R M #228

NON-COOPERATIVE PRICE TAKING

IN LARGE DYNAMIC MARKETS

Edward J. Green*
Economics Department
Princeton University

July 1978
Revised January 1979

* Financial support for this research was provided through a grant from the Sloan Foundation to the Princeton University Economics Department. I would like to thank Robert Anderson and Hugo Sonnenschein for valuable discussions. Special thanks are owed to Edward Prescott, who originally stated as conjectures several of the results presented in this paper.

NON-COOPERATIVE PRICE TAKING
IN LARGE DYNAMIC MARKETS

by

Edward J. Green
Department of Economics
Princeton University

1. Introduction

The question of why agents in large economies exhibit price-taking behavior has been important throughout the history of mathematical economics. Recent work has treated this question within the formal context of the Arrow-Debreu economy, in which the set of agents is either finite or a continuum. This work has shown roughly that, for some cooperative (e.g., the core) or non-cooperative (Cournot-Nash) equilibrium concept, if standard convexity and continuity conditions hold, then

(a) (Inclusion principle for non-atomic economies.) If the set of agents is a continuum, then every equilibrium allocation is a Walras allocation, and

(b) (Limit principle for finite economies.) If a sequence of increasingly large finite economies converges in a natural way, and if their equilibrium allocations have a limit, then the limit is a Walras allocation of the limit economy.

Actually, although the theory of the core is well developed[1/], the current results about non-cooperative equilibria of large Arrow-Debreu economies is less complete and less unified. One explanation for this situation is that, while a cooperative game can be defined from the feasible allocations of an economy and its sub-economies, the Arrow-Debreu model must be augmented to define a non-cooperative game and it is not clear how this ought to be done. Nevertheless, for an exchange economy with a continuum of traders, Mas Colell [5] has proved an inclusion theorem of type (a) which will apply to any non-cooperative game satisfying some apparently reasonable conditions. Limit theorems in the spirit of (b) for finite economies are due to Gabszewicz and Vial [3], Novshek and Sonnenschein [6], Postlewaite and Schmeidler [8], and Roberts and Postlewaite [9].[2/]

The present paper also is a study of the non-cooperative equilibria of large economies. However, an economy will be represented not as a set of Arrow-Debreu markets which clear simultaneously, but as an infinite sequence of markets which clear consecutively. This latter representation will be called a dynamic economy, and is to be distinguished from the corresponding Arrow-Debreu economy with infinitely many forward markets. In addition to the Cournot equilibria of their Arrow-Debreu counterparts, dynamic economies typically have other non-cooperative equilibria in which, by means of strategic threats, firms induce one another to act collusively. Because collusion can be non-cooperatively enforced even in some large dynamic economies, price-taking behavior does not necessarily characterize dynamic equilibrium. However, positive results along the lines of both (a) and (b) can be obtained. Interestingly, collusive equilibria are ruled out when firms are inconspicuous to one another, but not necessarily when they are insignificant in the usual sense of having negligible productive capacity relative to the market. This explicit dependence on considerations of market information distinguishes the dynamic theory from the results for Arrow-Debreu economies.

In section 3 a sequence of replica dynamic economies will be defined, each of which has a strategic non-cooperative equilibrium which supports the monopoly allocation at every time. This example refutes the limit principle for dynamic economies. An example is also given in which, if each firm can observe the output level of every other firm, then such a collusive non-cooperative equilibrium exists for a dynamic economy with a continuum of firms. Thus the most general version of the inclusion principle is refuted as well.

In the counterexample to the inclusion principle, firms are required to have exhaustive non-price information about their competitors. In section 4, it is proved that all equilibria which violate the principle require such information. This is done by interpreting dynamic economies within an abstract theory of repeated games, and providing a sufficient condition for all Nash equilibria of a repeated game to be constructible from Nash equilbria of the static game on which it is based. This condition yields an inclusion theorem for dynamic economies in which agents have information only about price and other market aggregates.

The inclusion principle derived in section 4 sheds light on the failure of the limit principle for dynamic economies. The principle must fail because there is information about individual producers which the price mechanism conveys in every finite economy, but not in a non-atomic economy. In fact, in a finite economy price information can reveal to a firm that a competitor has disregarded its strategic threat, although the identity of the violator is indeterminate. If the price varies because of random demand, its informativeness about the supply side of the market should be reduced. As the scale of an individual producer becomes arbitrarily small relative to market demand in a sequence of replica dynamic economies, any stochastic demand disturbance should make strategic threats unenforceable in the limit. This is proven in section 5, again in an abstract game-theoretic context. It yields a limit theorem for dynamic economies with demand uncertainty.

In section 6 some remarks are made about the relevance of these results to anti-trust policy.

2. Non-cooperative equilibria which support the monopoly allocation.

The first inclusion and limit theorems were derived by Cournot in a static partial equilibrium setting. These will be reviewed in this section, and counterexamples to their dynamic versions will be given in the next. These results may be reformulated in a general equilibrium setting, building on the work of Gabszewicz and Vial [3]. The positive results of this paper are immediate consequences of game theoretic results which apply straightforwardly to the general equilibrium version.

First, static and dynamic markets are defined. Intuitively, a static market consists of a set of firms, each of which supplies the market good at a total cost which depends on the quantity it produces, and an inverse demand function which determines the price as a function of mean supply. Implicitly, this inverse demand function is determined by the actions of price-taking consumers. The question to be investigated is: are profit-maximizing firms also price-takers when there are many firms?

A (stationary) dynamic market is a static market which operates repeatedly, its times of operation being indexed by the natural numbers. Firms maximize the discounted present value of their profits.[3/] These definitions are now presented formally:

Definition: A cost function is an element of the set $Y = \{f \mid f: R_+ \to R_+ \cup \{\infty\},$ f is continuous$\}$. A normal inverse demand function is a continuous function $D: R_+ \cup \{\infty\} \to R_+$ with $D(\infty) = 0$. A static market is an ordered 4-tuple $\langle K, \mu, Y, D \rangle$

where K is an abstract set of firms, μ is a probability measure on K, $y:K\to Y$ specifies the cost function of each firm, and D is a normal inverse demand function. A (stationary) dynamic market is an ordered 5-tuple $\langle K,\mu,y,D,\beta\rangle$, of which the first four components specify a static market and $\beta\epsilon(0,1)$ is the market discount factor. Let M be a market (static or dynamic) with measure space $\langle K,\mu\rangle$ of firms. A supply vector for M is a bounded μ-measurable function $q:K\to R_+$.

Given a supply vector q for M, mean supply will be $\int_K q(k)d\mu$ and the market price will $D(\int_K q(k)d\mu)$. If a firm $k\epsilon K$ has cost function $y(k)=f$, then a supply vector determines a net profit level for k . This net profit will be denoted by $\pi_k(q)$, and satisfies the equation

$$(1) \quad \pi_k(q) = q(k)D(\int_K q(i)d\mu) - f(q(k)) \ .$$

In selecting its output level, either a firm may accept the market price as an exogenous parameter of its decision problem, or it may recognize its own influence on the market price. A price-taking equilibrium results when all firms act in the former way, and a non-cooperative equilibrium results when they act in the latter way.

Definition: A price-taking equilibrium of the static market $\langle K,\mu,y,D\rangle$ is a supply vector q such that, for almost all (w.r.t. μ) $k\epsilon K$,

$$(2) \quad \pi_k(q) = \max_{r\epsilon R_+}[rD(\int_K q(i)d\mu) - (y(k))(r)] \ .$$

A price-taking equilibrium of the dynamic market $\langle K,\mu,y,D,\beta\rangle$ is a sequence $\langle q_t\rangle_{t\epsilon N}$ of supply vectors such that, for almost all $k\epsilon K$,

$$(3) \quad \sum_{t\epsilon N} \beta^t \pi_k(q_t) = \max\{ \sum_{t\epsilon N} \beta^t [r_t D(\int_K q_t(i)d\mu) - (y(k))(r_t)] \,|\, \langle r_t\rangle_{t\epsilon N} \epsilon R_+^N\} \ .$$

A <u>non-cooperative equilibrium</u> of the static market $\langle K,\mu,y,D\rangle$ is a supply vector q^* such that, for almost all $k\in K$,

$$(4) \quad \pi_k(q^*) = \max_{q\in Q}[q(k)D(\textstyle\int_K q(i)d\mu) - (y(k))(q(k))],$$

where $Q = \{q \mid q$ is a supply vector and $q(i) = q^*(i)$ for all $i \neq k\}$.

To define non-cooperative equilibrium for a dynamic market, the notion of a strategy must be introduced. A strategy of a firm is a rule which, at each time, determines an output level for the firm as a function of information which is available to the firm at that time. A strategy vector is an assignment of strategies to firms. Thus, if S is the space of strategies, a strategy vector is a measurable function $f:K\to S$. A non-cooperative equilibrium of a dynamic market is a strategy vector f such that

(a) for almost all $k\in K$, $f(k)$ yields as high a discounted present value of returns (given that the other firms are using the strategies assigned to them by f) as would any other strategy $s\in S$, and

(b) Clause (a) will continue to hold at every future time for almost all $k\in K$, regardless of firms' information at that time.[4]

This rather informal definition of dynamic non-cooperative equilibrium will be adequate to verify that the strategy vectors to be discussed in this section are equilibria. A more explicit definition will be given in section 4. In the remainder of this section, straightforward generalizations of Cournot's theorems for static markets will be stated, and counterexamples to their dynamic versions will be constructed. Cournot's inclusion theorem can be stated immediately:

Theorem 1: Every non-cooperative equilibrium of a non-atomic static market is a price-taking equilibrium.

