Selection of Equilibrium in a DynamicOligopoly

with Cost-Reducing Investments

Milan Horniaček[1]

CERGE-EI, Charles University and the Academy of Sciences of the Czech Republic, Prague

In this paper we analyze an infinite horizon dynamic oligopoly, producing a homogeneous good, with costly changes of output between the periods, and investments affecting marginal costs. The requirement of continuity of strategies and the weakest possible criterion of renegotiation-proofness, called renegotiation-quasi-proofness, are used to select a subset of Markov perfect equilibria with a common limit of continuation equilibrium paths. In each renegotiation-quasi-proof continuous strategy Markov perfect equilibrium, each firm's price and marginal costs converge to common levels that would maximize net profit of each firm if they were infinitely repeated.
Keywords: continuous Markov strategies, costs reducing investments, dynamic oligopoly, equilibrium selection, renegotiation-proofness.
JEL Classification Numbers: C73, D43, L13.

Článek analyzuje model dynamického oligopolu s rovnorodým produktem, nákladmi změn výstupu mezi jednotlivými obdobími a investicemi, které ovlivňují mezní náklady. Autor používá požadavek spojitosti strategií a nejslabší možné kriterium odolnosti vůči renegociaci, které nazývá "kváziodolností vůči renegociaci," na výběr podmnožiny Markovových dokonalých rovnováh, které mají společnou limitu rovnovážných trajektorií v podřízených hrách. Ve všech Markovově dokonalých rovnováhách se spojitými strategiemi, kváziodolných vůči renegociaci, mezní náklady a cena produktu každé firmy konvergují k společným hodnotám, které by, kdyby byly nekonečně opakovány, maximalizovali čistý zisk každé firmy.

I. Introduction

Many infinite horizon, discrete time oligopoly models involve physical links between periods, i.e., they are oligopolistic difference games. These links can stem, for example, from investments, from advertising, or from the costs of changing outputs or prices. In difference games, a current state, which is payoff relevant, should be taken into account by rational players when deciding on a current period action. If strategies depend only on a current state

[1] A CERGE ESC Grant is acknowledged as a partial source of financial support.

(which is payoff relevant), they are called Markov.[2] By applying the requirement of subgame perfection to Markov strategies which form a Nash equilibrium, we are led to the concept of a Markov perfect equilibrium.[3] In many oligopolistic infinite horizon difference games, the Markov perfect equilibrium is non-collusive.[4] Maskin and Tirole's (1988) model of a dynamic Bertrand duopoly is a notable exception. In their paper, the existence of payoff relevant states stems from the sequential naming of a price by duopolists (i.e., one of them names a price in odd periods, the other in even periods). The price set in a period t also remains in effect in a period t + 1. Unfortunately, it is hard to imagine how to modify that model to include more than two firms.

A firm's costs can change over time. Thus, the parameter(s) of the current period cost function is (are) a natural component of a payoff relevant state vector. The current period expenditures affecting the next period cost function are then included among the firm's choice of variables. We have adopted this approach for this paper. The model analyzed here is an infinite horizon dynamic oligopoly, composed of firms which have the following characteristics: 1) they produce to order a single homogeneous non-durable good, 2) changes of their output between periods are costly, 3) their investments affect the next period's marginal costs (constant over the range of possible outputs), 4) they discount future profits.[5]

The costs of changes in output can stem, for example, from additional charges when the quantities of inputs ordered are changed shortly before delivery. The costs of a reduction

---

[2] A justification for focusing on Markov strategies is given below.

[3] For a thorough characterization of a Markov perfect equilibrium see Maskin and Tirole (1994).

[4] See Maskin and Tirole (1987) for a typical example. This differs strikingly from results of the research on subgame perfect equilibria in supergames with discounting of payoffs, known as folk theorems, which imply the existence of a continuum of collusive subgame perfect equilibria. See Fudenberg and Maskin (1986) for study of folk theorems.

[5] This paper is a sequel to Horniaček (1996). The model analyzed in that paper is extended here by allowing for changes in costs.

in output can also arise from labour-related legal restrictions. Thus, these costs are, in their economic substance, different from investments into enlarging or maintaining capacity. In a period t, firms, taking into account their output in a period t - 1, name prices, on the basis of which customers place their orders and choose levels of investments which affect the next period's marginal costs. Orders are rationed if more than one firm charges the lowest price. Paying attention to the costs of changes in output, the firms then decide which portion of the orders they will confirm. The unconfirmed portion of the orders then becomes available to the other firms. Outputs equal to the confirmed orders are produced, then sold, with each firm incurring the costs of a change in output compared with the previous period. A state in a period t is identified with a vector of outputs in a period t - 1 and a vector of marginal costs in a period t. (Outputs in a period t - 1 are payoff relevant, because changes in output are costly.)

Focusing on Markov perfect equilibria imposes three limitations on equilibrium strategies. The first limitation is that repetitions of a certain state vector cannot be counted. Therefore, after a profitable unilateral deviation (i.e., a unilateral deviation that would increase the deviator's continuation average discounted net profit if the other firms ignored it), the play generally passes through several different state vectors lying on a punishment path. The second limitation concerns vectors of prices charged and investment expenditures, here called action vectors. An action vector prescribed by an equilibrium strategy profile for the first period must be the same as an action vector prescribed when the initial state reappears after a deviation. Thus, at the beginning of the game, a collusive action vector can be only gradually approached, unless the initial state cannot result from a profitable unilateral deviation from any continuation equilibrium and does not lie on a punishment path. The third limitation concerns the punishment path along which (or a subset of which) the play proceeds after a unilateral deviation. The punishment path must be the same for all firms. Otherwise, it would not always be possible to determine from which of the punishment paths a deviation took place on the basis of only a previous period's output vector and a vector of the current period's marginal costs.

