding the relationships of formulations of renegotiation-proofness in finite and infinite horizon supergames. Several papers, most notably Benoit and

Krishna (1988) have explored the implications of "Pareto perfection" (Bernheim and Ray (1989)) in finitely repeated games.¹⁶ The definition is recursive. In the final period, the credible Nash equilibria are those that are Pareto efficient among all Nash equilibria of G . In the last two periods, the equilibria that are efficient among those whose continuations are credible, are deemed credible, and so on. In technical terms, if W^t is the set of credible average values in the final t -period subgame, then $W^{t+1} = \text{eff } B(W^t)$. So as Bergin and MacLeod point out, the WRP concept is really not analogous to the commonly accepted finite horizon definition. Actually, Bernheim and Ray's internally renegotiation-proof solution is a closer analogue to the latter. Bergin and MacLeod also propose a new solution called recursive efficiency; see their paper for a discussion of its relationship to DeMarzo's point of view.

I cannot end this catalogue without mentioning an intriguing paper by Matsushima (1990). His idea is that just as an equilibrium specifies what will happen if its "instructions" are not obeyed, societies have metacodes indicating what happens when a social convention (equilibrium) is breached, that is, when the group as a whole sets aside an equilibrium. The ensuing analysis is highly ingenious; to my astonishment, Matsushima emerges from a jungle of infinite sequences of social conventions and breaching rules, with an existence result. I will not try to explain the motivation for the solution concept; on that score, despite some enjoyable discussions with the author, I remain mystified.

How are we to choose among the multitude of theories of renegotiation? I don't think this can be done purely on logical grounds: each theory is consistent on its own terms and respects all the relevant intertemporal incentive constraints. Many strategic situations (as traditionally described) are fundamentally underdetermined, even if one imposes the questionable restriction of equilibrium. What people will believe after observing a particular history of play, and what weight they give to verbal agreements, are partly questions of psychology. Why not leave the problem to the psychologists, then? I find this abdication of responsibility unattractive: the psychological aspects of the puzzle are inextricably interwoven with complicated considerations of sequential rationality and bargaining power. Thus, while guidance certainly should be sought from other disciplines, the skills of game theorists and economists are highly relevant in the construction of *educated guesses* about cooperation in supergames.

¹⁶Abreu and Pearce (in press) suggest that the logic of renegotiation is rather different in finitely, as opposed to infinitely, repeated games. If verbal agreements are influential because of their *prospective* usefulness, what weight can they carry in the final period of $G^T(\delta)$, where there is no future for the players to consider?

To guess is unavoidable¹⁷ if we are to make any contribution to many of the most important areas of the social sciences (beyond asserting that nothing can be said with much certainty, sometimes a useful remark in itself). The danger is that the guesses may be taken too seriously. "Equilibrium in dominant strategies" and "weakly renegotiation-proof equilibrium" are worlds apart in the immediacy of their links to basic principles of rationality, and correspondingly in the degree of confidence they ought to command. Yet one term, "solution concept," is used to describe them both (along with scores of other notions). We need to develop a means of communicating the intended interpretations of our various solution concepts and of distinguishing what is relatively solid in our analyses from what is of a more speculative nature. Such a language is needed not only for studying repeated games, but for game theory in general. Its absence is a stumbling block for the useful fusion of noncooperative and cooperative theories of strategic behavior.

5. CONCLUSION

Study of the equilibria of repeated games has been intense over the past decade. The results have been rewarding. While progress occurred largely in models chosen for their tractability, a picture has emerged that seems likely to be broadly representative of more general environments. Among other things, we have some sense of what optimal self-enforcing agreements look like, when they are likely to approach the first best, which theoretical simplifications are fairly innocuous, how implicit and explicit contracts compare, and which techniques extend to dynamic and stochastic games. But beneath our understanding of the mechanics of supergame equilibria lie foundational issues of the most immediate relevance.