To facilitate the statement of Cournot's limit theorem, the notion of a sequence of replica markets is introduced. Intuitively, such a series is formed by starting with a finite market, and at each stage adding a new "clone" of each of the original market agents (firms and consumers). This is done explicitly for firms, but must be done implicitly for consumers. Since there are n times as many consumers of each type in the n^{th} market as in the first, aggregate demand in the n^{th} market at any price is n times what it is in the first. However, since the inverse demand function is defined formally in terms of mean quantity rather than of aggregate quantity, and since there are n firms in the n^{th} market for every one in the first, all markets in the sequence should share the same inverse demand function. As replication continues, the members of the sequence should resemble the non-atomic representation of the market in their equilibrium behavior. The sequences are now defined formally, after which Cournot's limit theorem is immediately stated.

Definition: A sequence of replica (static or dynamic) markets is an infinite sequence $\langle M_n \rangle_{n \varepsilon N}$ of markets such that the following hold for some finite set $K, y{:}K \to Y, \beta \varepsilon (0,1)$ and normal inverse demand function D. For all n, $M_n = \langle K_n, \mu_n, y_n, D, \beta \rangle$, where

(5) $K_n = K \times \{0 \ldots n\}$

(6) $\mu_n(B) = \#B / \#K_n$ for all $B \subseteq K_n$.

(I.e., μ_n is normalized counting measure.)

(7) $y_n(\langle k,m\rangle) = y(k)$ for $k\varepsilon K$, $m\leq n$.

The market $M_\infty = \langle [0,1]\times K, \mu_\infty, y_\infty, D, \beta\rangle$ is the <u>non-atomic representation</u> of the sequence if

(8) $\mu_\infty(B) = \sum_{k\varepsilon K} \lambda(\{r\varepsilon[0,1] \mid \langle r,k\rangle\varepsilon B\})$ for all $B\subseteq[0,1]^\times K$.

(I.e., μ_∞ is Lebesgue measure on the copies of $[0,1]$ in K_∞.)

(9) $y_\infty(\langle r,k\rangle) = y(k)$ for all $\langle r,k\rangle\varepsilon K_\infty$.

<u>Theorem 2</u>: Suppose that $\langle M_n\rangle_{n\in N}$ is a sequence of replica static markets, that the supply vector q_n is a non-cooperative equilibrium of M_n for every n, and that some subsequence of $\langle\langle q_n, y_n\rangle\rangle_{n\in N}$ converges in distribution.[5/] Then the subsequence has a limit $\langle q_\infty, y_\infty\rangle$, and q_∞ is a price-taking equilibrium of the non-atomic representation M_∞.

3. <u>Non-cooperative equilibria which support the monopoly allocation (continued)</u>.

Counterexamples to the analogues of theorems 1 and 2 for dynamic markets are constructed by exhibiting non-cooperative equilibria of large markets, in which the monopoly price is maintained by mutual strategic threats among the firms. The existence of these equilibria in markets with finitely many firms was first pointed out by James Friedman [2]. A counterexample is provided to the analogue of theorem 2, simply by noting that such an equilibrium may exist for every market in a sequence of replica dynamic markets. Such a sequence is easy to construct. It requires only one type of firm (K is a singleton $\{k\}$) with a production function $y(k) = f$ described by

(10) $f(r) = \begin{cases} 0, & \text{if } r \leq 2/3 \\ r-2/3, & \text{if } r>2/3 \end{cases}$

Price is specified by the inverse demand function

(11) $D(r) = \begin{cases} 1-r, & \text{if } r \leq 1 \\ 0, & \text{if } > 1 \end{cases}$

The discount factor for the sequence is $\beta = 0.9$. The counterexample to the limit principle is now presented.

<u>Theorem 3</u>: There exist a sequence of replica dynamic economies $\langle M_n \rangle_{n\varepsilon N}$ and a double sequence of supply vectors $\langle q_{nt} \rangle_{n,t\varepsilon N}$ such that

(a) For each n, the sequence of supply vectors $\langle q_{nt} \rangle_{t\varepsilon N}$ is determined (in the sense that q_{nt} specifies firms' output levels at time t) by a non-cooperative equilibrium of M_n, and

(b) For each t there is a supply vector $q_{\infty t}$ such that $\lim_{n\to\infty} q_{nt} = q_{\infty t}$ in distribution, but

(c) $\langle q_{\infty t} \rangle_{t\varepsilon N}$ is not a price-taking equilibrium of the non-atomic representation M_∞ of the sequence.

Proof: Let $\langle M_n \rangle_{n\varepsilon N}$ be determined by the conditions that $K = \{k\}$, $y(k)=f$ defined by (10), D is determined by equation (11), and $\beta = 0.9$. Recall that a strategy for a firm in a dynamic market is a rule which, at each time, determines an output level on the basis of information then available to the firm. The double sequence $\langle q_{nt} \rangle_{n,t\varepsilon N}$ will specify the output decisions of firms which follow the strategy

(d) Produce at level 1/2 (the monopoly level) at time 0.

(e) If $D(\int_{K_n} q_{n0} d\mu_n) = \ldots = D(\int_{K_n} q_{n(t-1)} d\mu_n) = \frac{1}{2}$. Then produce at level 1/2 at time t.

(f) If the condition of (e) is not satisfied at t, then produce at level 2/3 (the output level corresponding to the unique non-cooperative equilibrium of the static market $\langle K_n, \mu_n, y_n, D\rangle$).

By induction on $t, q_{nt} \equiv 1/2$ for all n and t. Therefore condition (b) is satisfied with $q_{\infty t} \equiv 1/2$ for all t. Condition (c) is also satisfied, since $D(\int_{K_\infty} q_{\infty t} d\mu_\infty) = 1/2$ while firms have zero marginal cost. Condition (a) is equivalent to the statement, proved below as lemma 1, that for all firms to follow the strategy defined by (d), (e), and (f) is a non-cooperative equilibrium of M_n. Q.E.D.

<u>Lemma 1</u>: For all firms to employ the strategy defined by (d), (e), and (f) above is a non-cooperative equilibrium of the dynamic market M_n.

Proof: Intuitively, each firm makes a strategic threat to the others. It announces: "Initially, we will participate in a cartel in which all firms produce equal shares of the monopoly output $n/2$. We will do our part to maintain this cartel, as long as there is reciprocity. However, if at any time a lowering of the market price indicates that another firm is exceeding its production limit, we will thereafter protect ourselves by acting as static Cournot oligopolists."

Consider the prospects of a potential deviant firm, if all other firms make this threat. Suppose that it were to adopt a strategy which would lead it, in this environment, to select an output sequence $\langle r_t\rangle_{t\in N}$. It will be shown that $r_t = 1/2$ for all t is optimal. Because the strategy (d), (e), (f) produces this sequence if all other firms follow (d), (e), (f), part (a) of the definition of non-cooperative equilibrium for a dynamic market is satisfied.

To show this, consider first an output sequence $\langle r_t \rangle_{t \in N}$ with $r_t \neq 1/2$ for some t. Let $t_o = \min\{t | r_t \neq 1/2\}$. If other firm's outputs are determined by (d), (e), and (f), then the market price at time t is

$$(12) \quad D(\int_{K_n} q_{nt} d\mu_n) = \begin{cases} 1/2 & \text{for } t<t_o \\ (n+1-2r_t)/2n & \text{for } t=t_o \\ (n+2-3r_t)/3n & \text{for } t_o<t \end{cases} .$$

Since the outputs of other firms are the same for all sequences $\langle r'_t \rangle_{t \in N}$ which satisfy $t_o = \min\{t | r'_t \neq 1/2\}$, r_t must be the output level which maximizes returns at time t for $t_o \leq t$. I.e., by (12),

$$(13) \quad r_{t_o}(n+1-2r_{t_o})/2n - f(r_{t_o}) = \max_{r \in R_+} r(n+1-2r)/2n - f(r) \text{ , and}$$

$$(14) \quad r_t(n+2-3r_t)/3n - f(r_t) = \max_{r \in R_+} r(n+2-3r)/3n - f(r) \text{ for } t_o<t .$$

Solving equations (13) and (14) yields $r_t = 2/3$ for all $t \geq t_o$. Thus discounted profits are

$$(15) \quad \sum_{t \in N} (0.9)^t [r_t D(\int_{K_n} q_{nt} d\mu_t) - f(r_t)] \leq \sum_{t<t_o} (0.9)^t [(1/2)(1/2)-0]$$

$$+ (0.9)^{t_o}[(2/3)(1/2)-0] + \sum_{t_o<t} (0,9)^t [(2/3)(1/3)-0]$$

$$= (\sum_{t<t_o} (0.9)^t)/4 + (0.9)^{t_o}(7/3) .$$

On the other hand, if $r'_t = 1/2$ for all t and if $\langle q'_{nt} \rangle_{t \in N}$ is the sequence of supply vectors which are realized when the firm in question supplies $\langle r'_t \rangle_{t \in N}$, then the firm's discounted profits are

$$(16)\quad \sum_{t \in N} (0.9)^t [r'_t D(\int_{K_n} q'_{nt} d\mu_n) - f(r'_t)] = \sum_{t<t_o} (0.9)^t [(1/2)(1/2)-0]$$

$$+ \sum_{t_o \leq t} (0.9)^t [1/2)(1/2)-0] = (\sum_{t<t_o} (0.9)^t)/4 + (0.9)^{t_o}(5/2) .$$

Together, (15) and (16) establish the optimality of the supply sequence $r'_t \langle r'_t \rangle_{t \in N}$ for the firm. Thus part (a) of the definition of non-cooperative equilibrium is established. To show that part (b) of the definition holds, two cases must be considered at any future time in question. One possibility is that the market price has always been 1/2 previously. In this case, the firm faces the same decision problem as it did in the initial period, so the earlier part of the proof is still relevant. The other possibility is that the firm has received information that some previous price has been other than 1/2 (N.B. Although this possibility will not be realized in equilibrium, it must be considered for the reason explained in footnote 4). Then, since other firms will have this information also, they will always produce the Cournot output 2/3 in the future. Thus it is optimal for the firm under consideration to produce this output in response. Q.E.D.

