

OPTIMAL POLICIES UNDER DIFFERENT PRICING STRATEGIES IN A PRODUCTION SYSTEM WITH MARKOV-MODULATED DEMAND

Production-Pricing Control

E. L. Örmeci* J. P. Gayon† I. Talay-Değirmenci ‡
F. Karaesmen §

Abstract

We study the effects of different pricing strategies available to a continuous review inventory system with capacitated supply, which operates in a fluctuating environment. The system has a single server with exponential processing time. The inventory holding cost is nondecreasing and convex in the inventory level, the production cost is linear with no set-up cost. The potential customer demand is generated by a Markov-Modulated (environment-dependent) Poisson process, while the actual demand rate depends on the offered price. For such systems, there are three possible pricing strategies: Static pricing, where only one price is used at all times, environment-dependent pricing, where the price changes with the environment, and dynamic pricing, where price depends on both the current environment and the stock level. The objective is to find an optimal replenishment policy under each of these strategies. This paper presents some structural properties of optimal replenishment policies, and a numerical study which compares the performances of these three pricing strategies.

Keywords: Inventory control, pricing, Markov Decision processes

*lormeci@ku.edu.tr, Koç University, İstanbul, TURKEY

†INPG, Grenoble, FRANCE

‡Duke University, Durham, USA

§Koç University, İstanbul, TURKEY

1 Introduction

During the last few decades, it is realized that the joint optimization of pricing and replenishment decisions results in significant improvements on the firm's profit (see e.g., [1]). The inspiring results obtained on this topic so far encouraged us to analyse an inventory pricing and replenishment problem. On the other hand, the environmental factors affect the density of the demand distribution unpredictably, and the focus in the recent studies of inventory control has been shifting to model the impact of fluctuating demand on the optimal replenishment policy. Hence, we consider an inventory system operating in a fluctuating demand environment, which controls the prices as well as the replenishment. As a result, our work stands at the junction of three main-stream research topics, inventory control, price control and the effects of environmental changes on the control policies.

We study a continuous review, infinite horizon inventory pricing and replenishment problem with capacitated supply. The system has a single server with exponential processing time. There is no set up cost, and the production cost is linear. The inventory holding cost, on the other hand, is nondecreasing and convex in the inventory level. In order to model a fluctuating environment, we assume that the potential customer demand is generated by a Markov-Modulated (environment-dependent) Poisson process. Moreover, the actual demand depends on the price offered at the time of the transaction, such that the actual demand rate decreases as the price increases. For a system operating in this environment, there are three possible pricing strategies: Static pricing, where only one price is used at all times, environment-dependent pricing, where the price is allowed to change with the environment, and dynamic pricing, where price depends on both the current environment and the stock level. In this paper, we use a Markov Decision Process framework to model this system as a make-to-stock queue operating under each of these strategies. Using this framework, we show that optimal replenishment policies are of environment-dependent base-stock level policies for these pricing strategies. We also compare the performances of these three strategies by an extensive numerical study.

The objective of inventory management is to reduce the losses caused by the mismatches that arise between supply and demand processes. With the advances in computers and communication technology, the role of inventory management has changed from cost control to value creation. Therefore, the issues inventory management studies now include both the traditional decisions such as inventory replenishment and the strategic decisions made by the firm such as pricing. In fact, there has been an increasing amount of

research on pricing with inventory/production considerations, see the excellent review papers [3], [4], and [5].

The widely known results in inventory control model the randomness of demand by using a random component with a well-known density in the definition of the demand process. However, the focus in the recent studies of inventory control has been shifting to model the impact of fluctuating demand on the optimal replenishment policy (see [6] and [7] among others). In particular, changes in the demand distribution might be caused by economic factors such as interest rates, or they might be caused by the changes in business environment conditions such as progress in the product-life-cycle or the consequences of rivals' actions on the market. The model we present below considers the effect of external factors on the demand distribution.

This paper is organized as follows: In the next section we introduce the models for the pricing strategies described above. Section 3 will present structural results for an optimal replenishment policy for each of the pricing strategies. In section 4, we will present our numerical results, which compare the performances of the three policies and provide insights, and point out possible directions of future research.

