Federal Reserve Bank of Minneapolis
Research Department Staff Report 57

Revised August 1981

A GENERALIZED EQUILIBRIUM SOLUTION FOR GAME THEORY

John Bryant

Federal Reserve Bank of Minneapolis

and Rice University

ABSTRACT

Iterated contraction by dominance produces a generalized equilibrium. This solution to game theory is motivated, generated, analyzed, and compared to Nash equilibrium. One implication drawn is that a realized event in a social situation need not be uniquely determined by simple individual choices, even though the preference orderings implying those choices are the appropriate primitive.

A Generalized Equilibrium Solution for Game Theory

by John Bryant

Game theory is both at the heart of economics and without a definitive solution. This paper proposes a solution. It is argued that a dominance criterion generates a, and indeed the, generalized equilibrium solution for game theory. First we provide a set theoretic perspective from which to view game theory, and then present and discuss the proposed solution.

1. The Game

First we describe a perspective from which to view game theory. Let A be the set of events, and let τ be the set of individuals. For each $t\varepsilon\tau$, let X_t be the subset of A for which individual t has her unique complete preordering (preference ordering) $\gtrsim_t$ on points in X_t. In game theory it is typically assumed that events are probability measures on outcomes and that $X_t = X \subset A$ for all t. Now we get to the conflict part of game theory. For each $t\varepsilon\tau$, there exists a collection of subsets of A, W_t, such that individual t can restrict events to belong to one of the sets in W_t, $S_t\varepsilon W_t$. The strategy of the individual t is choice of a member of W_t. In game theory it is typically assumed that for any collection of sets

$$\left\{S_t\right\}_{t\varepsilon\tau}$$

such that $S_t\varepsilon W_t$ for all t, $\bigcap S_t$ is a point in X, and that for all t, $S_t\varepsilon W_t$ implies $S_t \subset X$.

The problem in game theory is to determine a preference ordering on the sets $S_t\varepsilon W_t$ which is generated by the preference ordering $\gtrsim_t$ on points in X.

1.1 An Example

As an example to clarify this formulation of a game, consider the following two person game.

	B1	B2
A_1	I	II
A_2	III	IV

(1)

Then:

$$X = \left\{(P_I, P_{II}, P_{III}, P_{IV},) \mid \sum_{i=I}^{IV} P_i = 1, P_i \in [0,1]\ i = I, II, III, IV\right\}$$

$\geq_A$ is determined by $\sum_{i=I}^{IV} P_i U_A(i)$ and similarly for B.

$$W_A = \left\{ S_A^P \mid P \in [0,1] \right\} \text{ where}$$

$$S_A^P = \left\{ (P_I, P_{II}, P_{III}, P_{IV}) \mid \sum_{i=I}^{IV} P_i = 1; P_i \in [0,1]\ i = I, II, III, IV;\ P_I + P_{II} = P = P_I/[P_I + P_{III}] \text{ (independence)} \right\}$$

$$W_B = \left\{ S_B^P \mid P \in [0,1] \right\} \text{ where}$$

$$S_B^P = \left\{ (P_I, P_{II}, P_{III}, P_{IV}) \mid \sum_{i=I}^{IV} p_i = 1; p_i \in [0,1]\ i = I, II, III, IV;\ P_I + P_{III} = P = P_I/[P_I + P_{II}] \right\}$$

$$S_A^\alpha \wedge S_B^\beta = [\alpha\beta, \alpha(1-\beta), (1-\alpha)\beta, (1-\alpha)(1-\beta)]$$

2. Dominance

The problem of game theory is the extension of preferences over events to preferences over sets of events.

There is one unambigious extension of preferences over events to preferences over sets of events, dominance. let $>_t$ be the strict preference of individual $t \varepsilon T$ defined in the obvious manner. For sets of events S, S^1 we define (strong strict) dominance $>>_t$ by $S >>_t S^1$ if and only if $s \in S$ and $s^1 \varepsilon S^1$ implies $s >_t s^1$. If the individual picks S over S^1 no matter what happens the individual is better off. If S, S^1 are not strictly ranked by dominance, they are indifferent.