The multiplicity of equilibria that causes problems in many areas of game theory arises in a dramatic way in repeated games: without multiplicity, self-enforcing cooperative agreements would be impossible. The accompanying conceptual puzzles were long ignored, but recent years have seen an explosion of research on supergame solution concepts, with particular attention devoted to the renegotiation of implicit agreements. After all the activity, there remain more questions than conclusions. Under what conditions are players likely to expect others to behave non-myopically? What do they think when someone departs from an agreed-upon plan? Which negotiation statements are credible? What is an effective precedent? Issues like these float tantalizingly in a multidisciplinary limbo, beyond the reach of purely mathematical analysis. They are too central to ignore, yet

¹⁷In his scintillating essay on the foundations of game theory, Binmore (1992) urges that predictions about strategic behavior be informed by careful study of "librations" (equilibrating processes in real time) and of the actual thought process of (boundedly rational) humans. This may well prove productive, but is unlikely to remove the need for guesswork in the foreseeable future. There will remain ample room for debate about what rules of thumb, models of the mind, updating processes, and so on, are reasonable or plausible.

too nebulous to have definitive answers. This is at once the most frustrating and the most alluring aspect of the subject.

REFERENCES

- Abreu, D., 1986. "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory* 39: 191-225.
- Abreu, D., 1988. "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica* 56: 383-396.
- Abreu, D., P. Milgrom and D. Pearce, In press. "Information and Timing in Repeated Partnerships," *Econometrica*.
- Abreu, D. and D. Pearce, In press. "A Perspective on Renegotiation in Repeated Games." In R. Selten (ed.), *Game Equilibrium Models*. Springer-Verlag.
- Abreu, D., D. Pearce and E. Stacchetti, 1986. "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory* 39: 251-269.
- Abreu, D., D. Pearce and E. Stacchetti, 1989. "Renegotiation and Symmetry in Repeated Games," Cowles Foundation Discussion Paper No. 920, Yale University.
- Abreu, D., D. Pearce and E. Stacchetti, 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58: 1041-1064.
- Asheim, G., 1988. "Renegotiation-Proofness in Finite and Infinite Stage Games through the Theory of Social Situations," Mimeo, The Norwegian School of Economics and Business Administration.
- Asheim, G., In press. "Extending Renegotiation-Proofness to Infinite Horizon Games," *Games and Economic Behavior*.
- Atkeson, A., 1988. "International Lending with Moral Hazard and Risk of Repudiation," Mimeo, Stanford University.
- Aumann, R. and L. Shapley, 1976. "Long-term Competition: A Game-theoretic Analysis," Mimeo.
- Barro, R. and D. Gordon, 1983. "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics* 12: 101-121.
- Benoit, J. P. and V. Krishna, 1985. "Finitely Repeated Games," *Econometrica* 53: 905-922.
- Benoit, J. P. and V. Krishna, 1988. "Renegotiation in Finitely Repeated Games," mimeo, Harvard Business School.
- Bergin, J. and W. B. MacLeod, 1989. "Efficiency and Renegotiation in Repeated Games," Mimeo, Queen's University at Kingston.
- Bernheim, B. D., B. Peleg and M. Whinston, 1987. "Coalition-Proof Nash Equilibrium 1. Concepts," *Journal of Economic Theory* 42: 1-12.
- Bernheim, B. D. and D. Ray, 1989. "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior* 1: 295-326.
- Berry, S. and H. Briggs, 1988. "A Non-Parametric Test of a First Order Markov Process for Regimes in a Non-Cooperatively Collusive Industry," *Economics Letters* 27: 73-77.

