Federal Reserve Bank of Minneapolis
Research Department

Dynamic Games With Hidden Actions and Hidden States*

Harold L. Cole and Narayana Kocherlakota

Working Paper 583
July 1997

ABSTRACT

We consider the large class of dynamic games in which each player's actions are unobservable to the other players, and each player's actions can influence a state variable that is unobservable to the other players. We develop an algorithm that solves for the subset of sequential equilibria in which equilibrium strategies are Markov in the privately observed state.

*We would like to thank V.V. Chari, Ed Green, and Andy McLennan for comments and suggestions. The views expressed herein are the those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. See the FRB of Minneapolis website http://research.mpls.frb.fed.us for further revisions, as well as other working papers.

1. Introduction

This paper is about the large class of dynamic games in which each player's actions are unobservable to the other players and each player's actions can influence a state variable that is unobservable to the other players. (A simple example of such a game is one in which two firms privately choose production quantities of a durable good.) As yet, no general methods have been developed with which to analyze these games. In this paper we show how the methods developed by Abreu, Pearce and Stacchetti (APS) (1986,1990) to analyze dynamic games with hidden actions can be extended to this more general class.

The key simplification that lies behind the work of APS is to note that since the set of possible continuation payoffs to a repeated game is independent of the history of play, one can employ a multi-player generalization of the principle of optimality of discounted dynamic programming in order to characterize the set of equilibrium payoffs. Specifically, they show that the vector of expected payoffs to the players, v, can be decomposed into the current payoff, u, and the discounted future payoff, δw (factorization), and that given the stationarity of the set of payoffs, V, it must be the case that both v and w are elements of V. Thus for any $v \in V$ there must exist an associated action vector, q, and map w from public outcomes, p, to elements of V, which delivers the expected payoffs v, and for which the actions in q are weakly optimal for every player given w (self-generation). These notions allow them to define a monotone "backwards" operator that maps a set of payoff vectors into another set of payoff vectors. They show that the largest fixed point of this operator is the set of sequential equilibrium payoff vectors and describe how to use this set of payoffs to construct the set of all equilibrium strategies.[1]

[1] Atkeson (1991) extends their results to dynamic games in which payoffs are time separable, period payoffs

Our goal is to extend these results of APS to games in which each player has a *privately* observed state. We assume that the transition rule for the private state is a deterministic function of the publicly observable variables and the private action of the individual, and, as in APS, we assume that any public signal has positive probability, regardless of the history of past play. We restrict attention to what we term Perfect Public, Markov Private (PPMP) equilibria: sequential equilibria in which individuals' strategies depend only on their current state and the public history, and in which initial beliefs about other players' states are degenerate. In these equilibria, given the public history and the initial value of the individual's state, the current value of an individual's state is perfectly predictable along the equilibrium path.

There are two key insights that help us solve for the set of PPMP equilibria. The first is to recognize that if all players (except agent n) have a common degenerate belief about agent n's state, the vector of these common beliefs can be used as an additional public state variable. The second is to understand that given the unobservability of individual states, the set of equilibrium payoffs is no longer a useful object. Instead, we need to focus on equilibrium payoff functions, which describe the payoff to any individual player n as a function of his actual private state, taking as given the strategies of the other players.

Given these two insights, we can generalize the notions of factorization and self-generation. The analog of factorization is that we can decompose an equilibrium payoff function (of the player's current true state) into a current payoff function (of his current true state) and a discounted future payoff function (his true state next period). We show that

and transition rules are stationary functions of the state and the actions of the agents, if the state is assumed to be observable and to take on only a finite set of values.

the set of possible vectors of equilibrium payoff functions depends on the history of past play only through the vector of common current beliefs about players' states. The correspondence from initial belief vectors into the set of vectors of equilibrium payoff functions then satisfies a property that corresponds to self-generation. Specifically, there must exist a specification of weakly optimal actions as a function of the individual's private state and a map from public outcomes p to continuation equilibrium payoff functions which yield the requisite expected payoff function, and for which the beliefs about individual's private states are updated in a manner consistent with the action functions.

Using these extended notions of factorization and self-generation, we can define an analog of APS' "backwards" operator. The range and domain of this operator are both equal to the set of all correspondences that map the set of vectors of (point-mass beliefs about) initial states into the set of vectors of payoff functions. By iterating the operator, we can find its largest fixed point; this largest fixed point is the correspondence that maps each possible vector of initial beliefs into a set of equilibrium payoff functions. Once we have this correspondence, it is straightforward to describe the set of all PPMP equilibrium strategies.