The analyzed dynamic oligopoly has a continuum of Markov perfect equilibria. We approach the problem of equilibrium selection using two additional restrictions imposed on Markov strategies. First, we require them to be continuous (functions of a current state). Second, we impose the weakest possible requirement of renegotiation-proofness, which we call renegotiation-quasi-proofness.

The requirement of the continuity of strategies, introduced into the analysis of infinite horizon games with the discounting of payoffs (but without restriction to Markov strategies) by J. W. Friedman and L. Samuelson (1994a, 1994b), is based on the view that punishments should "fit the crime." The two main arguments made by Friedman and Samuelson (1994b) in favour of continuous strategies are: 1) after a very small deviation, they are more appealing to real human players than a Draconian punishment; and, 2) following a deviation, the convergence of continuous strategies to the original action profile (or, in a Markov setting, to the limit of the original sequence of action profiles) reflects an intuitively appealing rebuilding of trust.

Continuous Markov strategies have the plausible property that large changes in payoff relevant variables have large effects on current actions, minor changes in payoff relevant variables have minor effects on current actions, and changes in variables that are not payoff relevant have no effect on current actions. This is an improvement over Markov strategies which do not require continuity. These strategies make a distinction only between effect and no effect, according to whether a variable that has changed is payoff relevant or not. Thus, minor changes in payoff relevant variables can have large effects on current actions.[6] Continuous Markov strategies are also an improvement in comparison with continuous strategies without the Markov property, which allow changes in variables that are not payoff relevant to affect current actions.

The requirement that strategies be Markov is an application of Harsanyi and Selten's (1992) principle of invariance of (selected) equilibrium strategies with respect to isomorphism of games.[7] The latter principle requires that strategically equivalent games have identical solutions (i.e. selected equilibrium or subset of equilibria). Applying this principle to a subgame perfect equilibrium of the analyzed difference oligopolistic game, we require that selected equilibrium strategy profiles prescribe the same play in all subgames that are strategically equivalent, i.e. in all subgames with the same initial state. The requirement of continuity of strategies is a strengthening of this principle. When two subgames are, from the strategic point of view, "close", i.e. the initial state of one of them is in a neighbourhood of the initial state of the other, a sequence of vectors of actions prescribed for the former should be in a neighbourhood (on the element-wise basis) of a sequence of vectors of actions prescribed for the latter.[8]

The requirement of the continuity of Markov strategies imposes two additional limitations on an equilibrium strategy profile. Firstly, the play cannot reach (or even approach) a cycle between different state vectors, because the play would have to approach it by switching between several paths, one for each element of the cycle. After certain

---

[6] See the discussion of minor causes and minor effects in Maskin and Tirole (1994).

[7] This fact is pointed out by Maskin and Tirole (1994).

[8] Let (q, c) and (q′, c′) be initial states in two different subgames and let $\{p(t), x(t)\}_{t=1}^{\infty}$ and $\{p'(t), x'(t)\}_{t=1}^{\infty}$ be sequences of price vectors and vectors of investment expenditures prescribed in them by a subgame perfect equilibrium profile of continuous Markov strategies. Then, when (q, c) converges to (q′, c′), (p(t), x(t)) converges to (p′(t), x′(t)) for each finite positive integer t.

profitable unilateral deviations, it would not be possible to determine, from a state vector alone, from which path the play deviated. Thus, for at least one of these paths, the set of states from which the play switches would have to be open. This would contradict the continuity of strategies. Also, an equilibrium strategy profile cannot generate a non-convergent sequence of state vectors. Therefore, since the punishment path is the same for all firms, all continuation equilibrium paths must converge to the same limit. Secondly, the convergence of the play to this limit can only have the form of approaching it, without reaching it in any finite time. Moreover, no matter to which vector of marginal costs it converges, the play cannot converge to an asymmetric price vector. (Proposition 1 in Section III makes this argument.) Thus, it must converge to a symmetric price vector.

Knowing this, intuition suggests that, since the costs affecting investments function is the same for all firms, the play should converge to the symmetric vector of marginal costs. Given the limit symmetric price vector, it is in the interest of each firm to approach the level of marginal costs at which its profit is maximized.

Nevertheless, there is still a continuum of symmetric price vectors and corresponding symmetric vectors of marginal costs, each of which is the limit of all continuation equilibrium paths of a certain continuous strategy Markov perfect equilibrium. Therefore, we impose on equilibrium strategies the weakest possible criterion of renegotiation-proofness, which we call renegotiation-quasi-proofness. A continuous strategy Markov perfect equilibrium is renegotiation-quasi-proof if no other continuous strategy Markov perfect equilibrium exists which gives all the firms a higher continuation average discounted net profit in every subgame. We show that, if all discount factors are close enough to one, in every renegotiation-quasi-proof continuous strategy Markov perfect equilibrium all continuation equilibrium paths converge to the same symmetric price and marginal cost vectors. These price and marginal cost vectors, if they are infinitely repeated, maximize each firm's average discounted net profit from an infinite repetition of the symmetric price and marginal cost vectors. That is, switching to infinite repetition of any other price and marginal cost vectors would not increase the average discounted net profit of any firm.

When applied to a continuous strategy Markov perfect equilibrium, the criterion of a renegotiation-quasi-proofness leads to the selection of the unique limit of all continuation equilibrium paths, although it does not lead to the unique equilibrium. Nevertheless, a symmetric limit of vectors of marginal costs does not imply that a vector of marginal costs in any finite period must be symmetric.[9]

The paper is organized as follows. In the next section we describe the analyzed dynamic oligopolistic game. Section III contains the statement and proof of the necessary conditions for the existence of a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium. In Section IV we prove sufficient conditions for its existence. Section V concludes.

