OR Spectrum
DOI 10.1007/s00291-006-0049-0

REGULAR ARTICLE

Jean-Philippe Gayon · Yves Dallery

Dynamic vs static pricing in a make-to-stock queue with partially controlled production

Abstract We consider a capacitated make-to-stock production system that offers a product to a market of price-sensitive users. The production process is partially controlled. On the one hand, the decision-maker controls the production of a single facility. On the other hand, an uncontrolled flow of items arrives at the stock. Such a situation occurs in several contexts; for example, when there is a return flow of products or a fixed delivery contract. We model the system as a make-to-stock queue with lost sales. We address the static pricing problem and the dynamic pricing problem with the objective of maximizing the average profit over an infinite horizon. For both problems, we characterize the optimal production and pricing policy. We also obtain analytical results for the static pricing problem. From numerical results, we show that dynamic pricing might be much more beneficial when the production is not totally controlled.

Keywords Make-to-stock queue · Static pricing · Dynamic pricing · Markov decision process

JEL Classification C00

1 Introduction

Recent years have seen an increased adoption of dynamic pricing (DP) strategies in retail and manufacturing companies. However, in industries where the sellers have the capability to store inventory and to replenish it, the benefits of DP with respect to static pricing (SP) are not always obvious. When pricing and replenishment

J.-P. Gayon (✉)
Laboratoire Gilco, ENSGI-INPG, 46, Avenue Félix Vialet, 38031 Grenoble, Cedex, France
E-mail: gayon@gilco.inpg.fr

Y. Dallery
Laboratoire Génie Industriel, École Centrale Paris, Grande Voie des Vignes, 92295 Chatenay-Malabry, Cedex, France
E-mail: yves.dallery@lgi.ecp.fr

decisions are jointly optimized, Chen et al. (2004) and Gayon et al. (2004) exhibit rather modest benefits of DP strategies with respect to SP strategies (when the demand is stationary). If the decision-maker loses, in part, control over the replenishment decisions, we might expect higher benefits from DP. In this paper, our goal is to investigate the benefits of DP when the replenishment process is not totally controlled. To our knowledge, there is currently no academic work studying this issue.

The assumption of a partially uncontrolled replenishment process is relevant in several contexts. In a reverse logistics setting, there exists a return flow of products that is not controlled. The literature makes a distinction between products which are in a good state and those which are not. The former can be put again in inventory without prior repair operations (after possible testing and repackaging). For example, mail-order companies and electronic retailers allow their customers to return products within a certain amount of days. These products might be sold in another market, remanufactured, recycled, or disposed of. For a complete discussion on inventory control with return flows, we refer the reader to Fleischmann et al. (1997). Several authors have proposed inventory control models with return flows, but none of the models we have encountered consider the price as a decision. Another context where replenishment is not totally controlled is when a firm has a fixed delivery contract. For example, Cheung and Yuan (2003) studied an infinite horizon inventory model of a buyer with a periodic order commitment. The commitment required the buyer to purchase, at least, a certain fixed amount on a periodic basis. For a survey on the different types of order commitments, refer to Anupindi and Bassok (1999). Finally, in the context of very high production set-up costs, it might be profitable not to stop the production for a long time. This is the case in the glass industry where the production is not stopped for several years. In the short term, this can be seen as an uncontrolled production.

The literature dealing with the coordination of pricing and inventory decisions can be broadly classified into two categories: systems with nonrenewable and perishable capacities (known as yield management) and systems where the sellers have the capability to store inventory and to replenish it. Writings on yield management have been appearing for more than 30 years and have been notably applied to airline and hotel industries. For an overview of this class of problems, see the papers of Weatherford and Bodily (1992), McGill and Van Ryzin (1999), and Bitran and Caldentey (2003). Recent years have seen many retail and manufacturing firms exploring new pricing strategies, and there has been, at the same time, a growing amount of literature on the coordination of pricing and inventory decisions. For a complete synthesis of the domain, we refer the reader to three recent surveys: those by Yano and Gilbert (2002), Elmaghraby and Keskinocak (2003), and Chan et al. (2004).

