COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 726

OPTIMAL CARTEL EQUILIBRIA WITH IMPERFECT MONITORING

by

Dilip Abreu, David Pearce and Ennio Stacchetti

October 1984

OPTIMAL CARTEL EQUILIBRIA WITH IMPERFECT MONITORING*

by

Dilip Abreu†, David Pearce‡ and Ennio Stacchetti‡

There exist optimal symmetric equilibria in the Green-Porter model [5],[8] having an elementary intertemporal structure. Such an equilibrium is described entirely by two subsets of price space and two quantities, the only production levels used by firms in any contingency. The central technique employed in the analysis is the reduction of the repeated game to a family of static games.

*This research was conducted mainly while the authors were visiting the Institute for Mathematics and its Applications of the University of Minnesota, and was supported by National Science Foundation Grant #MCS 81-20789. We would like to acknowledge helpful conversations with Peter Doyle, Motty Perry, Rob Porter and Hans Weinberger.

†Harvard University.

‡Yale University.

‡Stanford University.

I. Introduction

A subject of major interest in oligopoly theory is the nature and degree of implicit collusion that can be sustained amongst quantity-setting firms, via strategies that make output levels at time t depend upon aspects of the history of the oligopoly prior to t. The scope of such strategies clearly is limited by the extent of the firms' knowledge of that history. By far the most attention has been given to the case in which the production levels of all firms in each previous period are common knowledge (see, for example, Friedman [4] and Abreu [1]). The consequences of invoking the opposite polar assumption about information are explored in two important papers by Porter [8] and Green and Porter [5]. In their setting, firms cannot observe one another's past production levels; thus the actions of a firm are functions only of past prices, and possibly of the firm's own quantity history. Information about the history of production is further obscured by the fact that there is a stochastic component to the market price; the level of aggregate production cannot be inferred precisely from the price. Our paper exhibits strategies for the firms in the Green-Porter model that form optimal symmetric sequential equilibria of the discounted oligopolistic supergame. The equilibria have an extremely simple intertemporal structure, and their optimality is established under very general conditions.

Porter [8] studies the problem of how to maximize discounted profits with strategies based on trigger prices and Cournot-Nash reversion. He assumes that firms produce some quantity q. less than the Cournot-Nash amount, until price drops below a trigger price $\tilde{p}$ in some period. Thereafter, all firms produce their Cournot-Nash quantities for T periods, after which they return to q (which is maintained until price again falls below $\tilde{p}$, and so on).

Porter characterizes the choice of q, $\tilde{p}$, and T (the latter may be infinite) that yields cartel members the greatest discounted profit, subject to the constraint that the resulting regime must be a sequential equilibrium.

The optimization described above is severely limited in that it restricts attention to a small subset of the strategies available to cartel members. Included among the restrictions implicitly imposed by Porter are the following:

(i) The critical set of prices that trigger T periods of "punishment" is of the form $[0,\tilde{p}]$, that is, a tail test is used.

(ii) The trigger price at time t is independent of prices prior to $t - 1$, so review strategies such as those used by Radner [9] are ruled out.

(iii) Only one punishment is used: a firm cannot increase its quantity by different amounts, or for a different number of periods, depending on how low the preceding price was.

(iv) A firm's output at t cannot be a function of its previous outputs.

(v) Punishments more severe than Cournot-Nash reversion are not permitted (see Abreu [1]).

Porter [8] understands that a global optimum may not be achieved by equilibria of the sort he considers, and explains that complexity poses a major problem: "Since we will later allow the cartel to maximize joint value subject to enforcement constraints, the restriction to simple trigger strategies may not be desirable, given that more general strategies may lead to better outcomes... Unfortunately, models with strategies such as these are extremely complicated for computational purposes, and it is difficult to obtain any interesting results." (page 316).

After modifying the model (Section 2) by replacing the continuum of possible production levels with a discrete production set for each player, we are able to dispense with all five restrictions listed above. We find that there exist equilibria of surprising simplicity which are optimal among all symmetric sequential equilibria[1]. In these equilibria, only two quantities are ever produced. To compute which quantity to produce in period t, a firm simply needs to remember the price in the previous period, and which quantity was specified by the equilibrium in that previous period. Thus the sequence of equilibrium production levels is a Markov chain. In fact, (ii), (iii), and (iv) above are satisfied by these optimal equilibria: once (i) and (v) are relaxed, Porter's remaining conditions are unrestrictive! These results require only mild regularity assumptions on demand, and the stochastic disturbance in price can enter in a rather general way.

The central technique employed is the reduction of the repeated game to a static structure from which can be extracted the optimal equilibria in question (and indeed, equilibria supporting <u>any</u> symmetric sequential equilibrium payoff vector). Notice that every symmetric sequential equilibrium (hereafter S.S.E.) must prescribe other "successor" S.S.E.'s to follow each one-period price history (we show that quantity histories can be ignored). One can imagine constructing a new game by truncating the discounted supergame as follows: after each first-period history, replace the S.S.E. successor by the payoffs associated with that successor. The first-period equilibrium quantities will still constitute an equilibrium of the new game, and the resulting total payoffs will also be the same. More generally (if less intuitively), for any bounded set W of real numbers, let $B(W) \subseteq R$ represent the total payoffs that players

---

[1] For the moment we restrict attention to equilibria in pure strategies. However the extension to mixed strategy equilibria is straightforward (see the note at the end of the paper).

could receive in pure strategy equilibria of truncated games in which each first-period price is followed by some symmetric payoff drawn from W. At the end of the single period of the truncated game, firms receive their conventional one-period profits, plus some element of W (the same for each player), depending on the price that arises; the expected value of this sum, discounted to the beginning of the period, will be one element of $B(W)$.

W is said to be self-generating (in the context of a particular repeated game) if $W \subseteq B(W)$. Let V be the set of payoffs a player can receive in the various S.S.E.'s of the repeated game. Section 3 establishes that every self-generating set is a subset of V. Moreover, V itself is self-generating. These properties yield an elementary proof that V is compact, so best and worst S.S.E.'s exist whenever V is nonempty.

In supporting a symmetric equilibrium with value $v \in V$ in a truncated game, only the best and worst elements of V need be used following the occurrence of any price. This result, proved in Section 4, is analogous to bang-bang theorems in optimal control theory (see, for example, Artstein [2]). The implication for the repeated game is a dramatic dimensional simplification. If v is the payoff associated with any S.S.E., v can be supported by an S.S.E. which in every contingency (except for the first period) looks like the first period of either the best or the worst S.S.E. The first period of the equilibrium yielding v is no longer an exception if v is the optimal element of V. These results are explained in detail in Section 5 , which presents in addition a result on the sensitivity of maximal cartel profits to changes in the discount rate.

The conclusion raises the possibility of applying the analytic approach taken here to other classes of repeated games.

2. The Model

The structure presented here is based entirely upon the model developed in Porter [8] and Green and Porter [5]. Departures from their assumptions are noted below. The oligopoly is modelled by an N-person infinitely-repeated game with discounting. The first step in defining the game is to specify the single-period component game G.