2 Model formulation

In this section we present a make-to-stock production system with three different pricing strategies: (1) the static pricing problem where a unique price has to be chosen for the whole time horizon regardless of the environment and the inventory level, (2) the environment-dependent pricing where the price can be changed over time depending on the environment, but not on the inventory level (3) the dynamic pricing where the price can be changed over time depending on both the inventory level and the environment. The production system should also decide on the replenishment of the items.

Consider a supplier who produces a single part at a single facility. The processing time is exponentially distributed with mean $1/\mu$ and the completed items are placed in a finished goods inventory. The unit variable production cost is c and the stock level is $X(t)$ at time t , where $X(t) \in \mathbb{IN} = \{0, 1, \dots\}$. We denote by h the induced inventory holding cost per unit time and h is assumed to be a convex function of the stock level.

The environment state evolves according to a continuous-time Markov Chain with state space $E = \{1, \dots, n\}$ and transition rates q_{ej} from state e to state $j \neq e$. We assume that this Markov chain is recurrent to avoid technicalities. For all environment states, the set of allowable prices \mathcal{P} is

identical. The customers arrive according to a Markov Modulated Poisson process (MMPP) with rate Λ_e when the state of the exogenous environment is e . We assume that the potential demand rates are bounded, i.e., $\max\{\Lambda_e\} < \infty$; a reasonable assumption which will be necessary to uniformize the Markov decision process. The customers decide to buy an item according to the posted price p , so that the actual demand rate in environment e is $\lambda_e(p)$ when a price of p is offered. Obviously, the actual demand rate is bounded by the potential demand rate so that $\lambda_e(p) \leq \Lambda_e$ for all e and for all p . We note that the domain of the prices, \mathcal{P} , may be either discrete or continuous. When \mathcal{P} is continuous, it is assumed to be a compact subset of the set of non-negative real numbers \mathbb{R}^+ .

For a fixed environment state e , we impose several mild assumptions on the demand function. First, we assume that $\lambda_e(p)$ is decreasing in p and we denote by $p_e(\lambda)$ its inverse. One can then alternatively view the rate λ as the decision variable, which is more convenient to work with from an analytical perspective. Thus the set of allowable demand rates is $\mathcal{L}_e = \lambda_e(\mathcal{P})$ in environment state e . Second, the revenue rate $r_e(\lambda) = \lambda p_e(\lambda)$ is bounded. Finally we assume that p_e is a continuous function of λ when the set of prices \mathcal{P} is continuous.

At any time, the decision maker has to decide whether to produce or not. The decision maker may also choose a price $p \in \mathcal{P}$, or equivalently a demand rate $\lambda \in \mathcal{L}_e$ at certain times specified by the pricing strategies described above. If we are in search of optimal replenishment policies for the pricing strategies described above, then the optimal policy is known to belong to the class of stationary Markovian policies, see [8]. Therefore the current state of the system is exhaustively described by the state variable (x, e) with x the stock level and e the environment state and (x, e) belongs to the state space $\mathbb{N} \times E$. Then, for dynamic pricing strategy $p(x, e)$ is the price of the item when the system operates in environment e with x units of item on inventory, for environment-dependent pricing policy $p(e)$ is the price chosen a priori for environment e so that $p(e)$ is charged whenever the system enters environment e regardless of the current inventory level, and p_s is the static price to be always offered regardless of the environment and the inventory level.

