

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 855

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 References in publications to Discussion Papers (other than acknowledgment that a writer had access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

RENEGOTIATION-PROOF EQUILIBRIA:

COLLECTIVE RATIONALITY AND INTERTEMPORAL COOPERATION

by

David G. Pearce*

December 1987

1. Introduction

Foremost among the reasons for studying infinitely repeated games is that they provide a setting in which individuals can cooperate with one another, without recourse to binding contracts. A player may be willing to forego the immediate gains realizable from cheating on an implicit agreement, in order to enjoy continued cooperation from his fellow players in return. In Friedman (1971), for example, each individual is persuaded to take an action beneficial to the group, by the threat that if he fails to do so, the others will revert to playing a (relatively unattractive) Cournot-Nash equilibrium in each subsequent period. The use of perfect equilibrium threats more severe than "Cournot-Nash reversion" was introduced by Aumann and Shapley (1976) and Rubinstein (1979) in their celebrated work on the folk theorem, and was developed in the more realistic context of discounted repeated games by Abreu (1983) as a means of achieving maximal collusion. Green and Porter (1983) and Porter (1984) showed that self-enforcing collusion may often be supported even in models in which individuals' actions are unobservable, provided those actions affect the distribution of some publicly observed random variable: players could "punish" one another after an inauspicious realization of the public signal, by following an unattractive equilibrium in the remaining game.

With few exceptions, the literature on repeated games ignores a critical problem concerning the intertemporal consistency of equilibria that employ threats to support cooperation. Following the observation that

* This is a preliminary draft; comments will be most welcome. I have benefitted greatly from the advice and assistance of Dilip Abreu and Ennio Stacchetti while this work progressed. I also enjoyed helpful conversations with Truman Bewley.

someone has cheated, or an "adverse" realization of the random variable determining the equilibrium path, players have an incentive to renegotiate their original implicit agreement, abandoning the prescribed punishment path in favor of a more attractive equilibrium. Because players can anticipate this incentive to renegotiate, it undermines the credibility of the original threat. The first formalization of this idea is due to Farrell (1983), who proposed the restriction that if two equilibria σ_1 and σ_2 are plausible according to a given theory, neither equilibrium should Pareto-dominate the other; in addition, all such plausible equilibria should, after any history of play, induce "continuation equilibria" also considered plausible by that theory. Farrell and Maskin (1986) have developed this approach, giving characterization results for the definition in two-person games with perfect monitoring. Other recent work on renegotiation in supergames includes Cave (1985) and van Damme (1986); readers should also be aware of work by Greenberg (1986) in a slightly different spirit.

In this paper I propose a definition quite distinct from those in the existing literature on renegotiation, although it addresses the same concerns. The following informal discussion is intended to motivate the solution concept to be studied here. Consider a group of players in a repeated partnership game with imperfect monitoring (see, for example, Radner, Myerson and Maskin (1986), Fudenberg and Maskin (1986), and Abreu, Pearce and Stacchetti (1986a)) who are playing an optimal symmetric sequential

equilibrium in pure strategies¹ (hereafter, S.S.E.). Suppose that the best and worst S.S.E.'s, respectively, are $\bar{\sigma}$ and $\underline{\sigma}$, with corresponding average present discounted values \bar{V} and \underline{V} . Assume for simplicity that there is a unique symmetric one-shot Cournot-Nash equilibrium with value c , where $\underline{V} < c < \bar{V}$. Typically the best S.S.E. will require that after some equilibrium histories, players follow a continuation equilibrium (say $\underline{\sigma}$) with value \underline{V} (see Abreu, Pearce and Stacchetti (1986a and 1986b)). Suppose that such a history has transpired, and that player 1, for example, proposes that $\underline{\sigma}$ be abandoned in favor of $\bar{\sigma}$ (why suffer \underline{V} if \bar{V} is available?). The suggestion is likely to be greeted with skepticism. Players have the incentive to make the unobservable choices dictated by $\bar{\sigma}$ only if they believe that after certain histories, $\underline{\sigma}$ (or some other S.S.E. yielding \underline{V}) will be followed; this entails their believing that it is not possible to renegotiate away from $\underline{\sigma}$. Consequently the proposal in question lacks internal consistency: in order to believe that it is possible to abandon $\underline{\sigma}$ in favor of $\bar{\sigma}$, it is necessary to believe that it is not possible to renegotiate away from $\underline{\sigma}$. (This presumes a certain stationarity in perceptions regarding the credibility of a threat of a given severity: I assume that players do not imagine at time t that, although they can escape a prescribed "punishment" with some value w , they will be unable to do so in the future).

¹For technical reasons I consider only pure strategies here. The restriction to symmetric equilibria is convenient for purposes of exposition, but is relaxed in Section 5.

Not easily discouraged, player 1 next suggests reversion to one-shot Cournot-Nash behavior for the duration of the game. The charge of inconsistency cannot be brought against this new proposal; participation in Cournot-Nash reversion does not require players to believe that the punishment V will be available in the future (nor any other punishment worse than Cournot-Nash reversion itself). While this argument convinces everyone in the group that nothing worse than Cournot-Nash reversion is credible, player 2 points out that renegotiation can be taken one step further. He directs attention to an S.S.E. γ with value $x > c$, and having the property that after every history, the continuation of γ has value at least x . Again, the ability to renegotiate to the value x does not conflict with a belief that γ will be adhered to: for no history does γ call for a punishment worse than x . Player 2 further remarks that no other S.S.E. has value $y > x$, with all continuation values at least as great as y . Consequently, renegotiation to more favorable equilibria than γ is not possible. If one accepts this line of reasoning, an S.S.E. in this game is renegotiation-proof if none of its continuation values, after an arbitrary history, is less than x (otherwise, it would be vulnerable, after some history, to renegotiation to γ).

The preceding discussion can be formalized in a variety of models. I concentrate here on discounted repeated games with imperfect monitoring. Section 2 specifies the class of games under consideration, and collects a number of results from Abreu, Pearce and Stacchetti (1986b) on which the analysis rests. While a general definition of renegotiation-proof

equilibrium in asymmetric games is considered in Section 5, I begin in Section 3 by studying symmetric equilibria. Apart from the expositional advantages offered by its simplicity, the symmetric theory has some attractions, including uniqueness, that justify a separate treatment. Section 4 is an aside on the formulation of the solution concept for games with perfect monitoring. As a simple illustration, a comparison is made of optimal punishments achievable in linear Cournot and Bertrand oligopolies. The emphasis throughout is on the general properties exhibited by the theory. Of particular concern are the effects of introducing the restriction that equilibria be renegotiation-proof, on the maximal collusion sustainable, the severity of optimal punishments, the structural features of equilibria, and the payoffs achievable in the limit as players are made extremely patient.