The Single-Period Game: N identical firms simultaneously choose quantities q_i, $i=1,\ldots,N$, of output to produce. Whereas in [8] all non-negative output levels are feasible, we assume that there is an indivisibility in production; only integer multiples of some fundamental unit of production (which of course may be arbitrarily small) can be produced. This assumption simplifies the mathematics involved in Section 4. While it is a departure from tradition, it would appear to represent an increase in realism. Firms incur a total production cost $c(q_i) > 0$ ([8] requires constant marginal cost). Market price p depends on aggregate production q and a stochastic variable θ, that is, $p = P(q;\theta)$. In [8], the inverse demand function is linear in q, and θ enters as an additive or multiplicative disturbance: $p = a + bq + \theta$ or $p = (a + bq)\theta$. We dispense with these assumptions. The payoff of firm i given θ is

$$\pi_i(q_1,\ldots,q_N;\theta) = q_i \cdot P(\sum_{j=1}^{N} q_j;\theta) - c(q_i),$$

where q_j is the quantity produced by firm j. Firms maximize expected profit. We will denote by $\bar{\pi}_i$ the expected value of π_i (assumption (A1) below guarantees that $\bar{\pi}_i$ is well defined), that is

$$\bar{\pi}_i(q_1,\ldots,q_N) = \int_{-\infty}^{\infty} \pi_i(q_1,\ldots,q_N;\theta) \cdot f(\theta;\sum_{j=1}^{N} q_j)\, d\theta,$$

where f is the density function of θ, which we allow to depend on the aggregate output. We do not need the assumption in [8] that f is an increasing function over its support. Thus the one-period game is given by $G = (S_1,\ldots,S_N;\bar{\pi}_1,\ldots,\bar{\pi}_N)$, where $S_i = \{0,1,2,\ldots\}$, $i=1,\ldots,N$.

We assume that:

(A1) There exists a constant K such that $q \cdot E_\theta P(q;\theta) < K$ for all $q \in S_1$. (Note that $S_1 = \sum_{i=1}^N S_i$.)

(A2) $P(q;\theta)$ is strictly monotone in θ.

(A3) The set

$$\Omega = \left\{ P(\sum_{j=1}^N q_j;\theta) \;\middle|\; \theta \in \mathbf{R} \text{ and } f(\theta;\sum_{j=1}^N q_j) > 0 \right\}$$

is contained in R_+ and is independent of $(q_1,\ldots,q_N) \in \Pi_{i=1}^N S_i$.

This assumption, which will later be relaxed slightly, ensures that under no circumstances will firms be able to infer from the price that with probability 1 , someone has deviated from equilibrium behaviour. Ability to make such inferences fundamentally changes the nature of the problem (see, for example, Abreu [1]).

(A4) $c(0) = 0$, $c(q) > 0$ for all $q \in S_1$ and there exist $c_0 > 0$ and $q_0 > 0$ such that $c(q) > c_0 \cdot q$ for all $q > q_0$.

(A5) G has a symmetric Nash equilibrium (in pure strategies).

<u>The Repeated Game $G^{\infty}(\delta)$</u>: $G^{\infty}(\delta)$ is the infinitely-repeated game defined by the component game G and the discount factor $\delta \in (0,1)$. A strategy σ_i for firm i specifies an output in each period $t=1,2,\ldots$, as a function of past prices $p^{t-1} = (p(1),\ldots,p(t-1))$ and the firm's own past quantities $q_i^{t-1} = (q_i(1),\ldots,q_i(t-1))$. Thus $\sigma_i = (\sigma_i(1),\sigma_i(2),\ldots)$, where $\sigma_i(1) \in S_i$ and $\sigma_i(t)\colon \Omega^{t-1} \times S_i{}^{t-1} \to S_i$ is a (Lebesgue) measurable function for each $t > 2$. A strategy profile σ has the form $\sigma = (\sigma_1,\ldots,\sigma_N)$, and for each $t > 2$, $\sigma(t)(p^{t-1};q^{t-1}) = (\sigma_1(t)(p^{t-1};q_1^{t-1}),\ldots,\sigma_N(t)(p^{t-1};q_N^{t-1}))$, where $p^{t-1} = (p(1),\ldots,p(t-1))$, $q^{t-1} = (q(1),\ldots,q(t-1))$, and $q(s) = (q_1(s),\ldots,q_N(s))$, $s=1,\ldots,t-1$.

Given a history $H(t) = (p^t;q^t)$ and a strategy profile σ, we will denote by $\sigma|_{H(t)}$ the strategy profile induced by σ on the subtree following H(t). Thus, for any sequence of prices $\rho(1),\ldots,\rho(s)$ and any sequence of quantity vectors $\gamma(1),\ldots,\gamma(s)$,

$$\sigma|_{H(t)}(s+1)(\rho^s;\gamma^s) = \sigma(s+t+1)(p(1),\ldots,p(t),\rho(1),\ldots,\rho(s);q(1),\ldots,q(t),\gamma(1),\ldots,\gamma(s))$$

Let θ^t denote the vector $(\theta(1),\ldots,\theta(t))$. Given a strategy profile σ, a path $(q(1),q(2)(\theta^1),q(3)(\theta^2),\ldots)$ for the game is generated in the following way. In the first period, firms produce $q(1) = (\sigma_1(1),\ldots,\sigma_N(1))$, a value $\theta(1)$ of the stochastic component arises from the density $f(\cdot;\sum_{j=1}^N q_j(1))$, and the market price is $p(1)(\theta^1) = P(\sum_{j=1}^N q_j(1);\theta(1))$. In period t, $t > 2$, firms produce $q(t)(\theta^{t-1}) = \sigma(t)(p(1)(\theta^1),\ldots,p(t-1)(\theta^{t-1});q(1),\ldots,q(t-1)(\theta^{t-2}))$, and the market price is $p(t)(\theta^t) = P(\sum_{j=1}^N q_j(t)(\theta^{t-1});\theta(t))$, where $\theta(t)$ is the stochastic component that arises from the density $f(\cdot;\sum_{j=1}^N q_j(t)(\theta^{t-1}))$. We will assume:

(A6) The stochastic variables $\theta(t)$ $t=1,2,\ldots$, are independently generated.

The value $v_i(\sigma)$ of the strategy profile σ for firm i is the expected discounted sum of the payoffs. Define

$$R_i(\sigma;t) = \int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty}\pi_i(q(t)(\theta^{t-1});\theta(t))\cdot f(\theta(1);\sum_{j=1}^{N}q_j(1))\cdots f(\theta(t);\sum_{j=1}^{N}q_j(t)(\theta^{t-1}))d\theta(1)\cdots d\theta(t).$$

Then

$$v_i(\sigma) = \sum_{t=1}^{\infty}\delta^t\cdot R_i(\sigma;t).$$

Single-period payoffs are received at the end of each period; $v_i(\sigma)$ is the value of the infinite stream discounted to the beginning of period 1.

We use sequential equilibrium (see Kreps and Wilson [7]) as the solution concept in the repeated game. A strategy profile σ is symmetric if $\sigma_1(t) = \sigma_2(t) = \ldots = \sigma_N(t)$ for all $t > 1$. (A5) guarantees the existence of a symmetric sequential equilibrium (S.S.E.), because the strategy profile specifying that in every period, independently of the history, each firm produces its Cournot-Nash output is easily shown to be a sequential equilibrium[2]. Therefore the set

$$V = \{\, v_1(\sigma) \mid \sigma \text{ is an S.S.E.} \,\}$$

is nonempty.