In a non-atomic market, the output decision of a single firm does not affect the price. Thus a firm cannot threaten its competitors by making its future actions contingent on the present price, because each competitor takes the price to be exogenous to its decision. Another description of this situation is that, if enforcement of a cartel is based only on price information, any firm can be a free rider on the cartel without being discovered. However, firms may be capable of observing the output levels of their competitors directly. Even in a large market, this output information might be supplied by a trade association or by the government. If the information is available, the limit principle may not hold.

Theorem 4: There exists a non-atomic dynamic market in which, if firms observe the output levels of their competitors directly, there is a non-cooperative equilibrium which is not a price taking equilibrium.

Proof: Let the market be the non-atomic representation of the sequence in theorem 3, and let each firm produce 1/2 unit of output as long as no competitor exceeds that output, but 2/3 unit if some competitor has produced more than 1/2 unit at some previous time. The proofs of theorem 3 and lemma 1 extend straightforwardly to show that these strategies are in non-cooperative equilibrium, but that each firm will produce 1/2 unit in every period, while price taking firms would produce 2/3 unit. Q.E.D.

4. An inclusion theorem for repeated games

The examples presented in the last section show that some justification is needed for applying static inclusion and limit theorems to actual markets. The question of when those results may legitimately be applied is the topic of this section and of the next. In these two sections, conditions will be given which insure that the opportunities for repeated trade in a large dynamic economy will be irrelevant to the decisions taken by agents at any particular time. Under these conditions, all dynamic non-cooperative equilibria will be sequences of static equilibria. Therefore, facts about a static market will be true of the corresponding dynamic market as well.

These conditions will be studied in an abstract game theoretic context. This approach has three virtues. First, it makes it easy to deal with issues about agent's information (which have already appeared in Theorem 4). Second, at no extra cost it yields results which extend most of the literature on static price taking behavior to dynamic situations. For instance, the condition under which the dynamic inclusion principle holds is one already imposed by Mas Colell in his theorem. Third, because the issues raised here are relevant whenever agents repeatedly face situations with "prisoner's dilemma" characteristics, the game theoretic results are of general interest. The theory of repeated games to be used here will now be presented. Then, the interpretation of a dynamic market as a repeated game will be explained.

Intuitively, a repeated game is simply a game which is replayed countably many times. A repeated game is defined in terms of three underlying spaces. The players are specified by a measure space $\langle K,\mu \rangle$. The measure μ will always be normalized to be a probability measure. For economy of notation, the σ-algebra on which μ is defined will not be referred to explicitly.

The actions available to a player are elements of a set A. Although it is assumed here that all players have the same set of feasible actions, this restriction could easily be removed.

The possible outcomes of the game at any time form a set X. An element of X specifies the information shared by all players about the results of a play of the game. For instance, in an economy the elements of X might be price vectors. It is important to note that this use of the term "outcome" is slightly different from the usual one, which is that an outcome completely specifies the result of a play (e.g., in an economy, X would be the set of feasible allocations) including aspects which may be unobservable to some of the players (e.g., other players' components of the allocation vector).

Throughout the paper, A and X will be assumed to be complete separable metric spaces, and mappings into these spaces from measurable spaces will be assumed to be Borel measurable.

Each player k is characterized by a return function $u_k: A \times X \to R$, and a discount factor $\beta_k \varepsilon (0,1)$. A player attempts to maximize the discounted value of his returns from repeated plays of the game.[6/]

A play q of the game is a measurable assignment of actions to players.[7/] A^K will denote the set of plays. The outcome of a play is determined by the outcome function $F: A^K \to X$.

A strategy s is a rule by which a player determines an action to take at each time, based on his knowledge at that time of the outcomes of past plays, and of his own past actions. If each past action has been generated as a deterministic function of prior outcomes and actions, the player can recover the information about his past actions recursively from his memory of past outcomes (this is easily proven by recursion on t). If the successive plays of the game are indexed by the natural numbers, a player's information at time t about prior history is specified by an element of X^t . Thus, s may be represented by a sequence $<s_t>_{t \varepsilon N}$, where $s_o \varepsilon A$, and $s_t: X^t \to A$ for $t \varepsilon N_+$. The space of strategies is then $A \times \prod_{t \varepsilon N^+} A^{X^t}$, which will be denoted by S .

A strategy vector is an assignment $<s^k>_{k \varepsilon K}$ of strategies to players such that, if $s^k = <s_t^k>_{t \varepsilon N}$, then

(a) s_o^k is a measurable function of k , and

(b) for $x_o \ldots x_{t-1} \varepsilon X$, $s_t^k(x_o, \ldots, x_{t-1})$ is a measurable function of k.

These conditions insure that players' strategies determine recursively a sequence of plays and outcomes. The initial play q_o is the vector $<s_o^k>_{k \varepsilon K}$ of initial actions specified by the strategies. The initial outcome is $x_o = F(q_o)$. The next play is $q_1 = <s_1^k(x_o)>_{k \varepsilon K}$, and $x_1 = F(q_1)$, and so forth. This process defines an infinite sequence $<<q_t, x_t>>_{t \varepsilon N}$ of plays and outcomes, called the history generated by the strategy vector. The sequence $<x_t>_{t \varepsilon N}$ determined by the history is called the path.

The interpretation of dynamic markets as repeated games is straightforward.

outcomes are prices. The return function of firm k specifies its profit as a function of its output level and the price. I.e., if f is the total cost function of firm k, $u_k(a,x) = xa - f(a)$. The market discount rate β is β_k for each firm. The outcome function is $F(q) = D(\int_K q d\mu)$.

Now it is possible to make precise the informal definition of non-cooperative equilibrium which was given in section 2. Let S denote the space of strategies of the game. I.e., $S = A \times \prod_{t \in N_+} A^{X^t}$. Also, let $y: K \to R^{(A \times X)} \times (0,1)$ specify the characteristics of players. I.e., $y(k) = \langle u_k, \beta_k \rangle$. Three related concepts will now be defined. A static non-cooperative equilibrium is a play which is a Nash equilibrium of the simple (non-repeated) game. A dynamic quasi-equilibrium is a strategy vector which satisfies the defining condition of Nash equilibrium with respect to discounted returns at the initial time. A dynamic non-cooperative equilibrium is a perfect (in the sense of footnote 4) dynamic quasi-equilibrium.

Definition: A static non-cooperative equilibrium of the repeated game $\langle K,A,X,F,\mu,y \rangle$ is a play $q^* \varepsilon A^K$ such that, for any other $q \varepsilon A^K$, the following implication holds for almost all $k \varepsilon K$: If $y(k) = \langle u,\beta \rangle$ and $q(i) = q^*(i)$ for all $i \neq k$, then

(17) $u(q(k),F(q)) \leq u(q^*(k),F(q^*))$.

A dynamic quasi-equilibrium is a strategy vector $v^* \varepsilon S^K$ such that, for almost all $k \varepsilon K$, the following implication holds for all $v \varepsilon S^K$: If $y(k) = \langle u,\beta \rangle, v(i) = v^*(i)$ for all $i \neq k$, and $\langle\langle q_t, x_t \rangle\rangle_{t \varepsilon N}, \langle\langle q_t^*, x_t^* \rangle\rangle_{t \varepsilon N}$ are the histories generated by v and v^* respectively, then

(18) $\sum_{t \varepsilon N} \beta^t u(q_t(k), x_t) \leq \sum_{t \varepsilon N} \beta^t u(q_t^*(k), x_t^*)$.

A dynamic non-cooperative equilibrium is a strategy vector v^* such that (a) v^* is a dynamic quasi-equilibrium, and (b) if $t \varepsilon N_+$, and $x_0^*, \ldots, x_{t-1}^* \varepsilon K$, and for each player k a new strategy is defined by

(19) $s_0 = s_t^*(x_0^*, \ldots, x_{t-1}^*)$

and, for $r \varepsilon N_+$,

(20) $s_r(x_o,\ldots,x_{r-1}) = s^*_{t+r}(x^*_o,\ldots,x^*_{t-1},\ldots,x_o,\ldots,x_{r-1})$,

then the vector of these new strategies is also a dynamic quasi-equilibrium.

It was asserted at the end of section 3 that an inclusion theorem holds for dynamic markets in which firms cannot discern changes in the output level of individual competitors. For this condition to be realized in a non-atomic market it is sufficient for the publicly observable outcome of the operation of the market to depend only on the distribution of firm's output level. This property of the outcome function can be generalized to abstract games.

