

Federal Reserve Bank of Minneapolis  
Research Department Staff Report 254

September 1998

## **Dynamic Games With Hidden Actions and Hidden States\***

Harold L. Cole

Federal Reserve Bank of Minneapolis

Narayana Kocherlakota

University of Minnesota  
and Federal Reserve Bank of Minneapolis

### **ABSTRACT**

We consider a 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 depend on private information only through the privately observed state.

\*We would like to thank V.V. Chari, Ed Green, Andy McLennan, Chris Phelan, and the participants at the 1998 SITE conference 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.

## 1. Introduction

It is natural in many economic environments to model state variables as being private information. For example, in principal-agent settings, it is hard for the principal to monitor the asset market activity of the agent (Fudenberg, Holmstrom, and Milgrom (1990); Cole and Kocherlakota (1998)). In Cournot oligopolies, the firms' quantity choices are often modelled as unobservable (Green and Porter (1984)). Given this assumption and the durability of the produced good, a firm's cumulated stock of production is an unobservable state variable.

Because of this consideration, the purpose of this paper is to extend the methods of Abreu, Pearce and Stacchetti (APS) (1986, 1990) to a class of dynamic games in which the players have hidden states as well as hidden actions. APS consider repeated games with hidden actions and show that the set of equilibrium payoffs for such games is the largest fixed point of a monotone set-valued operator. They prove that the largest fixed point is the limit of iterating this operator on a sufficiently large initial set of payoffs.

We derive a similar result for a class of dynamic games with hidden states. We focus on Markov-private equilibria, in which an individual's strategy is constrained to depend on his past private information only through his current private state. Given such an equilibrium, each individual's payoff is a function of his true initial state. We consider the set of pairs  $(v, \mu)$  that are consistent with Markov-private equilibria, where  $v$  is a vector of equilibrium payoff functions and  $\mu$  is the initial common prior over vectors of individual states. We show that this set is the largest fixed point of a monotone set-valued operator, and we demonstrate that the largest fixed point is the limit of iterating this operator on a sufficiently large initial set of  $(v, \mu)$  pairs.

The main difficulty that we face is formulating exactly what class of dynamic games

is susceptible to our approach. The key restriction turns out to be that beliefs must be Markov: that is, if all players follow Markov-private strategies, player  $i$ 's beliefs about the other players' states must depend on his private information only through his current state, both on and off the equilibrium path. (Otherwise, private information other than his state would typically affect his best response, meaning that it is not Markov-private.) We provide a simple procedure to check whether beliefs are Markov, and we present examples of such games.

The restriction that beliefs be Markov still allows for a wide range of behavior on the part of beliefs. On the equilibrium path, a player knows his own state, while others do not. Moreover, the beliefs of the less informed players evolve in nontrivial ways.

## 2. A Class of Dynamic Games with Hidden States

There are  $N$  players who play an infinite horizon dynamic game. The structure of the game is as follows. Define  $S_i$  to be player  $i$ 's finite state space, and let  $S = \prod_{i=1}^N S_i$ . In period 0, Nature draws a vector  $s_0$  from  $S$  according to the probability density function (p.d.f.)  $\mu_0$ , which is common knowledge. Player  $i$  observes  $s_{i0}$ .

Each player  $i$  begins period  $t > 0$  with state  $s_{i,t-1}$ . The play in period  $t$ ,  $t > 0$ , proceeds as follows. First, player  $i$  chooses  $q_{it}$  from the finite set  $Q_i$ . Let  $q_t = (q_{it})_{i=1}^N$ . Nature then draws a public shock  $z_t$  from the finite set  $Z$  and a vector of private shocks from the set  $\Theta = \prod_{n=1}^N \Theta_n$ , according to the joint p.d.f.  $\pi(z_t, \theta_t | q_t, s_{t-1})$ . Here,  $q_t = (q_{it})_{i=1}^N$  and  $s_t = (s_{it})_{i=1}^N$ . Player  $i$  observes  $(z_t, \theta_{it})$ . At the end of the period, agent  $i$  realizes payoff  $u_i(\theta_{it}, z_t, q_{it}, s_{i,t-1})$ . The players have common discount factor  $\delta$ . Player  $i$ 's state evolves according to the law of motion  $s_{it} = g_i(\theta_{it}, z_t, q_{it}, s_{i,t-1})$ .

