Available online at www.sciencedirect.com

SCIENCE DIRECT®

European Journal of Operational Research 157 (2004) 486–500

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/dsw

O.R. Applications

Timing of investments in oligopoly under uncertainty: A framework for numerical analysis

Pauli Murto [a,*], Erkka Näsäkkälä [a], Jussi Keppo [b]

[a] *Systems Analysis Laboratory, Helsinki University of Technology, P.O. Box 1100, FIN-02015 HUT, Finland*

[b] *Department of Industrial and Operations Engineering, University of Michigan, 1205 Beal Avenue, Ann Arbor, MI 48109-2117, USA*

Received 20 December 2001; accepted 18 March 2003

Abstract

We present a modeling framework for the analysis of investments in an oligopoly market for a homogenous commodity. The demand evolves stochastically and the firms carry out investment projects in order to adjust their production cost functions or production capacities. The model is formulated as a discrete time state-space game where the firms use feedback strategies. The firms are assumed to move sequentially to ensure a unique Markov-perfect Nash equilibrium. Once the equilibrium has been solved, Monte Carlo simulation is used to form probability distributions for the firms' cash flow patterns and accomplished investments. Such information can be used to value firms operating in an oligopoly market. An example of the model is given in a duopoly market. The example illustrates the trade-off between the value of flexibility and economies of scale under competitive interaction.

Keywords: Investment analysis; Game theory; Uncertainty modeling; Real options

1. Introduction

In fast developing industries, the timing of investment projects is an important part of firms' strategies. Typically, investment decisions are made under high uncertainties and they are irreversible. If an individual investment does not affect other projects, real options theory can be directly applied in choosing the optimal timing of investment. Examples of influential theoretical papers in this literature stream are, e.g., McDonald and Siegel (1986) and Pindyck (1988), while the main theory is summarized in Dixit and Pindyck (1994) and Trigeorgis (1996).

However, another characteristic feature of many industries is the competitive interaction between the firms. An investment of a firm may have a significant effect on market prices, which affects the profitability of the past and future investment projects of all firms. This makes the use of standard real options models questionable in such contexts.

Building of telecommunications capacity is a good example of an industry with the mentioned characteristics. Technologies and demand change

* Corresponding author. Address: Department of Economics, Helsinki School of Economics, P.O. Box 1210, 00101 Helsinki, Finland. Tel.: +358-9-43138384; fax: +358-9-4313 8738.

E-mail addresses: pauli.murto@hkkk.fi (P. Murto), erkka.nasakkala@hut.fi (E. Näsäkkälä), keppo@umich.edu (J. Keppo).

0377-2217/$ - see front matter
doi:10.1016/S0377-2217(03)00234-0

constantly and, therefore, network investments must be carried out regularly. In a given geographical area there usually exist only few large telecommunications companies capable of carrying out large-scale investments. A new capacity investment usually increases significantly the supply of the market, which means that in the investment planning the effects of new investments on the capacity price have to be considered.

The evident drawback of most of the existing literature of real options in this context is that the effects of competitors' actions are ignored. There are other streams of literature, however, that account for competitive interaction. A related literature considers investment in equilibrium under perfect competition and rational expectations. A seminal paper is Lucas and Prescott (1971), where the social optimality of the equilibrium in a discrete-time Markov chain model is established. Leahy (1993) discovered that the equilibrium entry time under free entry is the same as the optimal entry time of a myopic firm who ignores future entry by competitors. Another contribution is by Baldursson and Karatzas (1997), who establish the links between social optimum, equilibrium, and optimum of a myopic investor under a general stochastic demand process utilizing singular stochastic control theory. Grenadir (1999) enriches the analyses by including construction delays. However, some assumptions in these models make their application restricted. First, the assumption of perfect competition is not realistic in many industries. As mentioned earlier, in the telecom industry there are typically few firms who have the required resources to carry out new capacity investments. Second, the papers assume that investments are incremental, which means that they can be undertaken in infinitely divisible units. In reality firms often face increasing returns to scale on the project level, thus new capacity is added in lumps. Examples of real options models that study the timing and size of lumpy investments are Dixit (1995), Bar-Ilan and Strange (1999), and Dangl (1999), but they ignore competition.

There is a new literature stream that combines game theory and real options theory. Examples of such papers in discrete time are Smit and Ankum (1993) and Kulatilaka and Perotti (1998). Papers with uncertainty modeled as a diffusion process are Grenadier (1996), Lambrecht (1999), Joaquin and Butler (1999), Mason and Weeds (2001), Moretto (1996), Murto (2003), Murto and Keppo (2002), and Weeds (2002). These models are, however, restricted to very simple settings. In all mentioned papers the game is played on a single project, so the full dynamics of the industry is not analyzed. Exceptions in this respect are Baldursson (1998) and Williams (1993), who analyse the dynamics of oligopolistic industries under uncertainty. However, the investments are incremental in these models. On the other hand, dynamic oligopoly markets with multiple lumpy investments have been studied by Gilbert and Harris (1984) and Mills (1990), but in a deterministic setting assuming perfect foresight by the firms. Other models of dynamic oligopoly in deterministic settings can be found in, e.g., Fudenberg and Tirole (1986).

In the present paper, we are interested in the timing of lumpy investment projects under uncertainty and oligopolistic competition. The important characteristic is that the output price is influenced by two factors: exogenous uncertainty and new capacity investments. However, modeling the full dynamics of such markets is analytically so difficult that one has to rely on numerical analysis. For this purpose, we present a systematic modeling framework to analyze a market with the following properties. There are several large firms producing a homogeneous and non-storable commodity, the demand for which evolves stochastically. The firms have a number of investment opportunities available, which they can use to adjust their production cost functions or maximum capacities. In choosing the timings for such investments, the firms take into account the demand uncertainty and the actions of the other firms. Thus, the market can be seen as a game with two opposite factors affecting the optimal investment timing. On one hand, uncertainty and irreversibility give an incentive to postpone the investment. On the other hand, the firms recognize the preemptive effect of the investment against the other investors, which increases the incentive to act quickly.

We assume that the firms behave non-cooperatively. Their strategies are assumed to depend on the state of the market, which consists of the value of the stochastic demand process and the accomplished projects of all firms. The equilibrium resulting from such strategies is a Markov-perfect Nash equilibrium. To ensure a unique equilibrium, we make an assumption that the firms move sequentially. We present an algorithm to compute the equilibrium using backward induction.

Our approach has two steps. First, the equilibrium strategies of the firms are solved, and second, the development of demand over time is Monte Carlo simulated. This allows one to estimate the probability distributions of the firms' discounted cash flows, and thus to derive quantitative measures on the values of the firms or some of their operations under risky market development and competitive interaction. The characteristic feature of our approach is that it combines the game theoretic equilibrium concept with the simulation method that is a standard tool in analyzing market values and risk exposures under uncertainty.

To illustrate the model framework, we give an example of a market, where the demand goes through a period of strong growth, but eventually stabilizes. This kind of dynamics are typical for emerging markets. The interesting issue in such a market is the way in which the uncertainty in the demand growth is transmitted through the competitive interaction to the firms' profit distributions. In the example, we illustrate the distributions of the firms' total payoffs that result from the simulation. We compare a symmetric case to a case, where one of the firms must make investments in larger lumps than the other firm. The asymmetric case is an illustration of the trade-off between the value of flexibility and scale economies. The firm with the possibility to make small investments is more flexible, but the firm that makes larger investments is more cost-efficient. The simulations demonstrate how the properties of one firm change the shape of the other firm's payoff distribution.