<u>Definition:</u> Denote the set of probability measures on A by $M(A)$. Let $m: A^K \to M(A)$ map plays to the distribution of actions which they determine as random variables. I.e., for $q \varepsilon A^K$ and $B \subseteq A$, $(m(q))(B) = \mu(\{k \mid q(k) \varepsilon B\})$. A game with outcome function F is <u>anonymous</u> if $F(q)$ depends on q only through $m(q)$ (i.e., if there is a function $G: M(A) \to X$ which makes the diagram

(21) $$\begin{array}{ccc} & M(A) & \\ {}^{m}\nearrow & & \searrow^{G} \\ A^K & \xrightarrow[F]{} & X \end{array}$$

commute.

The preceding definitions make it easy to state and prove the result described at the beginning of this section, relating static and dynamic non-cooperative equilibria. The inclusion theorem for dynamic markets with price information is an immediate corollary of this result.

<u>Theorem 5:</u> If $v^* \varepsilon S^K$ is a dynamic non-cooperative equilibrium of the anonymous repeated game $<K,A,X,F,\mu,y>$, μ is non-atomic, and $<<q^*_t x^*_t>>_{t \varepsilon N}$ is the history

generated by v^*, then every q^*_t is a static non-cooperative equilibrium of the game.

Proof: It will be shown that, if some q^*_r is not a static NCE, then v^* is not a dynamic NCE. If q^*_r is not a static NCE, there is a subset $B \subseteq K$ with $\mu(B) > 0$, such that for every $k \in B$ (with $y(k) = < u,\beta>$) there exists a play $q \in A^K$ with, $q(i) = q^*_r(i)$ for $i \neq k$, such that (17) fails to hold. I.e., $u(q^*_r(k), F(q^*)) < u(q(k), F(q))$. Now, for some arbitrary player $k \in B$, define a new strategy as follows. Let his initial action s_o be the initial action $q^*_o(k)$ specified by $v^*(k)$, if $r \neq 0$, but let it be $q_o(k)$ if $r = 0$. For every $t>0$, the function s_t will be a constant function. Everywhere on its domain X^t, the function will take the value $q^*_t(k)$ if $t \neq r$, or $q(k)$ if $t=r$. Define the strategy vector $v \in S^K$ by $v(k) = <s_t>_{t \in N}$ and $v(i) = v^*(i)$ for $i \neq k$. It is claimed that (18) fails to hold for v and v^*, and that v^* is therefore not a dynamic NCE because this failure of (18) can be exhibited anywhere on the set B which has positive measure. This claim is substantiated by noting that, if $<x_t>_{t \in N}$ is the path generated by v, $x_t = x^*_t$ (this fact is proved below as lemma 2). By construction of $v(k)$, $q_t(k) = q^*_t(k)$ except for $t=r$. Therefore, for $t \neq r$, the t^{th} summands of the left and right hand sides of (18) are identical. Consequently, (18) holds only if $u(q_r(k), x_r) \leq u(q^*_r(k), x^*_r)$. But this last inequality is equivalent to (17), which fails by assumption. Q.E.D.

<u>Lemma 2</u>: Let F be anonymous, $v^* \in S^K <x^*_t>_{t \in N}$ be the path generated by v^*, $r \in N$, and $a \in A$. For some $k \in K$ with $\mu\{k\} = 0$, define the strategy $<s_t>_{t \in N}$ as follows: Define $s_o = q^*_o(k)$ if $0 \neq r$, but $s_o = a$ if $0 = r$. For $t>0$, define s_t to be the constant function on X^t which is equal everywhere to $q^*_t(k)$ if $t \neq r$, but to a if $t = r$. Define $v \in S^K$ by $v(k) = <s_t>_{t \in N}$ and $v(i) = v^*(i)$ for $i \neq k$. Let $<x_t>$ be the path

generated by v. Then for all $t\epsilon N$, $i\epsilon K-\{k\}$, $q_t(i) = q_t^*(i)$ and $x_t = x_t^*$.

Proof: By induction on t. For $t = 0$, $q_t(i)$ is determined directly by $v^*(i)$ for all $i\neq k$. Since $\mu(\{k\}) = 0$, $m(q_o) = m(q_o^*)$, so $x_o = x_o^*$ because F is anonymous. Now suppose that the lemma holds for $0,\ldots,t$. Then for $i\neq k$, if $v^*(i) = \quad ,\langle s_t^*\rangle_{t\epsilon N}$

$$q_{t+1}(i) = s_{t+1}^*(x_o,\ldots,x_t) = s_{t+1}^*(x_o^*,\ldots,x_t^*) = q^*{}_{t+1}(i) .$$

As in the case $t=0$, $m(q_{t+1}) = m(q_{t+1}^*)$, so $x_{t+1} = x_{t+1}^*$. This completes the induction. Q.E.D.

Theorem 5 is an inclusion theorem, if this term is interpreted broadly to mean a characterization of the equilibria of non-atomic games. It describes the dynamic non-cooperative equilibria of a non-atomic anonymous iterated game in terms of the equilibria of the static game which is iterated. If the equilibria of the static game have some property, theorem 5 entails that the equilibria of the iterated game inherit the property. Thus, the theorem proved by Mas Colell in [5] extends immediately to dynamic economies. Theorem 1 of the present paper extends in a similar way.

Suppose that firms in a non-atomic dynamic market receive only price information about the operation of the market. Because price is a function of mean supply, the market is an anonymous game. By theorem 5, if $\langle\langle q_t,x_t\rangle\rangle_{t\epsilon N}$ is a dynamic NCE of the market, every q_t is a static NCE. By theorem 1, then, every q_t is a price taking equilibrium of the static market. It is evident from the definition of price taking equilibrium that the sequence $\langle q_t\rangle_{t\epsilon N}$ is therefore a price taking equilibrium of the dynamic market. This proves the inclusion theorem for dynamic markets with price information only. In fact, the theorem continues to hold if firms receive other information about market aggregates (e.g., the variance of output levels in the market). The significance of this observation will be considered in section 6.

5. A limit theorem for anonymous repeated games with random outcomes.

Although the inclusion principle is valid for dynamic markets under the hypothesis of anonymity, the limit principle remains invalid. This fact is evident from the counter-example given in the proof of theorem 3. The reason why the implications of anonymity are different for the two principles is obvious. In a non-atomic market, anonymity guarantees that no firm can have any effect at all on the information which its competitors receive about the market. If inverse demand is strictly downward sloping, though, a unilateral change of output level by a firm in a finite market must have a perceptible effect on the market price. Although the magnitude of this price change becomes small in large replica markets, the change remains perceptible. The limit principle fails because the individual firm's effect on price information does not diminish in large markets, although its effect on the price level shrinks. In order for the limit principle to be valid for a dynamic market, a situation must occur in which price information about supply is unreliable.

One such situation is the existence of random fluctuations, not directly observed by firms, in consumers' demand schedules. Suppose that firms in this situation attempt to maintain a strategic non-cooperative equilibrium like that described in lemma 1. When the market price falls, each firm must decide whether the decline reflects a spontaneous downward shift of the inverse demand function, or whether it was caused by a competitor having exceeded his output quota. In a large market where the scale of an individual firm is insignificant relative to aggregate demand uncertainty, price information cannot provide evidence about the firm's level of output. Just as in the non-atomic market, the firm has incentive to break the cartel agreement.

In this section, a limit theorem will be proven which applies to the situation just discussed. The theorem applies only to equilibria of the type studied in lemma 1, rather than to all dynamic non-cooperative equilibria. This is a preliminary theorem. A much more general one should be provable.

The limit theorem will be stated for anonymous repeated games with random outcomes. A game with random outcomes is one for which, at each play, the outcome is a random vector taking values in X, rather than a determinate value $x \in X$. The random vector which is the outcome depends only on the current play. Players' return functions are now interpreted as von Neumann-Morgenstern utility functions, and players are assumed to maximize expected discounted returns.

The definitions of a repeated game with random outcomes and of an anonymous game with random outcomes are the same as those given in the last section, except that now $F: A^K \to M(X)$ and $G: M(A) \to M(X)$. When plays have random outcomes, a strategy vector will not generate a unique path in X^N. Rather, it will determine a probability distribution over paths. If $B \subseteq X, x_o, \ldots, x_{t-1} \in X$, $\langle s^k \rangle_{k \in K}$ is a strategy vector, q is the play defined by $q(k) = s^k_t(x_o, \ldots, x_{t-1})$, and $\nu = F(q)$, then $\nu(B)$ is the probability that $x_t \in B$ conditional on the outcomes of the first t plays having been $x_o, \ldots, x_{t-1}$. In order for this conditional probability to be well defined, $\nu(B)$ must be a measurable function of $x_o, \ldots, x_{t-1}$. A sufficient condition for measurability of this function in the case of an anonymous game is that (a) G is continuous when M(A) and M(X) are endowed with the weak and total variation norm[8/] topologies, respectively, and (b) $m(q)$ is measurable as a function of $x_o, \ldots, x_{t-1}$, when M(A) has the weak topology. These conditions will be incorporated in the definitions of anonymity and of strategy vector.

Definition: A repeated game with random outcomes $\langle K,A,X,F,\mu,y\rangle$ is defined exactly as was a repeated game in section 4, except that $F:A^K \to M(X)$. The shorter term repeated random game will be used synonymously with this. A random game is anonymous if the diagram analogous to that which defines an anonymous determistic game (but with $M(X)$ instead of X) commutes, and if G is a continuous function from $M(A)$ under the weak topology into $M(X)$ under the total variation norm topology. A strategy for a repeated random game is defined as in the deterministic case. An assignment $\langle s^k\rangle_{k\varepsilon K}$ of strategies to players is a strategy vector for a repeated random game if it satisfies the conditions specified in the deterministic case and if, in addition, when for $t\varepsilon N_+$ the function $\theta_t:X^t \to M(A)$ is defined by

$$(22)\quad \theta_t(x_o,\ldots,x_{t-1}) = m(\langle s_t^k(x_o,\ldots,x_{t-1})\rangle_{k\varepsilon K}) ,$$

each θ_t is measurable ($M(A)$ having the weak topology).