The motivating discussion given earlier assumes that after all histories, players are free to talk to one another. Whether this matters is an intriguing question: if it is clear that renegotiation would occur following some history and some group discussion, is there any need for the discussion to take place? If only the obvious is said, is it redundant? Van Damme (1986) answers the latter question in the affirmative, but others could be found holding the opposite view. In any case, if players can (with or without communicating) renegotiate implicit agreements, perhaps they can also contemplate coalitional deviations, use correlated strategies, and so on. I ignore these possibilities, keeping the analysis as traditional as possible while taking into account intertemporal consistency for the group of players. The use of public randomizing devices or correlated strategies

could be allowed for, but permitting deviations by arbitrary coalitions would lead to problems that are as yet not fully resolved even for static games.² Consequently the theory proposed here seems most satisfying when regarded as a solution concept for two-person games.

2. The Model and Some Preliminaries

With the exception of Section 4, the class of games considered is that analyzed in Abreu, Pearce and Stacchetti (hereafter, APS). In each period $t = 1, 2, \dots$, each player $i = 1, \dots, N$ selects an action $s_i(t)$ from the finite set S_i . The choice $s_i(t)$ is unobservable to $j \neq i$, but the realization of a random variable $p(t)$ is publicly observed at the end of period t . The signal $p(t)$ has density function $f(p; s_1(t), \dots, s_N(t))$ and support $\Omega \subseteq R^k$ independent of $s(t) = (s_1(t), \dots, s_N(t))$. Player i 's payoff $\Pi_i(s(t))$ in the component game in period t is the expectation of his realized payoff $\pi_i(s_i(t), p(t))$. Thus a player cares about the unobserved actions of others insofar as these determine the distribution of the payoff-relevant signal. I assume that the component game, labelled G , has at least one Nash equilibrium in pure strategies.

The choice made by i in period t can depend upon his own past actions, and on the realizations $p(1), \dots, p(t-1)$. Therefore a strategy

²Significant contributions to the literature on coalitional deviations have been made recently by Bernheim, Peleg and Whinston (1987) and Bernheim and Whinston (1987).

σ_i for i in the infinitely repeated game is a sequence of measurable functions $\{\sigma_i(t)\}_{t=1}^{\infty}$ where for $t \geq 2$, $\sigma_i(t) : [S_i \times \Omega]^{t-1} \rightarrow S_i$, and $\sigma_i(1) \in S_i$. The discount factor is $\delta \in (0,1)$. Period t payoffs are received at the end of period t , and are discounted to the beginning of period 1. Denote by $\psi(\sigma) \in \mathbb{R}^N$ the vector of present discounted expected values for the players when the supergame strategy profile σ is employed; the corresponding average (discounted) payoff vector is $v(\sigma) = [(1-\delta)/\delta]\psi(\sigma)$. (This is known as the average payoff because if it were received in every period, its (total) present discounted value would be $\psi(\sigma)$.) The infinitely repeated game with discount factor δ is denoted $G^{\infty}(\delta)$.

The definition of sequential equilibrium given by Kreps and Wilson (1982) is easily generalized to this setting, where there are an infinity of information sets. I abuse terminology in applying the term "sequential equilibrium" to a profile, rather than to an assessment. There is no loss of generality in considering only equilibria in which each strategy σ_i conditions choices on signal realizations only (see APS), that is, $\sigma_i(t) : \Omega^{t-1} \rightarrow S_i$, $t = 2, 3, \dots$. A sequential equilibrium σ in which for each i , σ_i is a pure strategy that conditions only on signal realizations, is denoted an S.E. Let $V = \{v(\sigma) \mid \sigma \text{ is an S.E.}\}$. When necessary I write $V(\delta)$ to make explicit the dependence of the equilibrium (average) value set on δ .

Some further definitions permit a summary of those results from APS that are needed in what follows. In the definition of an admissible pair

$(s, u(\cdot))$, one can think of the vector s as being the action profile played in the first period of some S.E., while $u(p)$ plays the role of the average continuation values for players in the equilibrium after the signal p arises in the first period. In APS, admissibility is expressed in terms of total present discounted values; here I work with averages, to facilitate comparison of value sets corresponding to different discount factors.

A pair $(s, u(\cdot)) \in S \times L^\infty(\Omega, \mathbb{R}^N)$ is admissible with respect to $W \subseteq \mathbb{R}^N$ if:

$$(i) \quad u(p) \in W \quad \forall p \in \Omega ,$$

and (ii) for all i and all $s'_i \in S_i$,

$$\Pi_i(s) + \frac{\delta}{1-\delta} \int_{\Omega} u(p) f(p; s) dp \geq \Pi_i(s_{-i}, s'_i) + \frac{\delta}{1-\delta} \int_{\Omega} u(p) f(p; s_{-i}, s'_i) dp .$$

The (average) value of an admissible pair is

$$E(s; u) = (1-\delta)\Pi(s) + \delta \int_{\Omega} u(p) f(p; s) dp .$$

A nonempty set $W \subseteq \mathbb{R}^N$ is self-generating if $W \subseteq B(W)$, where

$$B(W) = \{E(s; u) \mid (s, u) \text{ is admissible w.r.t. } W\} .$$

Note that $B(\cdot)$ is monotonic: $W_1 \subseteq W_2 \Rightarrow B(W_1) \subseteq B(W_2)$.

(R1)-(R6) below are proved in APS. They are helpful here because much of the analysis is conducted in terms of self-generating sets and equilibrium value sets. The second part of (R6) is not stated explicitly in APS, but is covered by a simple relabelling in Proposition 6 of that paper.

- (R1) (Factorization) $V = B(V)$.
- (R2) (Self-generation) $W \subseteq \mathbb{R}^N$ bounded and $W \subseteq B(W) \Rightarrow B(W) \subseteq V$.
- (R3) W self-generating \Rightarrow $\text{cl } W$ (the closure of W) is self-generating.
- (R4) V is compact.
- (R5) For any compact $W \subseteq \mathbb{R}^N$, $B(\text{ext } W) = B(W)$, where $\text{ext } X$ denotes the set of extreme points of $X \subseteq \mathbb{R}^N$.
- (R6) (Monotonicity) If $0 < \delta_1 < \delta_2 < 1$, then $V(\delta_1) \subseteq V(\delta_2)$. More generally, if W is self-generating for δ_1 , then W is self-generating for $\delta_2 > \delta_1$.

For any S.E. σ and any t -period history p , the "successor S.E." induced by σ on the subtree following that history is denoted $\sigma|_p$. If σ is an S.E., the set of continuation values of σ is $C(\sigma) = \{v(\sigma)\} \cup \{v(\sigma|_p) \mid p \text{ is some signal history}\}$. It will often be useful to work with the following simple characterization of a self-generating set.