The paper is organized as follows. The model framework and the main notation are presented in Section 2. In Section 3, the strategies and the equilibrium are defined. The algorithm to compute the equilibrium strategies is given in Section 4. Section 5 contains an illustrative example and Section 6 concludes.

2. Model framework

In this section we give a formal description of the model. We explore a market for a homogenous non-storable commodity within a finite time horizon. Let $\Gamma = \{0, 1, \dots, T\}$ be the set of time periods. By $t \in \Gamma$ we refer to an arbitrary time period and by $t+1 \in \Gamma$ to the period directly after period t.

The set of firms is denoted by $W = \{1, \dots, n\}$. For each $w \in W$, there exists a countable set of investment opportunities $I_w = \{I_w^1, I_w^2, \dots\}$. If firm w decides to accomplish the investment project I_w^i at period t, it has to give up the corresponding sunk cost $\delta_w^i(t)$ at t. The project will be finished by the period $t+1$. Investments cannot be reversed. Let $[I_w]$ denote the collection of all subsets of I_w including the empty set and the set I_w itself. By $\tilde{I}_w^t \subset I_w$, we refer to the set of investment projects firm w has accomplished by the time period t. We call the set $\tilde{I}_w^t$ firm w's capacity at t.

The firms produce a given homogenous commodity. The production cost is a function $C_w : \mathbf{R}_+ \times [I_w] \to \mathbf{R}_+$, so that $C_w^t = C_w(q_w^t, \tilde{I}_w^t)$ is the cost at period t with output level q_w^t and capacity $\tilde{I}_w^t$. If there is an upper limit $M(\tilde{I}_w^t)$ for output with a given capacity, we define the production cost to be infinite for higher outputs, that is, $C_w(q_w, \tilde{I}_w^t) = \infty$ for $q_w > M(\tilde{I}_w^t)$. For instance, with the telecommunications capacity, $C_w(q_w, \tilde{I}_w^t)$ is typically close to zero if $q_w < M(\tilde{I}_w^t)$ and otherwise it is equal to infinity. Thus, a firm can alter its production cost and/or increase its capacity by carrying out costly investment projects.

The price of the commodity is defined by an inverse demand function, which is subject to a stochastic demand shock variable X. The value of the demand shock at period t is $X^t \in \Theta$, where Θ is a finite subset of $\mathbf{R}$. The inverse demand function is then a mapping $P : \mathbf{R}_+ \times \Theta \to \mathbf{R}$, such that $P(Q^t, X^t)$ is the market price at period t with given industry output $Q^t = \sum_{w=1}^{p} q_w^t$ and demand shock value X^t. The shock variable follows a Markov process. Thus, there is a transition probability function $f : \Theta \times \Theta \times \Gamma \to [0, 1]$ such that the conditional transition probability is $P\{X^{t+1} = x^{t+1} | X^t = x^t\} = f(x^t, x^{t+1}, t)$.

Since the state variable X^t is stochastic, the treatment of the risk faced by the firms is crucial.

We assume that the firms discount their cash flows using a fixed discount factor γ, which is a standard assumption in theoretical papers combining real options analysis and game theory (e.g., Baldursson, 1998; Lambrecht, 1999; Mason and Weeds, 2001; Murto, 2003; Weeds, 2002; Williams, 1993). This allows the interpretation of the model using the equivalent risk-neutral valuation principle. If there are financial assets traded in a complete and arbitrage-free financial market that perfectly correlate with X, all the uncertainties in X can be hedged using these assets. This corresponds to the framework of Cox et al. (1985), which extends the basic risk-neutral valuation argument of Cox and Ross (1976) to cover non-traded assets. More precisely, in a complete market, all derivative assets can be replicated with financial portfolios, and due to the law of one price, the value of a derivative asset must be the same as the value of the replicating portfolio. This means that there exists a unique martingale measure such that the risky cash flows can be valued by discounting them with the risk-free interest rate, and taking the expectation under the martingale measure. Therefore, the model should be interpreted so that f gives the 'risk-neutral' transition probabilities, i.e., those under the martingale measure, and γ corresponds to the risk-free interest rate (which is for simplicity assumed to be constant). These risk-neutral probabilities can be calculated by using the tradable assets that are perfectly correlating with X over the discrete time intervals (see e.g., Jackwerth, 1999; Rubinstein and Jackwerth, 2001; Cakici and Foster, 2002). These assets can be, for instance, forwards on X and forwards on a risk factor that is perfectly correlated with X. Because they are traded in the market, their pricing function implicitly gives the needed risk-neutral probabilities. In order to see this let us denote the price of an asset that depends on X at time t by $\theta(t,X)$, i.e., for each time t we know the relationship between the asset price and the state variable. This means, e.g., that the cost of carry parameter in the forward pricing formula can be calculated from the existing market prices (see e.g., Hull, 2002). On the other hand, the tradable assets satisfy the risk-neutral pricing equation (see e.g., Duffie, 1996)

$$\begin{aligned}\theta(t,x^t) &= \gamma E\lfloor \theta(t+1,X^{t+1})|X^t = x^t\rfloor \\ &= \gamma \sum_{x^t\in\Theta} f(x^t,x^{t+1},t)\theta(t+1,x^{t+1})\end{aligned}$$

for all $t \in \{0,1,\dots,T-1\}$ and $x^t \in \Theta$, where the expectation is with respect to the risk-neutral probabilities. As an example, $\theta(t+1,x^{t+1})$ could be the payoff of a forward contract and $\theta(t,x^t)$ its value at time t. Thus, according to the above equation, if there are as many tradable assets as there are elements in Θ, and their pricing functions are known, we can solve for each $t \in \{0,1,\dots,T-1\}$ and $x^t \in \Theta$ the risk neutral probabilities $f(x^t,\cdot,t) : \Theta \to [0,1]$ from $|\Theta|$ linear equations (one for each tradable asset), where $|\Theta|$ is the number of elements in Θ.

On the other hand, if X is not spanned by the financial markets, the model could be interpreted so that f gives the 'true' transition probabilities, and γ corresponds to some subjective discounting by the firms, possibly adjusted to account for the risk preferences of the firms. Theoretically, the discount factor is then derived from the utility functions of the firms (see, e.g., Cochrane, 2002). Because utility functions are firm specific, the discount factors could actually be firm specific, as well as dependent on time and the state variable in this case. However, for simplicity, we retain a common constant γ in the notation throughout the paper (which is accurate when spanning conditions hold and the risk-free interest rate is constant). It should be noted that if firm-specific utility functions are used, the determination of the discount factors is in practice difficult. Moreover, it is known that the optimal investment policy may be very sensitive to the choice of the discount factor, thus a proper sensitivity analysis would be appropriate in such a case. [1]

[1] The book by Dixit and Pindyck (1994) is very illustrative concerning the matter. In that book, the contingent claims analysis (applicable when spanning holds) and dynamic programming approach (applicable without spanning, but with no sound theory for the choosing of the discounting) are used in parallel. A constant discount rate is applied throughout the book.