Let S^K denote the set of strategy vectors. To define the random path generated by a vector $v\varepsilon S^K$, let $\Omega = X^N$ and let $\mathcal{B}$ be the Borel σ-field generated by the product topology on Ω. Define the projection functions $\underset{\sim}{x}_t:\Omega\to X$ for $t\varepsilon N$ by $\underset{\sim}{x}_t(\langle x_o,x_1,\ldots,x_t,\ldots\rangle) = x_t$. Let $\mathcal{B}_t$ be the smallest σ-field with respect to which $\underset{\sim}{x}_o,\ldots,\underset{\sim}{x}_t$ are all measurable (N.B. $\mathcal{B}_t \subseteq \mathcal{B}$, and $\mathcal{B}_t$ is isomorphic as an algebra to the Borel σ-field on X^{t+1}, by [7],p.6)

Recall that, for a probability measure π defined on $\mathcal{B}$, a regular conditional probability relative to $\mathcal{B}_t$ is a function $P_t:\mathcal{B}\times\Omega\to[0,1]$ such that

(a) for fixed $\omega\varepsilon\Omega$, $P_t(\cdot|\omega)$ is a probability measure defined on Ω,

(b) for fixed $B\subseteq\Omega$, $P_t(B|\cdot)$ is $\mathcal{B}_t$-measurable, and

(c) for all $B\varepsilon\mathcal{B}$ and $C\varepsilon\mathcal{B}_t$,

$$(23)\quad \pi(B\cap C) = \int_C P_t(B|\omega)\,d\pi(\omega)$$

For every probability measure π defined on $\mathcal{B}$, and every $t\varepsilon N$, a regular conditional probability relative to $\mathcal{B}_t$ exists, and any two of these differ only on a set of π measure zero (i.e., if P and P' are regular conditional probabilities relative to $\mathcal{B}_t$, then $\pi(\{\omega|\exists B\varepsilon\mathcal{B}(P(B|\omega) \neq P'(B|\omega)\}) = 0)$. ([7], p. 147.) Thus, regular conditional probabilities may be used unambiguously to define random paths.

Definition: Let $\Omega, \mathcal{B}, \underset{\sim}{x}_t, \mathcal{B}_t$, and P_t be as above. Let $v\varepsilon S^K$, where $v = \langle s^k\rangle_{k\varepsilon K}$ and $s^k = \langle s^k_t\rangle_{t\varepsilon N}$. The random path generated by v is the probability space $\langle\Omega,\mathcal{B},\pi\rangle$, which satisfies, for every Borel subset B of X,

$$(24)\quad \pi(\underset{\sim}{x}_o(\omega)\varepsilon B) = \nu(B),$$

where $\nu = F(\langle s^k_o\rangle_{k\varepsilon K})$, and for $t\varepsilon N_+$,

$$(25)\quad P_t(\underset{\sim}{x}_t(\omega)\varepsilon B|\omega) = \nu(B),$$

where $\nu = F(\langle s^k_t(\underset{\sim}{x}_o(\omega),\ldots,\underset{\sim}{x}_{t-1}(\omega))\rangle_{k\varepsilon K})$, for all ω belonging to some $C_t\varepsilon\mathcal{B}_{t-1}$ where $\pi(C_t) = 1$.

It is proven in the appendix that the random path generated by a strategy vector exists and is unique. Definitions of static and dynamic NCE for a random game differ from the definitions in the deterministic case only in that expectations must be taken with respect to distributions of random outcomes and paths.

Definition: A static non-cooperative equilibrium of the repeated game $\langle K,A,X,F,\mu,y\rangle$ with random outcomes is defined as for a deterministic game, except that (17) becomes 9/

$$(26)\quad E_{F(q)}u(q(k),x) \leq E_{F(q^*)}u(q^*(k),x).$$

The definition of a dynamic quasi-equilibrium is the same as in the deterministic case, except that now the random paths $\langle\Omega,\mathcal{B},\pi\rangle$ and $\langle\Omega,\mathcal{B},\pi^*\rangle$ of v and v^*, respectively, are treated rather than histories in the deterministic case. Equation (18) is changed to

$$(27)\quad E_{\pi}(u(s_o^k, \underset{\sim}{x}_o) + \sum_{t\in N_+} \beta^t u(s_t^k(\underset{\sim}{x}_o,\ldots,\underset{\sim}{x}_{t-1}),\underset{\sim}{x}_t))$$

$$\leq E_{\pi^*}(u(s^{*k}_o, \underset{\sim}{x}_o) + \sum_{t\in N_+} \beta^t u(s^{*k}_t(\underset{\sim}{x}_o,\ldots,\underset{\sim}{x}_{t-1}),\underset{\sim}{x}_t)).$$

The definition of a <u>dynamic non-cooperative equilibrium</u> is exactly the same as in the deterministic case.

An explicit definition is now given of the type of dynamic non-cooperative equilibrium which was studied in lemma 1. Recall that in that type of equilibrium, there is a play q_o to which players are committed at each repetition of the game. This commitment is enforced by mutual threats that, if there is evidence that some player k has not taken action $q_o(k)$, then the other players will thereafter take the actions prescribed by another play q_1. This threat is self-enforcing, because q_1 is a static non-cooperative equilibrium. Evidence that a player has departed from q_o consists of an outcome which lies outside some subset $B \subseteq X$. That is, the players agree to play q_o as long as outcomes in B have been observed in the past, but always to play q_1 if some outcome outside of B has occurred in the past.

<u>Definition</u>: A <u>strategic non-cooperative equilibrium</u> of the game $\langle K,A,X,F,\mu,y\rangle$ with random outcomes is a dynamic NCE $v\in S^K$ such that, for some triple $\langle q_o,q_1,B\rangle \in A^K \times A^K \times 2^X$, the following conditions are satisfied for every $k\in K$ (with $v(k) = \langle s_t\rangle_{t\in N}$):

(a) q_1 is a static NCE.

(b) $a_o = q_o(k)$

(c) If, for any $t \varepsilon N_+$, $x_o,\ldots,x_{t-1} \varepsilon B$, then $s_t(x_o,\ldots,x_{t-1}) = q_o(k)$

and

(d) If, for any $t \varepsilon N_+$ and $r<t$, $x_r \notin B$, then $s_t(x_o,\ldots,x_r,\ldots,x_{t-1})$ $= q_1(k)$ for all $x_o,\ldots,x_{r-1},x_{r+1},\ldots,x_{t-1} \varepsilon X$.

By abuse of notation, the strategy vector defined by (b) - (d) above will also be referred to as $\langle q_o, q_1, B\rangle$.

There is a condition which, if the expectations referred to in its statement are well-defined, is necessary and sufficient for a strategy vector $\langle q_o, q_1, B\rangle$ to be a strategic NCE. Some notation is introduced now to facilitate the statement of this condition. Let $k \varepsilon K$ be fixed. For $a \varepsilon A$, define q_a to be the play which results from q_o when the action of k is changed to a. I.e., $q_a(k) = a$, and $q_a(i) = q_o(i)$ for $i \neq k$. Define $\nu_j = F(q_j)$ for $j = 1,2$, and $\nu_a = F(q_a)$ for all $a \varepsilon A$, and let E_j and E_a be the expectation operators (on functions defined over X) with respect to ν_j and ν_a, respectively.

<u>Lemma 3</u>: Let $q_o, q_1 \varepsilon A^K$, let q_1 be a static NCE, and suppose that for all $k \varepsilon K$ (with $y(k) = \langle u,\beta\rangle$), $E_o u(q_1(k),x)$ is finite and $E_a u(a,x)$ is finite for every $a \varepsilon A$. Let $B \subseteq X$, and define

(28) $V_k(a) = [E_a u(a,x) + \beta(1-\beta)^{-1}(1-\nu_a(B))E_1 u(q_1(k),x)]/[1-\beta\nu_a(B)]$,

for all $a \varepsilon A$. Then $\langle q_o, q_1, B\rangle$ is a strategic NCE if and only if, for almost all $k \varepsilon K$,

(29) $V_k(q_o(k)) = \max_{a\epsilon A} V_k(a)$.

Proof: By definition, $<q_o,q_1,B>$ is a strategic NCE if and only if it is a dynamic NCE. Furthermore, by the reasoning explained in lemma 1, the fact that q_1 is a static NCE entails that $<q_o,q_1,B>$ is a dynamic NCE if and only if it is a dynamic quasi-equilibrium. Thus, it is sufficient to prove that (27) holds with $v^* = <q_o,q_1,B>$ for all other $v\epsilon S^K$ with $v(i) = v^*(i)$ for $i\neq k$, if and only if (29) holds.This proof consists of two parts. First, it is argued that one need only consider (27) with respect to v^* and strategy vectors $v = <q_a,q_1,B>$ for $a\epsilon A$. Second, a dynamic programming argument is used to evaluate the expected discounted returns in (27) for paths generated by these strategy vectors.