Assumption (A1) and the fact that supergame payoffs are discounted, imply the existence of a bounded set $\tilde{S} \subseteq S_1$ such that any strategy for firm i specifying (in any contingency) an action $q_i \notin \tilde{S}$, is strongly dominated. By (A1)

---

[2] We abuse terminology throughout by referring to a profile σ as a sequential equilibrium. Technically, we mean that there exists a system of beliefs μ such that (μ,σ) is a sequential equilibrium. Extensions to repeated games of the ideas of consistency and sequential rationality used in defining a sequential equilibrium are immediate.

there exists $K > 0$ such that

$$q \cdot E_\theta P(q;\theta) < K \quad \text{for all} \quad q \in S_1.$$

Note that for any $q, \gamma \in S_1$,

$$q \cdot E_\theta P(q+\gamma;\theta) < (q+\gamma) \cdot E_\theta P(q+\gamma;\theta) < K.$$

Since $c(q) \geq 0$ for all $q \in S_1$, the maximum single-period payoff for player i is bounded by K, and consequently the maximum supergame payoff for player i is bounded by $\delta K/(1-\delta)$. For any $q > q^* = \max \{ q_0, K/(c_0(1-\delta)) \}$, the maximum supergame payoff that firm i faces if it produces the quantity q, for any aggregate production $\gamma \in S_1$ for the rest of the firms, is bounded by

$$\delta \cdot \{q \cdot E_\theta P(q+\gamma;\theta) - c(q) + \delta K/(1-\delta)\} < \delta \cdot \{K - c_0 \cdot q + \delta K/(1-\delta)\} < 0,$$

and since firm i can always choose to produce $q = 0$ (and therefore get a supergame payoff of 0), firm i will only consider production quantities in $\tilde{S} = [0,q^*]$. This fact can be used to slightly relax assumption (A3) by requiring only that the set Ω defined there be contained in R_+ and independent of $(q_1,\ldots,q_N) \in \tilde{S}^N$.

In the course of the previous argument we have shown that $V \subseteq [0,\delta K/(1-\delta)]$, hence V is bounded.

3. Reduction to a One-Period Problem

For each repeated game $G^\infty(\delta)$, we define below a function $B: 2^R \to 2^R$, where 2^R is the power set of the real numbers. From this function one can recover V, the set of all payoffs received by a player in various S.S.E.'s (symmetric sequential equilibria). B provides a simple method of constructing an S.S.E. giving any desired payoff $v \in V$ to each firm. The function B is related to a class of truncated games associated with $G^\infty(\delta)$ as described in the introduction. This relationship is developed in what follows, to give intuitive motivation for the definitions and propositions; it is not needed for any of the formal results.

The following definitions implicitly depend upon a particular game $G^\infty(\delta)$. By (A2) the function $p = P(\sum_{j=1}^N q_j;\theta)$ has an inverse $\theta = \Theta(p;\sum_{j=1}^N q_j)$. We perform a change of variables to get the density $g(\cdot;\sum_{j=1}^N q_j)$ on price space.

Definition: Given a set $W \subseteq R$, a pair (q,u) with $q \in \tilde{S}$ and $u: \Omega \to W$ (Lebesgue) measurable, is admissible with respect to W if it satisfies the following incentive compatibility constraints:

$$\bar{\pi}_1(q\cdot e_N) + \int_\Omega u(p)\cdot g(p;N\cdot q)\,dp \geq \bar{\pi}_1(\tilde{q},q\cdot e_{N-1}) + \int_\Omega u(p)\cdot g(p;(N-1)\cdot q + \tilde{q})\,dp \quad \text{for all } \tilde{q} \in \tilde{S},$$

where $e_N = (1,\ldots,1) \in R^N$.

Hereafter, for a vector $(q_1,\ldots,q_N) \in \Pi_{j=1}^N S_j$ and a measurable function $u: R_+ \to R$, we use the notation

$$E_i(q_1,\ldots,q_N;u) = \delta\left\{\bar{\pi}_i(q_1,\ldots,q_N) + \int_\Omega u(p)\cdot g(p;\sum_{j=1}^N q_j)\,dp\right\} \quad \text{for all } i=1,\ldots,N.$$

<u>Definition</u>: For each $W \subseteq R$ define

$$B(W) = \{ E_1(q \cdot e_N;u) \mid (q,u) \text{ is admissible with respect to } W \}.$$

Let Q and U be functions with domain B(W) such that for every $w \in B(W)$, $(Q(w),U(w))$ is admissible with respect to W, and $w = E_1(Q(w)\cdot e_N;U(w))$. Such a pair of functions exists (there may be many; if so, choose one pair) by the definition of B(W).

<u>Definition</u>: $W \subseteq R$ is <u>self-generating</u> if $W \subseteq B(W)$.

Notice that $B(W) \subseteq \delta \{ co(W) + \overline{\pi}_i(\tilde{S},\ldots,\tilde{S}) \}$, so B(W) is bounded when W is bounded. (Here, co(W) denotes the convex hull of W.)

An alternative definition of B(W) can be given via a class of one-period games that are modifications of the component game G. For any measurable function $u\colon R_+ \to W$, define $G(u;\delta) = (\tilde{S},\ldots,\tilde{S};E_1(\cdot;u),\ldots,E_N(\cdot;u))$. This is the same game as G, except that players' payoffs are the sum of their payoffs in G and the expected value of the reward function u, all discounted by δ. If $(q,\ldots,q)$ is a Nash equilibrium of $G(u;\delta)$, then by definition (q,u) is admissible with respect to W and $E_1(q\cdot e_N,u) \in B(W)$. Conversely, any $w \in B(W)$ is the payoff associated with some symmetric pure strategy equilibrium of $G(u;\delta)$, for some $u\colon \Omega \to W$. In other words, the incentive compatibility constraints invoked in the definition of the admissibility of (q,u), precisely express the fact that for any player with payoff function $E_i(\cdot;u)$, q is a best response to $(q,\ldots,q)$.

<u>Proposition 1</u>: <u>(Self-generation)</u>. <u>For any bounded set</u> $W \subseteq R$, $W \subseteq B(W)$ <u>implies</u> $B(W) \subseteq V$.