Given the demand shock value and the output levels of all firms, the surplus for firm w at period t is given by

$$\pi_w^t = \pi_w(X^t, Q^t, \tilde{I}_w^t) = q_w^t P(Q^t, X^t) - C_w(q_w^t, \tilde{I}_w^t). \tag{1}$$

The firms want to maximize the discounted surplus minus the investment costs over the time horizon. Given a realization of the demand shock parameter $x = [x^0, \ldots, x^T]$, and given some fixed capacity and output levels for each period, the total payoff for firm w is

$$V_w(x, Q, \tilde{I}_1, \ldots, \tilde{I}_n) = \sum_{t=0}^{T} \gamma^t \left\{ \pi_w(x^t, Q^t, \tilde{I}_w^t) - \sum_{i:I_w^i \in (\tilde{I}_w^{t+1} - \tilde{I}_w^t)} \delta_w^i(t) \right\} + \gamma^T R_w(x^T, \tilde{I}_1^T, \ldots, \tilde{I}_n^T), \tag{2}$$

where $R_w(x^T, \tilde{I}_1^T, \ldots, \tilde{I}_n^T)$ is the salvage value of w with a given demand shock value and capacities at the end of the time horizon, $Q = [Q^0, \ldots, Q^T]$ is a vector containing the output variables for all firms at all time periods, and $\tilde{I}_w = [\tilde{I}_w^0, \ldots, \tilde{I}_w^T]$ contains the evolution of w's capacity. Because investments are irreversible, it is required that $\tilde{I}_w^t \subseteq \tilde{I}_w^{t+1}$ for all w and all $t = 0, 1, \ldots, T-1$. The set $\tilde{I}_w^{t+1} - \tilde{I}_w^t$ contains projects that w has decided to build at period t, and which are then accomplished at period $t+1$. [2] Therefore, the sunk costs for them have to be paid at period t.

The firms choose their output and investment decisions in order to maximize the expectation of (2). However, as the model is dynamic, and the firms are able to observe the state of the stochastic variable X during the time horizon, they can adapt their decisions accordingly. Therefore, the output variables and capacity sets are stochastic processes adapted to the process X. In (2), the vectors Q and $\tilde{I}_w$ correspond to one given realization of these processes. Moreover, the output decisions of the firms affect each other's profits and accordingly optimal behavior. Therefore, we have to analyze the equilibrium in the firms' strategies. In the next section we define formally the strategies of the firms and the resulting Markov-perfect Nash equilibrium of the game.

[2] We have used notation $A - B = A \cap B^C$, where B^C is the complement of B.

3. Strategies and equilibrium

Let the value of the demand shock variable and the investment projects that the firms have completed define the state of the system at period t. Denoting the state space by $\Omega = \Theta \times [I_1] \times \cdots \times [I_n]$, the state vector is $Y^t = [X^t, \tilde{I}_1^t, \ldots, \tilde{I}_n^t] \in \Omega$, where $\tilde{I}_w^t$ is the set of projects completed by w. Later in this section we augment this definition with an additional variable.

It should be remarked that the two decisions variables the firms have, namely the output and the investment entry times, are different in nature. The accomplishment of an investment project has direct and permanent influences over time for all firms, while the output choice only affects the present period. Therefore, the output choice at a given period can be made without considering the investments. In this model, the emphasis is on the investment decisions, and to clarify the formulation, we make the following assumption.

Assumption 3.1. The outputs with given capacities are defined by a mapping $\Lambda : \Omega \to Q_1 \times \cdots \times Q_n$. This gives the output levels $q_1(Y^t), \ldots, q_n(Y^t)$ for all firms given the state Y^t.

This mapping can be understood as an implicit function determined by the Nash equilibrium for the short-term price or quantity competition. The total output to the market is $Q(Y^t) = \sum_{w=1}^{n} q_w(Y^t)$. According to Assumption 3.1 the output variables are functions of the state, and thus, the surplus for firm w at state Y^t can be written as

$$\pi_w^t = \pi_w(Y^t) = q_w(Y^t) \cdot P(Q(Y^t), X^t) - C_w(q_w(Y^t), \tilde{I}_w^t). \tag{3}$$

Next we start the formulation of the firms' strategies. Since the output of a firm is defined by Λ, the strategy only defines the investment decisions. As a first step, consider defining the strategy of firm w as a mapping $S_w : \Omega \times \Gamma \to [I_w]$ such that

$S_w(Y^t, t)$ contains the investment projects that firm w carries out at time t given state value Y^t. This definition for the strategies does not guarantee a unique Nash equilibrium, because the optimal action of a firm within a given period is conditional on the actions of its competitors' actions at the same period. Thus, there can be several equilibrium actions for the firms within the period.

To ensure a unique equilibrium, we make an assumption that only one of the firms can make an investment at each time period. Thus we have a sequential game where the decision makers move one at a time. It is straightforward to show that such a game has a unique sub-game perfect equilibrium and it can be solved using backward induction (see e.g., Fudenberg and Tirole, 1991).

Assumption 3.2. At each time period, only one of the firms gets the opportunity to decide whether it will undertake one or more of its investment opportunities.

This assumption seems restrictive from the point of view of actual firm management, because it sets an exogenous constraint on when it is possible to act. However, it can be argued that such a sequential decision structure approximates the actual strategic behavior, where delay and short-term commitment effects characterize firms' actions. Besides, it should be noted that the assumption allows different ways to determine which firm gets to invest at a specific period. Two natural choices are possible. First, we can assume that there is a deterministic fixed order in which the firms act. Maskin and Tirole (1988) use a similar assumption in a theoretical duopoly model, and they justify it by the notion of short-term commitment. It means that once a firm has decided to act, it is committed to the action for some time and the other firm will have a chance to react. It seems that as the number of periods is increased, the alternating moves of the firms result in a reasonable approximation of the actual decision structure.

In the second possible formulation the firm who gets the investment opportunity at a given period is drawn randomly. This can be interpreted as a short-term commitment of an uncertain duration. Thus, the strategy of every firm that may get the investment opportunity at period t must determine what the firm would do in such a case. Also, the firms' strategies have to take into account the probabilities in which they and their competitors get the chance to invest in the following periods. The symmetric situation between the competitors can be modeled easily with this kind of a formulation. Namely, the probabilities are set so that at each period all the firms have equal probabilities to act before the others. The interpretation of the setting is that the firms do not know who will be the fastest to invest if the conditions turn favorable. Therefore, when a firm has a chance to invest, it has to take into account the fact that before it gets another chance, another firm may already have invested.

Next, we augment the model so that it allows the formulation of the strategies so that Assumption 3.2 is taken into account. We do that by adding one variable in the state vector. This new variable, denoted by U, gets values in W, and $U^t = w$ means that at period t, firm w has the opportunity to make an investment.

The variable U^t, defined for all t, is either a deterministic or a stochastic process. To retain as much generality as possible, we define the evolution of the process in the same way as in the case of the demand shock variable. Thus, there is a transition probability function $g : W \times W \times \Gamma \to [0, 1]$ such that the conditional probability $P\{U^{t+1} = u^{t+1} | U^t = u^t\} = g(u^t, u^{t+1}, t)$. If every firm has the same probability at each period to get the chance to invest, then we have $g(w', w, t) = 1/n$ for all w, w', and t.

Let $\Psi = \Theta \times [I_1] \times \cdots \times [I_n] \times W$ be the augmented state space and define the new state vector as

$$Z^t = [X^t, \tilde{I}_1^t, \ldots, \tilde{I}_n^t, U^t] \in \Psi. \quad (4)$$

Now we can formally define the strategies of the firms.