While this use of dominance is standard, it is worth noting that it differs from standard game theoretic use. Consider, for example, the two person game [3,p.123]

	B1	B2
A1	(10,10)	(-10, 11)
A2	(11,-10)	(1,1)

(2)

In standard game theory use A2 dominates A1 as $11 > 10$ and $1 > -10$, but not by our use as $1 < 10$. The intuitive notion behind the standard game theory use is that whatever individual B does A is better off playing A2. This is correct, but it does not imply that playing A2 necessarily makes A better off. The individual A is not choosing 11 over 10 or 1 or -10, but between the set $\{11, 1\}$ and the set $\{10, -10\}$. The "natural" pairing is irrelevant. To further convince oneself of this, consider the game where the pairing is reversed. **(2) Illustrates the problem with the standard game theory use of dominance. Players achieving (10,10) can hardly be viewed as making an error!**

Our use of dominance is the natural extension of preference over events to preferences over sets of events when there is no measure inherent to the set of events. It is in this sense on a par with expected utility as the natural extension of preferences when there is an intensity of preference and an inherent measure over sets, probability measure.

3. Iterative Contraction by Dominance

Dominance generates a preference ordering over the sets of events between which individuals choose. However, the dominance criterion can be pushed a little harder to generate an equilibrium solution concept for the game. Consider the simple game.

	B1	B2
A1	(0,0)	(-5,1)
A2	(1,0)	(-6,1)

(3)

B2 dominates B1, but A1 and A2 are indifferent, not ranked by dominance. However, as individual B will not play B1 individual A can consider the contracted game

	B2
A1	(-5,1)
A2	(-6,1)

in which A1 dominates A2. Dominance can be used to iteratively contract a game.

Let us formalize this iterative contraction by dominance. To do so, we must introduce some additional notation. Let $W = \underset{t\varepsilon\tau}{X} W_t$ be the cross product of the collections of sets of events to which individuals can restrict the world. For any cross product of collections of sets $C = \left\{ \underset{t\varepsilon\tau}{X} C_t \mid C_t \subseteq W_t \right\}$ define $R(C) = \left\{ s \middle| s = \bigcap_{\iota\varepsilon\tau} S_t,\ S_t \varepsilon C_t \right\}$. R (C) is then a set of events in X. For each t we define a deletion function as follows: $F_t(C) = \bigcap_{C_t^1 \varepsilon C_t^*} C_t^1$ where

$$C_t^* = \left\{ C_t^1 \;\middle|\; C_t^1 \subseteq C_t,\ R[C_t^1\ X(\underset{\substack{v\varepsilon\tau \\ v\neq t}}{X} C_v)] >>_t R[C_t \setminus C_t^1 X(\underset{\substack{v\varepsilon\tau \\ v\neq t}}{X} C_v)] \right\}.$$

We define the deletion function to be $F(C) = \underset{t\varepsilon\tau}{X} F_t(C)$. Then $F(C) \subseteq C$. Notice that if $F_t(C) \subset C_t$ then $F(C) >>_t [C \backslash F(C)]$. Now we are ready to formalize the contraction process. We define the collection of cross products of collection of sets $\tilde{C}$ using the following three properties: (1) $W \varepsilon \tilde{C}$ (2) $C \varepsilon \tilde{C}$ implies $F(C) \varepsilon \tilde{C}$ (3) $\{C\} \subseteq \tilde{C}$ implies $\cap\{C\} \varepsilon \tilde{C}$. Then $C \varepsilon \tilde{C}$ if and only if $C \not\varepsilon \tilde{C}$ implies $\tilde{C}$ does not satisfy properties (1), (2) and (3). Notice that $\tilde{C}$ is a nested collection and if C', $C'' \varepsilon \tilde{C}$ and $C'_t \subset C''_t$ then $C' >>_t (C'' \backslash C')$. The solution to our iterative contraction process is $V = \underset{C\varepsilon\tilde{C}}{\cap} C$. Notice that $V = F(V) \varepsilon \tilde{C}$.