<u>Proof</u>: The idea of the proof is to exhibit, for each $w \in B(W)$, an S.S.E. $\hat{\sigma}(w)$ with $v_i(\hat{\sigma}(w)) = w$ for all $i=1,\ldots,N$. With Q and U as defined above, for each

$w \in B(W)$, let $U^0(w) = w$, $U^1(w) = U(w)$, and

$$U^t(w)(p^t) = U(U^{t-1}(w)(p^{t-1}))(p(t)) \quad \text{for all} \quad p^t \in \Omega^t \quad \text{and} \quad t=2,3,\ldots$$

The functions U^t are well defined because W, the range of U, is contained in $B(W)$, the domain of U. For each $w \in B(W)$, define the symmetric strategy profile $\hat{\sigma}(w)$ by

$$\hat{\sigma}_i(w)(1) = Q(w) \quad \text{and} \quad \hat{\sigma}_i(w)(t)(p^{t-1};q_i^{t-1}) = Q(U^{t-1}(w)(p^{t-1})) \quad t=2,3,\ldots$$

Notice that $\hat{\sigma}(w)(t)$ does not depend on the quantity history $q^t = (q(1),\ldots,q(t))$; consequently when referring to $\hat{\sigma}\big|_{H(t)}$, we will write $\hat{\sigma}\big|_{p^t}$, suppressing the redundant quantity arguments of $H(t) = (p^t;q^t)$. We first show that the value of $\hat{\sigma}(w)$ for each player is w. Since

$$Q(U^{t-1}(w)(p^{t-1})) = Q(U^{t-2}(U(w)(p(1)))(p(2),\ldots,p(t-1)))^3,$$

one can check that

$$\hat{\sigma}(w)\big|_{p(1)} = \hat{\sigma}(U(w)(p(1))).$$

By definition of the functions Q and U, we have:

$$w = \delta \left\{ \overline{\pi}_i(Q(w)\cdot e_N) + \int_\Omega U(w)(p)\cdot g(p;N\cdot Q(w))\, dp \right\}.$$

Also by definition:

$$v_i(\hat{\sigma}(w)) = \delta \left\{ \overline{\pi}_i(\hat{\sigma}(w)(1)) + \int_\Omega v_i(\hat{\sigma}(w)\big|_p)\cdot g(p;N\cdot\hat{\sigma}(w)(1))\, dp \right\} \quad \text{for all } i=1,\ldots,N.$$

But $\hat{\sigma}(w)(1) = Q(w)\cdot e_N$ and $\hat{\sigma}(w)\big|_p = \hat{\sigma}(U(w)(p))$. Therefore

---

[3] It is understood that when $t = 2$, the right side of this equality becomes $Q(U^0(U(w)p(1))) = Q(U(w)(p(1)))$.

$$w - v_i(\hat{\sigma}(w)) = \delta \int_\Omega \{ U(w)(p) - v_i(\hat{\sigma}(U(w)(p)))\} \cdot g(p;N \cdot Q(w))\, dp,$$

so

$$|w - v_i(\hat{\sigma}(w))| < \delta \int_\Omega \sup_{x \in B(W)} |x - v_i(\hat{\sigma}(x))| \cdot g(p;N \cdot Q(w))\, dp.$$

Since g is a probability density, and the above inequality holds for every $w \in B(W)$,

$$\sup_{w \in B(W)} |w - v_i(\hat{\sigma}(w))| < \delta \cdot \sup_{x \in B(W)} |x - v_i(\hat{\sigma}(x))| .$$

However, since $\tilde{S}$ is compact and B(W) is bounded, the last inequality implies that

$$\sup_{w \in B(W)} |w - v_i(\hat{\sigma}(w))| = 0.$$

Hence $v_i(\hat{\sigma}(w)) = w$ for each $w \in B(W)$ and $i=1,\ldots,N$.

Next we show that $\hat{\sigma}(w)$ is a Nash equilibrium for each $w \in B(W)$. Because $G^\infty(\delta)$ and $\hat{\sigma}(w)$ are symmetric in players, it is sufficient to work only with player 1. Define the subsets of strategies for the first player $\Sigma(t,x)$, $t=1,2,\ldots$, $x \in W$, in the following way: $\tau \in \Sigma(t,x)$ if $\tau(s) = \hat{\sigma}_1(x)(s)$ for all $s > t$. Also let $(\tau,\hat{\sigma}_{-1}(x)) = (\tau,\hat{\sigma}_2(x),\ldots,\hat{\sigma}_N(x))$. We first show that for $t > 1$, if $v_1(\tau,\hat{\sigma}_{-1}(x)) < v_1(\hat{\sigma}(x))$ for all $\tau \in \Sigma(t,x)$ and all $x \in W$, then $v_1(\tau,\hat{\sigma}_{-1}(x)) < v_1(\hat{\sigma}(x))$ for all $\tau \in \Sigma(t+1,x)$ and all $x \in W$. Let $x \in W$ and $\tau \in \Sigma(t+1,x)$. Notice that

$$\tau|_{(p,\tau(1))} \in \Sigma(t,U(x)(p)) \quad \text{for all} \quad p > 0.$$

Then

$$v_i(\tau,\hat{\sigma}_{-1}(x)) = \delta \{ \overline{\pi}_1(\tau(1),Q(x)\cdot e_{N-1}) + \int_\Omega v_1(\tau|_{(p,\tau(1))},\hat{\sigma}_{-1}(U(x)(p)))\cdot g(p;\tau(1) + (N-1)\cdot Q(x))\, dp \}$$

$$< \delta \; \{ \; \overline{\pi}_1(\tau(1),Q(x)\cdot e_{N-1}) + \int_\Omega v_1(\hat{\sigma}(U(x)(p)))\cdot g(p;\tau(1) + (N-1)\cdot Q(x)) \; dp \; \}$$

$$= \delta \; \{ \; \overline{\pi}_1(\tau(1),Q(x)\cdot e_{N-1}) + \int_\Omega U(x)(p)\cdot g(p;\tau(1) + (N-1)\cdot Q(x)) \; dp \; \}$$

$$< \delta \; \{ \; \overline{\pi}_1(Q(x)\cdot e_N) + \int_\Omega U(x)(p)\cdot g(p;N\cdot Q(x)) \; dp \; \} \quad \text{(since } (Q(x),U(x)) \text{ is admissible w.r.t. } W)$$

$$= x = v_1(\hat{\sigma}(x)).$$

For $t=1$, $\tau \in \Sigma(1,x)$ only if $\tau = \hat{\sigma}(x)$. Thus $v_1(\tau,\hat{\sigma}_{-1}(x)) = v_1(\hat{\sigma}(x))$. We have then shown, by induction, that for every $t > 1$, no player wishes to deviate from $\hat{\sigma}(x)$ in the first t periods given that he must conform with $\hat{\sigma}(x)$ thereafter.

Finally, assume that there exist $w \in W$, a strategy σ_1 for player 1 and $\varepsilon > 0$ such that

$$v_1(\sigma_1,\hat{\sigma}_{-1}(w)) = v_1(\hat{\sigma}(w)) + \varepsilon = w + \varepsilon.$$

But for T sufficiently large,

$$\sum_{t > T} \delta^t \cdot \left| R_1(\sigma_1,\hat{\sigma}_{-1}(w);t) - R_1(\hat{\sigma}(w);t) \right| < \varepsilon/2.$$

Therefore the strategy $\sigma_1^* \in \Sigma(T,w)$, defined by $\sigma_1^*(t) = \sigma_1(t)$ for all $t=1,\ldots,T-1$, satisfies:

$$v_1(\sigma_1^*,\hat{\sigma}_{-1}(w)) > w + \varepsilon/2.$$

This is a contradiction. Therefore $\hat{\sigma}(w)$ is a Nash equilibrium.