In the remainder of the paper, we define the class of games under study, give a couple of examples, define a natural refinement of sequential equilibrium for these games, and then present a self-generation algorithm to find the set of equilibria. We close with some concluding comments.

2. Definition of the Class of Games

Define an infinite horizon N player game as follows. The N players have common discount factor δ. At the beginning of play in period $(t+1)$, $t \geq 0$, player n's state is denoted by $s_{nt} \in S$, where S is finite. The players' states are private information. Assume that for any player i, all of the other players' beliefs as to player i's initial state are common and degenerate: they believe that his state is $\hat{s}_{0i}$. Let $\hat{s}_0$ denote the vector of these beliefs.

In period t, each player n chooses an unobservable action q_{nt} from a finite action set Q. Nature then chooses an element p_t from a finite set P according to the density $\pi(p;q_t, s_{t-1})$, where q_t is the vector of player actions and s_{t-1} is the vector of player states.. The probability density assigns positive probability to any element of P for all vectors of player actions and states. The realization of the variable p_t is common knowledge.

Agent n then realizes period t payoff $u_n(q_{nt}, p_t, s_{n,t-1})$, where $u_n : Q \times P \times S \rightarrow [0, (1-\delta)]$. His state is updated according to the law of motion

$$s_{nt} = g_n(q_{nt}, p_t, s_{n,t-1})$$

where $g_n : Q \times P \times S \rightarrow S$. We assume that for all s'_n in S, p in P, and s_n in S, there exists $q \in Q$ such that $s'_n = g_n(q, p, s_n)$; this means that regardless what player n's state is, it is physically possible for the state to be any element of S. Note that the updating of the state variable is deterministic.

The specification of the game is fairly general, but there are two key limitations. First, the support of the observable public signal p_t is independent of the unobservable actions of the individuals. This restriction, combined with the assumption that individuals' states are common knowledge at the beginning of play, means that seeing a particular realization of p_t does not alter an individual's prior about the other players' period $(t-1)$ states. Second, all signals received by the agents are public information.

3. Examples

The following are examples of the general type of game described above.

A. Pollution Game

Consider two countries. In period t, they simultaneously choose production levels y_{1t} and y_{2t} from the finite set Y; country 1 never observes country 2's choice, and vice-versa. After the production choice is made, country 1 experiences an observable level of pollution p_{1t}, which is drawn from the finite set P according to the positive density $\pi(p; y_{1t}, y_{2,t-1})$ and country 2 experiences an observable level of pollution p_{2t}, which is drawn from the same set P according to the positive density $\pi(p; y_{2t}, y_{1,t-1})$. In words, country 1's level of pollution is affected by its current level of production and country 2's level of production last period; intuitively, it takes time for the two countries' level of production to affect each other's pollution levels. Country n gets momentary utility $u(y_{nt}, p_{nt})$, where u is increasing in its first argument and decreasing in its second.

B. Cournot Duopoly with Durable Goods

Consider two firms which both produce the same good. In period t, firm i chooses an quantity of production q_{it} from the quantity set Q; this choice is not observable to the other firm. This amount of production affects the stock s_{it} of firm i's goods, according to the law of motion $s_{it} = g(q_{it}, s_{i,t-1})$. Nature then draws a publicly observable price from a set of prices P;the good is durable, so the probability density $\pi(p; s_{1t} + s_{2t})$ depends on the total stock of goods produced by the two firms. The objective function of firm i is given by $u(q_{it}, p_t, s_{i,t-1}) = p_t q_{it}$.

4. Equilibrium

When S is a singleton, APS (1990) look at a subset of sequential equilibria (Kreps and Wilson (1982)) that Fudenberg and Tirole (1991) call perfect public equilibria. In these equilibria, agents' strategies are restricted to be functions only of the history of public signals, and not their past private actions. Fudenberg and Tirole (1991) argue that the set of sequential equilibrium payoffs is equal to the set of perfect public equilibrium payoffs; in this sense, there is no loss of generality in concentrating on the latter.

Similarly, we focus on a subset of sequential equilibria in which agent strategies are restricted to be functions only of the history of public signals and their current private states. More precisely, given a game, define a strategy σ_n for agent n to be a collection of functions $(\sigma_{nt})_{t=1}^{\infty}$ such that σ_{nt} maps agent n's information set $(s_{n,0}, p_1, ...p_{t-1}, q_{n1}, ..., q_{n,t-1})$ into an action choice q_{nt} for period t. We define a strategy σ_n to be *public and Markov private* if for all t, σ_{nt} maps any information sets with the same values for $(p_1, ..., p_{t-1}, s_{n,t-1})$ into the same action. Then, we use the following definition of equilibrium.