In this paper, we consider the problem of jointly coordinating price and replenishment decisions in a make-to-stock queue with partially controlled production. In a make-to-stock queue, the replenishment lead-times are load-dependent and are affected by the number of outstanding orders due to limitations in production capacity. The framework of our model is inspired by Li (1988) and Gayon et al. (2004). Li considered an $M/M/1$ make-to-stock queue with lost sales where the demand rate is a continuous function of the price. In addition to the production control, the manager sets the price dynamically over time to maximize average profit over an infinite horizon. Because of the memory-less property of the system, the price depends only on the stock level. Li proved the optimality of a base-stock policy. Moreover, he showed that the optimal sequence of prices is a nonincreasing function of the inventory level. Gayon et al. considered an extension of this model where the demand fluctuates over time. The rate of the Poisson process depends on the price, but also on the economical environment, which evolves according to a continuous-time Markov chain. The environment state is known by the decision-maker and the optimal policy consists of base-stock levels, one for each environment state. For a given environment state, the sequence of optimal prices is nonincreasing. Gayon et al. also carried out a numerical study to assess the potential benefits of DP in different situations. They showed that the potential benefits of DP over SP are rather modest when the demand rate of the Poisson process does not depend on the environment state. DP appears to be more effective when the demand rate depends on the environment state.

The rather modest benefits of DP observed by Gayon et al. are in part due to the totally controlled production process. When the production is not totally controlled, we might expect higher benefits. We take the results of Li and Gayon et al. as a starting point in investigating the effect of partially controlled replenishment on the benefit of DP over SP. Our model differs from the one of Li in the following way: In addition to the controlled replenishment process, there is an uncontrolled flow of items arriving in the inventory, modeled by a Poisson process. In the DP case, we extend the results of Li by showing that the optimal policy is of base-stock type and that the optimal prices are nonincreasing in the stock level. In the SP case, we extend the results of Gayon et al. by showing that the optimal policy is also of base-stock type. The numerical computation of the optimal SP strategy appears to be more difficult than the one of the optimal DP strategy. The SP strategy computation indeed requires many value iteration programs to run (one for each possible price) instead of one for the DP problem. To address this issue, we analytically compute, in a SP setting, the average profit for a given price and a given base-stock level. The problem then turns into a two-dimensional optimization. Furthermore, when the production is fully uncontrolled, we obtain explicit expressions for the optimal price and the optimal profit when the demand function is either linear or exponential. Based on previous results, we carry out a numerical study on the benefits of DP with respect to SP. Our main insight is that DP is potentially much more beneficial when the replenishment process is not totally controlled.

The rest of the paper is organized as follows: In Section 2, we present the formulation of the models and the problems. In Section 3, we characterize the structure of the optimal policy for the DP problem. Section 4 identifies, for the SP problem, the optimal policy and establishes a certain number of analytical results. Section 5 provides numerical results, which we use to derive insights. In Section 6, we offer a summary of main contributions.

2 Models formulation

In this section, we present the DP problem and the SP problem. We formulate the DP problem as a continuous-time review model where production and pricing decisions can be made at any point of time. The SP formulation is identical except that a unique price has to be chosen for the whole time horizon.

2.1 Dynamic pricing problem

Consider a supplier who partially controls his production. On the one hand, he has a single facility and can decide to produce or not. The processing time of this facility is exponentially distributed with rate μ_1 and the completed items are placed in a finished goods inventory. The unit variable production cost is c_1. On the other hand, there is an uncontrolled production process and the inventory receives items according to a Poisson process with rate μ_2, with unit variable production cost c_2. The stock level x belongs to $\mathbb{N}$, the set of nonnegative integers. The induced inventory holding cost is $h(x)$, and is convex in x.

Demand for items in stock arrive according to a Poisson process with rate $\lambda(p)$ depending on the posted price p. A demand that cannot be met from the stock is definitely lost. Partly following Gallego and van Ryzin (1994), we impose several assumptions on the demand function. First, we assume that $\lambda(p)$ is decreasing in p. Therefore, there is a one-to-one correspondence between prices and demand rates so that $\lambda(p)$ has an inverse denoted $p(\lambda)$. One can then alternatively view the rate λ as the decision variable, which is sometimes more convenient to work with from an analytical perspective. Second, we assume that the set of allowable demand rates is a bounded interval of the form $[0, \Lambda]$, where Λ is the maximum demand rate. Third, we assume the revenue rate $r(\lambda) = \lambda p(\lambda)$ is a continuous and strictly concave function. Concavity of $r(\lambda)$ stems from the standard economic assumption that the marginal revenue is decreasing in output. Finally, to ensure the stability of the stock level, we assume that the uncontrolled production rate is smaller than the maximum possible value of the demand rate, that is, $\mu_2 < \Lambda$.