II. The Analyzed Dynamic Oligopoly

The analyzed industry consists of a finite number of firms, indexed by a subscript $i \in I = \{1, 2, ..., n\}$, $n \geq 2$, producing to order a single homogeneous non-durable good. We analyze this industry in an infinite horizon model with discrete time.

The demand function $D:[0, \infty) \rightarrow [0, D(0)]$, with $D(0)$ finite, is continuous on its domain. There is $\rho > 0$ such that $D(\lambda) > 0$ for all $\lambda \in [0, \rho)$ and $D(\lambda) = 0$ for all $\lambda \geq \rho$. The

[9] Therefore, our model is not inconsistent with (empirically observed) differences between the costs of firms in the same industry.

function D is strictly decreasing, concave and twice differentiable on the interval $[0, \rho]$ (with the right-hand side derivative at 0 and the left-hand side derivative at ρ).

For each firm $i \in I$ and at each period $t \in \{1, 2, ...\}$, the costs of producing output $q_i(t)$ are $c_i(t)q_i$, where $c_i(t) \in (0, \infty)$. We let $c(t) = (c_1(t), ..., c_n(t))$. These costs depend (for $t > 1$) on firm i's investments $x_i(t-1)$ in period $t - 1$. (The marginal costs of each firm in period one are given.) The costs affecting investments function g is the same for all firms in all periods. If the firm's current period marginal costs are $c_i(t)$ and it wants to have marginal costs $c_i(t+1)$ in the next period, it must invest $g(c_i(t), c_i(t+1))$. Thus, $g:(0, \infty)^2 \to (0, \infty)$.[10] (Negative investments, in the sense that firms receive payment if their marginal costs increase enough, are not possible.) We assume that g is continuous, twice differentiable, strictly increasing in its first argument, strictly decreasing in its second argument, strictly convex in the vector of arguments, and that, for each $c_i \in (0, \infty)$ and each $p_i \in [0, \rho)$, the function $n^{-1}(p_i - c_i)D(p_i) - g(c_i, c_i)$ is strictly concave.[11] There is $p_i \in [0, \rho]$ and $c_i \in (0, \infty)$ such that $n^{-1}(p_i - c_i)D(p_i) - g(c_i, c_i) > 0$. For each $c_i \in (0, \infty)$ and each $c_i' \in (0, c_i)$, $g(c_i, c_i') + g(c_i', c_i) > 2g(c_i, c_i)$.

We also assume that

$$(1) \qquad \lim_{c_i(t+1)\to 0} g(c_i(t),\ c_i(t+1)) = \infty, \qquad c_i(t) \geq c_i(t+1),$$

and

$$(2) \qquad \lim_{c_i(t+1)\to\infty} g(c_i(t), c_i(t+1)) = 0, \qquad c_i(t) \leq c_i(t+1).$$

[10] A similar investment function was used by Flaherty (1980).

[11] Given the assumptions already made (namely, concavity of D and strict convexity of g), sufficient condition for this is that $-n^{-1}(\partial D(p_i)/\partial p_i) - 2(\partial^2 g(c_i, c_i)/\partial c_i(t)\partial c_i(t+1)) < (\partial^2 g(c_i, c_i)/\partial c_i(t)^2) + (\partial^2 g(c_i, c_i)/\partial c_i(t+1)^2)$ for each $c_i \in (0, \infty)$ and each $p_i \in [0, \rho)$.

In addition, we assume that function g is continuously invertible with respect to its second argument and we denote this inverse by $g^{-1}[c_i(t), x_i(t)]$.

Besides production costs and investment expenditures affecting the next period marginal costs, there are also the costs of changing an output in comparison with the previous period, which are a function of the absolute value of a difference between the current and the previous period output. For each firm $i \in I$, these costs are expressed by a function $\alpha_i:[0, D(0)]^2 \rightarrow [0, \infty)$. We have $\alpha_i[q_i(t - 1), q_i(t)] = \gamma_i[|q_i(t - 1) - q_i(t)|]$, where the function $\gamma_i:[0, D(0)] \rightarrow [0, \infty)$ is continuous, twice differentiable, and strictly convex[12] on its domain, strictly increasing at each point of its domain, except at 0 where it has a zero derivative, and $\gamma_i(0) = 0$. The costs of changes in an output can stem, for example, from additional charges incurred when ordered quantities of inputs are changed shortly before delivery or, in the case of a reduction of an output, from labour-related legal restrictions.[13] Thus, these costs are, in their economic substance, different from investments into enlarging or maintaining capacity.

Firm $i \in I$ discounts revenue and all costs with a discount factor $\beta_i \in (0, 1)$ and we set $\beta = (\beta_1, ..., \beta_n)$.

In a period $t \in \{1, 2, . . . \}$, each firm, taking into account its output in the previous period, names a price $p_i(t) \in P_i = [0, \rho]$. We let $P = X_{i \in I} P_i$.

For each $i \in I$, the initial output level $q_i(0) \in [0, D(0)]$ is given. This can be explained by the fact that the analyzed oligopoly existed before period one, but we started to observe and analyze it only as of period one.

---

[12] The assumption that γ_i is strictly convex realistically implies that the firm's competitive supply $y_i(p_i, z_i, c_i)$ (see below) is finite for all feasible p_i and z_i.

[13] Our qualitative results would not change if we assumed that only increases in an output are costly. Nor they would change if we assumed asymmetry between costs of increasing and reducing an output.