It remains only to show that for each $w \in B(W)$, $\hat{\sigma}(w)$ is a sequential equilibrium. Let player i's beliefs at an information set following price history p^t, when i has followed some strategy σ_i, be generated by the belief that the profile $(\hat{\sigma}_1(w),\ldots,\hat{\sigma}_{i-1}(w),\overline{\sigma}_i,\hat{\sigma}_{i+1}(w),\ldots,\hat{\sigma}_N(w))$ is being followed (with certainty), where $\overline{\sigma}_i = (\sigma_i(1),\ldots,\sigma_i(t),\hat{\sigma}_i(t+1),\hat{\sigma}_i(t+2),\ldots)$. The problem of consistency (see [7]) does not arise, because (A3) ensures that <u>every</u> price observation is compatible with the above beliefs. Note that actions prior to period $(t + 1)$ do not affect payoffs from $(t + 1)$ onward, and $\hat{\sigma}(w)$ is independent of quantity histories. Therefore beliefs about the past are irrelevant, and to demonstrate sequential rationality, it is sufficient to show that after every price history, the profile induced by $\hat{\sigma}(w)$ on the "remainder supergame" is a Nash equilibrium. Now, for any price history p^t, $t \geq 1$,

$$\begin{aligned}\hat{\sigma}(w)\Big|_{p^t}(s)(\rho^s) &= \hat{\sigma}(w)(s+t)(p(1),\ldots,p(t),\rho(1),\ldots,\rho(s)) \\ &= \hat{\sigma}(U^{t+1}(p^t))(s)(\rho(1),\ldots,\rho(s)).\end{aligned}$$

Therefore $\hat{\sigma}(w)\Big|_{p^t} = \hat{\sigma}(w^*)$, where $w^* = U^{t+1}(p^t) \in W$, and we have just shown that $\hat{\sigma}(w^*)$ is a Nash equilibrium for all $w^* \in B(W)$. Q.E.D.

In the context of "simple strategy profiles", Abreu [1] has shown that in order to check for subgame perfection, it is only necessary to verify that in no contingency is a "one-shot" deviation profitable. In the course of proving Proposition 1, we have established the equivalent principle for sequentiality of strategy profiles in $G^{\infty}(\delta)$. The method of proof is much more general than that of [1]; it also appears in a recent paper of Harris [6].

The idea of the proof of Proposition 1 can be explained informally using the modified component games described earlier. Suppose that $w \in B(W)$ for some self-generating set W. Any S.S.E. with value w has to specify some q for all firms to produce in the first period, and "successor" S.S.E.'s for each price p(1) that might arise. These successor profiles implicitly generate a "future reward function" u(p). Begin to construct $\hat{\sigma}(w)$ (of Proposition 1) by choosing a pair $(q;u)$ that is admissible with respect to W, and satisfies $w = E_1(q,\ldots,q;u)$. $\hat{\sigma}(w)$ specifies that each firm produce q in period 1, and implicitly "promises" future rewards $u(p)$ for each $p(1)$ that might arise. (For the moment, postpone worrying about how $\hat{\sigma}(w)$ will provide S.S.E's that deliver the promised utilities.) If each firm takes these future rewards as given, essentially the modified game $G(u;\delta)$ is being played in the first period of the supergame, and $(q,\ldots,q)$ is an equilibrium of that game. Since for all p, $u(p)$ is in W, it is possible, following any first-period price $p(1)$, for $\hat{\sigma}(w)$ to "deliver" $u(p(1))$ by choosing another admissible pair (r,v) with $u(p(1)) = E_1(r,\ldots,r;v)$, specifying that r be produced in period 2, and "promising" future rewards $v(p(2))$ beyond period 2. Again, if firms believe that $v(p(2))$ will occur if $p(2)$ is the realization of the price in the second period, it is optimal for them to produce $(r,\ldots,r)$, an equilibrium of the game $G(v;\delta)$. Repeating this process, one can generate, for any given t, quantities to be produced by obedient firms in the first t periods, for every possible history; in fact, the proof of the Proposition uses an inductive step to define $\hat{\sigma}(w)(t)$ for <u>all</u> t. This produces a strategy profile from which no one wishes to deviate in only one period (admissibility of the (q,u) pair invoked in each contingency guarantees this). Finally, the principle alluded to in the paragraph immediately following the proof, implies that no firm will deviate from $\hat{\sigma}(w)$ even if it is free to do so in infinitely many periods.

<u>Proposition 2</u>: (<u>Factorization</u>) $V = B(V)$.

<u>Proof</u>: Consider $w \in V$. We show that $w \in B(V)$. By definition, there exists an S.S.E. σ such that $v_i(\sigma) = w$, $i = 1,\ldots,N$. Define $q = \sigma_1(1)$ and $u\colon \Omega \to R$ by $u(p) = v_1(\sigma|_{(p,\,\sigma(1))})$ for every $p \in \Omega$. Since σ is a symmetric strategy profile, $\sigma_i(1) = \sigma_1(1)$ for all $i = 2,\ldots,N$, and $\sigma|_{(p,\sigma(1))}$ is also symmetric.

By (A3) the information sets $(p,\sigma_i(1))$, $i = 1,\ldots,N$ are reached in equilibrium for all $p \in \Omega$. Because σ is a sequential equilibrium, this implies that $\sigma\big|_{(p,\sigma(1))}$ is a sequential equilibrium too. Hence $\sigma\big|_{(p,\sigma(1))}$ is an S.S.E. and $u(p) = v_1(\sigma\big|_{(p,\sigma(1))}) \in V$ for all $p \in \Omega$.

By definition,

$$E_1(q\cdot e_N;u) = \delta\ \{\ \overline{\pi}_1(q\cdot e_N) + \int_\Omega u(p)\cdot g(p;N\cdot q)\ dp\ \}$$

$$= v_1(\sigma) = w,$$

and to complete the proof we need only check that the pair (q,u) satisfies the incentive compatibility constraints of admissibility. For any $\gamma \in \tilde{S}$, define σ_1^* by $\sigma_1^*(1) = \gamma$ and $\sigma_1^*\big|_{(p,r)} = \sigma_1\big|_{(p,q)}$ for all $p \in \Omega$ and $r \in \tilde{S}$. Then

$$v_1(\sigma_1^*,\sigma_{-1}) = \delta\ \{\ \overline{\pi}_1(\gamma,q\cdot e_{N-1}) + \int_\Omega v_1(\sigma_1^*\big|_{(p,\gamma)},\ \sigma_{-1}\big|_{(p,q\cdot e_{N-1})})\cdot g(p;(N-1)\cdot q + \gamma)$$

Since $\sigma_1^*\big|_{(p,\gamma)} = \sigma_1\big|_{(p,q)}$ and $u(p) = v_1(\sigma\big|_{(p,q\cdot e_N)})$,

$$v_1(\sigma_1^*,\sigma_{-1}) = \delta\ \{\ \overline{\pi}_1(\gamma,q\cdot e_{N-1}) + \int_\Omega u(p)\cdot g(p;(N-1)\cdot q + \gamma)\ dp\ \}.$$

σ is an N.E., therefore $v_1(\sigma) > v_1(\sigma_1^*,\sigma_{-1})$, that is

$$\delta\ \{\ \overline{\pi}_1(q\cdot e_N) + \int_\Omega u(p)\cdot g(p;N\cdot q)\ dp\ \}$$

$$> \delta\ \{\ \overline{\pi}_1(\gamma,q\cdot e_{N-1}) + \int_\Omega u(p)\cdot g(p;(N-1)\cdot q + \gamma)\ dp\ \}$$

as required. This establishes that $V \subseteq B(V)$. By Proposition 1, $B(V) \subseteq V$.

