Existence of Markov Perfect Equilibria (MPE) in Undiscounted Infinite Horizon Dynamic Games

Alfredo Garcia and Robert L. Smith
Department of Industrial and Operations Engineering
University of Michigan, Ann Arbor MI 48109

February 22, 1998

Abstract

We prove existence of *MPE* in undiscounted infinite horizon dynamic games, by exploiting an structural property (*Uniformly Bounded Reachability*) of the state dynamics. This allows to identify a suitable finite horizon equilibrium relaxation, the ending state *Constrained MPE,* that captures the relevant features of an infinite horizon *MPE* for a long enough horizon. An application to an asynchronous dynamic duopoly is presented.

1 Introduction

The traditional approach to prove existence of equilibria of infinite horizon games relies heavily on the continuity of payoff functionals. The procedure begins by proving existence of finite horizon equilibria and follows by taking limits as the horizon diverges. Compactness of the infinite horizon strategy space or the history space is usually required to ensure the existence of a limit point (see for example, Fudenberg and Levine [1983], [1986] and Borgers [1989]) which by continuity will inherit the desired properties.

In dynamic games with undiscounted or "average" reward payoffs, this approach fails since, the infinite horizon payoff functionals are not continuous. Moreover, since future rewards are as valuable as present rewards, "end of horizon" effects are magnified, thus, there may exist infinite horizon and finite horizon equilibria of a substantially different nature. In other words, there are infinite horizon equilibrium strategies that are not the limit of finite horizon equilibrium strategies.

In this paper, we provide a new proof of existence of *Markov Perfect Equilibria (MPE)* in the context of infinite horizon nonstationary undiscounted dynamic games. The proof relies on a new method to overcome "end of horizon" effects as in Schochetman and Smith [1997]. The idea is to restrict the deviation possibilities for players by forcing an ending target state for every finite horizon. This relaxation leads to the definition of a *Constrained MPE.* A *reachability* assumption, which essentially requires the "cooperative" controllability of the state dynamics, ensures that play in *early* periods (as opposed to play in *late* periods) is more relevant in identifying profitable deviations in the long run.

We apply our results to an asynchronous dynamic duopoly (see Maskin and Tirole [1988]). Interestingly enough, a full sequential characterization of infinite horizon *MPE,* as limits of finite horizon *Constrained MPE*, is possible in this setting. This result in turn, may help in proving the sustainability of first best outcomes as equilibrium play of the infinite horizon game. This issue, as carefully exposed by Dutta [1995], has been recently examined by Wallner [1997]. Of further research is the application of the techniques here introduced to linear quadratic dynamic games (see Engwerda [1996]) where, hopefully, an analytical representation of the *Constrained MPE* is possible.

2 Nonstationary Dynamic Games.

A $T-$horizon, $N-$player nonstationary dynamic game can be represented by means of the collection $G^T = \{A_k^i(\cdot) \subset A_k^i; r_k^i(\cdot,\cdot); f_k(\cdot,\cdot)\}_{k=0}^T$ with state space denoted by $S \subset \mathcal{R}$ and given initial state $s_0 \in S$, where :

- $A_k^i \subset \mathcal{R}$ is the set of feasible actions at time period k for player $i \in \mathcal{I} = \{1, 2, 3, ...N\}$ and for $s \in S$, $A_k^i(s)$ denotes the subset of feasible actions given state s.
- $r_k^i : S \times A_k^1 \times A_k^2 \times \cdots \times A_k^N \to \mathcal{R}$ is the player's i, $k-$th period reward function.
- $f_k : S \times A_k^1 \times A_k^2 \times \cdots \times A_k^N \to S$ is the state transition function for period k, i.e given state $s \in S$ and actions $a^i \in A_k^i(s)$ we shall denote by a, the action $N-$tuple $(a^1, a^2, ..., a^N)$. With this notation, the state at the beginning of period $k+1$ will be given by :