Definition 1. *A* ***Perfect Public, Markov Private Equilibrium (PPMPE)*** *is a collection of N public and Markov private strategies such that agent n's strategy prescribes a weakly optimal action at every information set, taking as given other agents' strategies and given that at any information set, agent n believes that all other agents have always followed their equilibrium strategies. (Note that if S is a singleton, then the set of PMPPE coincides with the set of PPE.)*

We have required that strategies be defined for every possible initial state of the players. Consequently, the strategies imply the appropriate updating rule for beliefs. Given the strategy of player n, and the initial belief as to his state $\hat{s}_{n0}$, the other players update their beliefs as to player n's private state according to

$$\hat{s}_{n,t} = g_n(\sigma_{n,t-1}(p_1, ..., p_{t-1}, \hat{s}_{n,t-1}), p_t, \hat{s}_{n,t-1}).$$

Note that since every possible public signal occurs with positive probability in every period, this defines beliefs over all possible public histories.

When the state space S is a singleton (that is, there are no hidden state variables), there is a vector of equilibrium payoffs associated with each PPMPE. However, we shall see that in a game with hidden state variables, it is more useful to consider the vector of *equilibrium payoff functions* associated with a given PPMPE. In particular, suppose N agents have initial believed state $\hat{s}$. Let $(\sigma_n^*)_{n=1}^N$ be a PPMPE. Define the equilibrium payoff function v_n to be:

$$\begin{aligned} v_n &: S \to [0,1] \\ v_n(s_n) &= \max_{\sigma} U_n(\sigma_n; s_n, \hat{s}, \sigma_{-n}^*) \end{aligned}$$

where U is the period zero payoff derived by agent n, given an arbitrary strategy σ_n, his current state s_n, the beliefs of the players as to each other's initial state, and the strategies of the other agents. Hence, the equilibrium payoff function v_n describes the agent's maximal utility in period 0 for each initial state s_n, given that all other agents play their equilibrium strategies.

5. Algorithm

Suppose there are no state variables (so S is a singleton). Then, every PPMPE has associated with it a vector of equilibrium payoffs. Let D be the set of all such vectors. We know that after any history of realizations of the public signal, the continuation strategies of the agents form a PPMPE; hence, the vector of continuation payoffs must lie in D. It is exactly this recursive structure that APS (1990) exploit in designing a solution algorithm for such games.

With state variables, the situation is complicated in two ways. First, in order to calculate the consequences of his actions, an individual needs to know not just his continuation equilibrium payoff, but also his continuation equilibrium payoff given that he starts with any possible state next period. Hence, as discussed above, given a PPMPE, we need to know the associated equilibrium payoff functions, not just the associated equilibrium payoffs. Second, the set of vectors of equilibrium payoff functions depends on the current beliefs of the agents. Hence, our goal is to solve for the set of equilibrium payoff functions given any possible vector of initial beliefs.

It is useful to formalize the problem as follows. Let $F(S, [0,1])$ be the set of functions from the state space S (of private states) to the set $[0,1]$; equilibrium payoff functions will lie in $F(S,[0,1])$. Let Γ be the set of all correspondences mapping the set of possible beliefs, S^N, into $(F(S,[0,1]))^N$; an element of Γ thus links initial beliefs about states to sets of payoff function vectors. Our goal is to find $\gamma^* \in \Gamma$ such that $\gamma^*(\hat{s})$ is the set of all PPMPE payoff functions when believed states are characterized by the vector $\hat{s}$ in S^N. We assume that for all for all $\hat{s} \in S^N$, $\gamma^*(\hat{s}) \neq \emptyset$.