At this point, the class of games under consideration is very large. Players can receive different signals about each other's actions or states. These signals may be public or not. The game may be repeated or dynamic.

We restrict the class of games somewhat by making a relatively standard assumption about the probabilistic structure of the game: the support of the signals/shocks is unaffected by the actions or states of the individuals:

For any  $(q, s)$ ,  $\pi(z, \theta|q, s) > 0$  for all  $\theta \in \Theta$  and  $z \in Z$ .

This assumption is similar to the assumption made in APS (1990) about the support of the public signal. Indeed, this kind of assumption is standard in unobservable action (i.e., moral hazard) settings.<sup>1</sup>

Throughout, it is helpful to keep in mind the following simple example of the general structure.

*Example 1.* Suppose there are two firms that produce the same good. Firm  $i$  begins period  $t$  with a state variable  $s_{i,t-1}$  that affects the firm's cost function. The initial states  $(s_{10}, s_{20})$  are drawn from the joint p.d.f.  $\mu_0$ . The two firms simultaneously choose quantities  $(q_{1t}, q_{2t})$ . Nature then draws the price of the good from a p.d.f.  $\pi_p(p_t|q_{1t}, q_{2t})$ . The price is publicly observable. Firm  $i$ 's period  $t$  payoff is given by  $p_t q_{it} - c(q_{it}, s_{i,t-1})$ , where  $c$  is a cost function. The state evolves exogenously according to the law of motion  $s_{it} = g_i(\theta_{it}, s_{i,t-1})$ , where  $(\theta_{1t}, \theta_{2t})$  are drawn from a joint p.d.f.  $\pi_\theta$ . The cost shocks  $(\theta_{1t}, \theta_{2t})$  are i.i.d. over time, and firm  $i$  only observes  $\theta_{it}$ .

---

<sup>1</sup>Note that because all elements of  $\Theta$  have positive probability, the intersection over  $i$  of the  $\sigma$ -algebras generated by  $(z, \theta_i)$  is the  $\sigma$ -algebra generated by  $z$ . Hence, the only information common to all players in period  $t$  is the history of public signals  $(z_1, \dots, z_t)$ .

This example is essentially the Green and Porter (1984) duopoly model modified to include Markov cost shocks that are private information.

### 3. Markov Beliefs and Markov-Private Equilibria

The APS (1990) algorithm applies to games in which  $S_i$  and  $\Theta_i$  are singletons for all  $i$ , so all shocks are public. We want to develop a similar algorithm for games in which some of these conditions are relaxed (especially the restrictions on the state space). But there are necessarily some limitations to our analysis. First, if we allow equilibrium strategies to depend on a player's entire private history. Then we have to keep track of other players' beliefs about that history. The set of such beliefs will have an arbitrarily high dimension as time passes. Thus, if we are to have a recursive algorithm that solves for the set of equilibrium payoffs, we must restrict attention to equilibria in which strategies are required to depend on some finite-dimensional summary statistic of private histories. A natural summary statistic is the physical state itself.

Given this reasoning, we define a strategy  $\sigma_i$  for agent  $i$  to be a collection of functions  $(\sigma_{it})_{t=1}^{\infty}$  such that  $\sigma_{it}$  maps agent  $i$ 's history  $(s_{i0}, \theta_i^{t-1}, q_i^{t-1})$  into an action choice  $q_{it}$  for period  $t$ . Let  $z^t = (z_{\tau})_{\tau=1}^t$  be the public history. We define a strategy  $\sigma_i$  to be *Markov-private* if for all  $t$ ,  $\sigma_{it}$  maps any histories with the same values for  $(z^{t-1}, s_{i,t-1})$  into the same action. Hence, in a Markov-private strategy, players can condition their actions arbitrarily on the public history, but their actions can only depend on their private history through their current private state.