Lemma 1. A nonempty bounded set W is self-generating if and only if for every $w \in W$, there exists an S.E. σ such that:

(a) $v(\sigma) = w$

and (b) $C(\sigma) \subseteq W$.

Proof: Suppose that for every $w \in W$, there exists an S.E. σ satisfying (a) and (b). Then $(\sigma(1), v(\sigma|_{p(1)}))$ is admissible with respect to W , and has value w . (Here, $p(1)$ is the realization of the signal in period 1.) Thus $W \subseteq B(W)$, and W is self-generating. Conversely, assume that W is self-generating. The proof of Proposition 2 in APS constructs, for each $w \in W$, an S.E. having value w and satisfying (a) and (b). Q.E.D.

Corollary: For each S.E. σ , $C(\sigma)$ is a self-generating set.

3. Symmetric Equilibria

This section studies symmetric equilibria of symmetric games. In this setting the definition of renegotiation-proofness that I suggest identifies uniquely the maximal severity of punishment that can be sustained in equilibrium, and the corresponding maximal attainable degree of "collusion." A degree of cooperation remains possible provided that Cournot-Nash reversion is capable of supporting some collusion, although Cournot-Nash reversion itself will typically not be renegotiation-proof. This is in strong contrast to the theory explored in Farrell and Maskin (1986), where restricting attention to symmetric equilibria would make all cooperation impossible.

An S.E. σ is symmetric if for every pair of players i and j , $\{\sigma_i(t)\}_{t=1}^{\infty} = \{\sigma_j(t)\}_{t=1}^{\infty}$. A symmetric S.E. is called an S.S.E. The definitions of Section 2 require slight modification for the symmetric analysis. Here $V \subseteq R$ will denote $\{v_1(\sigma) | \sigma \text{ is an S.S.E.}\}$, the set of symmetric equilibrium values. For any S.S.E. σ , $C(\sigma) = \{v_1(\sigma)\} \cup \{v_1(\sigma|_p) | p \text{ is some history of finite length}\}$. In the definition of an admissible pair

(s,u) , only symmetric vectors s and symmetric vectors $u(p)$ with components in $W \subseteq R$ are considered; $B(W) = \{E_1(s,u) \mid (s,u) \text{ is an admissible pair w.r.t. } W\}$. With this notational reinterpretation, the results presented in Section 2 all hold for the symmetric case. For any nonempty compact set $W \subseteq R$, \underline{W} and \bar{W} are the minimal and maximal elements of W , respectively.

I turn now to the formal definition of a renegotiation-proof equilibrium, which is followed immediately by a lemma giving two convenient restatements of the definition.

Definition: An S.S.E. σ is renegotiation-proof if for all $w \in C(\sigma)$, there is no S.S.E. γ satisfying:

$$(i) \quad v(\gamma) > w$$

and (ii) $x \geq v(\gamma)$ for all $x \in C(\gamma)$.

Lemma 2: Let σ be an S.S.E. The following three conditions are equivalent:

(i) σ is renegotiation-proof.

(ii) There exists a compact self-generating set W such that $C(\sigma) \subseteq W$ and for all compact self-generating sets Y ,

$$\underline{Y} \leq \underline{W} .$$

(iii) For each S.S.E. γ , $\inf C(\gamma) \leq \inf C(\sigma)$.

Proof: By the Corollary to Lemma 1, σ is renegotiation-proof if and only if $\exists W_1$ bounded and self-generating with $\inf W_1 > \inf C(\sigma)$. But by (R3), $\text{cl } W_1$ is also self-generating, so (i) and (ii) are equivalent. Further, (iii) clearly implies (i). Finally, if (iii) is violated, $\exists \gamma$ an S.S.E. with $\inf C(\gamma) > \inf C(\sigma)$. But $\text{cl } C(\gamma)$ is a compact self-generating set with minimum strictly exceeding $\inf C(\sigma)$, therefore (ii) cannot hold. Thus (ii) \Rightarrow (iii). Q.E.D.

Suppose that the continuation value of an S.S.E. σ after some history is $w \in \mathbb{R}$. If σ is renegotiation-proof, no one can claim that "punishments" as harsh as w are not needed, and can be abandoned by renegotiating to some $x > w$: every S.S.E. γ uses (after some histories) continuation values strictly worse than x . Hence, any assertion that renegotiation should allow x to be attained is internally inconsistent.

The following definition of a renegotiation-proof set is linked to the definition for S.S.E.'s by Lemma 2. Working directly with renegotiation-proof sets allows the most economical use of the results summarized in Section 2.

Definition: A compact self-generating set $W \subseteq \mathbb{R}$ is renegotiation-proof if

$$\underline{W} = \max\{\underline{X} \mid X \text{ is a compact self-generating set}\} .$$

Proposition 1: A renegotiation-proof set exists.

Proof: Let $M = \{X | X \text{ is a compact self-generating set}\}$. Note that $M \subseteq V$, so M is bounded. For each $x \in M$, let

$$W(x) = \cup\{Z | Z \text{ is a compact self-generating set with } \underline{Z} \geq x\} .$$

As a union of self-generating sets, $W(x)$ is self-generating, and by (R3), $\text{cl } W(x)$ is self-generating. Choose an increasing sequence $\{x_t\}$ converging to $\sup M$. The sequence $\{\text{cl } W(x_t)\}$ is a decreasing sequence of compact self-generating sets, therefore $W = \bigcap_t \text{cl } W(x_t)$ is a compact self-generating set (see Appendix) with $\min \text{cl } W = \sup M$, hence W is a renegotiation-proof set. Q.E.D.

Definition: $R = \cup\{W \subseteq R | W \text{ is renegotiation-proof}\} .$

Lemma 3 asserts that R is itself a renegotiation-proof set. Maximal credible cooperation, when the threat of renegotiation is accounted for, involves players following some S.S.E. σ with $v(\sigma) = \bar{R}$, and $C(\sigma) \subseteq R$. The most severe punishment that can be invoked is \underline{R} .

Lemma 3: R is a (compact) renegotiation-proof set.

Proof: $\text{cl } R$ is the closure of a union of self-generating sets, therefore $\text{cl } R$ is a compact self-generating set. Each of the sets over which the union is taken has minimum $\max M$ (as defined in the proof of Proposition 1), so $\min \text{cl } R = \max M$. Hence $\text{cl } R$ is renegotiation proof, $R = \text{cl } R$, and R is renegotiation-proof. Q.E.D.

The next proposition identifies a basic comparative statics result for self-generating sets that is applied repeatedly to establish the structural properties of renegotiation-proof sets. The result appears, considerably disguised, in the proof of Proposition 6 of APS, but an explicit proof is included here.