The problem of solving for the set of equilibria can be solved by using an analogue of

APS (1990)'s backwards operator. Define B to be a mapping from the set of all subsets of Γ into itself as follows. Let γ be a subset of Γ. Then, $(v_1, v_2, ..., v_N) \in B(\gamma)(\hat{s})$ if for all n and all p in P, there exists $q_n : S \to Q$ and $(w_1, ..., w_N)(p) \in \gamma(\hat{s}'(p))$, such that for all n and all s_n in S:

$$v_n(s_n) = \int_P \begin{bmatrix} u_n(q_n(s_n), p, s_n)+ \\ \delta w_n(p)(s_n'(p, s_n)) \end{bmatrix} \pi(p; q_n(s_n), q^*_{-n}, s_n, s_{-n})dp \tag{1}$$

$$s_n'(p, s_n) = g_n(q_n(s_n), p, s_n) \tag{2}$$

$$\hat{s}_n'(p) = g_n(q_n^*, p, \hat{s}_n) \tag{3}$$

$$q_n(s_n) \in \arg\max_{q \in Q} \int_P \begin{bmatrix} u_n(q, p, s_n)+ \\ \delta w_n(p)(g_n(q, p, s_n)) \end{bmatrix} \pi(p; q, q^*_{-n}, s_n, s_{-n})dp \tag{4}$$

$$q_n^* = q_n(\hat{s}_n) \text{ and } q^* = (q_1^*, ..., q_n^*) \tag{5}$$

Equation (1) insures that the actual expected payoff to player n when his state is s_n is $v_n(s_n)$, given the actions of the other players, the law of motion for his state tomorrow $s_n'(p, s_n)$, and the continuation payoff function conditional on p, $w_n(p)$. Equation (2) ensures that the law of motion for the individual's state be consistent with his state contingent action choice, $q_n(s_n)$. Equation (3) requires the updated belief about individual n's state, $\hat{s}_n'(p)$, be consistent with the transition rule if individual n's state is what the other agents believe it to, $\hat{s}_n$, and his action is q_n^*. Equation (4) requires that the individuals state contingent action plan be weakly optimal. Equation (5) requires that beliefs about individuals actions, q^*, be consistent with the beliefs about their states, $\hat{s}$, and the individuals' state contingent action plans.

It is straightforward to prove the following lemma.

LEMMA 1. If $\gamma(\widehat{s})$ is closed for all $\widehat{s} \in S$, $B(\gamma)(\widehat{s})$ is closed for all $\widehat{s} \in S$.

Proof. Let $(v^j)_{j=1}^{\infty}$ be a convergent sequence of elements of $B(\gamma)(\widehat{s})$ with limit v^{∞}. We need to prove that $v^{\infty} \in B(\gamma)(\widehat{s})$. For each j and for all p, there exists q^j and $w^j(p) \in \gamma((g^n(q_n^j(\widehat{s}_n), p, \widehat{s}_n))_{n=1}^N)$ such that (v^j, q^j, w^j) satisfy (1-5), given $\widehat{s}$. The sequence (q^j, w^j) has a convergent subsequence, so let's assume without loss of generality that the sequence itself converges to (q^{∞}, w^{∞}). Because Q and S are both finite, there exists J such that $q_n^j(s) = q_n^{\infty}(s)$ for all n, for all s, and for all $j \geq J$. It follows that $w^j(p) \in \gamma((g^n(q_n^{\infty}(\widehat{s}_n), p, \widehat{s}_n))_{n=1}^N)$ for all $j \geq J$; because $\gamma(\widehat{s})$ is closed for all $\widehat{s}$, we conclude that $w^{\infty}(p) \in \gamma((g^n(q_n^{\infty}(\widehat{s}_n), p, \widehat{s}_n))_{n=1}^N)$. In a similar fashion, we can prove that $(v^{\infty}, q^{\infty}, w^{\infty})$ satisfies equations (1-5), given $\widehat{s}$. The lemma follows. ■

The main proposition for dynamic games is:

PROPOSITION 1. *Let* $\gamma_0 = \Gamma$; *that is,* $\gamma_0(\hat{s}) = F(S, [0,1])^N$ *for all* $\hat{s} \in S^N$; *and let* $\gamma_j = B(\gamma_{j-1})$, $1 \leq j \leq \infty$. *Then, if* $\gamma^*(\hat{s})$ *denotes the set of all PPMPE payoff functions given initial beliefs* $\widehat{s}$, *the following hold:*

i.) $B(\gamma^*) = \gamma^*$

ii.) For all $\hat{s} \in S^N$, $\gamma^*(\hat{s}) \subseteq \gamma_j(\hat{s}) \subseteq \gamma_{j-1}(\hat{s})$.

iii.) Define $\gamma_{\infty} : S^N \Rightarrow F(S, [0,1])^N$ *so that for all* $\hat{s}$ *in* S^N, $\gamma_{\infty}(\hat{s}) \equiv \cap_{j=1}^{\infty} \gamma_j(\hat{s})$. *Then,* $\gamma_{\infty} = \gamma^*$.