Restricting attention to equilibria in which players use Markov-private strategies is only sensible under relatively strong conditions on the evolution of beliefs. Suppose, for

example, that player  $i$ 's beliefs about the other players' states depend arbitrarily on player  $i$ 's private history. Then, generically, player  $i$ 's optimal response would depend on his entire private history. Thus, if we want to restrict attention to equilibria in which players use Markov-private strategies, we need to focus on games in which player  $i$ 's beliefs about other players' current states are restricted to depend on player  $i$ 's private history only through his current state.

This consideration motivates the following definition. The definition uses the notation that given an initial p.d.f.  $\mu_0$  over  $S$  and strategies  $\sigma$ ,  $\Pr(z^t, \theta^t, s_0; \mu_0, \sigma)$  represents the implied joint p.d.f. over realizations of  $(z^t, \theta^t, s_0)$ .

DEFINITION 1. *Let  $\Delta$  be some set of p.d.f.'s with support  $S$ . Beliefs are Markov over  $\Delta$  in a game if for any  $\mu_0$  in  $\Delta$ , any collection of Markov-private strategies  $\sigma$ , and any strategy  $\sigma'_i$  :*

$$(1) \quad \Pr(s_{-i,t} | s_{i0}, z^t, \theta_i^t; \mu_0, \sigma_{-i}, \sigma'_i) = \Pr(s_{-i,t} | s_{it}, z^t; \mu_0, \sigma), \text{ for all } (s_{i0}, z^t, \theta_i^t).$$

$$(2) \quad \text{The p.d.f. of } s_t \text{ conditional on } z^t, \Pr(s_t | z^t; \mu_0, \sigma), \text{ lies in } \Delta \text{ for all } z^t.$$

This definition imposes two requirements on the evolution of beliefs. The first requirement is that player  $i$ 's beliefs about the other players' states (which are given by the p.d.f. of  $s_{-i,t}$ , conditional on player  $i$ 's information) depend only on his current private state and on the public history and that these beliefs are uninfluenced by the particular strategy that player  $i$  has followed. The second requirement is that the probability of the players' states, conditional on the public history, always remains in  $\Delta$ . It is important to note that even if beliefs are Markov, players need not be symmetrically informed after all histories, and players' preferences over action profiles need not be common knowledge (as is true in Fudenberg,

Holmstrom and Milgrom (1990)).

Because of the stationary structure of the game, Definition 1 can be verified by simply checking whether it holds over a single period. In particular, let  $\Delta$  be a set of p.d.f.'s with support  $S$ , and suppose that  $\mu$  lies in  $\Delta$ . Suppose that  $\alpha = (\alpha_i)_{i=1}^N$ , where  $\alpha_i : S_i \rightarrow Q_i$  is an arbitrary strategy. For any  $(\alpha, \mu)$ , we can define the p.d.f.  $\phi(\cdot; \alpha, \mu)$  as follows:

$$\phi(z, \theta, s; \alpha, \mu) = \pi(z, \theta | \alpha(s), s) \mu(s)$$

over  $Z \times \Theta \times S$ . Then, it is easy to see that beliefs are Markov over  $\Delta$  if and only if for any strategy  $\alpha$ , any  $\mu$  in  $\Delta$ , and any  $\alpha'_i$ ,  $\phi$  satisfies the following:

$$(3) \quad \text{For all } z, \mu'(z) \in \Delta$$

$$\text{where } \mu'(z)(s') \equiv \phi(g(\theta, z, q(s), s) = s' | z; \alpha, \mu)$$

$$(4) \quad \phi(g_j(\theta_j, z, q_j(s_j), s_j) = s'_j \text{ for } j \neq i | z, \theta_i, s_i; \alpha_{-i}, \alpha'_i, \mu)$$

$$= \phi(g_j(\theta_j, z, q_j(s_j), s_j) | z, s'_i; \alpha, \mu) \text{ for all } (\theta_i, s_i).$$