Definition 3.1. The strategy of the firm $w \in W$ is a mapping $S_w : \Psi \times \Gamma \to [I_w]$ such that the set $S_w(Z^t, t)$ contains the investment projects that firm w carries out at time t given state value Z^t. $S = [S_1, \ldots, S_n]$ is called a strategy profile. Let Φ_w

denote the set of admissible strategies of firm w. The strategy space Φ_w contains all mappings $S_w : \Psi \times \Gamma \to [I_w]$ satisfying the two conditions:

(1) $S_w(Z^t, t) \subset (\tilde{I}^t_w)^C$ for all t and all Z^t,
(2) $S_w(Z^t, t) = \emptyset$ for all such Z^t that $U^t \neq w$.

Condition (1) means that w can only invest in projects that it has not yet accomplished. Condition (2) simply formalizes the assumption of sequential investment: w cannot invest if some other firm gets the opportunity to invest.

As the projects that firm w decides to carry out at period t are accomplished at period $t+1$, the set of projects $\tilde{I}_w$ completed by w evolves therefore according to equation $\tilde{I}^{t+1}_w = \tilde{I}^t_w \cup S_w(Z^t, t)$. Since the strategies are functions of the stochastic variables X and U, the sets of completed projects are stochastic processes adapted to X and U.

The state transition rule of the model is summarized in the following three equations:

$$\tilde{I}^{t+1}_w = \tilde{I}^t_w \cup S_w(Z^t, t), \tag{5}$$

$$P\{X^{t+1} = x^{t+1} | X^t = x^t\} = f(x^t, x^{t+1}, t), \tag{6}$$

$$P\{U^{t+1} = u^{t+1} | U^t = u^t\} = g(u^t, u^{t+1}, t). \tag{7}$$

Eqs. (6) and (7) describe the evolution of the processes X and U, and (5) describes how the firms' capacities evolve according to their strategies given a realization of the processes X and U.

The firm w chooses the optimal strategy in terms of maximal expected profits. Let $V_w(Z^t, S, t)$ be the expected net present value of all cash flows from period t to the last period for firm w given strategies $S = [S_1, \ldots, S_n]$. This satisfies the recursive equations

$$\begin{aligned} &V_w(Z^t, S, t) \\ &= \pi_w(Z^t) - \sum_{i: I^i_w \in S_w(Z^t, t)} \delta^i_w(t) \\ &\quad + \gamma \left\{ E_{X,U}\left[V_w(X^{t+1}, \tilde{I}^{t+1}_1, \ldots, \tilde{I}^{t+1}_n, U^{t+1}, S_1, \ldots, S_n, t+1) \right] \right\}, \end{aligned} \tag{8}$$

where the state evolves according to (5)–(7). At all possible states, the firms choose strategies that give them the highest possible expected payoffs. The equilibrium of the game is defined as:

Definition 3.2. The Nash equilibrium is a strategy profile $S^* = [S^*_1, \ldots, S^*_n] \in \Phi_1 \times \cdots \times \Phi_n$ such that

$$\begin{aligned} &V_w(Z^t, S^*_1, \ldots, S^*_n, t) \\ &\quad \geqslant V_w(Z^t, S^*_1, \ldots, S_w, \ldots, S^*_n, t) \\ &\quad \forall w \in W, \ \forall t \in \Gamma, \ \forall Z^t \in \Psi, \ \forall S_w \in \Phi_w. \end{aligned} \tag{9}$$

In a stochastic game, state dependent strategies as given in Definition 3.1 are called Markov strategies. A profile of strategies giving Nash equilibrium in every proper sub-game, as the Definition 3.2 requires, is called a Markov-perfect equilibrium (MPE). Since only one firm can act at each period, there is certainly one unique MPE (see, e.g., Fudenberg and Tirole, 1991).

4. Computing the equilibrium

Let $V^*_w(Z^t, t)$ denote the expected discounted value of the cash flows for w at Z^t when all the firms use their MPE strategies from there on. These strategies can be computed using backward induction as follows:

1. Start from period T and calculate the present value of surplus $\pi_w(Z^T)$ at all possible states for all w using Eq. (3). The value function at the last period is given value $V^*_w(Z^t, T) = \pi_w(Z^T) + R(X^T, \tilde{I}^T_1, \ldots, \tilde{I}^T_n)$. Set $S^*_w(Z^T, T) = \emptyset$ for all w and Z^T. Set the time index at $t = T - 1$.
2. Using recursion (8), calculate $V^*_w(Z^t, t)$ for all w at all states at period t. In computing $V^*_w(Z^t, t)$ for the firm w, which has the opportunity to invest, assign the set $S^*(Z^t, t)$ that gives the largest $V^*_w(Z^{t+1}, t+1)$ when Z^t evolves according to (5)–(7). For the other firms, set $S^*_w(Z^t, t) = \emptyset$.
3. If $t > 0$, set $t := t - 1$ and go back to step 2. Otherwise the equilibrium strategies $S^*(Z^t, t)$ and values $V^*_w(Z^t, t)$ have been calculated for all w and Z^t.

After the equilibrium strategies have been calculated, the market can be simulated. The state is initiated at $t = 0$ by giving initial values for $\tilde{I}^0_w$, X^0, and U^0. The stochastic variable is given a value for the next period, and the sets $\tilde{I}^t_w$ are updated according to (5) using the previously computed strategies. This is continued until $t = T$. X^t and U^t can be chosen randomly using Eqs. (6) and (7), or different scenarios can be investigated by deterministically assigning certain paths for the variables.

5. Example case

In this section we present an example case, where a number of specifications are introduced in the general model developed in Sections 2 and 3. The purpose of this example is to illustrate the modeling framework rather than to derive conclusions concerning the economic theory of dynamic oligopoly. To carry out the computations, we have written a computer program for this specific purpose in C programming language. This program implements the algorithm described in Section 4.

The example considers a market, where the demand shock process goes through three phases. In the first phase the demand is stable. In the second phase the demand grows strongly, until it again stabilizes in the third phase. Such dynamics take often place in an emerging market. The firms know in advance that in the future the demand will grow rapidly and eventually stabilize. In any case, there are high uncertainties on the exact path that the demand development will follow. The different phases of the demand process are modeled with different parameterizations of the stochastic process.

We make the following assumptions to specify the example:

(1) There are two firms $W = \{1, 2\}$ with the following features:
   (a) Each firm has investment opportunities of one single type.
   (b) The number of investment opportunities is unlimited.
   (c) Each investment costs δ_w units of money.
   (d) Each investment increases the maximum capacity by an amount K_w^+.
   (e) The initial capacity is 0 for both firms.
   (f) The variable production cost is zero up to the capacity limit. In the notation of Section 2, we have:

$$C_w(q_w, \tilde{I}_w^t) = \begin{cases} 0, & \text{when } q_w \leqslant mK_w^+ \\ \infty, & \text{when } q_w > mK_w^+ \end{cases},$$

where m is the number of elements in $\tilde{I}_w^t$.

(2) The number of periods is 101, i.e., $\Gamma = \{0, 1, \ldots, 100\}$.