Our solution of the game does not generally produce a single event, as can be seen by inspection of game (2) above. Between the strategies of our solution V individuals are indifferent, and anything can happen. There is no prediction of a probability measure over the outcomes included in V. All that can be predicted is that one of them will occur. This we refer to as unsystematic or chaotic behavior. It is the position of this author that this prediction of chaotic behavior should not be viewed as a flaw in this solution concept, but as an important feature of the solution to game theory. A single outcome of a social situation is not implied by the solutions to simple maximization problems. We now turn to an analysis of the properties of this solution, and a comparison of it with Nash equilibrium.

4. Properties of the Solution

In the analysis of the properties of the solution of iterative contraction by dominance we proceed a bit more formally with a set of propositions.

Proposition 1: If $W \neq \phi$ then $V \neq \phi$.

Proof: $C \neq \phi$ implies $F(C) \neq \phi$.

A prediction that nothing occurs is meaningless, and our solution never is this prediction.

A most important characterization of a solution is how it relates to the standard solution concept of Nash equilibrium. First we define Nash equilibrium in the above notation, then discuss the relation of Nash equilibrium and the proposed solutions, and finally demonstrate that the proposed solution satisfies a generalized equilibrium property.

4.1 Nash Equilibrium

First we must define Nash equilibrium. $s^e = \bigcap_{t \varepsilon \tau} S_t^e$ for $S_t^e \in W_t$ for all t is a Nash equilibrium if and only if for all t $s^e \geq_t (\bigcap_{\substack{v \varepsilon \tau \\ v \neq t}} S_v^e) \bigcap S_t$ for all $S_t \varepsilon A_t \backslash S_t^e$. Strict preference defines strong Nash equilibrium. The relation of Nash equilibrium to the solution of iterated contraction by dominance is provided by the following propositions.

Proposition 2: If s^e is a Nash equilibrium $s^e \varepsilon R(V)$.

Proof: Consider a Nash equilibrium $s^e = \bigcap_{t \varepsilon \tau} S_t^e$.

Consider the collection of cross products of collections of sets $\hat{C} = \left\{C' \mid C' \varepsilon \tilde{C} \text{ and } s^e \varepsilon R(C')\right\}$. Then for all $\bar{C} \subseteq \hat{C}$ $F(\bigcap_{C' \varepsilon \bar{C}} \underline{C}') \varepsilon \hat{C}$. Therefore $\hat{C} = \tilde{C}$.

All Nash equilibrium are included in our solution.

Proposition 3: If $R(V) \varepsilon X$, V is a strong Nash equilibrium.

Proof: Suppose V is not a strong Nash equilibrium. Then there exists for some $t \varepsilon \tau$ $S_t \varepsilon W_t$, $S_t \neq V_t$ such that $X^0 = S_t \cap (\bigcap_{\substack{j\varepsilon\tau \\ j\neq t}} V_j) \geq_t R(V)$. Consider $\hat{C} = \{C^1 / C^1 \varepsilon \tilde{C} \text{ and } S_t \varepsilon C^1_t\}$. As in the proof of proposition 2 $\hat{C} = \tilde{C}$, so $V \supset V$-- a contradiction.

If our solution is a single event, that event is a strong Nash equilibrium.

Proposition 4: If s^e is the unique (strong or otherwise) Nash equilibrium of a game $R(V) = s^e$ does not necessarily hold.

Proof: Consider the following two person game

	B1	B2	B3
A1	(10,10)	(5,11)	(1,1)
A2	(9,7)	(6,6)	(1,1)
A3	(1,1)	(1,1)	(2,2)

(4)

$A3 \cap B3$ is the unique (strong) Nash equilibrium but V is the entire game.