Both of these arguments are transparent when it is realized that the decision problem of player k is equivalent to a very simple Markov dynamic programming problem. In this equivalent problem there are two states: either all previous outcomes of the game have been in B (state 0), or some outcome has not been in B (state 1). In state 0, taking an action $a\epsilon A$ yields a present return (the expected return $E_a u(a,x)$ from the play q_a) and incurs a risk (the probability $1-\nu_a(B)$ of producing an outcome not in B) of causing a change to state 1 in the next period. State 1 is an absorbing state, and in that state the player's best action is $q_1(k)$ which yields return $E_1 u(q_1(k),x)$. A basic theorem of dynamic programming states that, if any optimal control rule exists for this problem, there is an optimal rule which specifies a fixed action $a\epsilon A$ to be taken as long as state 0 persists. If player k follows this rule while other players use the strategies assigned by $<q_o,q_1,B>$ the strategy vector $<q_a,q_1,B>$ results. Thus, if any strategy of k leads to a violation of (27), the equation fails to hold for some strategy vector of this special form.

It remains to calculate the expected discounted returns to player k from paths of strategy vectors $\langle q_a . q_1, B\rangle$ (N.B. If $a=q_o(k)$, then $\langle q_a, q_1, B\rangle = \langle q_o, q_1, B\rangle$). Denote his expected discounted return from the path of $\langle q_a, q_1, B\rangle$ by $V_k(a)$. With this interpretation of $V_k(a)$, (29) holds if and only if $\langle q_o, q_1, B\rangle$ is a dynamic quasi-equilibrium. Thus it is only necessary to verify that (28) is true under this interpretation, in order to complete the proof.

To do this, note that $V_k(a)$ must equal the expected discounted return in the equivalent dynamic program under the control rule: perform action a in state 0, and action $q_1(k)$ in state 1. From the first time that state 1 occurs, the expected discounted return is $(1-\beta)^{-1} E_1 u(q_1(k), x)$. Thus, if $W: \{0,1\} \to R$ is the value function of the control rule, $V_k(a) = W(0)$ and $(1-\beta)^{-1} E_1 u(q_1(k), x) = W(1)$. $W(0)$ is defined by the functional equation

$$(30)\quad W(0) = E_a u(a,x) + \beta[P(\text{the state next period will again be } 0)\cdot W(0) + P(\text{the state next period will change to } 1)\cdot W(1)].$$

Given that the transition probabilities are $\nu_a(B)$ and $(1-\nu_a(B))$, respectively, (30) is equivalent to (28). Q.E.D.

Lemma 3 will be used to prove a limit theorem for anonymous repeated games with random outcomes. To state this theorem, a sequence of replica repeated games must be defined. This definition is a straightforward generalization of that of a sequence of replica dynamic markets.

Definition: Let $H = \langle K, A, X, F, \mu, y\rangle$ be an anonymous repeated game with random outcomes. In particular, let the diagram

(31) 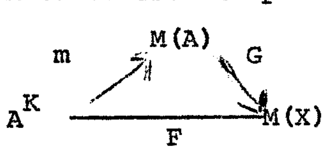

commute. Suppose that K is a finite set. Then, a <u>sequence of replica repeated</u>

<u>games</u> $\langle H_n \rangle_{n\in N}$ with random outcomes is defined as follows: For every $n\in N$, $H_n = \langle K_n, A, X, F_n, \mu_n, y_n \rangle$, where $K_n = K\times\{0,\ldots n\}$, μ_n is determined by $\mu_n(\{\langle k,r\rangle\}) = \mu(\{k\})/(n+1)$ for all $\langle k,r\rangle\in K_n$, and $y_n(\langle k,r\rangle) = y(k)$ for all $\langle k,r\rangle\in K_n$. It remains to define the outcome function F_n. Let $m_n{:}A^{K_n}\to M(A)$ map plays of H_n to the distributions of actions which they determine as random variables on $\langle K_n,\mu_n\rangle$. Then define $F_n{:}A^{K_n}\to M(X)$ by $F_n(q) = G(m_n(q))$ for all $q\in A^{K_n}$. Define the <u>non-atomic representation</u> H_∞ of the sequence analogously, where $K_\infty = K\times[0,1]$, and, for $B\subseteq K_\infty$, $\mu_\infty(B) = \sum_{k\in K} \mu(\{k\})\cdot\lambda(\{r\in[0,1]\,|\,\langle k,r\rangle\in B\})$.

The limit theorem to be proved here states that, under appropriate hypotheses, the limit of strategic non-cooperative equilibria is trivial in the sense that the "collusive" play supported in the limit must itself be a static non-cooperative equilibrium. Formally, if $\langle H_n\rangle_{n\in N}$ is a sequence of replica repeated games with random outcomes, if $\langle q_o^n, q_1^n, B_n\rangle$ is a strategic NCE of H_n for each n, and if $\langle\langle y_n, q_o^n, q_1^n\rangle\rangle_{n\in N}\to\langle y_o, q_o^\infty, q_1^\infty\rangle$ in distribution, then q_o^∞ is a static NCE of H_∞

The statement and proof of the limit theorem are now given.

Theorem 6 (limit theorem for anonymous repeated games with random outcomes):

Suppose that $\langle H_n \rangle_{n \in N}$ is a sequence of replica repeated games with random outcomes, having H_∞ as non-atomic representation. Let $H = \langle K, A, X, F, \mu, y \rangle$, and let G make (31) commute. Suppose that $\langle q_o^n, q_1^n, B_n \rangle$ is a strategic NCE of H_n for every $n \in N$, and that the following hypotheses are satisfied.

(a) The countable subset $D \subseteq A$ is dense in A.

(b) If $y(k) = \langle u, \beta \rangle$ for some $k \in K$, then $E_{G(\eta)} u(a,x)$ is a finite-valued, continuous function on $A \times M(A)$, where $M(A)$ has the weak topology.

(c) The sequence $\langle \langle y_n, q_o^n, q_1^n \rangle \rangle_{n \in N}$ converges in distribution.

Then there is a random vector $\langle y_\infty, q_o^\infty, q_1^\infty \rangle$ defined on K_∞ such that

(32) $\langle \langle y_n, q_o^n, q_1^n \rangle \rangle_{n \in N} \to \langle y_\infty, q_o^\infty, q_1^\infty \rangle$

in distribution, and q_o^∞ is a static NCE of H_∞.

Proof: The two ideas behind this proof are that the analogue (with expected utilities) of theorem 5 must be true for non-atomic games with random outcomes, and that the functions V_k defined by (28) are continuous in the actions and measures involved in their definition. If q_o^∞ is not a static NCE of H_∞, the first assertion entails that $\langle q_o^\infty, q_1^\infty, B_\infty \rangle$ cannot be a dynamic quasi-equilibrium for any $B_\infty \subseteq X$. Therefore (29) cannot be satisfied almost everywhere for H_∞ by any $\langle q_o^\infty, q_1^\infty, B_\infty \rangle$. By the second assertion, neither can (29) hold almost everywhere on K_n, for the near-by games H_n and strategy vectors $\langle q_o^n, q_1^n, B_n \rangle$ when n is large.

Formally, the limiting random vector $\langle y_\infty, q_o^\infty, q_1^\infty \rangle$ exists by Skorokhod's theorem ([4], p.50). Suppose that q_o^∞ is not a static NCE. Then there is a set $J \subseteq K_\infty$, with $\mu_\infty(J)>0$, on which (26) does not hold for all necessary q when $q^* = q_o^\infty$. Because K is finite, for some $k \varepsilon K$ (N.B. k will refer throughout the proof to this fixed element of K), $\mu_\infty(J \cap (\{k\} \times [0,1]))>0$. Without loss of generality, it may be assumed that $J \subseteq \{k\} \times [0,1]$.

Because H_∞ is non-atomic and anonymous, and because D is dense in A, the failure of (26) is equivalent by (b) to the existence of some $d \varepsilon D$ such that

$$(33)\quad E_{F_\infty(q_o^\infty)} u(q_o^\infty(\langle k,r\rangle),x) \leq E_{F_\infty(q_o^\infty)} u(d,x).$$

As above, the countability of D permits (33) to be assumed without loss of generality to hold everywhere on J for some fixed action d of D (N.B. d also be held fixed throughout the proof). Similarly, by the separability of R_+, it may be assumed that a stronger version

$$(34)\quad E_{F_\infty(q_o^\infty)} u(d,x) - E_{F_\infty(q_o^\infty)} u(q_o^\infty(\langle k,r\rangle),x)>\delta$$

of (33) holds uniformly on J for some $\delta>0$, and that by (b),

$$(35)\quad |E_{F(q_1^\infty)} u(q_1^\infty(\langle k,r\rangle),x)|<\varepsilon$$

holds uniformly on J for some $\varepsilon>0$.

By (b), (34) and (35), for each r with $\langle k,r\rangle \varepsilon J$, there exist open sets $U^*_{or}, U^*_{1r} \subseteq A$ and $W_{or}, W_{1r} \subseteq M(A)$ with $q_j^\infty(\langle k,r\rangle) \varepsilon U^*_{jr}$ and $m_\infty(q_j^\infty) \varepsilon W_{jr}$ for j=0,1, which satisfy

$$(36)\quad E_{G(\eta)} u(d,x) - E_{G(\theta)} u(a,x)>\delta$$

for $\eta,\theta \varepsilon W_{or}$ and $a \varepsilon U^*_{or}$, and

$$(37)\quad |E_{G(\eta)} u(a,x)|<\varepsilon$$

for $\eta \varepsilon W_{1r}$ and $a \varepsilon U^*_{1r}$.