(The notation  $\phi(g(\theta, z, q(s), s) = s' | z; \alpha, \mu)$  stands for “the probability of the updated state vector  $g(\theta, z, q(s), s)$  equalling  $s'$ , conditional on the public signal's equalling  $z$ .”) On the one hand, these are simply the one-period analogs of the conditions in Definition 1, so they must be satisfied in any game in which beliefs are Markov over  $\Delta$ . On the other hand, the conditional probabilities in Definition 1 can be constructed by doing a series of one-period Bayesian updates; hence, (3) and (4) imply (1) and (2).

A key feature of games in which beliefs are Markov over  $\Delta$  is that it is easy to keep

track of players' beliefs over each other's states. Let  $A = \{(\alpha_1, \dots, \alpha_N) | \alpha_i : S_i \rightarrow Q_i\}$ . Then, associated with any game in which beliefs are Markov over  $\Delta$ , we can define a *belief updating function*  $T : \Delta \times A \times Z \rightarrow \Delta$  by:

$$\mu' = T(\mu, \alpha, z)$$

$$\text{where } \mu'(s') = \phi(s' = g(\theta, z, \alpha(s), s) | z; \alpha, \mu)$$

where  $\phi$  is defined as above. The function  $T$  tells us how beliefs over  $s'$  evolve according to realizations of the public signal  $z_t$ .

At this point, it is helpful to return to the duopoly example and think about when beliefs are Markov in that context. First, suppose that  $\pi_\theta(\theta_1, \theta_2) = \pi_{\theta_1}(\theta_1)\pi_{\theta_2}(\theta_2)$ , so that cost shocks are independent. Define  $\Delta = \{\mu_0(s_1, s_2) | \mu_0(s_1, s_2) = \mu_{01}(s_1)\mu_{02}(s_2)\}$ , so that initial states are guaranteed to be independent. Then, beliefs are Markov over  $\Delta$ . The laws of motion of the two firms' states are entirely independent of each other; thus, the price is the only signal that firm  $i$  receives about firm  $j$ 's state, and that information is public. In contrast, if  $\pi_\theta(\theta_1, \theta_2)$  does not exhibit independence, then, in general, beliefs are not Markov.

Having defined the concept of Markov beliefs, we can now describe our notion of equilibrium.

DEFINITION 2. Consider any game in which beliefs are Markov over  $\Delta$  and  $\mu_0$  lies in  $\Delta$ .

A **Markov-Private Equilibrium (MPE)** is a collection of Markov-private strategies such that player  $i$ 's strategy prescribes a weakly optimal action at every information set, given that at any information set, player  $i$ 's beliefs about the other agents' states are described by

$$\Pr_t(s_{-i,t} | z^t, s_{it}; \sigma, \mu_0).$$



Note that if  $S_i$  is a singleton for all  $i$ , MPE are equivalent to perfect public equilibria as defined by Fudenberg and Tirole (1991).

For repeated games with public signals and unobservable actions, APS (1990) demonstrate that there is a recursive representation for the set of equilibrium payoffs. In our setting, given the possibility of hidden states, there is no such recursive representation for the set of equilibrium payoffs. In making action choices in a given period, a player has to know how his continuation payoff depends on his updated state. With hidden states there is a recursive representation for the set of equilibrium payoff functions.

This discussion makes clear that we have to express payoffs in terms of equilibrium payoff functions. Specifically, let  $\sigma^*$  be an MPE. We define the equilibrium payoff function  $v_i$  to be

$$v_i : S \rightarrow [0, 1]$$

$$v_i(s_i) = \max_{\sigma_i} U_i(\sigma_i; s_i, \mu_0, \sigma_{-i}^*)$$

where  $U_i$  is the period zero payoff derived by agent  $i$ , given an arbitrary strategy  $\sigma_i$ , his current state  $s_i$ , the initial prior over states, and the strategies of the other agents. Hence, the equilibrium payoff function  $v_i$  describes the agent's maximal utility in period 0 for each initial state  $s_i$ , given that all other agents play their equilibrium strategies.