We summarize now the previous notations:

x	Stock level
p	Posted price
μ_1	Controlled production rate
μ_2	Uncontrolled production rate
c_1	Unit controlled production cost
c_2	Unit uncontrolled production cost
$h(x)$	Unit holding cost when the stock level is x
$\lambda(p)$	Demand rate when the price is p
Λ	Maximum demand rate
$p(\lambda)$	Inverse of $\lambda(p)$
$r(\lambda)$	Revenue rate

In particular, we will consider two classes of demand rate functions frequently used in the pricing literature and which satisfy the above conditions. Let a and b be two positive real numbers with $\mu_2 < a$. We then define the linear demand function, and its associated revenue rate, by

$$\lambda^{lin}(p) = a - bp\,,\ p \in [0,\, 1/b],$$

$$r^{lin}(\lambda) = \frac{\lambda}{b}(a - \lambda),\ \lambda \in [0,\, a]$$

The second demand function we consider is the exponential

$$\lambda^{exp}(p) = ae^{-bp}\ ,\ p \geq 0,$$
$$r^{exp}(\lambda) = -\frac{\lambda}{b}\ln\left(\frac{\lambda}{a}\right),\ \lambda \in [0,\, a]$$

The problem is to decide, at any time, whether to produce or not and to choose a price p, or, equivalently, a demand rate λ, to maximize the average profit over an infinite horizon. We can formulate this problem as a Markov decision process. We define the state x as the inventory level. We let $v_d(x)$ be the relative value function of being in state x and g_d be the optimal average profit. Let $\beta = \Lambda + \mu_1 + \mu_2$. Using uniformization (Serfozo 1979), we can transform the continuous-time Markov decision process into an equivalent discrete time Markov decision process, where the following optimality equations are:

$$v_d(x) + \frac{g_d}{\beta} = \frac{1}{\beta}\{-h(x) + \mu_1 T_0 v_d(x) + T_1 v_d(x) + \mu_2[v_d(x+1) + c_2]\}$$

with

$$T_0 v_d(x) = \max[v_d(x), v_d(x+1) + c_1],$$
$$T_1 v_d(x) = \begin{cases} \max_\lambda\{r(\lambda) + \lambda v_d(x-1) + (\Lambda - \lambda)v_d(x)\} & \text{if } x > 0 \\ \Lambda v_d(x) & \text{if } x = 0 \end{cases}$$

The operator T_0 corresponds to the production decision, while T_1 corresponds to the arrival rate decision, or, equivalently, the price decision. Notice that the maximum in T_1 is well defined due to the assumptions on the revenue rate, $r(\lambda)$, which is strictly concave and defined on a bounded interval. The quantity μ_2 $(v_d(x+1) + c_2)$ corresponds to the uncontrolled arrival of an item in the inventory. Finally, we define the operator T, such that

$$v_d(x) + \frac{g_d}{\beta} = Tv_d(x)$$

and the operator Δ such that $\Delta v(x) = v(x+1) - v(x)$.

2.2 Static pricing problem

The setting of the SP problem is similar, except that the price cannot be changed over time. Let us consider first the problem with a given price p (or, equivalently, a given demand rate λ). The problem is then to find the optimal production policy maximizing the average profit over an infinite horizon. We can formulate again this problem as a Markov decision process. When the demand rate is λ, we let $v_s(x, \lambda)$

be the relative value function of being in state x and let $g_s(\lambda)$ denote the optimal average profit. Then, we can write the optimality equations:

$$v_s(x,\lambda)+\frac{g_s(\lambda)}{\beta}=\frac{1}{\beta}\{-h(x)+\mu_1 T_0 v_d(x,\lambda)+\tilde{T}_1 v_d(x,\lambda)+\mu_2[v_s(x+1,\lambda)+c_2]\}$$

where $\tilde{T}_0 = T_0$, and

$$\tilde{T}_1 v_s(x)=\begin{cases} r(\lambda)+\lambda v_s(x-1)+(\Lambda-\lambda)v_s(x) & \text{if}\, x>0 \\ \Lambda v_s(x) & \text{if}\, x=0\end{cases}$$