- Binmore, K., 1992. "Foundations of Game Theory." In J. J. Laffont (ed.), *Advances in Economic Theory: Sixth World Congress*. Cambridge: Cambridge University Press.
- Bresnahan, T., 1989. "Empirical Studies of Industries with Market Power." In R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organization*, Vol. 2, Chap. 17. Amsterdam: North Holland.
- Cave, J., 1987. "Long-Term Competition in a Dynamic Game: The Cold Fish War," *Rand Journal of Economics* 18: 596-610.
- Cooper, R. and A. Daughety, 1988. "Cooperation in Infinite Horizon Games with Finitely Lived Players," Mimeo, University of Iowa.
- Chou, C. and J. Geanakoplos, 1987. "On Finitely Repeated Games and Pseudo-Nash Equilibria," Mimeo, Yale University.
- Cr mer, J., 1986. "Cooperation in Ongoing Organizations," *Quarterly Journal of Economics* 101: 33-49.
- Cronshaw, M. and D. Luenberger, 1990. "Subgame Perfect Equilibria in Infinitely Repeated Games with Perfect Monitoring and Discounting," Mimeo, Stanford University.
- Damme, E. van, 1989. "Renegotiation-Proof Equilibria in Repeated Prisoners' Dilemma," *Journal of Economic Theory* 47: 206-217.
- DeMarzo, P., 1989. "Coalitions, Leadership and Social Norms: The Power of Suggestion in Games," Dept. of Finance Working Paper No. 85, Kellogg School, Northwestern University.
- Demougin, D. and A. Fishman, 1988. "Efficient Budget Balancing Cartel Equilibria with Imperfect Monitoring," Mimeo, University of Toronto.
- Evans, G. and E. Maskin, 1989. "Efficient Renegotiation-Proof Equilibria in Repeated Games," *Games and Economic Behavior* 1: 361-369.
- Farrell, J., 1984. "Credible Repeated Game Equilibria," Mimeo.
- Farrell, J. and E. Maskin, 1989. "Renegotiation in Repeated Games," *Games and Economic Behavior* 1: 327-360.
- Forges, F., J. F. Mertens and A. Neyman, 1986. "A Counterexample to the Folk Theorem with Discounting," *Economics Letters* 20: 7, 8.
- Friedman, J., 1971. "A Noncooperative Equilibrium for Supergames," *Review of Economic Studies* 38: 1-12.
- Friedman, J., 1985. "Cooperative Equilibria in Finite Horizon Noncooperative Supergames," *Journal of Economic Theory* 35: 390-398.
- Fudenberg, D., 1992. "Explaining Cooperation and Commitment in Repeated Games." In J. J. Laffont (ed.), *Advances in Economic Theory: Sixth World Congress*. Cambridge: Cambridge University Press.
- Fudenberg, D. and D. Levine, 1983. "Subgame Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory* 31: 251-263.
- Fudenberg, D. and D. Levine, 1986. "Limit Games and Limit Equilibria," *Journal of Economic Theory* 38: 261-279.

- Fudenberg, D. and D. Levine, 1989. "Equilibrium Payoffs with Long-Run and Short-Run Players," Mimeo, M.I.T. and U.C.L.A.
- Fudenberg, D. and D. Levine, In press. "An Approximate Folk Theorem with Imperfect Private Information," *Journal of Economic Theory*.
- Fudenberg, D., D. Levine and E. Maskin, 1989. "The Folk Theorem with Imperfect Public Information," Mimeo, M.I.T. and U.C.L.A.
- Fudenberg, D. and E. Maskin, 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* 54: 533-556.
- Fudenberg, D. and J. Tirole, 1990. "A Course in Game Theory," Manuscript, M.I.T.
- Green, E. and R. Porter, 1984. "Noncooperative Collusion under Imperfect Price Information," *Econometrica* 52: 87-100.
- Greenberg, J., 1990. *The Theory of Social Situations: An Alternative Game-Theoretic Approach*. Cambridge: Cambridge University Press.
- Greif, A., 1989. "Reputation and Coalitions in Medieval Trade: Maghribi Traders," *Journal of Economic History* 59: 857-882.
- Greif, A., P. Milgrom and B. Weingast, 1990. "The Merchant Guild as a Nexus of Contracts," Mimeo, Stanford University.
- Güth, W., W. Leininger and G. Stephan, 1988. "On Supergames and Folk Theorems: A Conceptual Discussion," ZiF Working Paper, University of Bielefeld.
- Hajivassiliou, V., 1989. "Testing Game-Theoretic Models of Price-Fixing Behavior," Cowles Foundation Discussion Paper No. 935, Yale University.
- Harris, C., 1985. "Perfect Equilibrium in Infinite Horizon Games," *Journal of Economic Theory* 37: 99-125.
- Harsanyi, J. and R. Selten, 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, Mass.: MIT Press.
- Hart, O., and B. Holmström, 1987. "The Theory of Contracts." Chapter 3 in T. Bewley (ed.), *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press.
- Howard, R., 1960. *Dynamic Programming and Markov Processes*. New York: M.I.T. and John Wiley and Sons.
- Kandori, M., 1988. "Monotonicity of Equilibrium Payoff Sets with Respect to Observability in Repeated Games with Imperfect Monitoring," Mimeo, Stanford University.
- Kandori, M., 1989a. "Repeated Games Played by Overlapping Generations of Players," Mimeo, University of Pennsylvania.
- Kandori, M., 1989b. "Social Norms and Community Enforcement," CARESS Working Paper No. 89-14, University of Pennsylvania.
- Jacquemin, A., and M. Slade, 1989. "Cartels, Collusion and Horizontal Merger." Chapter 7 in R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organization*, Vol. 1. Amsterdam: North Holland.