The following point is worth keeping in mind. As APS (1990) emphasize, when  $S_i$  is a singleton for all  $i$ , the set of Markov-private equilibrium payoffs is equivalent to the set of pure strategy sequential equilibrium payoffs (indeed, it is equivalent to the set of pure

strategy Nash equilibrium payoffs). This equivalence does not carry over to our setting with hidden states. Once  $S_i$  has multiple elements, then the set of Markov-private equilibrium payoffs may be smaller than the set of pure strategy sequential equilibrium payoffs.

## 4. Algorithm

Consider a game in which beliefs are Markov over  $\Delta$ , where  $\Delta$  is a closed set of pdf's with support  $S$ . Denote the game's belief updating rule by  $T$ . Let  $F_i$  be the set of functions from  $S_i$  into  $[0, 1]$ . This is the set of possible payoff functions for player  $i$ . Let  $\Gamma = \prod_{n=1}^N F_n \times \Delta$ . Then, an element  $(v, \mu) \in \Gamma$  is a specification of a payoff function for each agent and a pdf  $\mu$  over  $S$ .

Let  $\gamma^*$  be the subset of  $\Gamma$ , with the property that  $(v, \mu)$  lies in  $\gamma^*$  if and only if  $v$  is a vector of payoff functions for a Markov-private equilibrium, given that player states are initially drawn from the pdf  $\mu$ . Our goal is to solve for  $\gamma^*$ , which we assume is non-empty.

In what follows, we define  $B$  to be a function from the set of all subsets of  $\Gamma$  into itself. Conceptually,  $B(\gamma)$  consists of payoff function vector-belief pairs that are consistent with some Markov-private equilibrium, given that  $\gamma$  contains all possible continuation payoff function vector-belief pairs. This is, of course, analogous to the  $B$  operator in APS (1990).

Mathematically, let  $\gamma \subset \Gamma$ . Then,  $(v, \mu) \in B(\gamma)$  if for all  $i$ , there exists  $q_i^*$  that maps  $S_i$  into  $Q_i$ , and for all  $z$  in  $Z$ , there exists  $(v', \mu')(z)$  in  $\gamma$  such that

$$(5) \quad \mu'(z) = T(\mu, q^*, z)$$

and for all  $s_i$  in  $S_i$

$$(6) \quad v_i(s_i) = \sum_{(s_{-i}, z, \theta)} [u_i(\theta_i, z, q_i^*(s_i), s_i) + \delta v'_i(z)(g_i(\theta_i, z, q_i^*(s_i), s_i))] \pi(z, \theta | q^*(s), s) \mu(s_{-i} | s_i)$$

$$(7) \quad q_i^*(s_i) \in \arg \max_{q_i \in Q_i} \sum_{(s_{-i}, z, \theta)} \left[ \begin{array}{c} u_i(\theta_i, z, q_i, s_i) + \\ \delta v_i'(z)(g_i(\theta_i, z, q_i, s_i)) \end{array} \right] \pi(z, \theta | q_{-i}^*(s_{-i}), q_i, s) \mu(s_{-i} | s_i).$$

The first condition (5) describes the evolution of beliefs on and off the equilibrium path. The second (promise-keeping) condition describes the payoff to agent  $i$ , depending on his initial state and assuming that next period's payoff function is  $v'(z)$  and that next period's state is given by  $g_i(\theta_i, z, q_i^*(s_i), s_i)$ . The final (incentive) condition requires the continuation equilibrium action to be weakly optimal, given any initial state.

It is straightforward to prove the following lemma.

LEMMA 1. If  $\gamma$  is closed,  $B(\gamma)$  is closed.