(3) The demand shock process is specified by the transition rule:

$$X^{t+1} = \begin{cases} u \cdot X^t & \text{with probability } p^t \\ (1/u) \cdot X^t & \text{with probability } 1 - p^t \end{cases},$$

$$\text{where } p^t = \begin{cases} 0.5 \text{ for } t \in \{0, 1, \ldots, 32\}, \\ 0.9 \text{ for } t \in \{33, \ldots, 66\}, \\ 0.5 \text{ for } t \in \{67, \ldots, 100\}. \end{cases}$$

(4) The inverse demand function is time-invariant and has constant elasticity: $P(Q^t, X^t, t) = X^t \cdot a \cdot (Q^t)^{-\frac{1}{\varepsilon}}$, where price elasticity satisfies $\varepsilon > 1$.

(5) Each firm has an equal probability to get an opportunity to invest at each period. Formally, $g(1,1,t) = g(1,2,t) = g(2,1,t) = g(2,2,t) = 1/2$ $\forall t \in \Gamma$.

(6) The salvage value is $R_w(X^T, \tilde{I}_1^T, \tilde{I}_2^T) = \sum_{t=T+1}^{\infty} \gamma^{t-T} \pi_w^T$.

We next discuss these assumptions. Assumptions 1(a)–(e) specify that the firms carry out investments in order to increase their production capacities in lumps of constant size. Each firm has an unlimited set of such investment opportunities, [3] but due to technological restrictions, each firm can undertake projects of one single type only. We allow firm specific investment costs and project sizes, which reflects possibly different technological resources possessed by different firms (for example patents or technological competence). In some contexts, e.g., in electricity production, this might also reflect specific natural resources or local conditions available to the firms.

The assumption of zero variable cost (Assumption 1(f)) is a plausible assumption for telecommunications capacity, where the cost of

[3] In the actual computations we have used fixed maximum capacities, but set these maximums so high that it is never relevant for a firm to use all of its investment opportunities.

transmitting a unit of data is negligible given that the sufficient transmission infrastructure is in place (for more on this, see e.g., Clark, 1997; MacKie-Mason and Varian, 1996). Thus, all costs associated with production arise through the fixed investment costs. However, this kind of a cost structure would not be realistic in some other industries of homogenous commodities. For example, in electricity power generation considerations of variable production costs are very important with some technologies, whereas with some other technologies, e.g., hydro power, wind power, or geothermal power, the zero variable cost assumption can be well justified. Note that in our model during the long optimization period the fixed investment costs can be viewed as variable costs because the companies choose their investment amounts.

The number of periods (Assumption 2) is chosen large enough to have a reasonable approximation of the decision structure, where the firms decide about the timing of investments in continuously evolving market. The demand shock process in Assumption 3 reflects the three phases of market dynamics, the stable beginning and end, and the high growth in the middle. [4] As discussed in Section 2, if there are enough financial instruments that perfectly correlate with X then these transition probabilities can be calculated from these instruments (for more on this kind of estimation see e.g., Jackwerth, 1999; Rubinstein and Jackwerth, 2001; Cakici and Foster, 2002). For instance, in the telecommunications market these transition probabilities can be calculated from bandwidth leasing contracts (see e.g., Keppo, 2002). The expected demand shock value with 68% confidence interval is illustrated in Fig. 1.

The inverse demand function with constant elasticity (Assumption 4) has been chosen for two reasons: first because it has the realistic property that the price gets always a positive value, and

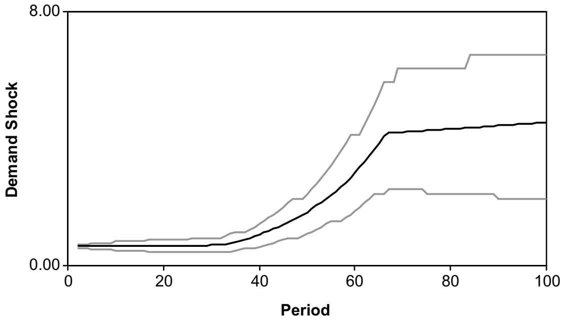

Fig. 1. Expected value of the demand shock process.

second because it is convenient for specifying the optimal outputs of the firms given fixed capacities. Namely, given zero marginal costs and given that the elasticity is high enough (more precisely when $\varepsilon > 1$), it is easy to show that the marginal revenue is positive for any firm, no matter what the outputs of the firms are. This implies that in equilibrium the firms want to utilize fully their capacities, that is, we have $q_w(Y^t) = mK_w^+ \ \forall Y^t \in \Omega$. Thus, we have an example of a realistic situation, where the strategic behavior concerns only the timing of investments, but not the output decisions. The assumption of positive marginal revenues has also been used in many influential theoretical models of dynamic oligopolies, e.g., Gilbert and Harris (1984), Ghemawat and Nalebuff (1990), and Mills (1990). Alternatively, in case of a lower elasticity, we could simply specify the mapping Λ so that the outputs are determined as the Cournot–Nash equilibrium of the short-term quantity game with fixed capacities.

Assumption 5 simply means that the turns in which the firms are allowed to move are randomized symmetrically. For simplicity, we assume that the agents are risk-neutral with respect to this uncertainty. Assumption 6 means that the salvage value for each firm at the last period is equal to the discounted stream of future cash flows assuming that the last period demand shock value and capacities will prevail forever.

The equilibrium strategies are calculated according to the algorithm given in Section 4. As given in Definition 3.1, the strategy determines which projects the firm undertakes at a given state. In the context of this example, where each firm can

---

[4] Varying the probabilities of going up actually also changes the volatility of the demand process. With the parameter value $u = 1.07$, which we use in the computations (see Table 1), the standard deviation of the percentage change of X within one period is about 6.8% at the beginning and end of the time horizon and about 4% during the high growth phase.

only undertake one type of projects, it is possible to express the strategies graphically. Namely, given some fixed capacities for both of the firms, it is easy to understand that the higher the demand shock value, the more attractive it must be for a firm to take the next investment project. Thus, the strategies can be expressed as "threshold curves" such that a threshold curve value with given capacities represents the lowest demand shock value at which the firm is willing to take the next project. This type of characterization of the optimal investment rules is typical in the real options literature (see, e.g., Dixit and Pindyck, 1994).

Fig. 2 shows the lowest threshold curves for firm 1 in a symmetric situation where $\delta_1 = \delta_2$ and $K_1^+ = K_2^+$ (in this symmetric situation the strategies for firm 2 are identical). The exact parameter values are given in Table 1. The two numbers associated with each threshold curve refer to the capacities of firms 1 and 2, respectively. For example, the curve 0–1 represents the lowest demand shock value at which firm 1 is willing to build one more capacity unit given that it has 0 units of capacity while firm 2 already has 1 unit of capacity. The gray borders represent the highest and lowest possible state values that can ever be reached. One can see that the threshold curves are not monotonic, which is due to the different phases of the demand growth. In the beginning, where the high growth phase is approaching, it becomes more and more attractive to invest, thus the thresholds are

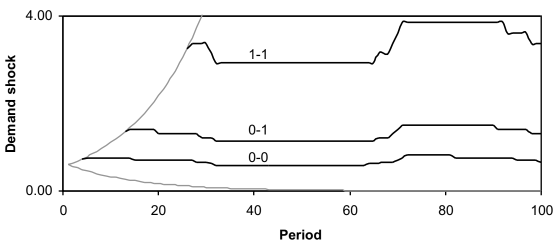

Fig. 2. Threshold curves.

decreasing. However, when the high growth period has started, and the eventual stabilization of the market is approaching, the investing becomes less attractive. Finally, the end transition with the high salvage value associated with the existing capacities and the absence of any further investment opportunities makes the investment increasingly attractive towards the end of the time horizon.