A vector of prices in a period t, $p(t) = (p_1(t),..., p_n(t))$, together with a vector of outputs in a period t - 1, $q(t-1) = (q_1(t-1),..., q_n(t-1))$, uniquely determines a vector of outputs in a period t, $q(t)$. Let $y_i(p_i, z_i, c_i)$ be firm i's competitive supply at price p_i when it has marginal costs c_i, and its output in the previous period was $z_i \geq 0$. That is, $y_i(p_i, z_i, c_i)$ is the output that firm i would produce if it were not constrained by supply of the other firms, or by market demand. The output $y_i(p_i, z_i, c_i)$ is defined by $y_i(p_i, z_i, c_i) = \text{argmax}\{(p_i - c_i)\lambda - \gamma_i(|\lambda - z_i|)|\lambda \in [0, \infty)\}$. Clearly, if $p_i > c_i$ (i.e., $p_i - c_i > 0 = \gamma_i'(0)$, where the apostrophe denotes the first derivative), then $y_i(p_i, z_i, c_i) > z_i$ and the quantity $y_i(p_i, z_i, c_i)$ satisfies the first order condition $p_i - c_i = \gamma_i'[y_i(p_i, z_i, c_i) - z_i]$. If $p_i = c_i$, then $y_i(p_i, z_i, c_i) = z_i$. If $p_i < c_i$ (i.e. $p_i - c_i < 0 = -\gamma_i'(0)$), then $y_i(p_i, z_i, c_i) < z_i$ and the quantity $y_i(p_i, z_i, c_i)$ satisfies the first order conditions $p_i - c_i \leq -\gamma_i'[z_i - y_i(p_i, z_i, c_i)]$ and $(p_i - c_i + \gamma_i'[z_i - y_i(p_i, z_i, c_i)])y_i(p_i, z_i, c_i) = 0$.

Consider $p \in P$ and $z \in [0, D(0)]^n$. The vector of outputs sold when the current price vector is p, the vector of immediately preceding period outputs is z, and the current vector of marginal costs is c, is denoted by $q(p, z, c)$, and is defined inductively. For price λ and price vector p set $J_0(\lambda, p) = \{i \in I|\ p_i = \lambda\}$. Let $\lambda_0 = \min_{i\in I} p_i$. For firm

$$i = \min\{j \in J_0(\lambda_0, p)|\ y_j(p_j, z_j, c_j) \leq y_k(p_k, z_k, c_k)\ \forall\ k \in J_0(\lambda_0, p)\}$$

the output sold is

$$q_i(p, z, c) = \min\{y_i(p_i, z_i, c_i),\ |J_0(\lambda_0, p)|^{-1}D(\lambda_0)\}^{14}$$

and we set $J_1(\lambda_0, p) = J_0(\lambda_0, p)\backslash\{i\}$ and $D_1(\lambda_0) = D(\lambda_0) - q_i(p, z, c)$. If $J_1(\lambda_0, p) \neq \emptyset$, we find firm

$$i = \min\{j \in J_1(\lambda_0, p)|\ y_j(p_j, z_j, c_j) \leq y_k(p_k, z_k, c_k)\ \forall\ k \in J_1(\lambda_0, p)\}.$$

The output sold by this firm is

$$q_i(p, z, c) = \min\{y_i(p_i, z_i, c_i),\ |J_1(\lambda_0, p)|^{-1}D_1(\lambda_0)\},$$

[14] For a finite set J, the symbol |J| denotes its cardinality.

and we set $J_2(\lambda_0, p) = J_1(\lambda_0, p)\backslash\{i\}$ and $D_2(\lambda_0) = D_1(\lambda_0) - q_i(p, z, c)$. We continue in this way with $J_k(\lambda_0, p)$ and $D_k(\lambda_0)$ for $k = 3, ..., |J_0(\lambda_0, p)| - 1$. Then, unless $J_0(\lambda_0, p) = I$, we set $\lambda_1 = \min\{p_i |\ i \in I\backslash J_0(\lambda_0, p)\}$ and, using the residual demand $D_r(\lambda_1)$ equal to $D(\lambda_1)$ minus the sum of outputs sold by the firms in $J_0(\lambda_0, p)$, instead of $D(\lambda_0)$, we repeat the above procedure for the firms in $J_0(\lambda_1, p)$. We continue in this way until we determine $q_i(p, z, c)$ for all $i \in I$. We let $q(p, z, c) = (q_i(p, z, c))_{i\in I}$. We also set $Q = [0, D(0)]^n$.

The changes in an output in a period t in comparison with an output in a period t - 1 are costly. Also, γ_i is strictly convex for all $i \in I$, which implies that $y_i(p_i, z_i, c_i)$ is finite and depends on z_i, so that firm i's preferences over various levels of its current period output depend on its output in a preceding period. So, an output vector $q(t - 1)$ is payoff relevant in a period t. The current vector of marginal costs is obviously payoff relevant. Therefore, the state vector in period t is $a(t) = (q(t - 1), c)$ and the set $A = [0, D(0)]^n \times (0, \infty)^n$ is the state space.[15] We assume that each firm in every period observes the vector of the outputs of the previous period and knows its own current marginal costs (but it need not know the current marginal costs of the other firms).

Firm i's ($i \in I$) profit in a period t, gross of costs of changing an output in comparison with a period t - 1 and gross of investment expenditures affecting the next period marginal costs, is

(3a) $$\pi_i[p(t), q(t - 1), c(t)] = (p_i(t) - c_i(t))\, q_i[p(t), q(t - 1), c(t)].$$

Firm i's net profit in a period t is

(3b) $$\pi_i[p(t), q(t - 1), c(t)] - \alpha_i[q_i(t - 1), q_i(p(t),$$

[15] This definition of the state space does not reflect the requirement that (as of period 2) any feasible $z \in Q$ satisfies $z = q(p, z', c(t))$ for some $p \in P$ and some feasible $z' \in Q$. The narrower definition of the state space, taking this into account, would be cumbersome and would not change the results of our analysis.

$q(t-1), c(t))] - x_i(t),$

where $x_i(t) = g(c_i(t), c_i(t+1))$.