There are open U_{jr} for $j = 0,1$ and $<k,r>\varepsilon J$, such that $q_j^\infty(<k,r>)\varepsilon U_{jr}$ and $\bar{U}_{jr}\subseteq U^*_{jr}$ (N.B. $\bar{U}$ is the closure of U). Let $\Gamma = \bigcup_{<k,r>\varepsilon J}(U_{or}\times U_{1r})$. By (a) and Lindelöf's theorem ([1], p. 12) there is a countable subset $Y\subseteq[0,1]$ such that $\{U_{or}\times U_{1r}|r\varepsilon Y\}$ is a cover of Γ. Since $\mu_\infty(J)>0$ and $J\subseteq\{<k,r>|<q_o^\infty(<k,r>),q_1^\infty(<k,r>)>\varepsilon\Gamma\}$, and since Y is countable, there is some $U = U_{or}\times U_{1r}$ such that

(38) $\mu_\infty(\{<k,r>|<q_o^\infty(<k,r>),q_1^\infty(<k,r>)>\varepsilon U\})>0.$

By [7], p. 4, there is a continuous function $f:A\times A\rightarrow[0,1]$ having support in $\bar{U}^*$ (where $U^* = U^*_{or} \times U^*_{1r}$ for the same r which defines U) such that $f\equiv 1$ on $\bar{U}$. By (38).

(39) $\int_{[0,1]}f(<q_o^\infty(<k,r>),q_1^\infty(<k,r>)>)d\lambda(r)>0.$

It is evident from (c) and the definition of convergence in distribution that asymptotically (i.e., for all sufficiently large $n\varepsilon N$) there exist $r_n\leq n$ for which

(40) $<q_o^n(<k,r_n>),q_1^n(<k,r_n>)>\varepsilon U$

For $n\varepsilon N$, define $a_j^n=q_j^n(<k,r_n>)$ for $j=0,1$. Also define (as in lemma 4) q_d^n to be the play of H_n which results when $<k,r_n>$ deviates from q_o^n by taking action d. I.e., $q_d^n(<k,r_n>) = d$ and $q_d^n(<k',r'>)= q_o^n(<k',r'>)$ if $k'\neq k$ or $r'\neq r_n$. Define $\eta_j^n = m_n(q_j^n)$ for $j=0,1,d$. By (c),

(41) $<\eta_j^n>_{n\varepsilon N}\rightarrow m_\infty(q_j^\infty)$

for $j=0,1$, and

(42) $<\eta_d^n>_{n\varepsilon N}\rightarrow m_\infty(q_o^\infty)$.

Thus asymptotically $\eta_o^n, \eta_d^n \varepsilon W_{or}$ and $\eta_1^n \varepsilon W_{1r}$, so that (36), (37) and (40) imply that for sufficiently large n,

(43) $E_{G(\eta_d^n)}u(d,x) - E_{G(\eta_o^n)}u(a_o^n,x) > \delta$

and

(44) $|E_{G(\eta_1^n)}u(a_1^n,x)| < \varepsilon$.

Equations (43) and (45) will be used now to establish that (29) fails asymptotically for $q_o = q_o^n$, so that $<q_o^n, q_1^n, B_n>$ cannot be a strategic NCE of H_n. Define $\nu_j^n = G(\eta_j^n)$ and let E_j^n, be expectation with respect to ν_j^n for j=0,1,d. It will be shown that asymptotically

(45) $[E_o^n u(a_o^n,x) + \beta(1-\beta)^{-1}(1-\nu_o^n(B_n))E_1^n u(a_1^n,x)]/[1-\beta\nu_o^n(B_n)]$

$< [E_d^n u(d,x) + \beta(1-\beta)^{-1}(1-\nu_d^n(B_n))E_1^n u(a_1^n,x)]/[1-\beta\nu_d^n(B_n)).$

By (28), (45) is equivalent for H_n to $V_{<k,r_n>}(q_n^o(<k,r_n>)) < V_{<k,r_n>}(d)$,

so that (45) contradicts (29) for H_n. Thus the theorem will hold by lemma 3.

Multiplying both sides of (45) by $[1-\nu_o^n(B_n)]$ yields

(46) $E_o^n u(a_o^n,x) + \beta(1-\beta)^{-1}(1-\nu_o^n(B_n))E_1^n u(a_1^n,x)$

$< [(1-\beta\nu_o^n(B_n))/(1-\beta\nu_d^n(B_n))][E_d^n u(d,x) + \beta(1-\beta)^{-1}(1-\nu_d^n(B_n)) E_1^n u(a_1^n,x)].$

For each $\xi > 0$ there is an open $W_\xi \subseteq M(A)$, with $m_\infty(q_o^\infty) \varepsilon W_\xi$, which satisfies

(47) $\sup_{B \subseteq X} |1-[[1-\beta(G(\eta))(B)]/[1-\beta(G(\theta))(B)]]| < \xi$

for $\eta, \theta \varepsilon W_\xi$. By (41), (42) and (47), (46) holds asymptotically if

(48) $E_o^n u(a_o^n,x)-E_d^n u(d,x)+\beta(1-\beta)^{-1}[\nu_d^n(B_n)-\nu_o^n(B_n)]E_1^n u(a_1^n,x) < 0$

asymptotically. By (d), (41) and (42), $\lim_{n\to\infty}[\nu_d^n(B_n)-\nu_o^n(B_n)] = 0$, so by (44), (48) holds asymptotically if

(49) $\overline{\lim_{n\to\infty}}[E_o^n u(a_o^n,x) - E_d^n u(d,x)]<0.$

This completes the proof, because (43) implies (49). Q.E.D.

From theorems 2 and 6, it is evident that the paths of strategic non-cooperative equilibria of a sequence of replica dynamic markets with unobserved random demand fluctuations converge to a price-taking equilibrium of the non-atomic representation of the sequence. One difficulty arises in the application of theorem 6 here. That is, that hypothesis (d) fails if firms have unbounded production sets. However, if marginal costs become sufficiently high as output levels rise, this difficulty may be handled by imposing prior bounds, uniform in n, on the output levels which firms would consider. Hypothesis (d) is satisfied in the resulting sequence of replica markets because inverse demand is continuous and production sets are compact.

6. Concluding Remarks

The results which have been presented here concern the extent to which price-taking behavior characterizes the non-cooperative equilibria of large economies. Theorems 1 and 2 of the present paper exemplify the positive results for this question which hold under the simplifying assumption that trade in each market occurs only once, during a single, initial trading period in which all markets clear simultaneously. Theorems 3 and 4 show that the validity of these results for actual markets, which continue to be active over time cannot be taken for granted. Finally, theorems 5 and 6 describe the extent to which results for static markets can legitimately be applied to temporal trade. In this section, the practical significance of these results will be examined.

Most economists are in rough agreement about which actual markets are paradigm cases which corroborate the static theory of competitive behavior in large markets, and which markets are accounted for less successfully by the theory. Theorems 5 and 6 justify the application of the static theory to the paradigm cases, and they also clarify the status of some of the marginal cases. Theorem 5 asserts (in conjunction with theorem 3) that, if firms in a market literally have no effect on the price, then the market must be competitive unless firms have access to extremely disaggregated information about their competitors. Theorem 6 states that, in the presence of imperfections which may reasonably be thought to exist in the price system, the assumption in theorem 5 that individual firms do not influence the price may be taken as a good approximation of the situation in large finite markets. Thus, when a market contains many un-coordinated sellers, the static theory may legitimately be used to predict that the market will be competitive.

Problems typically arise in deciding what the words "many" and "un-coordinated" mean, when the applicability of the theory of large markets to a particular market is in question. Thus one class of marginal cases for the static theory consists of moderately concentrated industries (typically having fewer than a dozen major suppliers), and another class consists of markets in which the actions of a trade or professional association might affect the behavior of sellers. One school of thought in the economics profession holds that collusive agreements of even the smallest size are very unstable and that non-cooperative equilibria converge rapidly to price-taking as the number of sellers is increased, so that there is a strong presumption that markets in both of these classes are competitive. An apparently conflicting view is held by a number of economists who have been persuaded of the monopolistic effects of professional associations in medicine, law, and so forth, on the basis of evidence which for the most part has been rather impressionistic.

Each of these views presupposes a set of beliefs about what determines whether firms in an industry would be able to detect and punish a competitor which attempted to violate a cartel agreement. It has long been recognized in the informal literature on industrial organization that this question, more than the question of how concentrated the industry is according to the usual quantitative measures, is important in assessing the competitiveness of an industry with restricted entry. By stating precise, formal criteria for cartel organization in an industry to be enforceable, theorems 5 and 6 point the way toward an embodiment of this informal tradition in a rigorous, empirically satisfactory theory.[10/] Industry studies based on such a theory would help to resolve long-standing empirical disputes about the extent of cartel organization, and might contribute to the determination of anti-trust policy on a more rational basis.

Appendix: Existence and uniqueness of a random path

Let $v = \langle s^k \rangle_{k\in K} \in S^K$, where $s^k = \langle s_t^k \rangle_{t\in N}$ for each t. It will be proven that there is a unique Borel measure on $\Omega = X^N$ which satisfies the definition of a random path of v of an anonymous repeated random game.