The second problem that the decision-maker faces is that of choosing the rate λ and is formulated as

$$\max_{\lambda} g_s(\lambda)$$

3 Characterization of the optimal policy for the dynamic pricing problem

To characterize the optimal policy, we first show that the relative value function $v_d(x)$ is concave in x. To that end, let us prove that the operator T preserves concavity. Assume that v is concave or, equivalently, that Δv is nonincreasing in x. Concavity is preserved under operator T_0 (Koole 1998), and $T_0 v$ is concave. Concavity is also preserved under operator T_1, and $T_1 v$ is concave. This last result is mentioned by Koole (1998), and a detailed proof is provided by Gayon et al. (2004). By assumption, $-h$ is also concave. Finally, the operator Tv, as a nonnegative linear combination of concave functions, is also concave. By the principle of value iteration, we finally obtain that v_d is also concave, and we obtain the following lemma:

Lemma 1 *The relative value function $v_d(x)$ is concave in x.*

As a result of Lemma 1, the difference $\Delta v_d(x) = v_d(x+1) - v_d(x)$ is nonincreasing in x, and there exists a level z_d such that:

$$\begin{cases} v_d(x) < v_d(x+1)+c_1 & \text{if } x < z_d \\ v_d(x) \geq v_d(x+1)+c_1 & \text{if } x \geq z_d\end{cases}$$

Relating these two equations with the production operator T_0, we conclude that it is optimal to produce when $x < z_d$ and to idle production otherwise, and we obtain

Property 1 *The optimal policy is of base-stock type: there exists a base-stock level z_d, such that it is optimal to produce if the stock level is smaller than z_d and to idle production otherwise.*

We also obtain results about the optimal prices, which are summarized in Property 2.

Property 2 *The sequence of optimal prices is unique and is nonincreasing in the stock level.*

Proof *Let p_x be the optimal price when the stock level is x. For $x > 0$, p_x is a maximizer of the following function:*

$$f_x(\lambda) = r(\lambda) - \lambda \Delta v_d(x-1)$$

$f_x(\lambda)$ is strictly concave in λ because $r(\lambda)$ is also strictly concave in λ. Consequently, there exists a unique maximizer, λ_x, on the interval $[0, \Lambda]$, which corresponds to a unique optimal price p_x because there is a one-to-one correspondence between a demand rate and a price.

We prove now, by contradiction, that the sequence of prices is nonincreasing in x. Let $x < y$ and denote by λ_x and λ_y the optimal rates in state x and y, respectively. Assume now that $\lambda_x > \lambda_y$. We can rewrite $f_y(\lambda_y)$ in the following way:

$$f_y(\lambda_y) = f_x(\lambda_y) + \lambda_y[\Delta v_d(x-1) - \Delta v_d(y-1)] + \Lambda(v(y) - v(x))$$

We have $f_x(\lambda_y) < f_x(\lambda_x)$ according to the uniqueness of the maximizer λ_y. Furthermore, from Lemma 1, v_d is concave, and thus, $\Delta v_d(x-1) - \Delta v_d(y-1) \geq 0$. As a result, we have

$$f_y(\lambda_y) < f_x(\lambda_x) + \lambda_x[\Delta v_d(x-1) - \Delta v_d(y-1)] + \Lambda[v_d(y) - v_d(x)] = f_y(\lambda_x)$$

Finally, we obtain that $f_y(\lambda_y) < f_y(\lambda_x)$, which is contradictory because λ_y is supposed to be the maximizer of f_y. Therefore, the assumption $\lambda_x > \lambda_y$ is false and we conclude that $\lambda_x \leq \lambda_y$, which is equivalent to $p_x \geq p_y$.

4 Characterization of the optimal policy for the static pricing problem

For the SP problem, proving the concavity of $v_s(x, \lambda)$ in x is similar to the DP problem, and it implies Property 3.