**Proof.** Let  $(v^m, \mu^m)_{m=1}^\infty$  be a convergent sequence of elements of  $B(\gamma)$  with limit  $(v^\infty, \mu^\infty)$ .

We need to prove that  $(v^\infty, \mu^\infty) \in B(\gamma)$ . For each  $m$  and for all  $z$ , there exists  $q^m$  and  $(v^{m'}(z), \mu^{m'}(z)) \in \gamma$  such that  $(v^m, \mu^m, q^m, v^{m'}(z), \mu^{m'}(z))$  satisfies (5)–(7).

The sequence  $(v^m, \mu^m, q^m, v^{m'}(z), \mu^{m'}(z))$  has a convergent subsequence, so let's assume without loss of generality that the sequence itself converges to  $(v^\infty, \mu^\infty, q^\infty, v^{\infty'}(z), \mu^{\infty'}(z))$ . Because  $\gamma$  is closed,  $(v^{\infty'}(z), \mu^{\infty'}(z)) \in \gamma$ .

Now, we need only verify that  $(v^\infty, \mu^\infty, q^\infty, v^{\infty'}(z), \mu^{\infty'}(z))$  satisfies (5)–(7). To see this, note that since  $Q_i$  is discrete, there exists  $M$  such that  $q^m(s) = q^\infty(s)$  for all  $s$  in  $S$  and all  $m \geq M$  and the right-hand sides of (5)–(7) are continuous in  $\mu$ . This completes the proof. ■

The main proposition for dynamic games is the following:

PROPOSITION 1. Let  $\gamma_0 = \Gamma$ , and let  $\gamma_j = B(\gamma_{j-1})$ ,  $1 \leq j < \infty$ . Then, if  $\gamma^*$  denotes the set

of  $(v, \mu)$  such that  $v$  is the payoff function vector for some Markov-private equilibrium when initial states are drawn according to  $\mu$ , the following conditions hold:

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

(ii.)  $\gamma^* \subseteq \gamma_j \subseteq \gamma_{j-1}$ .

(iii.) Define  $\gamma_\infty$  so that  $\gamma_\infty \equiv \bigcap_{j=1}^\infty \gamma_j$ . Then,  $\gamma_\infty = \gamma^*$ .

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

(i.)  $B(\gamma^*) = \gamma^*$ : First, we prove that  $\gamma^* \subseteq B(\gamma^*)$ . Consider an MPE with an equilibrium payoff function vector  $v$  and initial beliefs  $\mu$ . In the first period, along the equilibrium path, agents have strategies  $(q_i^*(s_i))_{i=1}^N$ . Conditional on Nature's draw of  $z$ , agents then play an MPE continuation equilibrium beginning in period two. Each continuation MPE has an associated equilibrium payoff function vector  $v'(z)$ . The equations defining  $B$  require  $q_i^*(s_i)$  to be the equilibrium action choice for agent  $i$ . Moreover, the equations defining  $B$  require  $v_i(s_i)$  to be the maximal utility agent  $i$  can get if he starts with arbitrary state  $s_i$  and given that the other agents play the MPE equilibrium strategies. It follows that  $(v, \mu) \in B(\gamma^*)$ . It is obvious that  $\gamma^* \supseteq B(\gamma^*)$ , which completes the proof.

(ii.)  $\gamma^* \subseteq \gamma_j \subseteq \gamma_{j-1}$ : It is obvious that if  $\gamma' \subseteq \gamma$ , then  $B(\gamma') \subseteq B(\gamma)$ . 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  $\gamma^* \subseteq \gamma_0$ , this implies that  $\gamma^* = B(\gamma^*) \subseteq \gamma_1$ , which implies that  $\gamma^* \subseteq \gamma_2$ , and so forth.

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

Part 1: If  $B(\gamma) = \gamma$ , then  $\gamma \subseteq \gamma^*$ . This is obvious from the definition of  $\gamma^*$ .