4.2 Equilibrium in Dominance

We define a generalized equilibrium with respect to dominance. Then it is shown that both strong Nash equilibrium and the solution of iterated contraction by dominance are such equilibria.

First let us define equilibrium with respect to dominance. $C \subseteq W$ is an equilibrium if and only if all $t \varepsilon \tau$: (1) $C_t \subseteq W_t$, (2) $R[C_t \times (\mathop{X}_{\substack{v\varepsilon\tau \\ v\neq t}} C_v)] \gg_t R\,[W_t \backslash C_t \times (\mathop{X}_{\substack{v\varepsilon\tau \\ v\neq t}} C_v)]$ and (3) C_t' satisfies (1) and (2) implies $C_t \subseteq C_t'$.

Proposition 5: For $C = \underset{t\varepsilon\tau}{X} S_t$, $S_t \varepsilon W_t$ C is an equilibrium if and only if C is a strong Nash equilibrium.

Proof: This follows immediately before from the definitions.

Proposition 6: V is an equilibrium.

Proof: (i) The first property of equilibrium that $V_t \subseteq W_t$ holds by construction.

(ii) Consider a violation of the second property of equilibrium: $R\,[V_t X(\underset{\substack{j\varepsilon\tau\\ j\neq t}}{X} V_j)] \not\succ_t R[W_t \setminus V_t X\,(\underset{\substack{j\varepsilon\tau\\ j\neq t}}{X} V_j)]$ for some t. By convention if $W_t \setminus V_t = \emptyset$ the strict preference does hold. For $V_t \neq W_t$ there exists $S_t \not\in V_t$, $S_t \varepsilon W_t$, such that for some sets $S_j^u \varepsilon V_j$ $j\varepsilon\tau$ and some $v\varepsilon R(V)$ $S_t \cap (\bigcap_{\substack{j\varepsilon\tau\\ j\neq t}} S_j^u) \geq_t v$. Consider $\hat{C} = \{C' \mid C' \varepsilon \tilde{C}$ and $S_t \varepsilon C_t'\}$. As in the proof of proposition 2 $\hat{C} = \tilde{C}$, so $V \supset V$ - - a contradiction.

(iii) Consider a violation of the third property of equilibrium: for some $C_t^1 \not\supseteq V_t$ $R[C_t^1\ X(\underset{\substack{j\varepsilon\tau\\ j\neq t}}{X} V_j)] \gg_t R[W_t \setminus C_t^1\ X(\underset{\substack{j\varepsilon\tau\\ j\neq t}}{X} V_j)]$. Then there exists $S_t \not\in V_t$, $S_t \varepsilon C_t^1 \subseteq W_t$ such that for some sets $S_j^u \varepsilon V_j$ $j\varepsilon\tau$ and some $v\varepsilon R(V)$ $S_t \cap (\bigcap_{\substack{j\varepsilon\tau\\ j\neq t}} S_j^u) \geq_t v$. Proceed as in (ii) above.

As many have observed, Nash equilibrium is an incomplete solution concept. However, Nash equilibrium is incomplete not because it isolates too many points as solution, but too few. There is not in general enough structure imposed in a game to imply Nash equilibrium as solution. We now turn to a specific criticism of Nash equilibrium to make this point.

4.3 A Criticism of Nash Equilibrium

Nash equilibrium is generated by a "logical process" that involves knowing the outcome of that process, clearly a non sequitur. Nash equilibrium is the result of a line of reasoning "If I knew they were going to do X then I would do Y," and then finding a fixed point of this process. However, it is part of the structure of the game that one does not know beforehand what the other payers are going to do. The alleged process starts with a conjecture which does not hold by construction. Nash equilibrium isolates events which are consistent with preferences, but not implied by preferences.