From the threshold curves, one can conclude the most likely order in which the firms carry out investments. Since the firms have identical strategies, it is equally likely for both of the firms to be the first to invest (under Assumption 5 that assigns equal probabilities for the investment turns). Once one of the firms has invested in one capacity unit, the question is who will undertake the next investment. Since the curve 0–1 is below that of 1–0 (which actually lies outside the figure area), it is more attractive for the firm with no capacity to build its first capacity unit than for the firm with 1 unit to build its second unit. Therefore, the second investment will be undertaken by the firm, which does not have any capacity yet. This same patterns will continue further, that is, one can see that the firms make investments sequentially so that the difference in capacities is at most one unit. This resembles the result in Baldursson (1998), where asymmetries in initial capacities are gradually evened out in a dynamic oligopoly, where firms increase capacity continuously.

Even if one can get some insight on how the market is likely to develop by looking at the threshold curves, the actual development of the market depends of course on the realizations of the stochastic processes. Fig. 3 illustrates the equilibrium investments for one randomly chosen realization of X and U. Also the price is presented in the figure. The parameters are given in Table 1. The first investment occurs at period 4 where firm 2 increases its capacity by one unit. Firm 1 follows at period 47. At period 65 firm 2 undertakes its second investment, and firm 1 follows at period 68.

Table 1
Parameter values with symmetric firms

Parameter	δ_1	δ_2	K_1^+	K_2^+	γ	u	a	ε	X^0
Value	10	10	1	1	0.97	1.07	0.5	1.1	0.6

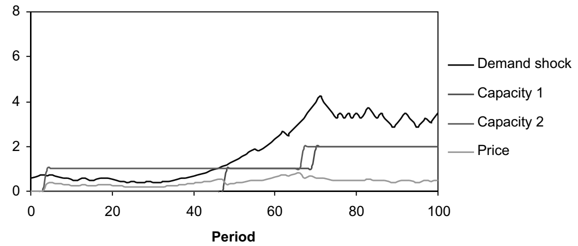

Fig. 3. Example with symmetric firms.

Therefore, the total capacity at the final period is 4 units. Each capacity investment has a decreasing effect on the market price, as seen in the figure.

Next, we change the parameters so that the firms are different with respect to the investment opportunities available. Firm 2 can now only make investments twice as large as firm 1, but on the other hand, investment cost per unit of capacity is considerably lower. Fig. 4 illustrates the equilibrium investments for one randomly chosen realization. The parameters used are summarized in Table 2. This time, firm 1 increases its capacity at period 13 by one unit, while firm 2 waits until the period 44 before increasing its capacity by two units. After that, firm 1 still makes two investments, at periods 63 and 82. This time, the total capacity at the final period is five units, where firm 1 has three units and firm 2 has two units.

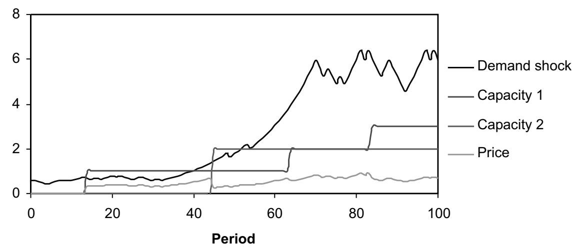

Fig. 4. Example with asymmetric firms.

It should be emphasized that the strategies as given in Definition 3.1 are deterministic decision rules that adapt to all market states. The optimal strategies are determined by taking into account all possible states of $Z = [X, \tilde{I}_1, \tilde{I}_2, U]$ and their probabilities. However, the actual investments and profits are stochastic, because they depend on the realization of the processes X and U. Therefore, Monte Carlo simulation can be used to analyze the probability distributions of the firms' cash flow patterns. This is done by repeatedly simulating X and U, and using the equilibrium strategy profile at each realization of the simulation to calculate the resulting cash flows. Fig. 5 shows the scatter plot of the firms' total discounted profits for 5000 realizations with the symmetric parameterization of Table 1. Fig. 6 presents the total profits for 5000 realizations with the asymmetric parameterization of Table 2.

As could be expected, the firms' payoff distributions in Fig. 5 are similar due to the identical cost structures. The outcomes are evenly spread around the dashed line, along which both of the firms have exactly the same payoffs. In Fig. 6 the firms have asymmetric cost structures, so the total profit distributions are not identical either. The scatter plot is curved and the outcomes are more distinctly split into different areas. When both firms get relatively low payoffs, it is typically firm 2 that profits more, which can be seen as an accumulation of points above the dashed line where the curvature of the distribution is highest. On the other hand, when both firms get high payoffs, which happens when the demand realization is high, it is typically firm 1 that profits more, which can be seen as an accumulation of points at the upper right hand corner below the dashed line. There are also some points along the line $y = 0$, which reflects the fact that in some realizations the demand growth is so weak that firm 2 does not make even a single investment.

Table 2
Parameter values with asymmetric firms

Parameter	δ_1	δ_2	K_1^+	K_2^+	γ	u	a	ε	X^0
Value	10	16	1	2	0.97	1.07	0.5	1.1	0.6

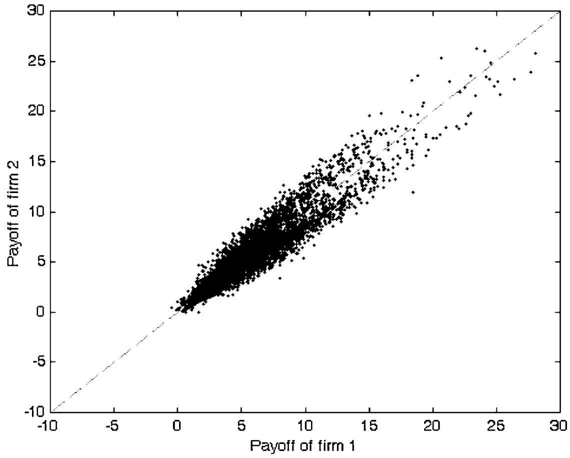

Fig. 5. Scatter plot of payoffs in symmetric case.

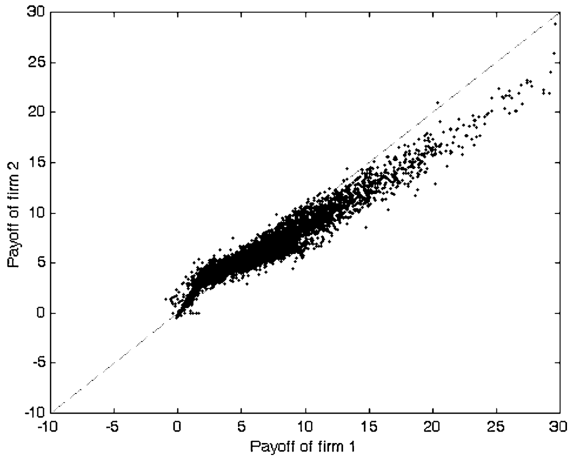

Fig. 6. Scatter plot of payoffs in asymmetric case.

This example is an interesting illustration of the trade-off between the value of flexibility and scale economies under competition. Firm 1 is more flexible with the possibility to make smaller investments, but firm 2 is more cost-efficient since the cost per unit of capacity is lower. It is interesting to note that with some realizations of the shock process flexibility is a more advantageous property, while for other realizations the converse is true.