Q.E.D.

In any S.S.E. σ, the value to firm i of the successor S.S.E. specified following a given first-period price p, must be independent of the quantity that i produced in period 1 (this is because no one else has observed that quantity, and hence i faces the same future environment regardless of his initial output level). Thus the value of σ for each player can be factorized into two terms: the profit from first-period production, and the discounted expected value of a reward function $u(p) = v_1(\sigma|_{(p(1),\sigma(1))})$. Since $\sigma|_{(p(1),\sigma(1))}$ is an S.S.E., this reward function is drawn from V. This, together with the constraints that the firms are willing to produce $\sigma(1) = (q,\ldots,q)$ in period 1, means precisely that (q,u) is admissible with respect to V. Thus the requirements for σ to be an S.S.E. are exactly those needed for $v(\sigma)$ to be in $B(V)$, and therefore $V = B(V)$.

Corollary 1: For every $w \in V$, there exists an S.S.E. σ such that $v_1(\sigma) = w$ and for every $t > 1$, $\sigma(t)$ is independent of the quantity components of the history $H(t) = (p(1),\ldots,p(t);q(1),\ldots,q(t))$.

Proof: For each $w \in V$, $\hat{\sigma}(w)$ (defined in Proposition 1) is such an S.S.E..

Proposition 3: Let $W \subseteq R$ be compact. Then $B(W)$ is compact.

Proof: For each $d \in \tilde{S} - \tilde{S}$ define the map $h_d\colon [\tilde{S} \cap (\tilde{S} - d)] \times L^\infty(\Omega;W) \to R$ by $h_d(q,u) = E_1(q+d,q\cdot e_{N-1};u)$. Endow $L^\infty(\Omega;W)$ with the weak-* topology. For each $q \in \tilde{S}$ define

$$\mu(q;W) = \{\, u \in L^\infty(\Omega;W) \mid (q,u) \text{ is admissible w.r.t. } W \,\}.$$

Then

$$\mu(q;W) = \{\, u \in L^\infty(\Omega;W) \mid h_d(q,u) \le h_0(q,u) \quad \text{for all} \quad d \in \tilde{S} - q \,\}.$$

For each $q \in \tilde{S}$ the maps $h_d(q,\cdot)\colon L^\infty(\Omega;W) \to R$, $d \in \tilde{S} - q$, are continuous, and using Alaoglu's Theorem [10] it is easy to check that $\mu(q;W)$ is compact.

By definition, $B(W) = \bigcup_{q \in \tilde{S}} h_0[\{q\} \times \mu(q;W)]$. Since $\mu(q;W)$ is compact for each $q \in \tilde{S}$, h_0 is continuous in u, and $\tilde{S}$ is finite, $B(W)$ is compact as required. Q.E.D.

Corollary 2: V is compact.

Proof: Notice that $W_1 \subseteq W_2$ implies $B(W_1) \subseteq B(W_2)$ (the operator B is monotone). Let cl(V) denote the closure of V. Since V is bounded (see Section 2), cl(V) is compact. B(cl(V)) contains B(V) = V, by monotonicity, and by Proposition 3 B(cl(V)) is compact. Hence cl(V) $\subseteq$ B(cl(V)), therefore Proposition 1 implies cl(V) $\subseteq$ V, that is, V is closed, and hence compact. Q.E.D.

Recall that (A5) implies V is nonempty.

Corollary 3: Let $\bar{v} = \max V$ and $\underline{v} = \min V$. Then

$$\bar{v} = \max \{ E_1(q \cdot e_N;u) \mid (q,u) \text{ is admissible w.r.t. } V \},$$

and

$$\underline{v} = \min \{ E_1(q \cdot e_N;u) \mid (q,u) \text{ is admissible w.r.t. } V \}.$$

Proof: $\bar{v} = \max B(V)$ and $\underline{v} = \min B(V)$ by Proposition 2. Q.E.D.

We have shown that V can be recovered from the function B. By Proposition 2, V is a fixed point of B; Proposition 1 implies V is the largest compact fixed point of B. Moreover, since V is self-generating, any element w of V is the payoff of $\hat{\sigma}(w)$, the S.S.E. constructed in Proposition 1. This supergame equilibrium is described entirely by two functions Q and U, and the number w. Section 4 proves that U may be chosen to have a "bang-bang" property that makes the intertemporal structure of the equilibrium entirely elementary.

4. Existence of a Bang-Bang (Optimal) Reward Function.

Proposition 4 shows that if the pair $(q,\bar{u})$ is admissible with respect to V, there exists a function u taking on only the two values $\bar{v} = \max V$ and $\underline{v} = \min V$, such that (q,u) is admissible and $E_1(q,\ldots,q;u) = E_1(q,\ldots,q;\bar{u})$, that is, the value is unchanged. As a consequence, each function U(w) used in constructing $\hat{\sigma}(w)$ can be chosen to have range $\{\underline{v},\bar{v}\}$.

Proposition 4: Let $(q,\bar{u})$ be an admissible pair with respect to a compact set $W \subseteq R$. Let $\underline{w} = \min W$ and $\bar{w} = \max W$. Then there exists a function $u: \Omega \to \{\underline{w},\bar{w}\}$ such that (q,u) is admissible with respect to W and $E_1(q,\ldots,q;u) = E_1(q,\ldots,q;\bar{u})$.

Proof: As in Proposition 3, let $h_d(q,u) = E_1(q + d, q\cdot e_{N-1};u)$, $d \in \tilde{S} - q$. We show that the set

$$F = \{ u: \Omega \to W \mid (q,u) \text{ is admissible w.r.t. } W \text{ and } h_0(q,u) = h_0(q,\bar{u})\}$$

contains an element with the bang-bang property. Clearly $F \subseteq \tilde{F}$, where

$$\begin{aligned}\tilde{F} &= \{ u: \Omega \to co(W) \mid (q,u) \text{ is admissible w.r.t. } co(W) \text{ and } h_0(q,u) = h_0(q,\bar{u})\} \\ &= \{ u \in L^\infty(\Omega;co(W)) \mid h_0(q,u) = h_0(q,\bar{u}) \text{ and } h_d(q,u) < h_0(q,u) \text{ for all } d \in \tilde{S} - q \}.\end{aligned}$$

It is easy to check that $\tilde{F}$ is convex, and as in Proposition 3, it may be verified that $\tilde{F}$ is compact when $L^\infty(\Omega;co(V))$ is endowed with the weak-* topology. Since $(q,\bar{u})$ is admissible with respect to W, F and therefore $\tilde{F}$ are nonempty. Thus $\tilde{F}$ is a nonempty compact convex set. By the Krein-Milman Theorem, $\tilde{F}$ has an extreme point. In what follows we show that each extreme point u of $\tilde{F}$ has the bang-bang property that $\bar{u}(p) \in \{\underline{w},w\}$ for all $p \in \Omega$.