We assume that binding contracts between firms are not possible. Each firm can, in any period, exit the industry without cost.[16] A Markov strategy of (the non-exiting) firm i, s_i, is a function that assigns to each element of the state space a price charged and investment expenditures affecting its next period marginal costs, i.e., $s_i{:}A \rightarrow P_i \times (0, \infty)$. Thus, a Markov strategy is a special case of a closed loop strategy.[17] The set of all feasible Markov strategies of firm i is denoted by S_i. We let $S = X_{i\in I}S_i$, $s = (s_1, ..., s_n) \in S$, and write s_{-i} for $(s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$.[18] We denote by s_{ip} the first component (price) and by s_{ix} the second component (investment expenditures) of the strategy s_i, i.e. $s_{ip}{:}A \rightarrow P_i$ and $s_{ix}{:}A \rightarrow (0, \infty)$. We also let $s_p = (s_{ip})_{i\in I}$ and $s_x = (s_{ix})_{i\in I}$.

For each firm $i \in I$, the average discounted net profit is its payoff function. The average discounted net profit of firm $i \in I$ in a subgame with an initial state (q, c), when the firms follow a Markov strategy profile s, is

$$(4) \qquad \Pi_i(s, q, c) = (1 - \beta_i) \sum_{t=1}^{\infty} \beta_i^{t-1}\{\pi_i[s_p(q(t-1), c(t)),$$

$$q(t-1), c(t)] - \alpha_i[q_i(t-1), q_i(t)] - s_{ix}(q(t-1), c(t))\},$$

where $q(0) = q$, $q(t) = q[s_p(q(t-1), c(t)), q(t-1), c(t)]$ for each positive integer t, and $c_i(t+1)$

[16] This is a natural assumption in an infinite horizon model. Since firms with a very low gross profit must also incur positive costs of investments affecting the next period marginal costs (bounded away from zero if their marginal costs are bounded from above), exit must be explicitly taken into account. Nevertheless, in order to avoid cumbersome notation, we do not add the possibility of exit to the firms' strategy spaces. Proceeds from selling the equipment of an exiting firm can be disregarded in our analysis.

[17] We restrict our attention to pure strategies. With respect to the costs of changing an output, a randomization between prices does not seem to be appealing. There is hardly any rationale for a randomization between investment expenditures.

[18] Other symbols with the subscript "−i" have an analogous meaning.

$= g^{-1}[c_i(t), s_{ix}(q(t\text{-}1), c(t))]$ for each integer $t > 1$ and every $i \in I$. (Without loss of generality, we can number the first period of a subgame by one.) We set $\Pi(s, q, c) = (\Pi_1(s, q, c), ..., \Pi_n(s, q, c))$. In what follows, "$G_O$," or "the game G_O," refers to the analyzed dynamic oligopoly.

A Markov perfect equilibrium is a profile of Markov strategies that yields a Nash equilibrium in every subgame of G_O. The following definition expresses this more formally:

Definition 1. A profile of Markov strategies $s \in S$ is a Markov perfect equilibrium of G_O if, for each state $a = (q, c) \in A$, for every firm $i \in I$, and for each strategy $s'_i \in S_i$, $\Pi_i(s, q, c) \geq \Pi_i((s'_i, s_{-i}), q, c)$.

If n - 1 firms use Markov strategies, then the best response of the remaining firm (chosen from the whole set of its closed loop strategies) is also a Markov strategy. If n - 1 firms use Markov strategies and the best response of the remaining firm is not unique, some of its best responses are also Markov strategies. Therefore, a Markov perfect equilibrium is still a subgame perfect equilibrium when the Markov restriction is not imposed.

A firm i's Markov strategy s_i is continuous if it is a continuous function from A to $P_i \times (0, \infty)$. The set of all feasible continuous Markov strategies of firm i is denoted by S^*_i. We let $S^* = X_{i\in I}S^*_i$.

A continuous strategy Markov perfect equilibrium is a Markov perfect equilibrium strategy profile in which all strategies are continuous functions. For the sake of completeness, we give a formal definition:

Definition 2. A continuous strategy Markov perfect equilibrium of G_O is a profile of continuous Markov strategies $s^* \in S^*$ such that for each state $a = (q, c) \in A$, for every firm $i \in I$, and for each strategy $s'_i \in S_i$, $\Pi_i(s^*, q, c) \geq \Pi_i((s'_i, s^*_{-i}), q, c)$.

Note that in Definition 2 we explicitly require that a strategy profile that is a continuous strategy Markov perfect equilibrium be immune to all unilateral deviations to all Markov strategies, including those that are not continuous.

We conclude this section by defining a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium:

Definition 3. A continuous strategy Markov perfect equilibrium $s^* \in S^*$ of G_O is renegotiation-quasi-proof if there is no other continuous strategy Markov perfect equilibrium of G_O, $s' \in S^* \backslash \{s^*\}$, such that $\Pi_i(s', q, c) > \Pi_i(s^*, q, c)$ for each $i \in I$ and all $(q, c) \in A$.

Thus, a continuous strategy Markov perfect equilibrium $s^* \in S^*$ is renegotiation-quasi-proof if firms, by (collectively) switching to another continuous strategy Markov perfect equilibrium, cannot increase their continuation equilibrium average discounted net profits in each subgame. This is the weakest possible concept of renegotiation-proofness. Switching to another equilibrium is assumed to take place only if it increases the continuation equilibrium average discounted payoff of each player in each subgame. Since it is such weak concept, we do not find it appropriate to call it simply "renegotiation-proofness." On the other hand, we cannot call it "weak renegotiation-proofness" because this term was already used by Farrell and Maskin (1989) for a different concept.