The following notation is used: for a function u and scalars λ and x , $\lambda u + (1-\lambda)x$ denotes the function g defined by $g(p) = \lambda u(p) + (1-\lambda)x$, for each p in the domain of u . The convex hull of a set W is written $\text{co } W$.

Proposition 2 (Radial Contraction): Suppose $0 < \delta_1 < \delta_2 < 1$, and let $W \subseteq R$ be a compact self-generating set for δ_1 . For each $x \in W$, the set $\{\lambda \underline{W} + (1-\lambda)x, \lambda \bar{W} + (1-\lambda)x\}$ is self-generating for δ_2 , where $\lambda = (1 - \delta_2)/(1 - \delta_1)$.

Proof: Because W is self-generating for δ_1 , (R5) implies the existence of pairs (\bar{s}, \bar{u}) and $(\underline{s}, \underline{u})$ admissible with respect to (\underline{W}, \bar{W}) , and having values \bar{W} and \underline{W} , respectively, for δ_1 . Let

$\gamma = (\delta_1/\delta_2) \cdot [(1 - \delta_2)/(1 - \delta_1)]$, and choose $x \in W$. I claim that the pairs $(\bar{s}, \gamma \bar{u} + (1-\gamma)x)$ and $(\underline{s}, \gamma \underline{u} + (1-\gamma)x)$ are admissible with respect to $\text{co}\{\lambda \underline{W} + (1-\lambda)x, \lambda \bar{W} + (1-\lambda)x\}$ for δ_2 . The reward functions (the second elements of each pair) have range

$$\{\gamma \underline{W} + (1-\gamma)x, \gamma \bar{W} + (1-\gamma)x\} \subseteq \text{co}\{\lambda \underline{W} + (1-\lambda)x, \lambda \bar{W} + (1-\lambda)x\}$$

(since $\gamma < \lambda$), and hence the first requirement of admissibility is

satisfied. Admissibility of (\bar{s}, \bar{u}) with respect to W for δ_1 yields, for each $s \in S_1$,

$$\begin{aligned} \Pi_1(s, \bar{s}, \dots, \bar{s}) - \Pi_1(\bar{s}, \dots, \bar{s}) &\leq \frac{\delta_1}{1 - \delta_1} \int_{\Omega} \bar{u}(p) \left[f(p; N\bar{s}) - f(p; s+(N-1)\bar{s}) \right] ds \\ &\quad - \frac{\delta_2}{1 - \delta_2} \int_{\Omega} \gamma \bar{u}(p) \left[f(p; N\bar{s}) - f(p; s+(N-1)\bar{s}) \right] dp, \end{aligned}$$

and adding the constant $(1-\gamma)x$ to the reward function has no effect on these incentive constraints. Thus the second condition for admissibility is also satisfied (the argument for $(\underline{s}, \underline{u})$ is identical). Straightforward algebra establishes that the values of $(\bar{s}, \gamma \bar{u} + (1-\gamma)x)$ and $(\underline{s}, \gamma \underline{u} + (1-\gamma)x)$ are $\lambda \bar{W} + (1-\lambda)x$ and $\lambda \underline{W} + (1-\lambda)x$, respectively. Thus, if B represents the function (defined in Section 2) associated with $G^\infty(\delta_2)$,

$$\begin{aligned} (\lambda \bar{W} + (1-\lambda)x, \lambda \underline{W} + (1-\lambda)x) &\subseteq B(\text{co}\{\lambda \bar{W} + (1-\lambda)x, \lambda \underline{W} + (1-\lambda)x\}) \\ &= B(\{\lambda \bar{W} + (1-\lambda)x, \lambda \underline{W} + (1-\lambda)x\}) \quad (\text{by (R5)}). \end{aligned}$$

Therefore $(\lambda \bar{W} + (1-\lambda)x, \lambda \underline{W} + (1-\lambda)x)$ is self-generating for δ_2 . Q.E.D.

Definition: Let s^* denote the best symmetric Nash equilibrium of the component game G .

In the traditional theory without renegotiation, punishments more severe than reversion to s^* are often sustainable. But $\Pi_1(s^*)$ is a lower bound on the value of any renegotiation-proof equilibrium: the

singleton set $\{\Pi_1(s^*)\}$ is self-generating, so $\underline{R} \geq \Pi_1(s^*)$. Also, in the traditional analysis higher discount factors allow more severe punishments: (R6) implies that $\underline{V}(\delta)$ is nonincreasing in δ . Considerations of renegotiation reverse this: $\underline{R}(\delta_1)$ is self-generating for δ_1 , and also for $\delta_2 > \delta_1$ (by (R6)). Thus $\underline{R}(\delta_2) \geq \underline{R}(\delta_1)$; the worst average value sustainable is a nondecreasing function of δ . Lemma 4 summarizes these results.

Lemma 4: (i) $\underline{R} \geq \Pi_1(s^*)$

(ii) $\underline{R}(\delta)$ is nondecreasing in δ .

A primary question is whether the theory proposed here is nontrivial: does it permit cooperation, and does it limit the degree of cooperation relative to the usual analysis without renegotiation? First, as long as the threat of Cournot-Nash reversion to s^* supports collusion, renegotiation-proof threats support some collusion also: both \underline{R} and \bar{R} usually exceed $\Pi_1(s^*)$ strictly.

Proposition 3: Suppose that the set $\Delta_a = \{\delta \in (0,1) \mid \exists W \text{ compact, self-generating s.t. } \bar{W} > \underline{W} \geq \Pi(s^*)\}$ is nonempty. Let $\delta_a = \inf \Delta_a$. Then for all $\delta > \delta_a$, $\underline{R} > \Pi_1(s^*)$.

Proof: Fix $\delta > \delta_a$, and choose δ_1 such that $\delta_a < \delta_1 < \delta$. $\underline{R}(\delta_1)$ is self-generating for δ_1 , therefore Proposition 2 implies that $\lambda \underline{R}(\delta_1) + (1-\lambda) \bar{R}(\delta_1)$ is self-generating for δ , where $\lambda = (1-\delta)/(1-\delta_1)$. Then the minimum of this set is a lower bound for $\underline{R}(\delta)$, and equals $\lambda \underline{R}(\delta_1) + (1-\lambda) \bar{R}(\delta_1) > \underline{R}(\delta_1) \geq \Pi(s^*)$, by Lemma 4. Q.E.D.