2.1 Optimal static pricing strategy

In static pricing, the decision maker has to choose a unique price in \mathcal{P} for the whole horizon. The static pricing problem can be viewed in two steps. First, we determine the optimal production policy, which depends on both

the environment and inventory level, for a given static price, p . Hence, let $v_s^\pi(x, e; p)$ be the expected total discounted reward when the replenishment control policy π is followed with $p_s = p$ over an infinite horizon starting from the state (x, e) . If we denote by α the discount rate, by $N(t)$ the number of demands accepted up to time t , and by $W(t)$ the number of items produced up to time t when the posted price is always p and the replenishment policy π is followed, then:

$$v_s^\pi(x, e; p) = E_{x,e}^\pi \left[\int_0^{+\infty} e^{-\alpha t} p dN(t) - \int_0^{+\infty} e^{-\alpha t} h(X(t)) dt - \int_0^{+\infty} e^{-\alpha t} c dW(t) \right],$$

where $X(t)$ is the inventory level at time t , as defined previously. We seek to find the policy π^* which maximizes $v_s^\pi(x, e; p)$ for a given price p . Let v_s^* be the optimal value function associated to π^* , so that:

$$v_s^*(x, e; p) = \max_{\pi} \{v_s^\pi(x, e; p)\}.$$

Now we can formulate this problem as a Markov Decision Process (MDP): Without loss of generality, we can rescale the time by taking $\mu + \sum \Lambda_e + \sum_e \sum_{j \neq e} q_{ej} + \alpha = 1$. After a uniformization, v_s^* satisfies the following optimality equations:

$$\begin{aligned} v_s^*(x, e; p) &= -h(x) + \mu T_0 v_s^*(x, e; p) + \lambda_e(p) v_s^*(x, e; p) + \sum_{j \neq e} q_{ej} v_s^*(x, j; p) \\ &\quad + \left(\sum_i \Lambda_i - \lambda_e(p) + \sum_{i \neq e} \sum_{j \neq i} q_{ij} \right) v_s^*(x, e; p), \end{aligned}$$

where the operator T_0 for any function $f(x, e)$ is defined as

$$T_0 f(x, e) = \max\{f(x, e), f(x+1, e) - c\}. \quad (1)$$

Hence, the operator T_0 corresponds to the production decision. We define $a_s(x, e)$ as the optimal replenishment decision in state (x, e) such that $a_s(x, e) = 1$ if it is optimal to produce the item, and $a_s(x, e) = 0$ otherwise. We also define also the operator T_s such that $v_s^* = T_s v_s^*$. Therefore, whenever a price p is given, we can find an optimal replenishment policy by solving an MDP.

The second step is to find the optimal price p_s^* in the set of prices \mathcal{P} , where there might exist potentially several optimal prices. Since we assume that the exogenous environment state follows a recurrent Markov chain, we choose the price p_s^* such that $p_s^* = \operatorname{argmax}\{v_s^*(0, 1; p) : p\}$ without loss of generality.

2.2 Optimal environment-dependent pricing strategy

The problem of environment-dependent pricing strategy is similar to the static pricing as it is also solved in two steps.

In the first step, the optimal production policy, π^* , is identified for a given set of prices $\bar{p}_{ed} = (p(1), \dots, p(N))$. Let v_{ed}^* be the optimal value function associated to π^* . Then:

$$v_{ed}^*(x, e; \bar{p}_{ed}) = \max_{\pi} \left\{ E_{x,e}^{\pi} \left[\int_0^{+\infty} e^{-\alpha t} p(E(t)) dN(t) - \int_0^{+\infty} e^{-\alpha t} h(X(t)) dt - \int_0^{+\infty} e^{-\alpha t} c dW(t), \right] \right\}$$

where $E(t)$ is the state of the exogeneous environment at time t , $p(E(t))$ is the posted price when the current environment is $E(t)$, and α , $X(t)$, $N(t)$ and $W(t)$ are defined as above. Optimal replenishment policy π^* can be determined by using uniformization as in the static pricing problem. Hence:

$$\begin{aligned} v_{ed}^*(x, e; \bar{p}_{ed}) &= -h(x) + \mu T_0 v_{ed}^*(x, e; \bar{p}_{ed}) \\ &\quad + \lambda_e(p) v_{ed}^*(x, e; \bar{p}_{ed}) + \sum_{j \neq e} q_{ej} v_{ed}^*(x, j; \bar{p}_{ed}) \\ &\quad + \left(\sum_i \Lambda_i - \lambda_e(p) + \sum_{i \neq e} \sum_{j \neq i} q_{ij} \right) v_{ed}^*(x, e; \bar{p}_{ed}), \end{aligned} \quad (2)$$

where the operator T_0 is defined as in (1). Now $a_{ed}(x, e)$ is the optimal replenishment decision in state (x, e) , so that $a_{ed}(x, e) = 1$ if it is optimal to produce the item, and $a_{ed}(x, e) = 0$ otherwise. We also define the operator T_{ed} such that $v_{ed}^* = T_{ed} v_{ed}^*$.