As is usually the case in the literature on renegotiation-proofness, the set of (subgame perfect) equilibria to which firms are allowed to renegotiate is restricted here. Since we assume, for reasons explained in the Introduction, that firms will coordinate on a continuous strategy Markov perfect equilibrium at the beginning of the game, there is no reason to assume that they will renegotiate to some other type of (subgame perfect) equilibrium.

When all discount factors are close to one, the concept of renegotiation-quasi-proofness is similar to the concept of renegotiation-proofness used in Maskin and Tirole's 1988 paper on dynamic price competition. In their paper, a renegotiation is based on the change of a price vector that is infinitely repeated (after a finite number of periods) in each continuation equilibrium. In our case it is based on the change of a common limit of all continuation equilibrium paths.

III. Necessary Conditions

In this section we show that every continuation equilibrium path (in price and marginal costs space) of each renegotiation-quasi-proof continuous strategy Markov perfect equilibrium must converge to the pair of symmetric vectors (p^*, c^*) with the following property: If the pair (p^*, c^*) is infinitely repeated, each firm's single period net profit, as well as its average discounted net profit, is maximized, subject to the constraint that infinitely repeated price and marginal cost vectors are symmetric. For each $i \in I$,

$$(5) \qquad (p_i^*, c_i^*) = \operatorname{argmax}\{n^{-1}(p_1 - c_1)D(p_1) - g(c_1, c_1) \mid p_1 \in P_1, c_1 \in (0, \infty)\}.$$

Note that the assumptions in the preceding section ensure that the maximum in (5) is unique.[19]

Proposition 1. Assume that the profile of continuous Markov strategies $s^* \in S^*$ is a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium of G_O (with all n firms active) for all vectors of discount factors $\beta \in X_{i\in I}[\beta_i^*, 1)$, where $\beta^* \in (0, 1)^n$. Then each continuation equilibrium path converges (in P x $(0, \infty)^n$) to the pair of symmetric price and marginal costs vectors (p^*, c^*) defined by (5).

[19] Obviously, we can restrict our attention to those marginal costs not exceeding ρ and such that $g(c_i, c_i) \leq \max\{p_1D(p_1) \mid p_1 \in P_1\}$.

Proof. Let Θ be the union of all continuation equilibrium paths in the state space (i.e. $\Theta \subset A$) with the following property: Following a unilateral deviation at any $a = (q, c) \in \Theta$, s^* prescribes a movement back along a continuation equilibrium path from which the deviation took place. (If a unilateral deviation at $a \in A$ cannot be punished without a decrease in prices charged by the firms that did not deviate, then $a \in \Theta$. This implies that $\Theta \neq \emptyset$.) Denote by $cl(\Theta)$ the closure of Θ. The set $cl(\Theta)$ must be a single connected compact curve in A, not containing loops. If $cl(\Theta)$ were a union of two or more disjoint (or any countable family of[20]) connected curves in A, or if it contained a loop, it would not be possible to determine, on the basis of the current state alone, from which of them (or from which part of the loop) a unilateral deviation took place. (Thus, an action profile prescribed by an equilibrium strategy profile would not depend on the state at which a unilateral deviation took place.) Therefore, for at least one of the connected subsets of $cl(\Theta)$, or for at least one part of the loop, the set of states from which the play switched to that subset of $cl(\Theta)$ (to that part of the loop) would have to be open, and this would contradict the continuity of strategies. The continuity of strategies implies that $cl(\Theta)$ must be connected. The fact that $cl(\Theta)$ is a single connected compact curve in A implies that all continuation equilibrium paths in A must converge to the same state vector. Denote this state vector by (q^+, c^+). Let p^+ be the corresponding limit of price vectors, i.e. $q^+ = q(p^+, q^+, c^+)$.

The equilibrium average discounted gross profit of a firm cannot tend to zero (or even a negative number) as the firm's discount factor tends to one (i.e. the limit of single period gross profits of no firm can be zero). The equilibrium average discounted net profit of such

[20] The set $cl(\Theta)$ cannot have a positive Lebesgue measure. In such a case, a direction of movement in the state space would have to depend on whether a unilateral deviation took place, and, if so, on the identity of the deviator. This approach is applied in Friedman and Samuelson (1994b). Such information cannot be obtained from a current state alone.

a firm would tend to a negative number (as its discount factor tends to one), so it would have to exit the industry. This implies that the price vector p^+ cannot be asymmetric. If the equilibrium average discounted gross profit of all firms with $p_i^+ > \min_{j\in I} p_j^+$ tended to a positive number, then at p^+, q^+, and c^+ the firms in $J_0(\min_{j\in I} p_j^+, p^+)$ could not produce outputs totalling $D(\min_{j\in I} p_j^+)$. This would imply $p_i^+ \leq c_i^+$ for all $i \in J_0(\min_{j\in I} p_j^+, p^+)$. Therefore, the equilibrium average discounted gross profits of these firms would tend to a non-positive number as their discount factors tended to one.

The fact that p^+ is symmetric implies that q^+ must be symmetric. If it were not, for a firm i with $q_i^+ < n^{-1}D(p_1^+)$ we would have $p_i^+ \leq c_i^+$. With respect to the requirement of a renegotiation-quasi-proofness, this implies that c^+ must also be symmetric and each c_i^+ must maximize $n^{-1}(p_i^+ - c_i)D(p_i^+) - g(c_i, c_i)$. Otherwise, a firm could, at the state (q^+, c^+), increase its continuation equilibrium average discounted net profit by unilaterally deviating only in investment expenditures. Punishing such a deviation would be impossible if the firms did not observe the current marginal costs of the other firms. If the firms were to observe the current marginal costs of the other firms, the punishment would not be renegotiation-quasi-proof. Since p^+, q^+, and c^+ are all symmetric, a renegotiation-quasi-proofness implies that $(p^+, c^+) = (p^*, c^*)$. If (p^+, c^+) were different from (p^*, c^*), a renegotiation to a continuous strategy Markov perfect equilibrium, in which all continuation equilibrium paths (in $P \times (0, \infty)^n$) converge to (p^*, c^*), would increase the continuation equilibrium average discounted net profit of each firm in every subgame.[21] Q.E.D.