Proof. We prove the three parts of the proposition in order.

i.) $B(\gamma^*) = \gamma^*$: First, we can prove that for any initial belief vector $\hat{s} \in S^N$, $\gamma^*(\hat{s}) \subseteq B(\gamma^*)(\hat{s})$. Consider a PPMPE with equilibrium payoff function vector $(v_1, ..., v_N)$. In the first period, along the equilibrium path, agents play a vector of actions $(q_n^*)_{n=1}^N$; conditional on

Nature's draw of p, agents then play a PPMPE continuation equilibrium beginning in period two. Each continuation PPMPE has associated some equilibrium payoff function vector $(w_1, ..., w_N)(p)$. The equations defining B require q_n^* to be the equilibrium action choice for agent n. Moreover, the equations defining B require $v_n(s)$ to be the maximal utility agent n can get if he starts with state s instead of state s_n, and given that the other agents play the PPMPE equilibrium strategies. It follows that $(v_1, ..., v_N) \in B(\gamma^*)(\hat{s})$. It is obvious that for any $\hat{s} \in S^N$, $\gamma^*(\hat{s}) \supseteq B(\gamma^*)(\hat{s})$, which completes the proof.

ii.) For all $\hat{s}$ in S^N, $\gamma^*(\hat{s}) \subseteq \gamma_j(\hat{s}) \subseteq \gamma_{j-1}(\hat{s})$: It is obvious that if $\gamma'(\hat{s}) \subseteq \gamma(\hat{s})$ for all $\hat{s}$ in S^N, then $B(\gamma')(\hat{s}) \subseteq B(\gamma)(\hat{s})$ for all $\hat{s}$ in S^N. Since $\gamma_0 = \Gamma$ and $\gamma_1 = B(\gamma_0) \subseteq \Gamma$, $\gamma_2 \subseteq \gamma_1$, which implies that $\gamma_3 \subseteq \gamma_2$, and so forth. Since $\gamma^*(\hat{s}) \subseteq \gamma_0(\hat{s})$ for all $\hat{s}$ in S^N, this implies that $\gamma^* = B(\gamma^*)(\hat{s}) \subseteq \gamma_1$, which implies that $\gamma^* \subseteq \gamma_2$, and so forth.

iii.) $\gamma_\infty(\hat{s}) = \gamma^*(\hat{s})$ for all $\hat{s}$ in S.: It is obvious from (ii) that $\gamma^*(\hat{s}) \subseteq \gamma_\infty(\hat{s})$. Now we need to prove that $\gamma^*(\hat{s}) \supseteq \gamma_\infty(\hat{s})$. We prove this in three parts:

Part 1: If $B(\gamma) = \gamma$, then $\gamma(\hat{s}) \subseteq \gamma^*(\hat{s})$ for all $\hat{s}$. This is obvious.

Part 2: $B(\gamma_\infty)(\hat{s}) \subseteq \gamma_\infty(\hat{s})$. This follows from $\gamma_j(\hat{s}) \supseteq \gamma_\infty(\hat{s})$ for all j, which means $\gamma_{j+1}(\hat{s}) \supseteq B(\gamma_\infty)(\hat{s})$ for all j and all $\hat{s}$, which proves the result.

Part 3: $B(\gamma_\infty)(\hat{s}) \supseteq \gamma_\infty(\hat{s})$. This is somewhat more difficult. Take any $(v_1, ..., v_N)$ such that $(v_1, ..., v_N) \in \gamma_j(\hat{s})$ for all j. Then, since this implies that $(v_1, ..., v_N) \in B(\gamma_j)(\hat{s})$ for all j, so for all j, there exists some q^j and for all $p, w^j(p) \in \gamma_j((g(q_n^j(s_n), p, s_n))_{n=1}^N)$ such that $[(v_1, ..., v_N), q^j, w^j]$ satisfies (1-5), given $\hat{s}$. Since we can always choose a convergent subsequence, assume, without loss of generality, that this sequence of $\{(q^j, w^j)\}$ pairs converges, and let the limit be (q^∞, w^∞). Since Q and S are finite, the convergent sequence must have the property that there exists J such that for all n, $q_n^j(s_n) = q_n^\infty(s_n)$ for all s_n and all $j \geq J$. It follows that for all p, and all $j \geq J$, $w^j(p) \in \gamma_j((g(q_n^\infty(\widehat{s}_n), p, \widehat{s}_n)_{n=1}^N)$; because $\gamma_j(\widehat{s})$ is closed for all $\widehat{s}$, $w^\infty(p) \in \gamma_\infty((g(q_n^\infty(\widehat{s}_n), p, \widehat{s}_n))_{n=1}^N)$. Similarly, it can be shown that $[(v_1, ..., v_N), q^\infty, w^\infty]$ satisfy (1-5), given $\hat{s}$. We conclude that $(v_1, ..., v_N) \in B(\gamma_\infty)(\hat{s})$ ■