Lemma 4: To prove the existence and uniqueness of a random path, it is sufficient to prove that there is a unique sequence $\langle \pi_t \rangle_{t\in N}$ such that

(a) π_t is a measure defined on $\mathcal{B}_t$ for each t,

(b) if $t < t'$, then π_t is the restriction of $\pi_{t'}$ to $\mathcal{B}_t$,

(c) the equation

(50) $\pi_o = F(\langle s_o^k \rangle_{k\in K})$

holds, and

(d) if $t\in N_+$ and P_t^* is regular conditional probability for π_t relative to $\mathcal{B}_{t-1}$, then for every Borel subset B of X,

(51) $P_t^*(\underset{\sim}{x}_t(\omega) \in B \mid \omega) = \nu(B)$,

where $\nu = F(\langle s_t^k(\underset{\sim}{x}_o(\omega), \ldots, \underset{\sim}{x}_{t-1}(\omega)) \rangle_{k\in K})$, for all ω belonging to some $C_t \in \mathcal{B}_{t-1}$ where $\pi_t(C_t) = 1$.

Proof: By the Kolmogorov consistency theorem ([7], p. 139), for any sequence satisfying (a) - (d) there is a unique Borel measure π on Ω such that π_t is the restriction of π to $\mathcal{B}_t$, for every t. Thus (50) and (51) are equivalent to (24) and (25), so $\langle \Omega, \mathcal{B}, \pi \rangle$ is a random path generated by v. If $\langle \Omega, \mathcal{B}, \pi' \rangle$ is any random path generated by v, and if π_t' is the restriction of π' to $\mathcal{B}_t$, then $\langle \pi_t' \rangle_{t\in N}$ satisfies (a)- (d). By the uniqueness of the

Kolmogorov extension, $\langle\pi_t\rangle_{t\in N} \neq \langle\pi'_t\rangle_{t\in N}$ if $\pi \neq \pi'$. Thus, uniqueness of the sequence satisfying (a) - (d) is sufficient for the uniqueness of the random path.

Q.E.D.

Lemma 5: A unique sequence $\langle\pi_t\rangle_{t\in N}$ exists which satisfies (a) - (d) of lemma 4.

Proof: A sequence satisfying (a) - (d) will be constructed recursively. At each stage the measure chosen will be determined uniquely. To begin with, (50) defines π_o . Supposing that π_t has been chosen, condition (d) will be used to uniquely define π_{t+1} on $\mathcal{B}_{t+1}$. Define a block in X^{t+2} to be a product set $C\times B$, where $C\in\mathcal{B}_t$[11] and B is a Borel subset of X. Define $\pi^*:\mathcal{B}_t\times\mathcal{B}_o\to[0,1]$ by

$$(52)\quad \pi^*(C\times B) = \int_C P^*_{t+1}(\underset{\sim}{x}_{t+1}(\omega)\in B|\omega)d\pi_t(\omega),$$

where equation (51) is taken as the definition of P^*_t. [N.B. The measurability of $P^*_{t+1}(\underset{\sim}{x}_{t+1}\in B|\cdot)$ with respect to $\mathcal{B}_t$ is guaranteed by the definitions of strategy vector and anonymous game.] Now define $\mathcal{B}^*\subseteq\mathcal{B}_{t+1}$ to be the class of finite unions of disjoint blocks. $\mathcal{B}^*$ is a Boolean algebra ([1], p. 185), and π^* determines a finitely additive measure on $\mathcal{B}^*$. In fact, it is evident from Lebesque's monotone convergence theorem that π^* is countably additive on $\mathcal{B}^*$. By the Caratheodory-Hahn extension theorem([1])p.136), π^* extends uniquely to a measure on the smallest σ-field containing $\mathcal{B}^*$. By [7], p. 6, this σ-field is $\mathcal{B}_{t+1}$. Let the extension measure be π_{t+1}. Conditions (a) and (b) are trivially verified, and it follows from (23) that P^*_{t+1} defined by (51) agrees almost surely with regular conditional probability for $B\in\mathcal{B}^*$, so that π_{t+1} satisfies (d) and is the unique extension of π_t which does so. This completes the induction.

Q.E.D.

The existence and uniqueness of the random path of v is an immediate consequence of lemmas 4 and 5:

<u>Theorem 7</u>: If H is an anonymous repeated game with random outcomes and v is a strategy vector for H, then v generates a unique random path $\langle\Omega,\mathfrak{B},\pi\rangle$.

Footnotes

1. The standard reference for this theory is Hildenbrand [4].

2. The bibliography of [5] contains further references in this area. William Thomson has presented (in [11]) a counter example to the non-cooperative limit principle for a static exchange economy.

3. The present formulation of dynamic markets is restrictive. For instance, it does not allow for the holding of inventories. The inclusion theorem (Theorem 5) to be proved here generalizes straightforwardly to apply to richer models. I am confident that the limit theorem (Theorem 6) is also robust.

4. Clause (b) insures that a firm will not base its initial output decision on an expectation that, at a later time, it will produce at a level which from the perspective of that time will seem suboptimal. In the next section, equilibria will be studied in which firms in a market maintain collusion by means of mutual strategic threats. Although these threats will not be exercised in equilibrium, they will not have a deterrent effect unless it would be in the firms' interest to exercise them in the event that deviation from the cartel did occur. This explains why (a) must hold even conditional on firms having received information which they in fact will not receive in equilibrium. Condition (b) was introduced (in [10]) by Selten, whose concept of perfect Nash equilibrium is equivalent to the dynamic NCE defined here.

5. Let Z be a metric space. A sequence $\langle\eta_n\rangle_{n\in N}$ of probability measures converges weakly to a probability measure η if, for every bounded continuous function $f:Z\to R$, $\lim_{n\to\infty}\int_Z f d\eta_n \to \int_Z f d\eta$. If $\langle\Omega_n,\mu_n\rangle_{n\in N}$ is a sequence of probability spaces and $\langle\Omega,\mu\rangle$ is a probability space, and if $z_n:\Omega_n\to Z$ for every n, then $\langle z_n\rangle_{n\in N}\to z:\Omega\to Z$ in distribution if, when the measures η_n are defined by $\eta_n(B)=\mu_n(\{z_n\in B\})$ for all Borel sets $B\subseteq Z$ (and η is defined analogously for z), then $\langle\eta_n\rangle\to\eta$ weakly.

6. It is assumed that the discounted sum of optimal returns converges. In applications, this assumption typically is satisfied in equilibrium. If necessary, the assumption could be avoided by using a discounted overtaking criterion.

7. It might be suggested that all (not necessarily measurable) functions from K to A ought to be included in the set of plays and that an outcome of the game should be assigned to every such function. Such a change would not significantly affect the results of this paper. In the case of the limit theorem (Theorem 6), plays of the finite-player games are trivially measurable, and the measurability of the limiting play is proven in the theorem. The inclusion theorem (Theorem 5) would continue to hold as stated, if the definition of anonymity used there were replaced by: a game is anonymous if any two plays differing only in the action of a single player have the same outcome.

8. If $\mathcal{B}$ is a σ-field and $f: \mathcal{B} \to R$ is a countably additive set function, then the total variation norm of f is defined by

$$|f| = \sup\{ \sum_{n \in N} |f(B_n)| \,\Big|\, \langle B_n \rangle_{n \in N} \text{ is a sequence of disjoint elements of } \mathcal{B}\} .$$

Since this norm induces a stronger topology than the weak topology, the continuity of G is not automatic. However, if the range of G is any of the usual parametrized families of distribution (e.g., normal, if $X=R^n$) and does not contain any degenerate (e.g., singular covariance normal) distributions, then weak continuity of the parameter as a function of the distribution of actions is sufficient to guarantee continuity of G.

9. The functions $u(q(k),x)$ and $u(s_t^k(\underset{\sim}{x}_o,\ldots,\underset{\sim}{x}_{t-1},\underset{\sim}{x}_t)$ are random variables on X and Ω in equations (26) and (27), respectively.

10. In particular, equation (29) suggests how the hypothesis that there is collusion in a particular industry might be tested on the basis of a time series of prices and market shares for that industry. Evidence from such a test would be more persuasive than that from cross-industry comparisons on which current attempts to measure the extent of collusion are based.

11. In this proof, $\mathcal{B}_t$ will indicate both the Borel subsets of X^{t+1} and the smallest subfield of X^N with respect to which $\underset{\sim}{x}_o, \ldots, \underset{\sim}{x}_t$ are Borel measurable. The intended reference will be clear from context.

References

1. N. Dunford and J. Schwartz, "Linear Operators" vol. 1, Interscience, New York, N. Y., 1958.

2. J. W. Friedman, A non-cooperative equilibrium for supergames, Review of Economic Studies 28(1971), 1-12.

3. J. Gabszewicz and J. P. Vial, Oligopoly à la Cournot in general equilibrium theory, Journal of Economic Theory 4(1972), 381-400.

4. W. Hildenbrand, "Core and Equilibria of a Large Economy," Princeton University Press, Princeton, N.J., 1974.

5. A. Mas Colell, An axiomatic approach to the decentralizability of non-cooperative equilibrium in an economy with a continuum of traders, University of California, Berkeley, 1978.

6. W. Novshek and H. Sonnenschein, Cournot and Walras, Princeton University, 1977.

7. K. R. Parthasarathy, "Probability Measures on Metric Spaces," Academic Press, New York, N.Y., 1967.

8. A. Postlewaite and D. Schmeidler, Approximate efficiency of non-Walrasian equilibria, Econometrica 46(1978) 127-135.

9. J. Roberts and A. Postlewaite, The incentives for price-taking behavior in large exchange economies, Econometrica 44(1976) 115-127.

10. R. Selten, A re-examination of the perfectness concept for equilibrium points in extensive games, International Journal of Game Theory 4(1975) 25-55.

11. W. Thomson, Nash equilibria of Walras and Cournot mechanisms, University of Minnesota, 1978.