A similar argument establishes that whenever $\bar{V} > \underline{V}$, the threat of renegotiation reduces the severity of punishments available: $\underline{R} > \underline{V}$ (see Proposition 4 for a precise statement). The fact that \underline{R} exceeds \underline{V} usually implies that $\bar{R} < \bar{V}$. Under regularity conditions that are not restated here (the exact conditions are prohibitively lengthy), the continuation values after the first period of any S.S.E. with value \bar{V} must lie in the two-point set (\underline{V}, \bar{V}) (see APS, Proposition 7). But a renegotiation-proof equilibrium is constrained to using continuation values at least as great as $\underline{R} > \underline{V}$, so either $\bar{R} < \bar{V}$, or only \bar{V} is used as a continuation value. In the latter case, first-period behavior is not collusive: A Nash equilibrium of the component game is played, and $\bar{R} = \Pi_1(s^*)$.

Proposition 4: Suppose $\Delta_b = \{\delta \in (0,1) \mid \bar{V}(\delta) > \underline{V}(\delta)\}$ is nonempty. Let $\delta_b = \inf \Delta_b$. For all $\delta > \delta_b$, $\underline{R}(\delta) > \underline{V}(\delta)$.

Proof: Choose any $\delta_1 \in (\delta_b, \delta)$, and note that by (R6), $\underline{V}(\delta) \leq \underline{V}(\delta_1)$. By an argument identical to that of Proposition 3 (in this case, perform a radial contraction toward $\bar{V}(\delta_1)$), $\underline{R}(\delta_1) < \underline{R}(\delta)$. Thus

$$\underline{V}(\delta) \leq \underline{V}(\delta_1) \leq \underline{R}(\delta_1) < \underline{R}(\delta) . \quad \text{Q.E.D.}$$

As δ increases, two opposing effects operate on $\bar{R}(\delta)$, the maximal renegotiation-proof payoff. First, the great weight placed by players on the future facilitates cooperation: present gains from contemplated deviations are given less relative importance. On the other hand, $\underline{R}(\delta)$ increases, eroding the "deterrence power" available. Interestingly, the

former effect dominates dramatically as players become arbitrarily patient. Asymptotically, taking renegotiation into account does not reduce maximally collusive payoffs: both $\bar{R}(\delta)$ and $\underline{R}(\delta)$ have the same limit as $\bar{V}(\delta)$.

Proposition 5: $\lim_{\delta \rightarrow 1} \underline{R}(\delta) = \lim_{\delta \rightarrow 1} \bar{V}(\delta)$.

Proof: For any $\varepsilon > 0$, we must find $\delta^* > 0$ such that for all $\delta_1 > \delta^*$, $\underline{R}(\delta_1) \geq \lim_{\delta \rightarrow 1} \bar{V}(\delta) - \varepsilon$. ($\underline{R}(\delta_1)$ certainly does not exceed $\lim_{\delta \rightarrow 1} \bar{V}(\delta)$, because $\underline{R}(\delta) \leq \bar{V}(\delta)$ for all δ , and $\bar{V}(\delta)$ is nondecreasing in δ .) Choose δ_2 sufficiently large that $\bar{V}(\delta_2) > \lim_{\delta \rightarrow 1} V(\delta) - \varepsilon/2$. Because $\lim_{\delta \rightarrow 1} [(1-\delta)/(1-\delta_2)] = 0$, we can choose δ^* close enough to 1 that for all $\delta_1 > \delta^*$, $[(1-\delta_1)/(1-\delta_2)](\bar{V}(\delta_2) - \underline{V}(\delta_2)) < \varepsilon/2$. Since $V(\delta_2)$ is self-generating for δ_2 , Proposition 2 (radial contraction) ensures that $(\bar{V}(\delta_2), \bar{V}(\delta_2) - [(1-\delta_1)/(1-\delta_2)](\bar{V}(\delta_2) - \underline{V}(\delta_2)))$ is self-generating for δ_1 . Then

$$\begin{aligned} \underline{R}(\delta_1) &\geq \bar{V}(\delta_2) - \left(\frac{1-\delta_1}{1-\delta_2} \right) (\bar{V}(\delta_2) - \underline{V}(\delta_2)) \\ &> \bar{V}(\delta_2) - \frac{\varepsilon}{2} \\ &> \lim_{\delta \rightarrow 1} \bar{V}(\delta) - \varepsilon. \end{aligned} \quad \text{Q.E.D.}$$

To summarize, the symmetric theory yields unique quantitative predictions for maximally collusive payoffs and for severest credible punishments. Whenever the availability of Cournot-Nash reversion would have sufficed, for a particular δ , to maintain some cooperation, a degree of cooperation is

also possible in renegotiation-proof equilibrium. But for all $\delta' > \delta$, threats as severe as Cournot-Nash reversion are not renegotiation-proof, and (given certain regularity conditions) the best attainable payoffs fall short of $\bar{V}(\delta')$ (the best payoff available when the problem of renegotiation is ignored). Thus the possibility of renegotiation exerts a moderating influence on extremal payoffs. The limiting behavior of maximal payoffs, however, is exactly the same whether or not renegotiation is considered:

$\lim_{\delta \rightarrow 1} \bar{R}(\delta) = \lim_{\delta \rightarrow 1} \bar{V}(\delta) = \lim_{\delta \rightarrow 1} \underline{R}(\delta)$. Fortunately, the structural simplicity of the S.S.E.'s studied in Abreu, Pearce and Stacchetti (1986a) is preserved in this theory. For example, there exists a renegotiation-proof equilibrium with value \bar{R} , in which only two actions are ever used, the alternation between them constituting a Markov chain. In this equilibrium, only the continuation values \bar{R} and \underline{R} ever arise.

4. An Extension to Perfect Monitoring

Before proceeding to asymmetric settings, I would like to indicate very briefly how the theory of Section 3 can be adapted to games with perfect monitoring. The component game G is now a standard, deterministic finite N -person game; in the repeated game $G^\infty(\delta)$, each player can condition his choice in period t on the past choices of all players. It is convenient to allow public randomization over continuation paths, for purposes of convexification.

First, it is necessary to sketch the theory of self-generation that is relevant for $G^\infty(\delta)$ (that such a theory is available was noted in Abreu,

Pearce and Stacchetti (1986a)). Here V is the set of (average) values of all pure strategy symmetric subgame perfect equilibria (S.P.E.'s), where symmetry of a profile σ means that for no $t-1$ period history does σ specify that two players take different actions in period t ; S is the set of pure strategies available to each player. A pair (s,u) with $s \in S$ and $u : S^N \rightarrow \mathbf{R}$ is admissible with respect to $W \subseteq \mathbf{R}$ if

$$(i) \quad u(q_1, \dots, q_N) \in \text{co } W \quad \forall (q_1, \dots, q_N) \in S^N$$

and (ii) for each $s' \in S$,

$$\Pi_1(s', s, \dots, s) - \Pi_1(s, \dots, s) \leq u(s, \dots, s) - u(s', s, \dots, s) .$$

The value of the pair (s,u) is $E(s;u) = (1-\delta)\Pi_1(s, \dots, s) + \delta u(s, \dots, s)$, and $B(W) = \text{co}\{E_1(s;u) \mid (s,u) \text{ is admissible w.r.t. } W\}$. With the above interpretation of symbols and terminology, (R1)-(R6) of Section 2 can be proved for the perfect monitoring model.