Part 2:  $B(\gamma_\infty) \subseteq \gamma_\infty$ . This follows from  $\gamma_j \supseteq \gamma_\infty$  for all  $j$ , which means that  $\gamma_{j+1} \supseteq$

$B(\gamma_\infty)$  for all  $j$ , which proves the result.

Part 3:  $B(\gamma_\infty) \supseteq \gamma_\infty$ . This is somewhat more difficult. Take any  $(v, \mu) \in \gamma_j$  for all  $j$ . This implies that  $(v, \mu) \in B(\gamma_j)$  for all  $j$ . So for all  $j$  and for all  $z$ , there exists  $q^j$  and  $(v^{j'}(z), \mu^{j'}(z)) \in \gamma_j$  such that  $[v, \mu, q^j, v^{j'}(z), \mu^{j'}(z)]$  satisfies (5)–(7). Since we can always choose a convergent subsequence, assume, without loss of generality, that this sequence converges, and let the limit be  $[v, \mu, q^\infty, v^{\infty'}(z), \mu^{\infty'}(z)]$ . Because  $\gamma_\infty$  is the intersection of closed sets,  $\gamma_\infty$  is closed, and so  $(v^{\infty'}(z), \mu^{\infty'}(z))$  is in  $\gamma_\infty$  for all  $z$ . To complete the proof, we need only verify that  $[v, \mu, q^\infty, v^{\infty'}(z), \mu^{\infty'}(z)]$  satisfies (5)–(7). But this follows from the same logic as in the proof of Lemma 1. ■

Hence, we can solve for  $\gamma^*$  by iterating the backwards operator  $B$ . Having solved for  $\gamma^*$ , it is then straightforward to actually construct the set of all MPE for a given game (given that  $\mu_0$  lies in  $\Delta$ ). Let  $\mu_0$  be the initial p.d.f. over states, and pick an arbitrary element of  $\gamma^*$  with the belief component  $\mu_0$ . This element of  $\gamma^*$  implies some vector of action functions  $(q_1^*, \dots, q_N^*)$ . It also implies, for each realization of  $z$ , a p.d.f.  $\mu'(z)$  and a vector of continuation equilibrium payoff functions  $v'(z)$ . Given  $(\mu', v')$ , we can construct the second period action functions, and so on.

The above discussion assumes that the horizon of the game is infinite. However, it should be clear how to extend the analysis to  $\tau$  period games,  $\tau$  finite. In the last period  $\tau$ , for any initial p.d.f. over states  $\mu$ , 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 set  $\gamma_0$  of vectors of equilibrium payoff functions and p.d.f.'s. Now iterate on this initial set using the backwards operator described above. The iterated set  $B^{\tau-t}(\gamma_0)$  equals the set of equilibrium payoff functions in period  $t$ .

## 5. 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 both actions and states are hidden. 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.

Following Spear and Srivastava's (1987) 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 mechanism design problems in which the agent's actions or current private information influences an unobservable state variable (see Fernandes and Phelan (1998) for some work along these lines).

## 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] Cole, H., and N. Kocherlakota. 1998. "Efficient Allocations with Hidden Income and Hidden Storage," manuscript, Federal Reserve Bank of Minneapolis.
- [5] Fernandes, A., and C. Phelan. 1998. "A Recursive Formulation for Repeated Agency with History Dependence," manuscript, Federal Reserve Bank of Minneapolis.
- [6] Fudenberg, D., B. Holmstrom, and P. Milgrom. 1990. "Short-Term Contracts and Long-Term Agency Relationships." *Journal of Economic Theory* 51: 1–31.
- [7] Fudenberg, D., and J. Tirole. 1991. *Game Theory*. M.I.T. Press, Cambridge, Mass.
- [8] Green, E., and R. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52: 87-100.
- [9] Spear, S., and S. Srivastava. 1987. "Repeated Moral Hazard with Discounting," *Review of Economic Studies* 54, 599-617.