The proof is by contradiction. Assume that u is an extreme point of $\tilde{F}$, but there exists a set $M \subseteq \Omega$ of positive (Lebesgue) measure such that $u(p) \notin \{\underline{w},\overline{w}\}$ for each $p \in M$. Without loss of generality, we can assume that there exists $\varepsilon > 0$ such that $u(M) \subseteq [\underline{w} + \varepsilon, \overline{w} - \varepsilon]$ (see, for example Proposition 14, pg.61 in Royden [10]).

Let m be the cardinality of $\tilde{S}$. Partition M into m+1 sets M(k), k=1,..,m+1, of positive measure. Define the matrix $A = (a_{d,k})$ in the following way:

$$a_{d,k} = \int_{M(k)} g(p;N\cdot q + d)\, dp \quad \text{for all} \quad d \in \tilde{S} - q \quad \text{and all} \quad k=1,\dots,m+1.$$

Let $x \in R^{m+1}$ be a non-zero solution of $A\cdot x = 0$ with $|x_k| < \varepsilon$ for all $k=1,\dots,m+1$. Such a solution exists because A has m rows and m+1 columns. Define the function $v\colon \Omega \to (-\varepsilon,\varepsilon)$ by $v = \sum_{k=1}^{m+1} x_k \cdot \chi_{M(k)}$, where $\chi_{M(k)}$ is the characteristic function of M(k). Since $(u + v) \in L^\infty(\Omega;co(W))$ and $h_d(q,u + v) = h_d(q,u)$ for all $d \in \tilde{S} - q$, $(u + v) \in \tilde{F}$. Similarly $(u - v) \in \tilde{F}$. However $u = (u + v)/2 + (u - v)/2$, contradicting the fact that u is extreme. Therefore $u\colon \Omega \to \{\underline{w},\overline{w}\}$ as asserted.

Finally, since $\{\underline{w},\overline{w}\} \subseteq W$, u belongs to F. Q.E.D.

Let W be any compact subset of R, and define $\underline{w} = \min W$ and $\overline{w} = \max W$. An immediate consequence of Proposition 4 is that

$$B(\{\underline{w},\overline{w}\}) = B(W) = B(co(W)).$$

5. The Elementary Intertemporal Structure of S.S.E.'s

Proposition 4 of the last section makes it possible to simplify further the equilibria that support elements of V. We show that for every $w \in V$, there exists an S.S.E. with value w, that is completely described by three quantities and three subsets of Ω. If $w = \overline{v}$ or $\underline{v}$, only two quantities and two subsets are required. These two quantities are the only ones produced in the equilibrium in question, the alternation between them constituting a Markov process. A further application of Proposition 4 occurs in the proof (given later in this section) of the monotonicity of maximal symmetric cartel profits in the discount rate δ.

Recall that the profile $\hat{\sigma}(w)$ was constructed in Section 3 (for any $w \in V$) by choosing, for each $\tilde{w} \in V$, a pair $(Q(\tilde{w}),U(\tilde{w}))$ that was admissible with respect to V, and satisfied $\tilde{w} = E_1(Q(\tilde{w}),\ldots,Q(\tilde{w});U(\tilde{w}))$. Proposition 4 shows that this pair can always be chosen such that the range of the function $U(\tilde{w})$ is simply $\{\underline{v},\overline{v}\}$. In what follows we assume that the functions Q and U were indeed chosen this way. A trivial argument establishes that for any $w \in V$, only two quantities are ever produced after period 1 in the equilibrium $\hat{\sigma}(w)$. In period $t \geq 2$, the production level is $Q(U^{t-1}(w)(p^{t-1}))$, which is either $Q(\overline{v}) = \overline{q}$ or $Q(\underline{v}) = \underline{q}$ (depending on p^{t-1}), since $\overline{v}$ and $\underline{v}$ are the only values that $U^{t-1}(w)(p^{t-1}) = U(U^{t-2}(w)(p^{t-2}))(p(t-1))$ can assume.

For any $\tilde{w}$, $U(\tilde{w})$ partitions price space into a "reward" region, in which $U(\tilde{w})(p) = \overline{v}$, and a "punishment" region, in which $U(\tilde{w})(p) = \underline{v}$. We denote the reward region and its complement as follows:

$$R(\tilde{w}) = \{ p \in \Omega \mid U(\tilde{w})(p) = \overline{v} \}$$
$$R^c(\tilde{w}) = \{ p \in \Omega \mid U(\tilde{w})(p) = \underline{v} \}$$

Let $\overline{R} = R(\overline{v})$ and $\underline{R} = R^c(\underline{v})$. In the first period of the S.S.E. $\hat{\sigma}(w)$, each firm produces $Q(w)$, and some price $p(1)$ is generated. If $p(1)$ falls in $R(w)$, second-period production is $\overline{q}$, and third-period production is determined by whether or not $p(2)$ belongs to the favourable set $\overline{R}$. If instead $p(1)$ falls in $R^c(w)$, second-period production is $\underline{q}$, and third-period production is $\underline{q}$ if $p(2)$ is in $\underline{R}$, or $\overline{q}$ otherwise. In any period $t > 2$, either the favourable regime is in effect (firms are producing $\overline{q}$ and enjoying relatively high expected profits) or an unfavourable regime is operative (firms produce $\underline{q}$, a higher, less profitable quantity except in the degenerate case where $\overline{q} = \underline{q}$). Associated with these states are the reward regions $\overline{R}$ and $\underline{R}^c$, respectively. The regime in period $t + 1$ is favourable if and only if $p(t)$ lies in the reward region associated with the regime of period t. Hence, $\hat{\sigma}(w)$ is described entirely by the quantities $Q(w)$, $\overline{q}$, and $\underline{q}$, and the reward regions $R(w)$, $\overline{R}$, and $\underline{R}^c$. Of course if w is $\overline{v}$ or $\underline{v}$, <u>only two quantities and regions are involved</u>. A probability measure is induced on price space by equilibrium behaviour (given some history); the probability measures of the sets $\overline{R}$ and $\underline{R}$ give the probabilities of remaining in the reward and punishment states, respectively.

The next result concerns the sensitivity of maximal symmetric profits to changes in the discount rate. Since this requires reference to two different sets V (one for each discount rate) and two functions B, we now make the dependence on δ explicit by writing $V(\delta)$, $\overline{v}(\delta)$, $\underline{v}(\delta)$, and $B(W;\delta)$.

<u>Proposition 5</u>: (Monotonicity in Discount Factor). Let δ_1 and δ_2 be two discount factors such that $0 < \delta_1 < \delta_2 < 1$. Then:

$$(1)\ \frac{1 - \delta_2}{\delta_2} \cdot \overline{v}(\delta_2) > \frac{1 - \delta_1}{\delta_1} \cdot \overline{v}(\delta_1)$$

$$(2)\ \frac{1-\delta_2}{\delta_2}\cdot(\overline{v}(\delta_2)-\underline{v}(\delta_2)) > \frac{1-\delta_1}{\delta_1}\cdot(\overline{v}(\delta_1)-\underline{v}(\delta_1))$$

That is, $\overline{v}(\delta)$ and the length of the interval $co(V(\delta))$ are increasing in δ, even when discounted profits are "normalized" by the factor $(1 - \delta)/\delta$.