$$s_{k+1} = f_k(s, a)$$

The set of T-long feasible sequences of actions profiles that players may exert is commonly referred to as the *history* space :

$$H(T) = \prod_{k=0}^{T-1} A_k^1 \times A_k^2 \times \cdots \times A_k^N$$

We shall denote by $h_T(s_0) \in H(T)$ a *feasible history* sequence if it is of the form :

$$h_T(s_0) = (a_0; a_1; ...; a_T)$$

where for $k = 0, 1, 2, ..., T-1$:

$$a_k = (a_k^1, a_k^2, ..., a_k^N) \text{ and } a_k^i \in A_k^i(s_k)$$
$$s_{k+1} = f_k(s_k, a_k)$$

Finally, the total sum of rewards per stage for feasible history $h_T(s_0) \in H(T)$ is given by :

$$P_T^i(h_T(s_0)) = \sum_{k=0}^{T-1} r_k^i(s_k, a_k^1, a_k^2, ..., a_k^N)$$

Finally, if one is only concerned with an intermediate stream of rewards, we shall denote by $P_N^i(h_T(s_0))$ the sum of the rewards induced by history $h_T(s_0)$ up to period N with $N < T$, i.e :

$$P_N^i(h_T(s_0)) = \sum_{k=0}^{N-1} r_k^i(s_k, a_k^1, a_k^2, ..., a_k^N)$$

2.1 Strategies and Markov Perfect Equilibria (MPE).

We now introduce the concept of Nash Equilibria in strategies that employ available information for dynamic games with fixed finite horizon T.

A *closed-loop strategy* for player $i \in \mathcal{I}$, say π_i^T, is a $T-$tuple of maps $\pi_k^i : S \to A_k^i$, so that π_i^T is of the form :

$$\pi_i^T = (\pi_0^i, \pi_1^i, ..., \pi_{T-1}^i)$$

We denote $\Pi^i(T)$ the set of all such strategies for player $i \in \mathcal{I}$. We refer to the N-tuple $\pi^T \in \Pi^1(T) \times \Pi^2(T) \times \cdots \times \Pi^N(T)$ as a *closed loop strategy combination* and denote $\Pi(T)$ the set of all such strategy combinations.

We shall denote by $h_T^{\pi^T}(s_0) \in H(T)$ the *feasible history* induced by strategy combination π^T. Similarly, we shall denote by $h_T^{\pi^T}(s_k) \in \prod_{k}^{T-1} A_k^1 \times A_k^2 \times \cdots \times A_k^N$ the *feasible history* of play induced by strategy combination π^T after intermediate state s_k at time period k, where $0 < k < T-1$. As above, the payoff obtained for this case will be denoted by :

$$P_T^i(h_T^{\pi^T}(s_k)) = \sum_{j=k}^{T-1} r_j^i(s_j, a_j^1, a_j^2, ..., a_j^N)$$

The extension of a dynamic game when there is an infinite number of stages to play follows straightforwardly by setting the *history* space to be the infinite cartesian product :

$$H = \prod_{k=0}^{\infty} A_k^1 \times A_k^2 \times \cdots \times A_k^N$$

Similarly, we shall denote by Π, the set of all infinite horizon feasible strategy combinations.

Finally, the total aggregated reward received by player i under infinite horizon strategy combination π is defined as follows :

$$P^i(h^{\pi}(s_0)) = \lim \inf_{T\to\infty} \frac{P_T^i(h_T^{\pi^T}(s_0))}{T}$$

where π^T stands for the $T-$horizon truncation of infinite horizon strategy combination π.