Property 3 *The optimal policy of the SP problem is of base-stock type.*

The optimal policy for the SP problem can then be specified by a base-stock level, z_s, and a static price, p_s. Notice that we do not ensure the uniqueness of z_s and p_s. The computation of z_s and p_s might be computationally difficult. It actually requires, for each possible price, a dynamic program to be run. In the following subsections, we seek to compute the optimal paramaters z_s and p_s. Explicit expressions for the average profit for a given base-stock z and price p are obtained by analyzing a continuous-time Markov chain. Furthermore, when the system is fully uncontrolled, we obtain analytical expressions for the optimal price and profit when the demand function is either linear or exponential.

4.1 Partially controlled production

Consider a policy with a base-stock level z and a price p, or, equivalently, a demand rate λ. Let $g_s(z, p)$ be the associated average profit. Then, the stock level evolves according to a continuous-time Markov chain with state space $\mathbb{N}$ and transition rates q_{ij}, from state i to state j, given by

$$\begin{cases} q_{,i+1} = \mu_1 + \mu_2 & \text{if } 0 \leq i \leq z-1 \\ q_{i,i+1} = \mu_2 & \text{if } i \geq z \\ q_{,i-1} = \lambda & \text{if } i \geq 1 \\ q_{ij} = 0 & \text{otherwise} \end{cases}$$

We introduce the additional notations: $\rho_1 = \lambda/(\mu_1 + \mu_2)$ and $\rho_2 = \mu_2/\lambda$. By assumption, we have $\rho_2 < 1$.

If π_i denotes the stationary probability to be in state i, then we have the following relations:

$$\begin{cases} \pi_i = \rho_1^{-i} \pi_0 & \text{if } 0 \leq i \leq z \\ \pi_i = \rho_1^{-z} \rho_2^{i-z} \pi_0 & \text{if } i \geq z \end{cases}$$

Using the normalization condition, $\sum_{i=0}^{\infty} \pi_i = 1$, we obtain

$$\pi_0 = \rho_1^z \left(\frac{1 - \rho_1^{z+1}}{1 - \rho_1} + \frac{\rho_2}{1 - \rho_2} \right)^{-1}$$

The average profit, $g_s(z, p)$, can be expressed as the difference between the average revenue and the average holding cost:

$$g_s(z, p) = \lambda p(1 - \pi_0) - h \sum_{i=0}^{\infty} i\pi_i$$

We finally obtain an analytical expression for the average profit:

$$g_s(z, p) = \lambda p(1 - \pi_0) - h\pi_0 \left(\frac{\rho_1^{z+1} - \rho_1 - \rho_1 z + z}{1 - \rho_1^2} + \frac{\rho_2}{1 - \rho_2} \right)$$

This result will be used in the numerical section to compute the optimal parameters z_s and p_s maximizing the two-dimensional function $g_s(z, p)$.

4.2 Fully uncontrolled production

For a fully uncontrolled system ($\mu_1 = 0$), the stationary probability is the one of a $M/M/1$ queue with an arrival rate μ_2 and a service rate λ. We have, therefore, in this case: $\pi_i = \rho_2^i(1-\rho_2)$. There is no base-stock level (as $\mu_1 = 0$), and we can simplify the notation $g_s(z,p)$ to $g_s(p)$. We obtain that

$$g_s(p) = \mu_2\left(p - \frac{h}{\lambda - \mu_2}\right)$$

By analyzing the derivative of $g_s(p)$, we compute explicitly the optimal prices p_s^{lin} and p_s^{exp} for a linear and exponential demand function, respectively,

$$p_s^{lin} = \max\left\{0, \frac{a - \mu_2 - \sqrt{ah}}{b}\right\}$$

$$p_s^{exp}\left\{0, \frac{1}{b}\ln\left[\frac{a(h + 2\mu_2 - \sqrt{h^2 + 4\mu_2 h}\ )}{2\mu_2^2}\right]\right\}$$

In the linear case, the optimal average profit, g_s^{lin}, turns out to be very simple

$$g_s^{lin} = \begin{cases} \dfrac{-\mu_2 h}{a - \mu_2} & \text{if } p_s^{lin} = 0 \\ \dfrac{\mu_2}{a}\left(a - \mu_2 - (a+b)\sqrt{\dfrac{h}{a}}\right) & \text{otherwise} \end{cases}$$

5 Numerical study

In this section, we want to examine the benefits of implementing a DP policy when the production is not totally controlled.