In the second step, an optimal price vector $\bar{p}_{ed}^* = (p^*(1), \dots, p^*(n))$ is chosen such that $\bar{p}_{ed}^* = \operatorname{argmax}\{v_{ed}^*(0, 1; \bar{p}_{ed}) : \bar{p}_{ed}\}$, without loss of generality due to the recurrent Markov chain governing the environment process.

2.3 Optimal dynamic pricing strategy

The system with dynamic pricing is an extension of Li (1988), who analyzes the same system operating in a stationary environment, to the one operating in a fluctuating environment. This problem is different from the static and environment-dependent pricing in the following way: Since both optimal replenishment and optimal pricing policies depend on the current inventory level as well as the environment, both policies are determined as a result of an MDP. We let $v_d^*(x, e)$ be the maximal expected total discounted reward when an optimal dynamic control policy π^* , which controls both the replenishment

decisions and prices, is followed over an infinite-horizon with initial state (x, e) . Then we have:

$$v_d^*(x, e) = \max_{\pi} \left\{ E_{x,e}^{\pi} \left[\int_0^{+\infty} e^{-\alpha t} p(X(t), E(t)) dN(t) - \int_0^{+\infty} e^{-\alpha t} h(X(t)) dt - \int_0^{+\infty} e^{-\alpha t} c dW(t) \right] \right\},$$

where $E(t)$, α , $X(t)$, $N(t)$ and $W(t)$ are defined as above. We can still use uniformization, so that v_d^* should satisfy the following optimality equations:

$$v_d^*(x, e) = -h(x) + \mu T_0 v_d^*(x, e) + T_e v_d^*(x, e) + \sum_{j \neq e} q_{ej} v_d^*(x, j) + \left(\sum_{i \neq e} \Lambda_i + \sum_{i \neq e} \sum_{j \neq i} q_{ij} \right) v_d^*(x, e),$$

where the operator T_0 is defined as in (1), and T_e is given by:

$$T_e v_d(x, e) = \max_{\lambda \in \mathcal{L}_e} g_{x,e}(\lambda),$$

and the function $g_{x,e}$ is defined for any λ in \mathcal{L}_e by:

$$g_{x,e}(\lambda) = \begin{cases} r_e(\lambda) + \lambda v_d(x-1, e) + (\Lambda_e - \lambda) v_d(x, e) & : \text{if } x > 0 \\ \Lambda_e v_d(x, e) & : \text{if } x = 0. \end{cases}$$

Therefore, the operator T_e corresponds to the arrival rate decision, or equivalently the price decision in environment e . Optimal replenishment decision in state (x, e) is denoted by $a_d(x, e)$, where $a_d(x, e) = 1$ if it is optimal to produce the item, and $a_d(x, e) = 0$ otherwise. Finally, we define the operator T_d such that $v_d^* = T_d v_d^*$.

2.4 Discussion on different pricing strategies

Before describing our results, we want to discuss the advantages and disadvantages of these three pricing strategies. Obviously, optimal dynamic pricing policies always generate more profit than optimal environment-dependent policies, which in turn generate more than optimal static policies. Now we turn to the “qualitative” effects of these policies: Static pricing represents the traditional pricing since the price remains fixed over time, regardless of the changes in the environment and in the stock level. This type of policies is easy to implement. In addition, consumers may prefer the transparency of a known price that is not subject to any changes. At the other extreme, we have dynamic pricing that leads to frequent price changes, since even

a change in the stock level may trigger a change in price. Therefore, dynamic pricing may create negative consumer reactions. Moreover, its implementation requires sophisticated information systems that can accurately track sales and inventory data in real time, and can be extremely difficult especially if price changes require a physical operation such as a label change. Environment-depending pricing, on the other hand, allows the price to change only with the environmental state. Hence, the associated system changes the prices, but not as frequently as the one with the dynamic pricing does. As a result, this policy is in between static and dynamic policies regarding to the practical problems and difficulties they bring.