2.2 Markov Perfect Equilibrium.

We are now ready to introduce the solution concept we shall be dealing with :

Definition 1 (Markov Perfect Equilibrium) : *We say that π^T is a Markov Perfect Equilibrium (MPE) in closed-loop strategies iff for every player $i \in \mathcal{I}$ who would like to deviate from π^T by playing $\gamma_i^T \in \Pi^i(T)$ from every state $s_k \in S$ with $0 \le k < T-1$, would find no incentive in doing so, i.e :*

$$P_T^i(h_T^{(\gamma_i^T, \pi_{-i}^T)}(s_k)) \le P_T^i(h_T^{\pi^T}(s_k))$$

where $(\gamma_i^T, \pi_{-i}^T) \in \Pi(T)$ stands for the strategy combination in which all players $j \in \mathcal{I}$ and $j \neq i$ follow π_j^T and player $i \in \mathcal{I}$ follows γ_i^T.

This definition carries over straighforwardly to the infinite horizon setting with the above introduced framework.

We denote $\Pi^*(T)$ and Π^* the set of all "*Markov Perfect Equilibrium*" strategies for the $T-$horizon and infinite horizon games respectively.

2.3 Topologies on the set Π .

Since our interest is to study convergence of finite horizon equilibrium strategies to infinite horizon equilibrium strategies, it is very important to carefully define relevant topologies on Π, and consequently the different notions of convergence they induce. For a complete study the interested reader is referred to Harris[1985b].

We will adopt the convention that any finite horizon strategy combination is trivially extended through any feasible choice of continuation sequence of strategies, so that its extension is an element of Π.
We first concentrate on a topology for H. Given $h = (a_0, a_1, a_2, ...,)$ and $h' = (a'_0, a'_1, a'_2, ...,)$ we define the metric $D : H \times H \to \mathcal{R}^+$ by :

$$D(h, h') = \sup_t [\frac{\min\{d_t(a_t, a'_t), 1\}}{t}]$$

where d_t is any metric on $A_k^1 \times A_k^2 \times \cdots \times A_k^N$.

The metric $D(\cdot,\cdot)$ induces the product topology on H (see Munkres[1975], p.123). With this metric in hand one can define varios topologies on the set Π :

Definition 2 : *$\mathcal{W}$ is the topology with basis consisting of the sets :*

$$\{\pi \in \Pi \mid D(h^{\pi}, h) < \varepsilon\}$$

where $h, h^{\pi} \in H$. The basis is then obtained as we vary ε, h varies over H.

In words, the notion of convergence related to the topology $\mathcal{W}$ is simply the fact that $\pi^T \to \gamma$ with respect to $\mathcal{W}$ if and only if for any given intermediate state s_k, the sequence of histories induced by π^T, namely $\{h_T^{\pi^T}(s_k)\}_T$ converges to $h^{\gamma}(s_k)$ in the product topology.

Definition 3 : *$\mathcal{L}$ is the topology with basis consisting of the sets :*

$$\{\pi \in \Pi \mid \pi^T = \gamma^T\}$$

with $\gamma \in \Pi$, obtained as γ varies over Π and T varies over all periods and where π^T, γ^T are the $T-$truncations of infinite horizon strategy combinations π and γ respectively.

In words, $\pi^T \to \gamma$ with respect to $\mathcal{L}$ if and only if for all intermediate states, simultaneously, the sequence of histories induced by π^T converge in the discrete topology (they fully agree) to the histories induced by γ. $\mathcal{L}$ is essentially a uniform version of $\mathcal{W}$, Clearly, for practical purposes it may be easier to prove convergence with respect to $\mathcal{W}$, since convergence need only be verified for representative subgames. On the contrary, $\mathcal{L}$ imposes more restrictive conditions on an approximating sequence, so it is generally more helpful in proving uniqueness.

When action sets are discrete, as it will be assumed throughout this paper, these topologies coincide (see Harris[1985b]).