The definition of a renegotiation-proof equilibrium is precisely as in Section 3, but the characterization in terms of self-generating sets is slightly complicated by the public randomization implicit in the new definition of $B(W)$. The proof of Lemma 5 is a straightforward modification of that of Lemma 2, and is omitted.

Lemma 5: Let σ be an S.P.E. The following three conditions are equivalent:

- (i) σ is renegotiation-proof.
- (ii) Let $F = \{W \subseteq \mathbb{R} \mid W \text{ is a compact self-generating set s.t. for some pair } (s,u) \text{ admissible w.r.t. } W, \underline{W} = E_1(s,u)\}$. There exists $W_1 \in F$ with $C(\sigma) \subseteq W_1$, such that for all $W_2 \in F$, $\underline{W}_1 \geq \underline{W}_2$.
- (iii) For every S.P.E. γ , $\inf C(\gamma) \leq \inf C(\sigma)$.

An example may be helpful in clarifying the reason for re-writing condition (ii) in the lemma. Suppose that $W = [9,10]$ and $B(W) = [6,10]$, and that there is no pair (s,u) admissible with respect to W with value 8 (8 is in $B(W)$ because it can be obtained as a convex combination of values of admissible pairs, say 6 and 10). Then although W is self-generating, 8 is not a lower bound on \mathbb{R} , because the value 8 cannot be obtained without using (with positive probability) continuation values less than 8.

In most respects the theory with perfect monitoring shares the essential properties of the theory with imperfect monitoring, including those stated in the lemmas and propositions of Section 3. The remark in that Section that, subject to regularity conditions, \bar{R} falls short of \bar{V} , is no longer appropriate: under perfect monitoring, continuation values of optimal equilibria need not be restricted to the extreme points of the value set.

Example: Severest Punishments in Repeated Cournot and Bertrand Oligopolies

Consider N firms, each with constant marginal cost $c > 0$, facing a downward-sloping linear market inverse demand curve $P(Q)$. If they are Cournot competitors, quantity is the strategic variable, and the component game is denoted G_C . Alternatively, if they are Bertrand competitors, with the usual tie-breaking rule of equal division), price is the strategic variable, and the component game is denoted G_β . In either case, the strategy set of the component game is R_+ , which violates the finiteness assumption made earlier; this presents no problems for the theory. Neglecting renegotiation, a zero-profit punishment, namely reversion to a one-shot Bertrand equilibrium with price equal to marginal cost, is always available in $G_\beta^\infty(\delta)$, but not necessarily in $G_C^\infty(\delta)$, for a particular δ . While renegotiation precludes using zero-profit punishments in either supergame, Lemma 6 shows that invariably, $R_C(\delta) \geq R_\beta(\delta)$. As in the standard theory, this does not allow one to rank the payoffs from maximal collusion in the two games.

The essential idea behind the proof of Lemma 6 is that for a given market price, a Bertrand oligopolist can deviate more effectively (capturing his rivals' market shares without significantly decreasing the price) than a Cournot oligopolist. The difficulty of the proof is that the preceding statement is true only if the deviation involves decreasing price. Throughout the proof, superscripts and subscripts β and C identify whether a function or set is associated with the Bertrand game or the Cournot game, respectively.

Lemma 6: $R_C \geq R_\beta$.

Proof: Suppose not. Then there exists an element W of F_β (recall the definition of F in Lemma 5) such that for all $W_1 \in F_C$, $W > W_1$. Let Π_1 be player 1's payoff function in G_C , and q_e be the quantity produced by each firm in the Cournot equilibrium of G_C ; the singleton $\{\Pi_1(q_e)\}$ is in F_C , therefore $W > \Pi_1(q_e)$. We can obtain a contradiction by establishing that $(\underline{W}, \bar{W}) \in F_C$. There exists a pair (s, u) admissible with respect to W , for the Bertrand game, and having value \underline{W} . Consider two cases.

Case 1: $s \geq P(Nq_e)$. Choose q to solve $s = P(Nq)$. Let $\hat{u}(q) = u(s)$, and $\hat{u}(r) = \underline{W}$ for $r \neq q$. Then the pair (q, \hat{u}) has value \underline{W} , and the incentive constraints for admissibility (in the Cournot game) with respect to (\underline{W}, \bar{W}) are satisfied; because $q \leq q_e$, and \hat{u} punishes deviations, only defections that increase output need be checked. If a defection to some $q' > q$ yields a payoff higher than $E_1^C(q, \hat{u})$, the defection in the price game leading to the same price $P((N-1)q + q')$ is even more profitable, but this contradicts admissibility of (s, u) in the Bertrand game.

Case 2: $s < P(Nq_e)$. Let $\tilde{u}(r) = z$ for all $r \in R_+$, where z is chosen so that $E_1^C(q_e, \tilde{u}) = \underline{W}$. It is simple to check that $z \in W$. Moreover (q_e, \tilde{u}) satisfies the incentive constraints for admissibility (in the Cournot game) with respect to (\underline{W}, \bar{W}) , because q_e is the Cournot-Nash quantity, and the reward function \tilde{u} is constant. Therefore (q_e, \tilde{u}) is admissible with respect to (\underline{W}, \bar{W}) in the Cournot game.

Similarly, there exists a pair (s^*, u^*) admissible with respect to W in the Bertrand game, having value \bar{W} . Because $\bar{W} > \Pi_1(q_e, \dots, q_e)$, and the value of (s^*, u^*) is an average of values in W and the value $\Pi_1(s^*, \dots, s^*)$, the latter must exceed Cournot profits $\Pi_1^C(q_e, \dots, q_e)$, and hence $s^* > P(Nq_e)$. Now repeat the argument of Case 1 above. Q.E.D.

5. Asymmetric Equilibria

This Section generalizes the solution concept studied in Section 3 to equilibria of asymmetric games. The equilibrium value sets are now multi-dimensional, and this necessitates a more complicated notion of severest credible punishments. Unfortunately, many different beliefs regarding what continuation values are renegotiation-proof are consistent with a particular game; it is generally not possible to isolate a unique renegotiation-proof set. (The same problem arises in other theories of renegotiation.)