3 Structural results

The MDP formulations of the replenishment problems given in Section 2 provide not only a tool to numerically solve the corresponding problem but also an effective methodology to establish certain structural properties of optimal policies. In particular, we will use these formulations to prove that there exists an optimal environment-dependent base-stock policy under each of the pricing strategies. We first present the definition of an environment-dependent base-stock policy:

Definition 1 *A replenishment policy which operates in a fluctuating demand environment, as described in Section 2, is an environment-dependent base-stock policy, if it always produces the item in environment e whenever the current inventory level is below a fixed number $S(e)$, i.e., $x < S(e)$, and it never produces in environment e whenever $x \geq S(e)$, where the numbers $\{S(1), \dots, S(N)\}$ are called the base stock levels with $S(e) \in \mathbb{N}$.*

Now we argue that each of the pricing strategies yields to an optimal environment-dependent base-stock policy, if the corresponding value function is concave. Hence assume that $v_\pi^*(x, e)$ is concave with respect to x for each environment e , i.e.:

$$v_\pi^*(x+1, e) - v_\pi^*(x, e) \leq v_\pi^*(x, e) - v_\pi^*(x-1, e). \quad (3)$$

If it is optimal to replenish in a state (x, e) , from equation (1) we have:

$$v_\pi^*(x, e) \leq v_\pi^*(x+1, e) - c, \iff c \leq v_\pi^*(x+1, e) - v_\pi^*(x, e).$$

Then, by concavity, we have:

$$c \leq v_\pi^*(x+1, e) - v_\pi^*(x, e) \leq v_\pi^*(x, e) - v_\pi^*(x-1, e),$$

implying that it has to be optimal to replenish in state $(x - 1, e)$ as well. Therefore, whenever an optimal policy replenishes in a state (x, e) , it replenishes in all states (k, e) with $k \leq x$. We can, similarly, show that if an optimal policy does not replenish in a state (x, e) , it continues not to replenish in all states (k, e) with $k \geq x$. These two statements together imply the existence of an optimal base-stock level in each environment e , $S_\pi^*(e)$:

$$S_\pi^*(e) = \min\{x : a_\pi(x, e) = 0\},$$

where $a_\pi(x, e)$ is the optimal replenishment decision in state (x, e) with policy π . Now we show that the corresponding value functions are concave for all pricing strategies we describe above:

Lemma 1 *For a fixed environment e , for all $\pi = s, ed, d$: If $v_\pi^*(x, e)$ is concave with respect to x , then $T_\pi v_\pi^*$ is also concave with respect to x .*

Proof. $\pi = s$ is a special case of $\pi = ed$ if we set $p(e) = p$ for all e , and we refer to [2] for the proof of $\pi = d$. Hence, we show the statement for $\pi = ed$. In this proof we denote $v_{ed}^*(x, e; \bar{p}_e)$ by $v_{ed}^*(x, e)$. Assume that v_{ed}^* is concave in x for each environment e .

Now we consider each term in equation (2) separately. By assumption $-h$ is concave. To prove that T_0 preserves concavity, we need to show:

$$\delta = T_0 v_{ed}^*(x + 1, e) - 2T_0 v_{ed}^*(x, e) + T_0 v_{ed}^*(x - 1, e) \leq 0$$

Now let $a' = a_{ed}(x + 1, e)$ and $a'' = a_{ed}(x - 1, e)$. By our observation above, there exists an optimal environment-dependent base-stock policy, so that $a' \leq a''$. Since $v_{ed}^*(x + a', e) \leq T_0 v_{ed}^*(x, e)$ and $v_{ed}^*(x + a'', e) \leq T_0 v_{ed}^*(x, e)$:

$$\begin{aligned} \delta &\leq v_{ed}^*(x + 1 + a', e) - ca' - v_{ed}^*(x + a', e) + ca' - \\ &\quad v_{ed}^*(x + a'', e) + ca'' + v_{ed}^*(x - 1 + a'', e) - ca'' \leq 0. \end{aligned}$$

If $a' = a''$, then the statement is true by the concavity of v_{ed}^* . If $a' = 0$ and $a'' = 1$, then the term in the second inequality is exactly 0. All other terms in (2) are concave by concavity of v_{ed}^* . Thus, $T_{ed} v_{ed}^*$ is concave in x for an environment e , whenever v_{ed}^* is concave.