3 Constrained Markov Perfect Equilibria.

Let us now briefly discuss the motivations for the solution concept relaxation that we will introduce shortly. The main difficulty for a sequential characterization of infinite horizon equilibria as limits of finite horizon equilibria is due to "end of horizon" effects (see Fudenberg and Levine[1983]). In words, for a fixed finite horizon, the final state attained for finite horizon equilibrium will generally be different from the state attained by the truncation of the infinite horizon equilibrium. Myopic behavior close to the fixed finite horizon is the explanation for this. We will try to overcome this effect by forcing equilibrium strategies to attain a certain "target" state.

Definition 4 (Constrained strategies) : *Let $s \in S$ be some feasible state, we denote by $\Pi(T, s)$ the set of closed-loop strategy combinations such that for every $s_k \in S$, $0 \leq k < T-1$, and the state s is reachable from s_k; the play to follow after state s_k must reach s_T. In other words, the history prescribed, i.e $h_T^{\pi^T}(s_k)$ reaches state s, at time period T, whenever state s is reachable from $s_k \in S$, $0 \leq k < T-1$.*

Note that the play prescribed by any $\pi^T \in \Pi(T, s)$ from some state $s_k \in S$ from which s is not reachable, is completely irrelevant to the definition.

Definition 5 (Constrained MPE) :*A strategy combination $\pi^T \in \Pi(T, s)$ is called a "Constrained MPE to state s" iff for every deviation $\gamma_i^T \in \Pi(T)$ such that $(\gamma_i^T, \pi_{-i}^T) \in \Pi(T, s)$ from every $s_k \in S$ with $0 \leq k < T-1$, such that state s is reachable from s_k we have :*

$$P_T^i(h_T^{(\gamma_i^T, \pi_{-i}^T)}(s_k)) \leq P_T^i(h_T^{\pi^T}(s_k))$$

We denote $\Pi^*(T, s)$ the set of all "*Constrained MPE to state s*"

4 Existence of Markov Perfect Equilibria.

4.1 Standing Assumptions.

We now present the standing assumptions for our analysis:

- **Assumption 1**: (**Non-Emptyness**) For every T, there exists some $s \in S$ such that $\Pi^*(T, s) \neq \emptyset$.

- **Assumption 2:**
  (a) **Discreteness** : Each set A_k^i is discrete and finite, hence, the *history* space H is compact in the product topology and Π is compact in $\mathcal{L}$.
  (b) **Reward Boundedness** : for every $i \in \mathcal{I}$ and for every time period k :

$$-\infty < -M \leq r_k^i(\cdot, \cdot) \leq M < \infty$$

- **Assumption 3 (Uniformly Bounded Reachability)** : For any infinite feasible sequence of states, say $\mathbf{s} = (s_0, s_1, s_2, ...)$ and any intermediate state off the sequence, say s'_k, at time period k; there exist some finite time period $T > k$ and sequence of action profiles $\{a_s\}_{k<s\leq T}$, so that state s_T in the sequence, is reached in $T = k + \triangle(s'_k, \mathbf{s}) > t$ periods. Moreover,

$$\sup_k \sup_{\mathbf{s}} \triangle(s'_k, \mathbf{s}) \leq L < \infty$$

4.2 The Existence Result.

The intuition for the next result lies in the fact that under the *Uniformly Bounded Reachability* assumption, a sequence of finite horizon *Constrained MPE* will encompass all posible deviations (and not just the "constrained" deviations) as the horizon diverges to infinity. Compactness of the strategy space ensures that every sequence of *Constrained MPE* has a converging subsequence and the limit strategy will be an *MPE* for the infinite horizon game, by the argument above.

Lemma 1 : *Let $\mathbf{s} = (s_0, s_1, s_2, ...)$ be an infinite feasible sequence of states and $\{\pi^T : \pi^T \in \Pi^*(T, s_T)\}_T$ a sequence of finite horizons Constrained MPE such that :*

$$\pi = \lim_{T\to\infty} \pi^T \text{ with respect to } \mathcal{L}$$

then under assumptions 2 and 3 $\pi \in \Pi^*$.