Hence, we can solve for γ^* by iterating the backwards operator B. Having solved for γ^*, it is then straightforward to actually construct the set of all PPMPE for a given game. Let $\hat{s}$ be the initial belief state vector, and pick an arbitrary element of $\gamma^*(\hat{s})$. This element of $\gamma^*(\hat{s})$ implies some vector of action functions $(q_1^*, ..., q_N^*)$; it also implies, for each realization of p, a vector of future belief states $\hat{s}'(p)$ and a vector of continuation equilibrium payoff functions.

The above discussion assumes that the horizon of the game is infinite. However, it should be clear how to extend the analysis to T period games, T finite. In the last period T,for any initial belief state vector $\hat{s}_{T-1}$,we can calculate the set of Nash equilibria; given any possible Nash equilibrium, it is straightforward to calculate the vector of equilibrium payoff functions. This gives us a mapping γ_0 from states at the beginning of period T into sets of vectors of equilibrium payoff functions. Call this mapping γ_0, and then iterate on this mapping using the backwards operator described above. The iterated set $B^{T-t}(\gamma_0)(\hat{s})$ equals the set of equilibrium payoff functions in period t, given that the initial state vector is $\hat{s}$.

6. Concluding Comments

Self-generation has proved to be a useful method of characterizing the set of possible equilibrium outcomes in repeated games (see for example APS 1986), and in dynamic games (see for example Atkeson 1991). This paper provides a way to extend the concept of self-generation to a rich collection of dynamic games in which agents are symmetrically informed about one another's preferences along the equilibrium path, but asymmetrically informed off the equilibrium path. We conjecture that the other results obtained by APS 1990 (the bang-bang characterizations of equilibrium payoffs and monotonicity in the discount factor) are also obtainable for this class of games.

Like dynamic programming, the algorithmic nature of self-generation is extremely suggestive about how to actually solve dynamic games. Admittedly, for many games, the current state of technology, combined with the high dimensionality of the objects of interest, may preclude a direct application of our approach. However, if agent action sets and state spaces are both continuous, it is possible that the dimensionality problem can be greatly ameliorated if the equilibrium payoff functions can be shown a priori to be concave. In particular, the papers of Chang (1996) and Phelan and Stacchetti (1997) suggest that in this situation, we need not keep track of equilibrium payoff functions evaluated at every point in the state space, but rather just their derivatives, evaluated at the vector of beliefs. We suspect that this "local" approach, as opposed to our "global" approach, may be of great value in an interesting set of applications.

Following Spear and Srivastava (1987)'s work, a large number of papers have exploited the basic APS idea of using utility as a summary statistic in solving dynamic mechanism design problems. Similarly, it should be straightforward to extend the basic approach of this paper to dynamic principal-agent problems in which the agent's unobservable action influences an unobservable state variable. An example of such a problem is in Cole and Kocherlakota (1997).

References

[1] Abreu, D., D. Pearce, and E. Stacchetti. 1986. "Optimal Cartel Equilibrium with Imperfect Monitoring," *Journal of Economic Theory* 39:251-69.

[2] Abreu, D., D. Pearce, and E. Stacchetti. 1990. "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58:1041-1064.

[3] Atkeson, A. 1991. "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica* 59:1069-1090.

[4] Chang, R. 1996. "Credible Monetary Policy with Long Lived Agents: Recursive Approaches," manuscript: Federal Reserve Bank of Atlanta.

[5] Cole, H.,and N. Kocherlakota. 1997. "Efficient Allocations with Hidden Income and Hidden Storage," Federal Reserve Bank of Minneapolis Working Paper 577.

[6] Phelan, C., and E. Stacchetti. 1997. "Subgame Perfect Equilibria in a Ramsey Taxes Model," manuscript, Northwestern University.

[7] Spear, S., and S. Srivastava. 1987. "Repeated Moral Hazard with Discounting," *Review of Economic Studies* 54, 599-617.