The fact that firm 1 is able to make the largest payoffs may seem counterintuitive. It could be expected that when demand growth is very high, firm 2 would be in an advantageous position with the possibility to make large investments. However, a closer examination of individual simulation results reveals that along those realizations where demand grows strongly, firm 1 gradually moves far ahead of firm 2 in capacity. Typically, in high demand realizations the total capacities at the last period are around 6–9 for firm 1, but only 4–6 for firm 2. Thus, it seems that it is easier for firm 1 to keep up with the growing demand, since the small investments allow gradual capacity expansion as the demand growth goes along, without having so much downside risk due to large individual investments. On the other hand, in the cases where demand growth is not so high, even one large investment made by firm 2 prevents firm 1 to fully benefit from the advantage of being able to make small investments.

Fig. 7 presents a histogram of the profits in 5000 simulations for both firms in the symmetric setting with the parameterization of Table 1. Fig. 8 contains the profit histograms for the firms in the asymmetric setting with the parameterization of Table 2.

It is interesting to note that the shapes of the payoff distributions change for both of the firms when moving from the symmetric case to the asymmetric one. For firm 1 the payoff distribution gets wider with more outcomes giving high

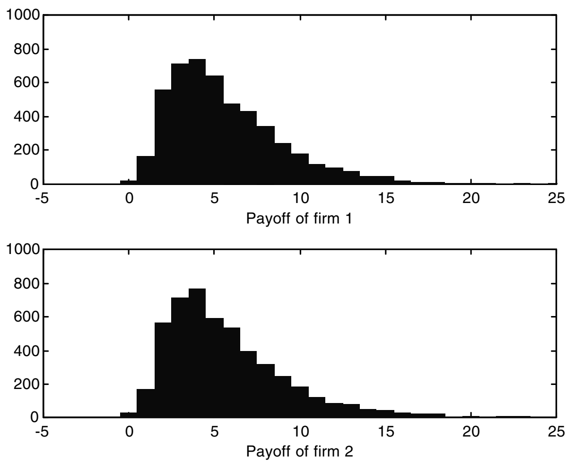

Fig. 7. Payoff histograms in symmetric case.

payoffs. An explanation is that when restricted to large lumpy investments, firm 2 is not able to adapt as well to demand changes, and thus, facing less competitive response, firm 1 can better utilize the realizations with high demand growth to its advantage. Indeed, in most of the realizations firm 2 makes 1–3 investments, whereas the number of investments undertaken by firm 1 ranges much more, typically from 1 to 9. As there is little variation in the investment behavior of firm 2, and on the other hand the investment behavior of firm 1 adapts more effectively to demand growth thus having a stabilizing effect on the price, there is little payoff variation for firm 2. This can be seen as a more concentrated payoff distribution. Overall, an interesting point to note is that the change in the parameters of one of the firms not only has an effect on its own payoff, but due to the strategic interaction it changes the shape of the payoff distribution for the other firm as well.

It is also evident from Figs. 7 and 8 that the profit distributions are not normally distributed. There is a considerable tail in the distributions at high total profits, while there is much less variation at the left hand side of the mean value. This is not surprising, because the distribution of the increment in the shock variable within one of the three separate growth phases approaches log-normal distribution as the number of periods is increased. Even though there is a non-linear relationship between the payoff per period and the shock variable, it can be expected that the total payoff is skewed to left.

Finally, it should be noted from Figs. 7 and 8 that the total payoffs for both of the firms are rather small when compared to the investment costs. This reflects the rent dissipation effect due to the preemptive competition. This can be further illustrated by comparing the duopoly profits to the monopoly case. Fig. 9 shows the histogram of the total industry payoffs (sum of the firms' payoffs) in the symmetric duopoly with parameters of Table 1. As a comparison, the figure shows also the histogram of a single firm with the same parameter values, but in the absence of competition. One can see that the competitive pressure drives the industry profits to relatively low levels even in the case of only two firms.

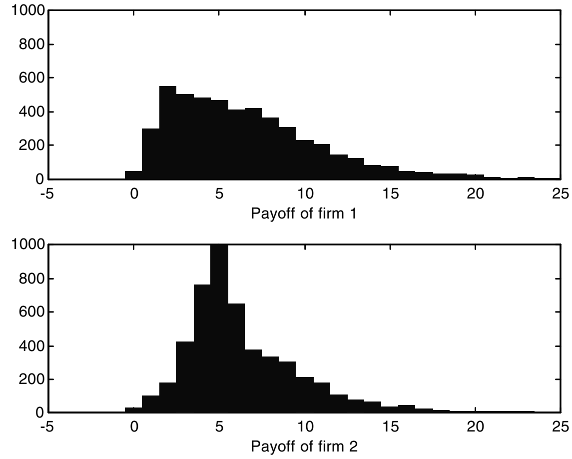

Fig. 8. Payoff histograms in asymmetric case.

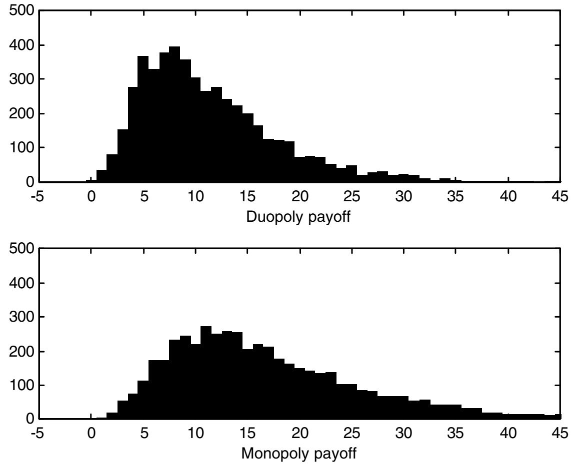

Fig. 9. Industry payoffs in duopoly and monopoly.

This example considers a duopoly, but an extension to more than two firms would be straightforward in principle. However, as the increase in the number of firms increases the dimension of the state space, such an extension increases significantly the required computing capacity. To extend our example to an oligopoly with three or more asymmetric firms would in practice require a reduction in the number of periods. Although it would be interesting to look at the effect of varying the number of firms, such an extension is not expected to add any conceptually important new feature in the model. It can be expected that the increased number of firms would drive the industry profits to even lower levels than in duopoly.

6. Conclusions

This paper considers the dynamics of a market, where there are several large producers of a homogenous commodity, and where each producer faces a set of discrete investment opportunities. By accomplishing investment projects, the firms lower their production costs or increase their maximum capacities (or, in general, alter their production cost functions). The set of investment opportunities can be seen as a portfolio of real options, and the firms have to decide the optimal timing of investments in an uncertain environment.

What makes the setting difficult to analyze is the competitive interaction between the firms. The profits of all firms are affected by the output choice of any individual firm through the inverse demand function. Therefore, in deciding the optimal investment timing the firms have to take into account the effects of their actions on the behavior of the other firms.