The logic behind Nash equilibrium seems to be as follows. If there is a solution to the game in the sense of a choice of a single event, then that event must be "consistent" with preferences. Therefore Nash equilibrium, the consistency criterion, isolates all candidates for solution. The problem is then viewed as isolating one of the typically multiple Nash equilibria as solution. For example, the requirement of perfection is used to delete some Nash equilibria in dynamic games. The problem with this whole logic is in the "If" part of "If there is a solution". The subsequent observations never erase the "If". The consistency criterion is that if I knew they were going to play a particular Nash equilibrium, then I would too, and this holds for everyone. However, that statement does not imply that therefore all individuals are going to play the Nash equilibrium. Pushing the arguement a bit farther, if I knew they were going to play a particular way that could only be because they would do so. In saying "I knew" the solution is already being asserted.

5. Weaker Preferences Over Sets of Events

Now let us consider whether the criterion for preferences over sets of events can be weakened. Such a weakening, by reducing the "area of indifference" between sets of events, might reduce the scope of unsystematic or chaotic behavior. Such a weakening might reduce the size of the solution.

4.1 Weaker Extensions of Preferences

Alternative extensions of preferences over events seem to be a reasonable place to start an analysis of alternative preferences over sets of events. Game theory is capturing the notion that individuals are ultimately concerned with realized parameters of the world (outcomes). It seems reasonable to maintain that notion. Preference orderings are a summary of individuals behavior with regard to gambols on those parameters. In particular note that we have not assumed preferences to be independent, just taken them as given.

We should start our analysis of alternative preferences over sets with alternative extensions of preferences over events. Individual behavior should, then, depend upon events in the manner that these are reflected in the preference orderings. There is no inherent measure on the sets of events making up strategies. Rankings of strategies should, then, depend only upon the preference orderings of the least and most preferred events in the sets (or appropriate limits). Clearly our dominance criterion is in this class. By the (strong, strict) dominance criterion one set is chosen over another if the least preferred points of the chosen set are strictly preferred to the most preferred points of the rejected set.

There are several such weaker extensions of preference. Consider, for example, weak strict dominance. A set S weakly strictly dominates a set S' if $s \in S$ and $s' \in S'$ implies $s \geq_t s'$ with strict preference holding for some s in S and some s' in S'. Weak strict dominance says that one set is chosen over another if the least preferred points of the chosen set are preferred to the most preferred points of the rejected set, and the most preferred points of the chosen set are strictly preferred to the least preferred points of the rejected set. Other common criteria are minimax which orders sets by their least preferred points and maximin which orders sets by their most preferred points.

A problem with such weaker preference criteria is that they do not tend to produce equilibrium solutions. Consider, for example, the simple two person zero sum game

	B1	B2
A1	1	-1
A2	1	1

(5)

With contraction by weak strict dominance the solutions is $A2 \cap B2$, but given A2 individual B should be indifferent between B1 and B2. Or consider the following two person game:

	B1	B2
A1	(-1,-1)	(2,0)
A2	(1,1)	(0,0)

(6)

By the minimax criterion $A2 \cap B2$ is the solution but individual A prefers $A1 \cap B2$ and individual B prefers $A2 \cap B1$. By the maximin criterion $A1 \cap B1$ is the solution but individual A prefers $A2 \cap B1$ and individual B prefers $A1 \cap B2$. Our basic point is simple. Unless sets are ranked by strong strict dominance deletions of portions of these sets from consideration can remove or reverse the preference ordering. One can use a modified contraction process of modified preference criteria in which deletions which would not subsequently be overturned are made.

4.2 Weaker Preferences Over Strategies

Having considered extensions of preferences over events, the remaining possibility is preferences over strategies themselves, rather than ultimately over outcomes. Yet more is added to our description of the individual psyche. We will consider only three such additional preferences which seem consistent with the preferences over events imposed in game theory.