Proof : Let us first show that;

$$P^i(h^{(\gamma_i,\pi_{-i})}(s_0)) \leq P^i(h^\pi(s_0))$$

for any player $i \in \mathcal{I}$, who would deviate by playing $\gamma_i \in \Pi$ from initial state s_0. We recall that $h_T^{(\gamma_i,\pi_{-i})}(s_0)$ and $h_T^\pi(s_0)$ stand for the $T-$truncations of the histories induced by strategies (γ_i, π_{-i}) and π, respectively.

By convergence in $\mathcal{L}$ there exists T^N such that for any π^T with $T > T^N$ the play prescribed by (γ_i^T, π_{-i}^T) and π^T coincide exactly with $h_T^{(\gamma_i,\pi_{-i})}(s_0)$ and $h_T^\pi(s_0)$ respectively, in the first $N < T$ periods. Moreover, the deviation for player i :

$$\bar{\gamma}_i^T = (\gamma_0^i, \gamma_1^i, .., \gamma_N^i, a_{N+1}^i, ..., a_{T-1}^i)$$

(in which we append from the $N-$period, the actions $(a_{N+1}^i, ..., a_{T-1}^i)$ as prescribed by π^T) is such that $(\bar{\gamma}_i^T, \pi_{-i}^T)$ "*reaches*" state s_T. Formally :

$$(\bar{\gamma}_i^T, \pi_{-i}^T) \in \Pi(T, s_T)$$

Hence, by hypothesis on π^T we have :

$$P_T^i(h_T^{(\bar{\gamma}_i^T,\pi_{-i}^T)}(s_0)) \leq P_T^i(h_T^{\pi^T}(s_0))$$

By cost boundedness and the choice of T^N, we have that total payoff accrued, up to period N satisfies :

$$\frac{P_N^i(h_T^{(\bar{\gamma}_i^T,\pi_{-i}^T)}(s_0))}{N} \leq \frac{P_N^i(h_T^{\pi^T}(s_0))}{N} + \frac{2M \cdot L}{N}$$

hence :

$$\frac{P_N^i(h_T^{(\gamma_i,\pi_{-i})}(s_0))}{N} \leq \frac{P_N^i(h_T^{\pi}(s_0))}{N} + \frac{2M \cdot L}{N}$$

Then, iterating on this construction, we have that :

$$P^i(h^{(\gamma_i,\pi_{-i})}(s_0)) = \lim\inf_{N\to\infty} \frac{P_N^i(h_T^{(\gamma_i,\pi_{-i})}(s_0))}{N} \leq \lim\inf_{N\to\infty} \frac{P_N^i(h_T^{\pi}(s_0))}{N} = P^i(h^{\pi}(s_0))$$

Thus, from the initial state, the proposed deviation is not profitable.

For a deviation from any other state $s_k \in S$ with $0 < k$ we use the same argument. ∎

By a standard compactness argument existence of *MPE* follows :

Theorem 1 : *Under Assumptions 1,2 and 3 there exists an MPE for the infinite horizon undiscounted game.*

Proof : By assumption 1 (non-emptyness) on can construct a sequence $\{\pi^T : \pi^T \in \Pi^*(T, s_T)\}_T$ of Constrained MPE for an infinite feasible sequence of states $\mathbf{s} = (s_0, s_1, s_2, ...)$. By compactness of the strategy space, there exists a converging subsequence, say $\{\pi^{T_k} : \pi^{T_k} \in \Pi^*(T_k, s_{T_k})\}_k$ and :

$$\pi = \lim_{k\to\infty} \pi^{T_k} \text{ with respect to } \mathcal{L}$$

Finally by Lemma 1, $\pi \in \Pi^*$. ∎

5 Application : Sequential Duopoly.

In this section we briefly illustrate all the definitions above introduced for the case of a duopoly competition in prices, as in Maskin and Tirole [1988].