Proof: Define

$$k = \frac{\delta_2 - \delta_1}{\delta_1(1-\delta_2)}, \quad \overline{\alpha} = k\cdot\overline{v}(\delta_1),\ \underline{\alpha} = k\cdot\underline{v}(\delta_1),$$

$$W = \{\underline{v}(\delta_1) + \underline{\alpha},\ \overline{v}(\delta_1) + \overline{\alpha}\},\text{ and } W_1 = co(W).$$

For each $\gamma \in R$, let $[\gamma]$ denote the constant function from Ω into $\{\gamma\}$. It is easy to check that if (q,u) is admissible with respect to $V(\delta_1)$, $(q,u + [\gamma])$ is admissible with respect to W_1 for all $\gamma \in [\underline{\alpha},\overline{\alpha}]$. Let $(\overline{q},\overline{u})$ be admissible with respect to $V(\delta_1)$ and satisfy $E_1(\overline{q}\cdot e_N;\overline{u},\delta_1) = \overline{v}(\delta_1)$. Then $(\overline{q},\overline{u} + [\overline{\alpha}])$ is admissible with respect to W_1 and

$$\begin{aligned} E_1(\overline{q}\cdot e_N;u + [\overline{\alpha}],\delta_2) &= \delta_2\left\{\overline{\pi}_1(\overline{q}\cdot e_N) + \int_\Omega \overline{u}(p)\cdot g(p;N\cdot\overline{q})\,dp + \overline{\alpha}\right\} \\ &= \delta_2\left(\frac{\overline{v}(\delta_1)}{\delta_1} + \overline{\alpha}\right) \\ &= \overline{v}(\delta_1) + \frac{\delta_2 - \delta_1}{\delta_1}\cdot\overline{v}(\delta_1) + \delta_2\cdot\overline{\alpha} = \overline{v}(\delta_1) + \overline{\alpha}. \end{aligned}$$

By a similar argument, there exists $(\underline{q},\underline{u})$ admissible with respect to $V(\delta_1)$ such that $E_1(\underline{q}\cdot e_N;\underline{u},\delta_1) = \underline{v}(\delta_1)$ and $E_1(\underline{q}\cdot e_N;\underline{u} + [\underline{\alpha}],\delta_2) = \underline{v}(\delta_1) + \underline{\alpha}$. Thus $W \subseteq B(W_1;\delta_2)$. By Proposition 4, $B(W_1;\delta_2) = B(W;\delta_2)$. Hence, by Proposition 1, $W \subseteq V(\delta_2)$, and

$$\overline{v}(\delta_2) > \overline{v}(\delta_1) + \overline{\alpha} \qquad \text{and} \qquad \underline{v}(\delta_2) < \underline{v}(\delta_1) + \underline{\alpha}\,.$$

The reader may verify that these inequalities imply (1) and (2) above. Q.E.D

6. Conclusion

This paper shows that every S.S.E. payoff in the Green-Porter model [5],[8] can be supported by sequential equilibria that are easily described, and are extremely simple in their intertemporal structure. In every period of the optimal S.S.E. (and in all but the first period of <u>any</u> S.S.E.), firms produce according to one of two "regimes", corresponding to the first periods of the best and worst S.S.E.'s, respectively. The alternation between the two regimes resulting from equilibrium behaviour is a Markov process. Potential complexities such as employing varying severities of "punishment" depending on how far a price realization is from its "normal" range, or making the set of prices that trigger a change from the favourable to the unfavourable regime depend on recent price history, <u>do not arise</u>. These and other results are derived by reducing the repeated game to a static structure that is far more accessible for the purpose of analysis. This technique also enables us to work with much weaker assumptions than those invoked in earlier papers on this subject. Apart from minor regularity conditions, the cost and demand functions (for a given realization of the random variable θ) are unrestricted, and the stochastic disturbance to demand is modelled in quite a general way.

While the paper has focused exclusively on the Green-Porter model, it demonstrates an approach to the study of repeated games that might be of broader interest. The factorization and self-generation properties of Section 3 in particular have analogues in many classes of supergame. Such results are available, for example, for the simple strategy profiles of Abreu [1], and yield an alternative proof of the existence of an optimal simple penal code. Adaptation of the reduction function B to particular strategic situations should permit productive analysis of many repeated games heretofor considered intractable.

Note on Mixed Strategies

Working with behaviour strategies rather than mixed strategies involves no loss of generality (see Aumann's extension [3] of Kuhn's Theorem). Consequently it suffices to study the pure strategy equilibria of the supergame $H^\infty(\delta)$, having component game $H = (M,\ldots,M;\overline{\pi}_1,\ldots,\overline{\pi}_N)$, where an element of M is a probability distribution on the set $\tilde{S}$ of Section 2, and the original functions $\overline{\pi}_i$, $i=1,\ldots,N$, are extended to M by an expected profit calculation. The self-generation and factorization propositions of Section 3 apply immediately to $H^\infty(\delta)$, and with some additional arguments Proposition 3 (and hence Corollary 2) can also be established. The proof of the bang-bang result depends upon there being only a finite number of incentive compatibility constraints defining the set $\mu(q;W)$ (see Proposition 3). Since a strategy in H can be no more profitable than the most lucrative strategy giving weight to only one quantity in $\tilde{S}$, there are still only a finite number of constraints. Thus Proposition 4 is unchanged by the introduction of randomizations over production levels. This implies that the values of all S.S.E.'s of $H^\infty(\delta)$ are generated by S.S.E.'s using only two elements of M after period 1. The two elements are those used in the best and worst S.S.E.'s of $H^\infty(\delta)$, respectively.

REFERENCES

[1] Abreu, Dilip: "Repeated Games with Discounting: A General Theory and an Application to Oligopoly", mimeo (1982), Princeton.

[2] Artstein, Zvi: "Discrete and Continuous Bang-Bang and Facial Spaces or: Look for the Extreme Points", SIAM Review, 22 (1980), 172-185.

[3] Aumann, Robert: "Mixed and Behavior Strategies in Infinite Extensive Games", in Advances in Game Theory, ed. by M. Dresher, L. S. Shapley and A. W. Tucker, Princeton: Princeton University Press, 1961.

[4] Friedman, J.W.: "A Non-cooperative Equilibrium for Supergames", Review of Economic Studies, 28 (1971), 1-12.

[5] Green, Edward J. and Robert H. Porter: "Noncooperative Collusion Under Imperfect Price Information", Econometrica, 52 (1984), 87-100.

[6] Harris, Christopher: "A Note on the Existence of Optimal Simple Penal Codes", mimeo (1984), Nuffield College, Oxford.

[7] Kreps, David M. and Robert Wilson: "Sequential Equilibria", Econometrica, 50 (1982), 863-894.

[8] Porter, Robert H.: "Optimal Cartel Trigger-Price Strategies", Journal of Economic Theory, 29 (1983), 313-338.

[9] Radner, Roy: "A Repeated Principal-Agent Game with Discounting", mimeo (1981), Bell Laboratories, Murray Hill, N.J.

[10] Royden,H.L.: Real Analysis, (1968), Macmillan, New York.