---

[21] Following the assumptions of Proposition 1, a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium exists. The argument that p^+ is symmetric does not use renegotiation-quasi-proofness. If a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium with all continuation equilibrium paths converging to $(p', c') \neq (p^*, c^*)$, where p' is symmetric, exists, then a condition analogous to (8a) and (8b) below, with p^* and c^* replaced by p' and c', holds. This implies that (8a) and (8b) hold. Thus, if a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium with all continuation equilibrium paths converging to (p', c') exists, then a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium with all continuation equilibrium paths converging to (p^*, c^*) also exists. (See Proposition 2.)

IV. Sufficient Conditions

Proposition 1 in the preceding section implies that, if G_O has a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium, then all its continuation equilibrium paths (in P x $(0, \infty)^n$) converge to the symmetric price vector p^* and the symmetric vector of marginal costs c^*, defined in (5). In this section we give the sufficient conditions under which such equilibrium exists.

Proposition 2. Assume that

(6) $\quad y_i(p_i^*, q_i^*, c_i^*) \geq n^{-1}D(p_i^*), \qquad \forall\ i \in I,$

where p^* is the symmetric price vector and c^* is the symmetric vector of marginal costs defined in (5), and q^* is the symmetric vector of outputs satisfying $q_1^* = n^{-1}D(p_1^*)$. Also assume that there is a symmetric price vector $p^0 \in P$ and a vector of marginal costs $c^0 \in (0, \infty)^n$ satisfying

(7) $\quad y_i(p_i^0, 0, c_i^0) \geq n^{-1}D(p_i^0), \qquad \forall\ i \in I,$

(8a) $\quad \max\{(p_i - c_i)D(p_i) - g(c_i', c_i) \mid p_i \in [0, p_i^0], (c_i', c_i) \in (0, \infty)^2\}$

$\quad < (p_i^* - c_i^*)q_i^* - g(c_i^*, c_i^*),$[22] $\qquad \forall\ i \in I,$

and

(8b) $\quad \max\{(p_i - c_i)[D(p_i) - \sum_{j\in I\setminus\{i\}} q_j((p_{-i}^0, p_i^0 + \varepsilon), 0, (c_{-i}^0, c_i))]$

$\quad - g(c_i', c_i) \mid p_i \in [p_i^0, \rho], (c_i', c_i) \in (0, \infty)^2\}$

[22] Note 19 also applies here.

$< (p^*_i - c^*_i)q^*_i - g(c^*_i, \quad c^*_i), \qquad \forall\, i \in I,$[23]

where ε is a small positive real number. Then, there exists a vector of discount factors $\beta^* \in (0, 1)^n$ such that for each $\beta \in X_{i\in I}[\beta^*_i, 1)$ there is a renegotiation-quasi-proof continuous strategy Markov perfect equilibrium of G_O in which, in each continuation equilibrium, the play converges to the price vector p^*, to the vector of marginal costs c^*, and to the output vector q^*.

Proof. Set $q^0_i = n^{-1}D(p^0_1)$ for each $i \in I$. Let $r \in (0, 1)$ and define a function ξ:A $\to [0, 1]$ by

$$(9) \quad \xi(q,c) = \max(r\frac{q_1^0 - \min(\max(\max_{i\in I}q_i,\ q_1^*),\ q_1^0)}{q_1^0 - q_1^*} + 1 - r$$

$$- \max_{i,j\in I}(q_i - q_j)^2,\ 0)$$

Description of the equilibrium strategy profile s^*.

The price part of the strategy profile s^* is defined by

$$(10) \qquad s^*_p(q, c) = \xi(q, c)p^* + [1 - \xi(q, c)]p^0.$$

The investments part of the strategy profile s^* is given, for each $i \in I$, by $s^*_{xi}(q, c) = g[c_i, c^+_i(q, c)]$, where, for each $i \in I$ and every $(q, c) \in A$, $c^+_i(q, c)$ is the second element of a sequence $\{c^+_i(\tau)\}^{\infty}_{\tau=1}$ maximizing

∞

[23] For a set of specific forms of functions D, g, and γ_i ($i \in I$), the conditions (7), (8a) and (8b) can be expressed solely in terms of restrictions on their parameters, i.e. solely in terms of restrictions imposed on primitives of the model.

$$(11a)\quad \sum_{\tau=1} \beta_i^{\tau-1} n^{-1}[p_i(\tau) - c_i^+(\tau)]D(p_i(\tau)) - g[c_i^+(\tau), c_i^+(\tau+1)]$$

subject to $c_i^{\prime}(\tau) \in (0, \infty)$ for all $\tau \in \{1, 2, ...\}$ and

$$(11b)\quad y_i[p_i(\tau), n^{-1}D(p_i(\tau\text{-}1)), c_i^+(\tau)] \geq n^{-1}D(p_i(\tau)),$$

for all $\tau \in \{2, 3, ...\}$,

where $p(1) = s_p^*(q, c)$, $c^+(1) = c$, and the sequence $\{p_i(\tau)\}_{\tau=2}^{\infty}$ is determined by (9) and (10) under the assumption that no firm deviates. With respect to (6), along any continuation equilibrium path price vectors converge to p^*, vectors of marginal costs converge to c^*, and output vectors converge to q^*. Also note that the solution of the dynamic programming problem (11) is unique as it follows from the strict convexity of g and it is a continuous function with the state space A as its domain. Therefore, the function s_x^* is continuous.