Now the above argument immediately implies the existence of optimal environment-dependent base-stock policies:

Theorem 1 *For all pricing strategies $\pi = s, ed, d$: The optimal replenishment policy is an environment-dependent base stock policy.*

Optimality of environment-dependent base stock policies shows that information about the environment in which a firm operates is crucial.

ϵ	$\max\{PG_{d,s}\}$	$\max\{PG_{ed,s}\}$	$\max\{PG_{d,ed}\}$
0.3	5.68%	2.25%	3.36%
0.6	10.19%	7.58%	2.87%
0.8	13.04%	11.25%	3.23%

Table 1: Maximum profit gain for different demand variability.

μ	0.11	0.21	0.31	0.41	0.51	0.61	0.71
$PG_{d,s}$	12.50%	8.67%	6.80%	5.72%	4.75%	3.80%	3.12%
$PG_{ed,s}$	10.90%	6.70%	4.16%	2.60%	1.47%	0.7%	0.5%
$PG_{d,ed}$	1.45%	1.85%	2.53%	3.04%	3.23%	3.07%	2.66%

Table 2: Profit gain of pricing policies for different service rates with $\epsilon = 0.8$.

4 Numerical results

In our model formulation, the system is controlled directly by the demand rate, defined as a function of the offered price. In this section we explicitly refer to the prices. We consider a linear demand rate function, which is frequently used in the pricing literature. Let p be the price offered. Then we define the linear demand function, and its associated revenue rate by:

$$\lambda_e(p) = \Lambda_e(1 - ap), p \in [0, 1/a], \quad (4)$$

where a is a positive real number.

ϵ	$S_s^*(L)$	$S_s^*(H)$	$S_{ed}^*(L)$	$S_{ed}^*(H)$	$S_d^*(L)$	$S_d^*(H)$
0.3	6	11	7	9	12	20
0.6	4	14	5	10	7	22
0.8	2	13	3	10	3	23

Table 3: The optimal base stock levels for different ϵ with $\mu = 0.11$.

For a given problem, let g_π^* be the optimal average profit using policy π , where discount rate is set to 0, i.e., $\alpha = 0$. We define the relative Profit Gain for using policy π instead of policy π' , $PG_{\pi,\pi'}$, by

$$PG_{\pi,\pi'} = \frac{g_\pi^* - g_{\pi'}^*}{g_{\pi'}^*}. \quad (5)$$

ϵ	p_s^*	$p_{ed}^*(L)$	$p_{ed}^*(H)$	$\bar{p}_d^*(L)$	$\underline{p}_d^*(L)$	$\bar{p}_d^*(H)$	$\underline{p}_d^*(H)$
0.3	0.78	0.74	0.82	0.82	0.42	0.87	0.51
0.6	0.75	0.65	0.84	0.75	0.33	0.88	0.51
0.8	0.78	0.57	0.84	0.65	0.19	0.88	0.51

Table 4: The optimal prices for different ϵ with $\mu = 0.11$, where $\bar{p}_d^*(e) = \max\{p_d^*(x, e)\}$, and $\underline{p}_d^*(e) = \min\{p_d^*(x, e)\}$.

As we know that $g_s \leq g_{ed} \leq g_d$, we will consider $PG_{d,ed}$, $PG_{d,s}$ and $PG_{ed,s}$.