5.1 Computational procedure

In all the numerical experiments, we consider a linear demand function of the form $\lambda(p) = a - bp$. Moreover, we take $a = b = 1$ without losing any generality, because it is similar to setting a monetary unit and a time unit. We assume for all problem instances that the holding costs are linear, of the form $h(x) = hx$, and that there is no production cost ($c_1 = c_2 = 0$).

For a given problem, let g_d be the optimal average profit using a DP policy and g_s the optimal average profit using a SP policy. Define *PG*, the relative profit gain for using a DP policy (instead of an SP policy), by

$$PG = \frac{g_d - g_s}{g_s}$$

We compute g_d by solving the dynamic programs corresponding to each problem instance using the value iteration method. The value iteration algorithm is terminated only when a five-digit accuracy is achieved. The size of the state space is increased until the average profit is no longer sensitive to the truncation level.

The procedure to evaluate g_s and the associated optimal policy is based on the analytical results of Section 4. When the production is totally uncontrolled ($\mu_1 = 0$), we compute directly the optimal price and profit because we have an exact formula. When the production is partially uncontrolled ($\mu_1 > 0$), we discretize the interval of prices $[0, 1]$ with increments of 0.001. For a given price p, we search the optimal base-stock $s_s(p)$ and optimal average profit $g_s(p)$. We repeat this operation for all prices in the discretized set, and the maximum average profit obtained is set as g_s and the corresponding optimal price as p_s.

5.2 The benefit of dynamic pricing

Gayon et al. (2004) investigate the benefit of DP with respect to SP in the case of a totally controlled production ($\mu_2 = 0$). Their main observation is that when the static price is chosen effectively, the potential impact of DP is limited. In particular, they show that the maximum profit gain is 3.81% for a linear demand function, whatever the problem parameters are. We will study now the impact of an uncontrolled production on the benefit of DP.

Let $\gamma = \mu_2/(\mu_1 + \mu_2)$ be the proportion of production capacity that is uncontrolled. Figure 1 represents the optimal SP and DP average profit vs γ when the maximum production rate $(\mu_1 + \mu_2)$ is fixed. As expected, the DP profit is always greater than the SP profit. We also observe that both curves decrease in γ. This is rather intuitive because, when γ increases, the decision-maker keeps the same production capacity but loses control over the production process. Moreover, the absolute profit gain of DP policy vs SP policy, $(g_d - g_s)$, increases in γ. This trend is confirmed on Fig. 2, which represents the relative profit gain, *PG*, vs γ for different values of the holding cost h. The relative profit also appears to be increasing in γ. For example, when $h = 0.01$, the profit gain is 15% when $\gamma = 1$ instead of 1.8% when $\gamma = 0$. The main insight provided by Figs. 1 and 2 is that when the production is less and less controlled, a flexibility on prices is a more and more efficient tool.

Let us now study the influence of γ on the optimal dynamic prices and static prices. Figure 3 presents the optimal dynamic prices $p(x)$ vs the stock level x for

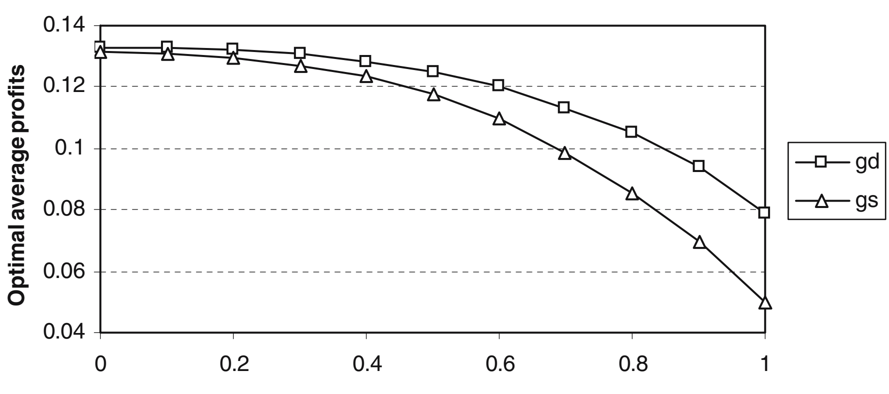

Fig. 1 The effect of γ on the optimal static and dynamic profits ($\mu_1 + \mu_2 = 0.5$, $h = 0.04$)