I continue to use compact self-generating sets to summarize players' attitudes concerning what values are renegotiation-proof. The notation used below is exactly that of Section 2. Let $R_+^N = \{(x_1, \dots, x_N) \in R^N \mid x_i \geq 0, i = 1, \dots, N\}$ and $R_{++}^N = \{(x_1, \dots, x_N) \in R^N \mid x_i > 0, i = 1, \dots, N\}$. For any $X \subseteq R^N$, define $X^+ = \{x+y \mid x \in X \text{ and } y \in R_+^N\}$. A renegotiation-proof set W will have the following interpretation. Players believe that continuation values will withstand renegotiation only if they are at least as attractive, in the Pareto sense, as points on the lower boundary of W , defined by

$$L(W) = \{x \in W \mid \exists (y, z) \in W \times \mathbb{R}_+^N \text{ with } x = y+z \text{ and } x \neq y\} .$$

It should not be possible to find another compact self-generating set $Z \subseteq \mathbb{R}^N$ with $Z^+ \subseteq W^+$, such that some point $w \in W$ is strictly Pareto-dominated by a point z on the lower boundary of Z . If such a Z did exist, it would be possible to renegotiate from the value w to the value z , by the following argument. Because Z is self-generating, there exists an S.E. σ with value z and continuation values all lying in Z . All points in Z are considered credible by players, since $Z^+ \subseteq W^+$ (notice that for W compact, $(L(W))^+ = W^+$). Also, the ability to renegotiate to z in the future does not threaten σ , none of whose continuation values is Pareto-dominated by z .

The development of a definition formalizing the preceding discussion proceeds by specifying a binary relation on the set Θ of all compact self-generating sets.

Definition: Let the binary relation $>$ on Θ be defined by: for

$W_1, W_2 \in \Theta$, $W_2 > W_1$ if

$$(i) \quad W_2^+ \subseteq W_1^+$$

and (ii) $w_2 - w_1 \in \mathbb{R}_{++}^N$ for some $(w_1, w_2) \in W_1 \times L(W_2)$.

Notice that $>$ partially orders Θ .

Definition: A set $R \in \Theta$ is renegotiation-proof if it is maximal with respect to the partial ordering $>$, that is, $\nexists X \in \Theta$ with $X > R$. An S.E. σ is renegotiation-proof if there exists a renegotiation-proof set X with $C(\sigma) \subseteq X$.

Because $>$ is not a total ordering of Θ , there may be many maximal elements of Θ , and hence many renegotiation-proof sets. Each one represents a consistent set of beliefs players could commonly hold about what values can survive the threat of renegotiation. Existence does not pose a problem. Indeed, for any $X \in \Theta$, there exists a renegotiation-proof set R such that $R^+ \subseteq X^+$.

Proposition 6: For any self-generating set $X \subseteq \Theta$, there exists a renegotiation-proof set R with $R^+ \subseteq X^+$.

Proof: If X is maximal with respect to $>$ in Θ , we are done. Otherwise, let Γ be a maximal linearly ordered subset of $\{T \in \Theta \mid T > X\}$. If Γ has a maximal element (with respect to $>$), it is a renegotiation-proof set meeting the required condition. If not, for each $W \in \Gamma$ define $>(W) = \{T \in \Gamma \mid T > W\}$ and $\tilde{W} = \cup\{T \mid T \in >(W)\}$. Because Γ has no maximal element, $>(W)$ and \tilde{W} are nonempty. Moreover, \tilde{W} is self-generating, and by (R3) of Section 2, so is $\text{cl } \tilde{W}$. One can check that $\text{cl } \tilde{W} \in >(W)$, therefore $\text{cl } \tilde{W} = \tilde{W}$, so $\tilde{W} \in >(W)$. Now let $Y = \cap\{\tilde{W} \mid W \in >(X)\}$. Since Y is the intersection of a nested family of compact sets having the finite intersection property, Y is a nonempty compact set. A straightforward topological argument establishes the existence of a sequence $\{Y_i\}_{i=1}^{\infty}$ of sets

from the collection $(\tilde{W} | W \in \succ(X))$ with $\bigcap_{i=1}^{\infty} Y_i = Y$. By the argument in the Appendix, Y is self-generating.

We need to show that Y is maximal in θ . If not, $\exists Z \in \theta$ such that $Z > Y$. But for every $W \in \Gamma$, $Z^+ \subseteq Y^+ \subseteq \tilde{W}^+ \subseteq W^+$ (this establishes the first of the conditions (i) and (ii) needed to show $Z > W$). Also there exists $z \in L(Z)$ and $y \in Y \subseteq \tilde{W}$ such that $z - y \in \mathbb{R}_{++}^N$. Thus $Z > \tilde{W} > W$. That this holds for all $W \in \theta$ contradicts the maintained hypothesis that Γ has no maximal element Q.E.D.

Readers familiar with Farrell and Maskin (1986) will recognize that any set W called renegotiation proof in that paper is self-generating; consequently, there is a set R satisfying the definition of renegotiation-proofness studied here, with $R^+ \subseteq W^+$.

As in the symmetric setting, it is not possible for a renegotiation-proof equilibrium to be Pareto-dominated by Cournot-Nash reversion; this is stated precisely in Lemma 7.

Lemma 7: Let $s \in S$ be a Nash equilibrium of G . For any renegotiation-proof set R , $\exists r \in R$ such that

$$\Pi(s) - r \in \mathbb{R}_{++}^N.$$

Proof: If s is a Nash equilibrium of G with $\Pi(s) - r \in \mathbb{R}_{++}^N$ for some $r \in R$, then $(\Pi(s))$ is in θ , $(\Pi(s))^+ \subseteq R^+$, and $(r, \Pi(s)) \in R \times L(\{\Pi(s)\})$. Hence $(\Pi(s)) > R$, and R is not renegotiation-proof. Q.E.D.

Lemma 8 (monotonicity): Suppose that R_1 is renegotiation-proof for the discount factor δ_1 , and $\delta_2 \in (\delta_1, 1)$. Then there exists a set R_2 renegotiation-proof for δ_2 such that $R_2^+ \subseteq R_1^+$.

Proof: R_1 is self-generating for δ_1 , and hence, by (R6) of Section 2, for δ_2 . The result follows from Proposition 6. Q.E.D.

The impossibility of resorting to punishments Pareto-inferior to Cournot-Nash reversion often reduces the collusive efficiency of renegotiation-proof equilibria, for fixed δ . But an analog to the limit theorem of Section 3 is easy to establish: when players are sufficiently patient, given any point w in the interior of the set of values attainable in the limit as $\delta \rightarrow 1$ neglecting the problem of renegotiation, a renegotiation-proof set can be found all of whose values are (for each player) almost as good as, or better than, w . In other words, patience ultimately overcomes the obstacles to cooperation posed by renegotiation.