We consider a system which operates in two environments, with low demand rate (L) and with high demand rate (H). The demand rates in these environments are $\Lambda_L = 1 - \epsilon$ and $\Lambda_H = 1 + \epsilon$. The factors that affect optimal policies are the ratios λ/μ and h/p , so we vary the service rate μ and the holding cost h , where we set $a = 1$, $c = 0$, and the average demand rate as 1. Moreover, here we only report $h = 0.01$ and $q_{LH} = q_{HL} = q = 0.01$, although we experimented with different h and q as well as asymmetric transitions rates. In the whole numerical study, we restrict our attention to the recurrent states of the Markov chain generated by an optimal policy.

As ϵ increases, the demand variability increases. We observe that optimal gain for each pricing policy decreases with ϵ . The profit gain of π_{ed} and π_d with respect to π_s also increases with ϵ (see Table 1), which shows the ability of these policies to adjust the highly uncertain environments. For small ϵ , on the other hand, $PG_{d,s} < 6\%$, suggesting that optimal static policy performs good enough with mild uncertainty. From Table 2, we observe that policy π_s performs the worst with capacitated supply ($\mu < 0.4$) and volatile demand with respect to π_{ed} and π_d . Optimal static prices are closer to the optimal environment-dependent prices in environment H , rather than those in environment L (see Table 4). Hence, the demand fluctuation hurts not only the firm by decreasing its average gain, but also the customers due to high prices, when static pricing strategy is followed.

We see that policy π_{ed} performs very closely to policy π_d with $\max\{PG_{d,ed}\} < 3.5\%$ (see Table 1). In fact, it brings most of the benefit of π_d , compare $PG_{d,s}$ with $PG_{ed,s}$ in Table 2. Moreover, policy π_{ed} has the advantage of lower inventory levels (see Table 3) and of smaller price differences (see Table 4). Hence, we can conclude that it is better to use π_{ed} , since it brings most of the benefit of π_d , while causing less reaction on the customer side with less variability in prices, and requiring a reasonable storage space with less variability in the stock levels.

Optimal pricing and replenishment policies may have further monotonicities: If we order the environment states with respect to the potential demand rates, i.e., $\Lambda_e \leq \Lambda_{e+1}$ for $e = 1, \dots, n - 1$, then, under certain conditions, we expect to have monotone base stock levels, i.e., $S_\pi^*(e) \leq S_\pi^*(e + 1)$ for all pricing strategies $\pi = s, ed, d$. The optimal environment-dependent prices should also be ordered with the potential demand rates. Finally, we expect that the effective demand rates are monotone with the potential demand rates, under certain conditions. Our future work will focus on finding sufficient conditions to observe these kind of monotonicities.

References

- [1] H. Chen, O. Wu, D.D. Yao. Optimal Pricing and Replenishment in a Single-Product Inventory System. Working Paper, Columbia University, Dept. of Industrial Engineering and Operations Research, 2004.
- [2] J.-P. Gayon, I. Talay-Değirmenci, F. Karaesmen, L. Ormeci. Dynamic Pricing and Replenishment in a Production/Inventory System with Markov-Modulated Demand. Working Paper, Koç University, Department of Industrial Engineering, 2004.
- [3] W. Elmaghraby, P. Keskinocak. Dynamic pricing in the presence of inventory considerations: research overview, current practices and future directions. *Management Science*. 49-10:1287-1309, 2003.
- [4] C.A. Yano, S.M. Gilbert. Coordinated pricing and production/procurement decisions: a review. In *Managing Business Interfaces*. Kluwer Academic Publishers, 2003.
- [5] L. M. A. Chan, Z. J. Max Shen, D. Simchi-Levi, J. L. Swann. Coordination of pricing and inventory decisions: a survey and classification. In *Handbook of quantitative supply chain analysis: modeling in the e-business era*. Kluwer, 2004.
- [6] S. Özekici, M. Parlar. Inventory models with unreliable suppliers in a random environment. *Annals of Operations Research*. 91:123-136, 1999.
- [7] F. Chen, J. S. Song. Optimal Policies for Multiechelon Inventory Problems with Markov-Modulated Demand. *Oper. Research*. 49-2: 226-234.
- [8] M. Puterman. *Markov Decision Processes*. John Wiley and Sons Inc, New York, 1994.