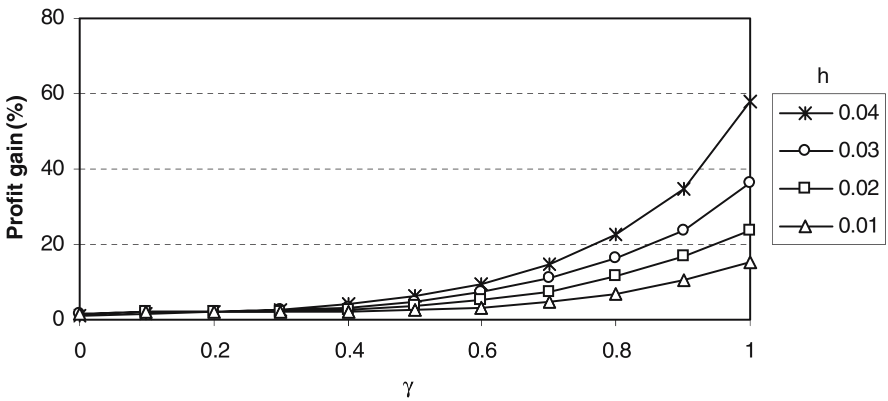

Fig. 2 The effect of γ on the relative profit gain ($\mu_1 + \mu_2 = 0.5$)

different values of γ. As proven in Section 3, we observe that the optimal dynamic prices are nonincreasing in the stock level. For each stock level x, we also notice that the optimal dynamic price $p(x)$ decreases in γ. When γ is increasing, the production is less and less controlled, and, to limit excess inventory, it is better to diminish the prices to increase the demand rate. We also notice that the influence of γ, on the prices, increases with the stock level. When the stock level increases, the incentive to decrease the price to get rid of the inventory surplus is also increasing. Finally, when the inventory level is high enough, the optimal dynamic price is 0 to limit excess inventory as much as possible. Figure 4 presents the optimal static prices vs γ for different values of the maximum production rate $(\mu_1 + \mu_2)$. As in the DP case, the optimal static price is decreasing in γ. The explanation is the same: when γ is increasing, it is better to decrease the price to raise the demand and thus to limit the inventory holding costs. We also observe that the optimal static price is decreasing in the maximum production rate $(\mu_1 + \mu_2)$. When the production capacity is decreasing, it becomes more and more difficult to satisfy the demand and it is better to raise the price to limit stockouts and to increase the profit at the same time.

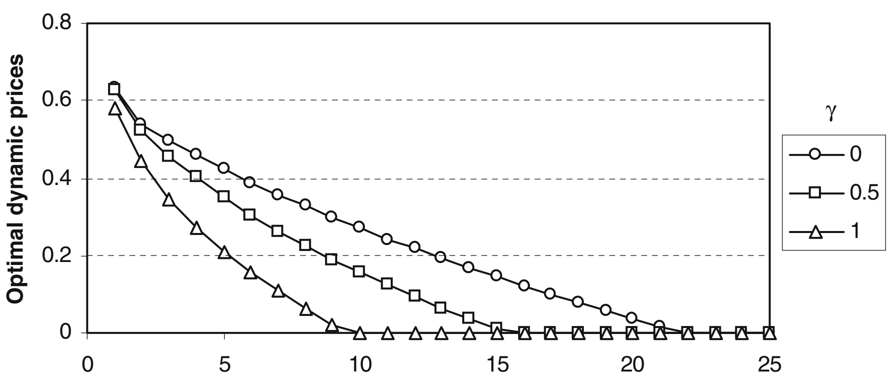

Fig. 3 Optimal dynamic prices vs stock level ($\mu_1 + \mu_2 = 0.5$, $h = 0.04$)