One example of such a consistent preference is Nash equilibrium. People may simply have a taste for a particular Nash equilibrium. By definition such a taste is consistent with preferences over events. It is, however, a single particular and peculiar taste, which should not be the basis of all game theory. Another, and to some readers perhaps more plausible, such preference is the use of the standard game theoretic dominance criterion in the iterative contraction procedure. The third and last such consistent preference which we consider is "joint dominance". $C \subseteq W$ is jointly dominant if (1) for all $t \epsilon \tau$ $R(C) \gg_t R[W \setminus C]$ and (2) C^1 satisfies (1) implies $C \subseteq C^1$. For example, in game (4) above this solution eliminates strategies A3 and B3. To some readers this solution may seem a plausible supplement to the iterative contraction by dominance solution of the original game, rather than an added preference over strategies. However, all of these examples share the same flaw as solutions to the original game. They all implicitly start with the conjecture "If I knew they were going to do X . . . ."

5. The Cooperative Game

Only noncooperative games have been treated. This is, nonetheless, a complete treatment of game theory. Cooperative games should be reformulated as noncooperative dynamic games. This idea goes back to Nash [6].

The game as a complete model is a complete description of the relevant environment. In a game at a point in time all players decide which $S_t \varepsilon W_t$ to restrict the world to. At this point in time all previous conversations, agreements, and so on, are irrelevant. Therefore, cooperation is impossible. The fact of having to make independent choices at a point in time, the basic structure of the game itself, rules out cooperation in that choice. The cooperative game violates a basic assumption of Utilitarianism, individual choice.

This arguement does not imply that game theory cannot confront the existence of coalitions. One can have a coalition if individuals can physically bind themselves to strategies prior to the decision point of a game. But that decision to bind oneself must come at a particular previous point, a point of time in which the decision to bind can be analyzed as another noncooperative game. Coalitions appear in a sequence of noncooperative games in which decisions in early games determine the W_t's of later games. The whole sequence of games should be analyzed at the initial point as a noncooperative super game. That there may be many possible sequences of noncooperative games corresponding to a single cooperative game does not invalidate this procedure. Quite the contrary. Rather, it implies that a search for a general solution to a cooperative game is fruitless as the cooperative game is not a well-posed problem. The cooperative game provides only a partial description of the relevant environment. Loosely speaking, in removing the key assumption of independent choice by allowing binding contracts, the cooperative game renders the model incomplete. Additional structure must be provided to replace the deleted assumption.

It is worth noting that the preceding arguement immediately implies that in dynamic games only perfect equilibria should be treated (although the equilibrium should be our solution, not Nash equilibrium). A person can no more cooperate with herself at a later date than can two individuals cooperate. This point is also made in [1].

6. Concluding Comments

In many textbook games our proposed solution is the entire game. All that can be said is that one of the outcomes will occur. This may seem to provide a very weak tool for economic analysis. We may have to face the fact that self interest does little to determine the outcome in a social situation. It is not obvious that this is correct, however.

In the first place, textbook games may be rarely observed pathologies. For example, the game theoretic foundation of competitive equilibrium provided in [2] uses Nash equilibrium. However with a slight strengthening of assumptions the iterative contraction by dominance solution is the same.[1/] This competitive solution is an example of the final offer [4,7,9,10], which is one approach to cooperative action. In the final offer a trader uses a technology of openly priorly binding herself to a strategy consisting of no trade and a set of trades. In contrast to the competitive situation in the simple bargaining problem the player making the final offer chooses the monopoly solution. Oligopolistic specifications do generate chaos as solutions because by our dominance criterion only technologies which generate the competitve or monopolistic solutions are systematically used. The properties of the environment which determine who must or can and does systematically use what final offer technology become an interesting subject for study.