The strategy profile s^* forms a Markov perfect equilibrium. Strategies are functions of a current payoff relevant state only, so they are Markov. Obviously, for each firm $i \in I$, we can restrict our attention to those marginal costs not exceeding $\max\{c_i(1), \rho\}$ and no lower than some level c_i^- satisfying $c_i^- \leq c_i(1)$ and $g(c_i^-, c_i^-) \leq \max\{p_iD(p_i)| \; p_i \in P_i$. Then G_O is continuous at infinity and it is enough to examine unilateral single period deviations.

Let Δ be a small positive real number such that (8a) and (8b) still hold when we replace p^0 by a symmetric price vector $p' \in P$ with each component equal to $p_1^0 + \Delta$. A deviation in price by firm i in a period t at a state $a = (q, c) \in A$ with $s_p^*(q) \notin [p^0, p']$ causes[24] the play to switch in a period t + 1 along the line segment $[p^0, p^*]$ towards p^0. (A common price prescribed by s_p^* after a deviation is lower than it would be if a deviation had not taken place.) A deviation in investment expenditures by firm i in a period t at a state a = $(q, c) \in A$ with $s_p^*(q) \notin [p^0, p']$, leading to an asymmetric vector of outputs in a period t

[24] For two n-dimensional vectors b and b′ the symbol [b, b′] denotes the set of all their convex combinations.

+ 1, causes the play to switch in a period t + 2 along the line segment $[p^0, p^*]$ towards p^0. (A deviation in investment expenditures that does not lead to an asymmetric price vector cannot increase the deviator's continuation average discounted net profit.) In the limit case $\beta_i = r = 1$ (fixing β_i and r to one beginning in the period in which a deviation took place) such a deviation strictly reduces firm i's continuation average discounted net profit. Due to the continuity of the latter in β_i and r, the same holds also for $\beta_i < 1$ and $r < 1$ close enough to one.

A deviation in price or a deviation in investment expenditures leading to an asymmetric vector of outputs, at a state $a \in A$ with $s^*_p(q, c) \in [p^0, p']$ restarts the movement along the line segment $[p^0, p^*]$, beginning with a vector in the set $[p^0, p']\backslash\{p'\}$. That is, it triggers a continuation equilibrium that gives a deviating firm a continuation equilibrium average discounted net profit no higher than the one it would earn in a subgame in the first period of which a deviation took place. The conditions (7), (8a) and (8b) imply that, when a deviator's discount factor is close enough to one, the resulting single period net profit is strictly lower than a deviator's continuation equilibrium average discounted net profit in a subgame in the first period of which the deviation took place. Therefore, such a deviation decreases a deviator's overall continuation average discounted net profit.

The continuity of s^* follows from the continuity of the functions maximum, minimum, ξ, g, $y_i(p_i, z_i, c_i)$ for all $i \in I$, and continuity of the function s^*_x.

Renegotiation-quasi-proofness of s^* follows from the definition of the price vector p^* and the vector of marginal costs c^* in (5), and the fact that $s^*_i(q^*, c^*) = (p^*_i, g(c^*_i, c^*_i))$ for all $i \in I$. Q.E.D.

So far we have only allowed an exit from the industry. Entry into the industry leads to a new renegotiation-quasi-proof continuous strategy Markov perfect equilibrium, based on new values of p^* and c^*, computed according to (5), but for a new n. This follows from the fact that the output vector and the vector of marginal costs created by entry can be viewed (for a new number of firms) as the initial state of G_O.[25]

V. Conclusions

In this paper we presented a method for the selection of the unique limit of all continuation (subgame perfect) equilibrium paths in a dynamic Bertrand oligopoly. The firms in this oligopoly produce a homogeneous good to order. Changes in their output are costly and investments affect their next period marginal costs. The method presented here is based on imposing three restrictions on strategies: Markov property, continuity, and renegotiation-quasi-proofness. In order to compute their equilibrium strategies, the firms need to know only the previous period vector of outputs and their own marginal costs. We (realistically) refrain from assuming that firms observe the marginal costs of their competitors. Nevertheless, our results continue to hold if marginal costs are publicly observable. Deviations in marginal costs (or in the investment expenditures determining them) that need to be punished, i.e. those that affect other firms, manifest themselves in deviations in outputs. We think that a similar decomposition of players' actions into the observable and the non-observable could be useful in other dynamic games in economics.

[25] Again, this new renegotiation-quasi-proof continuous strategy Markov perfect equilibrium is not uniquely determined by our approach, but the limit of all its continuation equilibrium paths is.

REFERENCES

Farrell, Joseph and Eric Maskin. 1989. "Renegotiation in Repeated Games," *Games and Economic Behavior* 1:327-360.

Flaherty, M.Thèrèse. 1980. "Industry Structure and Cost-reducing Investment," *Econometrica* 48:1187-1209.

Fudenberg, Drew and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* 54:533-554.

Friedman, James W. and Larry Samuelson. 1994a. "Continuous Reaction Functions in Duopolies," *Games and Economic Behavior* 6:55-82.

Friedman, James W. and Larry Samuelson. 1994b. "An Extension of the 'Folk Theorem' with Continuous Reaction Functions," *Games and Economic Behavior* 6:83-96.

Harsanyi, John C. and Reinhard Selten. 1992. *A General Theory of Equilibrium Selection in Games*. Cambridge, Mass: MIT Press.

Horniaček, Milan. 1997. "Selection of Markov Equilibrium in a Dynamic Oligopoly with Production to Order," CERGE Working Paper 113.

Maskin, Eric and Jean Tirole. 1987. "A Theory of Dynamic Oligopoly III: Cournot Competition," *European Economic Review* 91: 947-68.

Maskin, Eric and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica* 56:571-99.

Maskin, Eric and Jean Tirole. 1994. "Markov Perfect Equilibrium," Harvard Institute for Economic Research Discussion Paper 1698.