Players move sequentially, so that in odd numbered periods k, firm 1 chooses its price which remains unchanged until period $k+2$.That is, $p_{k+1}^1 = p_k^1$ if k is odd. Similarly, firm 2 chooses prices only in even numbered periods, $p_{k+1}^2 = p_k^2$ if k is even. Hence, at time period k , firm's i instantaneous reward $r_k^i(.)$ is a function of the "state", i.e the price that firm's j set on period $k-1$, say p_k^j , and the "action", i.e the price that firm's i will establish p_k^i. Feasible price set, say $p \in P$ are discrete and bounded, goods are perfect substitutes, that is, firms share the market equally whenever they charge the same price. Firms have the same unit cost c. Let $D_k(.)$ denote the market demand function at time period k. The total reward at time period k is given by :

$$r_k(p) = (p-c)D_k(p) \; p \in P$$

Then :

$$r_k^i(p^1,p^2) = \begin{cases} r_k(p^i) & \text{if } p^i < p^j \\ \frac{r_k(p_k^i)}{2} & \text{if } p^i = p^j \\ 0 & \text{if } p^i > p^j \end{cases}$$

Strategies are "Markovian" in that they depend on the current "state", i.e last period rival's action. Hence, the set of all histories is the same as the set of all feasible sequences of states.
Consider the infinite history $h = \{(p_k^1, p_k^2)\}_k$ then firm i undiscounted payoff is :

$$P^i(h) = \lim \inf_{T\to\infty} \frac{1}{T} \sum_{k=0}^{T-1} r_k^i(p_k^1, p_k^2)$$

Now, let us assume that p_T^1 is a feasible price decision for firm 1 at odd time period T. Then, $\Pi(T, p_T^1)$ stands for the set of all markovian strategy combinations for horizon T in which player 1 is constrained to play p_T^1 at time period T .Similarly, $\Pi^*(T, p_T^1)$ is the set of T-long horizon "constrained" MPE strategy combinations to "state" p_T^1. Notice that under the assumptions by a backwards induction argument one can easily see that $\Pi^*(T, p_T^1) \neq \emptyset$ and that the *Uniformly Bounded Reachability* assumption holds.

Moreover, a "converse" like result to Theorem 1 holds in this particular setting, i.e every infinite horizon Markov Perfect Equilibrium is the limit of a sequence of finite horizon *Constrained Markov Perfect Equilibrium* strategies.

Theorem 2 : *For every* $\pi \in \Pi^*$ *there exists an infinite feasible sequence of "states"* $(p_0, p_1, p_2, ...)$ *and a sequence of finite horizon Constrained Markov Perfect Equilibrium* $\{\hat{\pi}^T : \hat{\pi}^T \in \Pi^*(T, p_T)\}_T$ *such that :*

$$\pi = \lim_{T\to\infty} \hat{\pi}^T \text{ with respect to } \mathcal{L}$$

Proof : Let π^T denote the T-truncation of π and p_T the "state" reached from initial state. Clearly, $\pi^T \notin \Pi(T, p_T)$, since, off- equilibrium play need not necessarily lead to "state" p_T. However, one can construct a strategy combination $\hat{\pi}$ that "resembles" π such that $\hat{\pi}^T \in \Pi(T, p_T)$ as follows :

For all intermediate "states" $p_k \in P$ with $0 \leq k < T$:

- If $h_T^{\pi^T}(p_k)$ reaches state p_T at time period T, we set $h_T^{\hat{\pi}^T}(p_k)$ to be exactly $h_T^{\pi^T}(p_k)$.
- Else, $h_T^{\hat{\pi}^T}(p_k) = h_T^{\pi^T}(p_k)$ except for the last period action which is set to be exactly p_T.