Indeed, one implication of game theory for economics is that technologies of exchange are part of the basic structure of a model like tastes, endowments and productions technologies. Typically economic models simply specify the latter plus one strategy: no trade. Each individual is assumed to be able to use her technologies on her endowment and consume the product. If that were the only strategy available, the problem would be well posed. However, it is implicitly assumed in such models that other unspecified strategies are

[1/] For example, all firms are potential entrants.

also available in that no trade is not taken as the solution. As the strategies are not specified the models are not complete, we do not have a well posed problem. The environments underlying economic models not only have the potential for indeterminacy implied by our proposed solution to game theory, but the indeterminacy of being incompletely specified.

In the face of the indeterminacy in their underlying environments economic models typically impose a particular event as "solution." A frequently chosen, and exhaustively studied, "solution" is competitive equilibrium. From a positive viewpoint, this approach is appropriate: the competitive equilibrium frequently exists and is unique. However, one cannot assert that the underlying environment generates the competitive equilibrium. Admittedly competitive equilibrium may be in a weak sense consistent with a complete underlying environment. If in an underlying environment the strategies involving trade do imply a single event as our contraction by dominance solution, any individual engaging in trade must strictly prefer that event to no trade. At the same time in that environment all individuals prefer (not necessarily strictly) the competitive equilibrium event to the no trade event. This is very far indeed from an assertion that the actual free market mimics competitive equilibrium. [8] and [12] insist on the importance of technologies of exchange. However, a complete model has to generate strategy sets endogenously, a non-trivial task. Strategy sets are determined by the environment.

The ability of individuals to alter the environment in reaction to the possibility of unsystematic behavior might seem an interesting subject for study. Economic agents may try to eliminate unsystematic behavior from the environment. However we must preduct that economic agents' only <u>systematic</u> behavior in this regard is to so alter the environment when the result is preferred to the best result of the chaotic game, and not to do so if the result is worse than the worst result of the chaotic game. Of course the environment may generate little unsystematic behavior if, for example, the realizations of outcomes in chaotic games involve the creation of institutions which eliminate future chaotic games from the environment.

In the second place, chaos may be the correct prediction.

References

* The author wishes to thank Gregory Ballentine, Gerard Debreu, Jack Gelfand, John Harsanyi, Charles Holt, Edward Prescott and Stephen Salant for valuable comment. The views expressed herein do not represent those of Ballentine, Debreu, Gelfand, Harsanyi, Holt, Prescott or Salant, and errors and oversights are my responsibility alone.

[1] Bryant, J., "Should Control Theory Be Used for Economic Stabilization?: A Comment" in K. Brunner and A. Meltzer, eds., Optimal Policies, Control Theory and Technology Experts, Journal of Monetary Economics, 7, Amsterdam, North-Holland, 1977.

[2] __________, "Competitive Equilibrium with Price Setting Firms and Stochastic Demand", International Economic Review, 21 (October 1980), 619-626.

[3] Harsanyi, J., Rational Behavior and Bargaining Equilibrium in Games and Social Situations, Cambridge, Cambridge University Press, 1977.

[4] Holt, C., "Competitive Bidding for Contracts Under Alternative Auction Procedures", Journal of Political Economy 88, 1980, 433-445.

[5] Luce, R. and H. Raiffa, Games and Decisions, New York, John Wiley & Sons, c. 1957.

[6] Nash, J., "Non-cooperative Games", Annals of Mathematics, 54 (1951), 286-295.

[7] Samuelson, W., "First Offer Bargains," Management Science (forthcoming).

[8] Stahl, I., Bargaining Theory, Stockholm, The Economic Research Institute, 1972.

[9] Thompson, E., "Characteristics of Worlds with Perfect Strategic Communication," Journal of Economic Theory, 23 (1980), 111-119.

[10] Townsend, R., "Equilibrium with Endogenous Marketeers," working paper, Carnegie-Mellon University, 1980.

[11] Von Neumann, J. and O. Morgenstern, Theory of Games and Economic Behavior, Princeton University Press, c. 1944.

[12] Zabel, E., "Competitive Price Adjustment without Market Clearing," Econometrica (forthcoming).