- Kreps, D., 1987. "Corporate Culture and Economic Theory," Mimeo, 1987.
- Kreps, D. and R. Wilson, 1982. "Reputation and Imperfect Information," *Journal of Economic Theory* 27: 253-279.
- Krishna, V., 1987. "The Folk Theorems for Repeated Games," Mimeo, Harvard Business School.
- Lambson, V., 1987. "Optimal Penal Codes in Price-Setting Supergames with Capacity Constraints," *Review of Economic Studies* 54: 385-398.
- Lee, L.-F. and R. Porter, 1984. "Switching Regression Models with Imperfect Sample Separation Information - with an Application on Cartel Stability," *Econometrica* 52: 391-418.
- Legros, P., 1989. *Efficiency and Stability in Partnerships*, Ph.D. dissertation, California Institute of Technology.
- Legros, P. and H. Matsushima, 1990. "Efficiency, Limited Liability and Neutrality in Partnerships," Mimeo, Cornell University.
- Levenstein, M., 1989. "The Feasibility and Stability of Collusion: A Study of the Pre-World War I Bromine Industry," Mimeo, Yale University.
- MacLeod, B. and J. Malcomson, 1988. "Reputation and Hierarchy in Dynamic Models of Employment," *Journal of Political Economy* 96: 832-854.
- MacLeod, B. and J. Malcomson, 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica* 57: 447-480.
- Matsushima, H., 1989. "Efficiency in Repeated Games with Imperfect Monitoring," *Journal of Economic Theory* 48: 428-442.
- Matsushima, H., 1990. "Structure of Renegotiation in Infinitely Repeated Games," Mimeo, Stanford University.
- Mertens, J. F., 1987. "Repeated Games," *Proceedings of the International Congress of Mathematicians* 86: 1528-1577, American Mathematical Society.
- Milgrom, P., D. North and B. Weingast, 1990. "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs," *Economics and Politics* 2: 1-24.
- Milgrom, P. and J. Roberts, 1982. "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory* 27: 280-312.
- Mirrlees, J., 1974. "Notes on Welfare Economics, Information, and Uncertainty." In Balch, McFadden and Wu (eds.), *Essays on Economic Behavior under Uncertainty*. Amsterdam: North Holland.
- Okuno-Fujiwara, M. and A. Postlewaite, 1989. "Social Norms and Random Matching Games," Mimeo, University of Pennsylvania.
- Pearce, D., 1987. "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," Cowles Foundation Discussion Paper No. 855, Yale University.
- Pearce, D. and E. Stacchetti, 1988. "The Interaction of Implicit and Explicit Contracts in Repeated Agency," Cowles Foundation Discussion Paper No. 892, Yale University.