We embed the collection $\{\hat{\pi}^T : \hat{\pi}^T \in \Pi^*(T, p_T)\}_T$ in the space Π, by appending the infinite tail of play prescribed by π after "state" p_T. It is clear that :

$$\pi = \lim_{T\to\infty} \hat{\pi}^T \text{ with respect to } \mathcal{L}$$

Reasoning by contradiction, let us now assume that :

$$\hat{\pi}^T \notin \Pi^*(T, p_T) \text{ for all } T$$

By definition, this implies the existence of profitable deviations, i.e for each T there is γ_i^T with $i \in \mathcal{I}$ such that $(\gamma_i^T, \hat{\pi}_{-i}^T) \in \Pi(T, p_T)$ and :

$$P_T^i(h_T^{(\gamma_i^T, \hat{\pi}_{-i}^T)}(p_0)) > P_T^i(h_T^{\hat{\pi}^T}(p_0))$$

But by construction :

$$P_T^i(h_T^{(\gamma_i^T, \hat{\pi}_{-i}^T)}(p_0)) = P_T^i(h^{(\gamma_i, \pi_{-i})}(p_0))$$

and :

$$P_T^i(h_T^{\hat{\pi}^T}(p_0)) = P_T^i(h^{\pi}(p_0))$$

which in turn will imply that :

$$P^i(h^{(\gamma_i, \pi_{-i})}(s_0)) > P^i(h^{\pi}(s_0))$$

In other words, $\pi \notin \Pi^*$, hence a contradiction. ∎

6 Conclusion.

Existence of equilibria in undiscounted games is a difficult issue. In this paper, we have presented a new approach to this problem that relies heavily on structural properties of the game (the so called *Uniformly Bounded Reachability* assumption). This assumption allows to define a solution concept (the *Constrained MPE)* for the finite horizon game, that captures the relevant features of an infinite horizon *MPE* for a long enough horizon.

An application to an asyncronous dynamic duopoly is presented. In this setting, not only the limit point of finite horizon *Constrained MPE* is an infinite horizon *MPE,* but every infinite horizon *MPE* is the limit of a sequence of finite horizon *Constrained MPE.* This sequential characterization helps identify the existence of efficient equilibria for the infinite horizon game.

References

[1] Börgers,Tilman "Perfect Equilibrium histories of Finite and Infinite Horizon Games" *Journal of Economic Theory* 47 (1989) pp 218-227

[2] Dutta P. "Collusion, Discounting and Dynamic Games" *Journal of Economic Theory* 66 (1995) pp 289-296

[3] Fudenberg D. and Levine D. "Subgame-Perfect equilibria of Finite and Infinite-Horizon Games" *Journal of Economic Theory* 31 (1983) 251-267

[4] Engwerda J. "Computational aspects of the Infinite planning horizon Open-loop Nash Equilibrium in LQ games" *Tilburg University Working Paper* (1996)

[5] Fudenberg D. and Levine D. "Limit Games and Limit Equilibria" *Journal of Economic Theory* 38 (1986) 261-279

[6] Fudenberg D. and Tirole J. "Game Theory" *MIT Press* 1994

[7] Harris,Chris "A characterization of perfect equilibria in Infinite Horizon Games" *Journal of Economic Theory* 37 (1985a) pp. 99-127

[8] Harris,Chris "Existence and Characterization of Perfect equilibrium in Games of Perfect Information" *Econometrica* 53 No 3 (1985b) pp. 613-628

[9] Maskin E. and Tirole J. "A Theory of Dynamic Oligopoly II : Price Competition, Kinked Demand curves and Edgeworth cycles" *Econometrica* 56 No.3 (1988) 571-599

[10] Munkres J. "Topology: A First Course" *Prentice Hall* 1975

[11] Schochetman I. and Smith R.L "Infinite Horizon Optimization with Average Optimal Criterion" (To appear on *Mathematics of Operations Research*)

[12] Wallner K. "Sequential Moves and Tacit Collusion : Reaction Function Cycles in a Finite Pricing Duopoly" *Stockholm School of Economics, Working Paper No. 206,* November 1997