We have not attempted to conduct an analytical treatment of the equilibrium in such a setting. Instead, we have proposed a systematic two-stage approach for analyzing numerically such a market. In the first stage, the equilibrium strategies of the firms are computed. The industry equilibrium can be understood as a set of such strategies, or decision rules, where no firm has an incentive to change her strategy at any possible state of the game. In other words, we refer to the Markov-perfect Nash equilibrium as the solution concept. In the second stage the equilibrium strategies are used to simulate the market. This makes it possible to analyze the statistical properties of the cash flows resulting from such a competitive market. In other words, the model allows one to make conclusions on how the primary uncertainty in the exogenous demand shock parameter is transmitted through the competitive interaction to the resulting profit flows.

We have illustrated the model with an example of two firms that compete against each other in building, e.g., telecommunications capacity. We have assumed that the market for such capacity grows slowly in the beginning, then goes through a period of strong growth, and eventually stabilizes. These three stages of growth are modeled by using different parameters in the stochastic process characterizing the demand growth. We have illustrated the total payoff distributions of the firms resulting from the Monte Carlo simulations. The payoff distributions are skewed to left. The asymmetric case, where one of the firms can only make investments in larger lumps than the other firm, illustrates the trade-off between flexibility and scale-economies. In our example, flexibility is a more advantageous property with some realizations of the stochastic demand growth, while the converse is true for others. We also demonstrate that the change in the parameters of one firm not only has an effect on its own payoff, but changes the statistical properties of the other firm's payoff distribution as well.

Acknowledgements

We would like to thank Fridrik Baldursson, Pierre-Olivier Pineau, Rune Stenbacka, Jon Thor Sturluson, and two anonymous referees for helpful comments on earlier versions of the paper. Partial financing by the Nordic Energy Research is gratefully acknowledged.

References

Baldursson, F., 1998. Irreversible investment under uncertainty in oligopoly. Journal of Economic Dynamics and Control 22, 627–644.

Baldursson, F., Karatzas, I., 1997. Irreversible investment and industry equilibrium. Finance and Stochastic 1, 69–89.

Bar-Ilan, A., Strange, W.C., 1999. The timing and intensity of investment. Journal of Macroeconomics 21 (1), 57–77.

Cakici, N., Foster, K.R., 2002. Implied binomial trees from the historical distribution, Risk Magazine, forthcoming.

Clark, D.D., 1997. A model for cost allocation and pricing in the internet. In: McKnight, L.W., Bailey, J.P. (Eds.), Internet Economics. MIT Press, Cambridge, MA.

Cochrane, J.H., 2002. Asset Pricing. Princeton University Press, Princeton, NJ.

Cox, J., Ross, S., 1976. The valuation of options for alternative stochastic processes. Journal of Financial Economics 3, 145–166.

Cox, J., Ingersoll, J., Ross, S., 1985. An intertemporal general equilibrium model of asset prices. Econometrica 53, 363–384.

Duffie, D., 1996. Dynamic Asset Pricing Theory, second ed. Princeton University Press, Princeton, NJ.

Dangl, T., 1999. Investment and capacity choice under uncertain demand. European Journal of Operational Research 117, 415–428.

Dixit, A., 1995. Irreversible investment with uncertainty and scale economies. Journal of Economic Dynamics and Control 19, 327–350.

Dixit, A., Pindyck, R.S., 1994. Investment under Uncertainty. Princeton University Press, Princeton, NJ.

Fudenberg, D., Tirole, J., 1986. Dynamic Models of Oligopoly. Harwood Academic Publishers, London.

Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press, Cambridge, MA.

Ghemawat, P., Nalebuff, B., 1990. Exit. Rand Journal of Economics 16, 184–194.

Gilbert, R.J., Harris, R.G., 1984. Competition with lumpy investment. Rand Journal of Economics 15 (2), 197–212.

Grenadier, S.R., 1996. The strategic exercise of options: Development cascades and overbuilding in real estate markets. Journal of Finance 51 (5), 1653–1679.

Grenadir, S.R., 1999. Equilibrium with time-to-build: A real options approach. In: Brennan, M.J., Trigeorgis, L. (Eds.), Project Flexibility, Agency, and Competition: New Developments in the Theory of Real Options. Oxford University Press, New York.

Hull, J., 2002. Options, Futures, and Other Derivatives, fifth ed. Prentice Hall, Englewood Cliffs, NJ.

Jackwerth, J., 1999. Option implied risk-neutral distributions and implied binomial trees: A literature review. Journal of Derivatives 7 (2), 66–82.

Joaquin, D.C., Butler, K.C., 1999. Competitive investment decisions: A synthesis. In: Brennan, M.J., Trigeorgis, L. (Eds.), Project Flexibility, Agency, and Competition: New Developments in the Theory of Real Options. Oxford University Press, New York.

Keppo, J., 2002. Pricing of bandwidth derivatives under network arbitrage condition. University of Michigan, Preprint.

Kulatilaka, N., Perotti, E.C., 1998. Strategic growth options. Management Science 44 (8), 1021–1031.

Lambrecht, B., 1999. Strategic sequential investments and sleeping patents. In: Brennan, M.J., Trigeorgis, L. (Eds.), Project Flexibility, Agency, and Competition: New Developments in the Theory of Real Options. Oxford University Press, New York.

Leahy, J.V., 1993. Investment in competitive equilibrium: The optimality of myopic behavior. Quarterly Journal of Economics CVII (4), 1105–1133.

Lucas, R.E., Prescott, E.C., 1971. Investment under uncertainty. Econometrica 39 (5), 659–681.

MacKie-Mason, J., Varian, H., 1996. Some economics of the internet. In: Sichel, W. (Ed.), Networks, Infrastructure and the New Task for Regulation. University of Michigan Press, Ann Arbor.

Maskin, E., Tirole, J., 1988. A theory of dynamic oligopoly I: Overview and quantity competition with large fixed costs. Econometrica 56 (3), 549–569.

Mason, R., Weeds, H., 2001. Irreversible Investment with Strategic Interaction. Discussion Papers in Industrial Organization 3013. Center for Economic Policy Research, London.

McDonald, R., Siegel, D., 1986. The value of waiting to invest. Quarterly Journal of Economics 101, 707–728.

Mills, D.E., 1990. Capacity expansion and the size of plants. RAND Journal of Economics 21 (4), 555–566.

Moretto, M., 1996. Option Games: Waiting vs. Preemption in the Adoption of New Technology, Nota Di Lavoro 55.96, University of Padua, Fondazione Eni Enrico Mattei.

Murto, P., 2003. Exit in Duopoly under Uncertainty, RAND Journal of Economics, forthcoming.

Murto, P., Keppo, J., 2002. A game model of irreversible investment under uncertainty. International Game Theory Review 4, 127–140.

Pindyck, R.S., 1988. Irreversible investment, capacity choice, and the value of the firm. American Economic Review 78 (5), 969–985.

Rubinstein, M., Jackwerth, J., 2001. Recovering probabilities and risk aversion from option prices and realized returns, In: Lehmann, B. (Ed.), Essays in Honor of Fisher Black. Oxford University Press, Oxford, forthcoming.

Smit, H.T.J., Ankum, L.A., 1993. A real options and game theoretic approach to corporate investment strategy under uncertainty. Financial Management (Autumn), 241–250.

Trigeorgis, L., 1996. Real Options: Managerial Flexibility and Strategy in Resource Allocation. MIT Press, Cambridge, MA.

Weeds, H., 2002. Sleeping patents and compulsory licencing: An options analysis. Review of Economic Studies 69, 729–747.

Williams, J.T., 1993. Equilibrium and options on real assets. The Review of Financial Studies 6 (4), 825–850.