- Phelan, C. and R. Townsend, 1989. "Computing Multiperiod, Information-Constrained Optima," Mimeo, University of Chicago.
- Porter, R., 1983a. "Optimal Cartel Trigger Price Strategies," *Journal of Economic Theory* 29: 313-338.
- Porter, R., 1983b. "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics* 14: 301-314.
- Radner, R., 1980. "Collusive Behavior in Noncooperative Epsilon-Equilibria in Oligopolies with Long But Finite Lives," *Journal of Economic Theory* 22: 136-154.
- Radner, R., 1981. "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica* 49: 1127-1148.
- Radner, R., 1985a. "Repeated Partnership Games with Imperfect Monitoring and No Discounting," *Review of Economic Studies* 53: 43-58.
- Radner, R., 1985b. "Repeated Principal-Agent Games with Discounting," *Econometrica* 53: 1173-1198.
- 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.
- Ray, D., 1989. "Internally Renegotiation-Proof Equilibrium Sets: Limit Behavior with Low Discounting," Mimeo, Indian Statistical Institute, New Delhi.
- Rey, P. and B. Salanie, 1990. "Long-Term, Short-Term and Renegotiation: On the Value of Commitment in Contracting," *Econometrica* 58: 597-620.
- Rogoff, K., 1989. "Reputation, Coordination and Monetary Policy." In R. Barro (ed.), *Modern Business Cycle Theory*. Cambridge, Mass.: Harvard University Press.
- Rotemberg, J. and G. Saloner, 1986. "A Supergame-Theoretic Model with Price Wars During Booms," *American Economic Review* 76: 390-407.
- Rotemberg, J. and G. Saloner, 1989. "Tariffs vs Quotas with Implicit Collusion," *Canadian Journal of Economics* 22: 237-244.
- Rubinstein, A., 1977. "Equilibrium in Supergames," Research Memorandum 25, Center for Research in Mathematical Economics and Game Theory, Hebrew University, Jerusalem.
- Rubinstein, A., 1979a. "Equilibrium in Supergames with the Overtaking Criterion," *Journal of Economic Theory* 21: 1-9.
- Rubinstein, A., 1979b. "Offenses That May Have Been Committed by Accident -- an Optimal Policy of Retribution." In S. Brams, A. Schotter and G. Schwödiauer (eds.), *Applied Game Theory*. Würzburg: Physica-Verlag.
- Rubinstein, A. and M. Yaari, 1983. "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory* 30: 74-97.
- Sabourian, H., 1989. "Repeated Games: A Survey." In F. Hahn (ed.), *The Economics of Missing Markets, Information, and Games*. Oxford: Clarendon Press.

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale University
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 983

NOTE: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than acknowledgment that a writer had access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

REPEATED GAMES: COOPERATION AND RATIONALITY

by

David G. Pearce

June 1991

- Salant, D., 1988. "A Repeated Game with Finitely Lived Overlapping Generations of Players," Mimeo, GTE Labs.
- Selten, R., 1965. "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die Gesamte Staatswissenschaft* 121: 301-324.
- Selten, R., 1973. "A Simple Model of Imperfect Competition Where 4 Are Few and 6 Are Many," *International Journal of Game Theory* 3: 141-201.
- Selten, R., 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory* 4: 22-55.
- Slade, M., 1986. "Conjectures, Firm Characteristics, and Market Structure: An Empirical Analysis," *International Journal of Industrial Organization* 4: 347-370.
- Slade, M., 1987. "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion," *Journal of Industrial Economics* 35: 499-516.
- Slade, M., 1989a. "Price Wars in Price-Setting Supergames," *Economica* 56: 295-310.
- Slade, M., 1989b. "Strategic Pricing Models and Interpretation of Price-War Data," Mimeo, University of British Columbia.
- Smith, L., 1990. "Folk Theorems in Overlapping Generations Games," Mimeo, University of Chicago.
- Sorin, S., 1986. "On Repeated Games with Complete Information," *Mathematics of Operations Research* 11: 147-160.
- Syropoulos, C., 1989. "Cooperative Outcomes in Non-Cooperative Supergames of Trade Policy and the (Non) Equivalence of Tariffs and Quotas," Working Paper No. 1-89-2, Department of Economics, Pennsylvania State University.
- von Neumann, J. and O. Morgenstern, 1947. *Theory of Games and Economic Behavior*, 2nd edition. Princeton: Princeton University Press.
- Weinberger, C., 1990. "Reputation and Strikes in an Ongoing Bargaining Relationship," Mimeo, U. C. Berkeley.
- Williams, S. and R. Radner, 1987. "Efficiency in Partnership When the Joint Output Is Uncertain," Mimeo, Northwestern University.