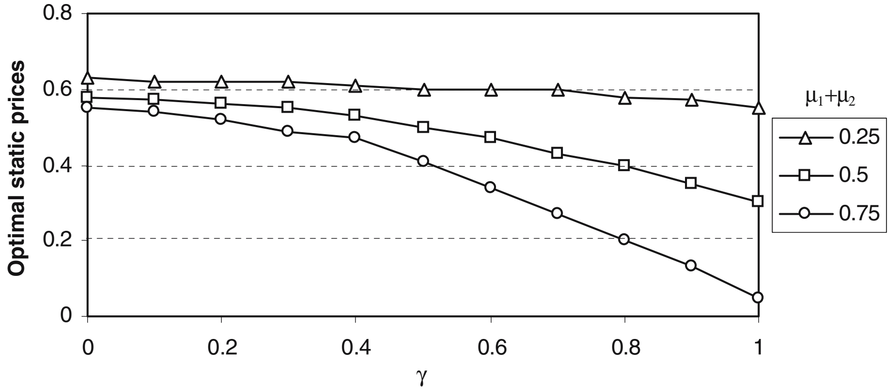

Fig. 4 The effect of γ on optimal static prices ($h = 0.04$)

6 Conclusion

In this paper, we have analyzed a SP problem and a DP problem in a make-to-stock queue with partially uncontrolled production. We have characterized the optimal policy for both problems. The optimal production policy is of base-stock type in both cases. In the DP case, the optimal pricing policy consists of a list-price with one price per stock level, the price being nonincreasing in the stock level. In the SP problem, we have obtained analytical results on the average profit. Furthermore, when the production is totally uncontrolled, we obtain analytical expressions of the optimal prices and profits for two classes of demand function (linear and exponential). Based on these results, we have carried out a numerical study on the potential benefits of DP with respect to SP when the demand function is linear. As observed by Chan et al. (2004) and Gayon et al. (2004), when the production is totally controlled, multiple price changes result in limited profit improvement over a single price. However, when some of the production is not controlled, we have shown that the potential benefits of DP are much more important. Throughout the paper, we have assumed that the uncontrolled flow of items was stochastic. This assumption was not realistic to model an order commitment, and it would be of interest to study the deterministic case. However, we believe that the insights we obtained should remain unchanged.

References

Anupindi R, Bassok Y (1999) Supply contracts with quantity commitments and stochastic demand. In: Tayur S, Ganeshan R, Magazine M (eds) Quantitative models for supply chain management. Kluwer, Dordrecht

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. Manuf Serv Oper Manag 5(3):203–229

Chan LMA, Shen ZJM, Simchi-Levi D, Swann J (2004) Coordination of pricing and inventory decisions: a survey and classification. In: Simchi-Levi D, Wu SD, Shen ZJM (eds) Handbook of quantitative supply chain analysis: modeling in the e-business. Kluwer, Dordrecht, pp 335–392

Chen H, Wu OQ, Yao DD (2004) Optimal pricing and replenishment in a single product inventory system. Working paper, IEOR Department, Columbia University

Cheung KL, Yuan XM (2003) An infinite horizon inventory model with periodic order commitment. Eur J Oper Res 146(1):52–66

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: research overview, current practices and future directions. Manage Sci 49 (10):1287–1309

Fleischmann M, Bloemhof-Ruwaard JM, Dekker R, van der Laan E, van Nunen JAEE, Van Wassenhove LN (1997) Quantitative models for reverse logistics: a review. Eur J Oper Res 103:1–17

Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. Manage Sci 40(8):999–1020

Gayon JP, Talay I, Karaesmen F, Örmeci EL (2004) Dynamic pricing and replenishment in a production/inventory system with markov-modulated demand. Working paper, Department of Industrial Engineering, Koç University

Koole G (1998) Structural results for the control of queueing systems using event-based dynamic programming. Queueing Syst 30:323–339

Li L (1988) A stochastic theory of the firm. Math Oper Res 13(3):447–466

McGill JI, Van Ryzin GJ (1999) Revenue management: research overview and prospects. Transp Sci 33(2):233–256

Serfozo RF (1979) An equivalence between continuous and discrete time markov decision processes. Oper Res 27(3):616–620

Weatherford LR, Bodily S (1992) A taxonomy and research overview of perishable-asset revenue management: yield management, overbooking, and pricing. Oper Res 40(5):831–844

Yano CA, Gilbert SM (2002) Coordinated pricing and production/procurement decisions: a review. In: Chakravarty A, Eliashberg J (eds) Managing business interfaces: marketing, engineering and manufacturing perspectives. Kluwer, Boston