Proposition 7 uses the following notation: for $W \subseteq \mathbb{R}^N$ and $(\lambda, x) \in \mathbb{R} \times \mathbb{R}^N$, $\lambda W + (1-\lambda)x = \{\lambda w + (1-\lambda)x \mid w \in W\}$.

Proposition 7: Let $V^* = \cup\{V(\delta) \mid \delta \in (0, 1)\}$. For any $w \in V^*$ and $\varepsilon > 0$, there exists δ^* such that for any $\delta \geq \delta^*$, there exists a renegotiation-proof set R for δ with $R \subseteq \{w - (\varepsilon, \dots, \varepsilon)\}^+$.

Proof: Choose δ_1 sufficiently close to 1 so that $V(\delta_1)$ contains some point x in $B_{\varepsilon/2}(w)$, the open ball of radius $\varepsilon/2$ around w . Now choose δ_2 large enough so that the closed set $W = \lambda V(\delta_1) + (1-\lambda)x$ is a subset of $B_{\varepsilon/2}(x)$, where $\lambda = (1 - \delta_2)/(1 - \delta_1)$. Then $W \subseteq B_\varepsilon(w)$. Proposition 2 guarantees that W is self-generating for δ_2 , therefore by Proposition

6, there exists a set R renegotiation-proof for δ_2 , satisfying $R^+ \subseteq W^+ \subseteq (w - (\varepsilon, \dots, \varepsilon))^+$. By Lemma 8 (monotonicity), we are done.

Q.E.D.

6. Conclusion

Cooperation in repeated games relies on the possibility that equilibrium play following some t -period history depends on more than simply the structure of the game remaining after the first t periods, that structure being always the same. In a nondegenerate theory of renegotiation, what a player expects, and the statements he finds credible at the end of period t must be affected by the history that has transpired, and perhaps by the implicit agreement that was in force. The solution concept proposed in this paper acknowledges both these influences, while imposing a certain stationarity on beliefs regarding what renegotiation options are available: renegotiation to an equilibrium σ will not take place if, after some history h , the continuation equilibrium $\sigma|_h$ is itself vulnerable to renegotiation to σ (in the sense that all players prefer σ to $\sigma|_h$). Even in repeated games having only symmetric equilibria, this leads to a substantial theory yielding specific predictions about (constrained) efficient cooperation and severest credible punishments. Although there is often a multiplicity of renegotiation-proof sets in asymmetric games, each of them typically shares most of the features of the symmetric theory. The possibility of renegotiation moderates the punishments available, as well as the collusive payoffs attainable, but in the special case when players are very

patient, the value of any renegotiation-proof equilibrium is near the Pareto frontier of the equilibrium value set in the traditional theory without renegotiation.

APPENDIX

Theorem: Let (W_n) be a nonincreasing sequence of compact self-generating sets, and let $W = \bigcap_n W_n$. Then W is self-generating.

Proof: For each $w \in W$ and $n = 1, 2, \dots$, because $w \in W_n$ we can choose a pair (s_n, u_n) admissible with respect to W_n such that $E(s_n; u_n) = w$. Since S is compact, without loss of generality assume $s_n \rightarrow s$. For every n , $\text{co } W_n \subseteq \text{co } W_1$ and W_1 is compact, $L^\infty(\Omega; \text{co } W_1)$ is weak*-compact and without loss of generality $u_n \xrightarrow{*} u$. For each $m = 1, 2, \dots$, $u_n(\Omega) \subseteq \text{co } W_n \subseteq \text{co } W_m$ for every $n > m$, therefore $u(\Omega) \subseteq \bigcap_m \text{co } W_m = \text{co } W$ (for the last equality, see, for example, APS, Lemma 4). Thus $u : \Omega \rightarrow \text{co } W$. Since $E : S \times L_\infty(\Omega; \mathbb{R}^N) \rightarrow \mathbb{R}^N$ is continuous when $L_\infty(\Omega; \mathbb{R}^N)$ is endowed with the weak-* topology (see Abreu, Pearce and Stacchetti (1986b)), it is easy to check that the incentive constraints for admissibility of (s, u) with respect to $\text{co } W$ are satisfied, and that $E(s; u) = w$. By (R5) of Section 2, there exists $\tilde{u} : \Omega \rightarrow \text{ext co } W$ such that $E(s; \tilde{u}) = w$ and (s, \tilde{u}) is admissible with respect to $\text{ext co } W$, and hence with respect to W , because $\text{ext co } W \subseteq W$. Q.E.D.

REFERENCES

- Abreu, D. (1983). "Repeated Games with Discounting: A General Theory and an Application to Oligopoly," Ph.D. Thesis, Princeton University.
- Abreu, D., D. Pearce and E. Stacchetti (1986a). "Optimal Cartel Equilibria with Imperfect Monitoring," Journal of Economic Theory, 39, 251-269.
- Abreu, D., D. Pearce and E. Stacchetti (1986b). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," Cowles Foundation Discussion Paper No. 791, Yale University.
- Aumann, R. and L. Shapley (1976). "Long Term Competition--A Game Theoretic Analysis," unpublished manuscript.
- Bernheim, D., B. Peleg and M. Whinston (1987). "Coalition-Proof Nash Equilibria. I. Concepts," Journal of Economic Theory, 42, 1-12.
- Bernheim, D. and M. Whinston (1987). "Coalition-Proof Nash Equilibria. II. Applications," Journal of Economic Theory, 42, 13-29.
- Cave, J. (1985). "The Cold Fish War: Long-Term Competition in a Dynamic Game," unpublished manuscript.
- van Damme, E. (1986). "Renegotiation-Proof Equilibria in Repeated Prisoners' Dilemma," University of Bonn Discussion Paper.
- Farrell, J. (1983). "Credible Repeated Game Equilibrium," unpublished manuscript.
- Farrell, J. and E. Maskin (1986). "Renegotiation-Proof Equilibrium," seminar notes, Harvard University.
- Friedman, J. (1971). "A Non-cooperative Equilibrium for Supergames," Review of Economic Studies, 38, 1-12.
- Fudenberg, D. and E. Maskin (1986). "Discounted Repeated Games with Unobservable Actions, II: Two-Sided Moral Hazard," unpublished manuscript.
- Green, E. and R. Porter (1983). "Noncooperative Collusion under Imperfect Price Information," Econometrica, 52, 87-100.

- Greenberg, J. (1986). "Perfect Equilibria Paths in Repeated Games: The Unique Maximal Stationary Stable Standard of Behavior," I.M.S.S.S. Technical Report No. 495, Stanford University.
- Porter, R. (1984). "Optimal Cartel Trigger-Price Strategies," Journal of Economic Theory, 29, 313-338.
- 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.
- Rubinstein, A. (1979). "Equilibrium in Supergames with the Overtaking Criterion," Journal of Economic Theory, 21, 1-9.