2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

Optimal Multiperiod Inventory Decisions with Partially Observed Markovian Supply Information

Haifeng Wang and Houmin Yan

Abstract— **This paper considers a multiperiod newsvendor problem with partially observed supply capacity information that evolves as a Markovian process. The supply capacity is fully observed by the buyer when the capacity is smaller than the buyer's ordering quantity. Otherwise, the buyer knows that the current-period supply capacity is greater than its ordering quantity. Based on these two observations, the buyer updates the future supply-capacity forecasting accordingly. With a dynamic programming formulation, we prove the existence of a unique optimal ordering policy.**

I. INTRODUCTION

Inventory control is one of the most studied topics in supply-chain management. In the classic newsvendor supply-chain model, the buyer satisfies its customers' requests by placing orders with its suppliers at the beginning of each selling season. The fundamental task of the buyer is to maintain a high customer-service rate without holding too many leftover inventories. Because neither the demand nor the supply capacities are always known to the buyer when orders are placed, both demand and supply capacity information is fundamental for any inventory-control decisions. Most studies in the supply-chain management literature assume that the supply capacity is either known to the buyer, or unlimited. Few of studies consider supply-capacity information uncertainty and its influence to the inventory policy. Yet the supply capacity is always unknown or only partially observable in practice. In such an environment, most well-known results may not hold. This paper aims to explore the structural properties of an inventory system with partially observed supply capacity information and to find if an optimal policy exists that can solve the above problem.

A. The Importance of Modeling the Partially Observed Capacity

The study of inventory systems with partially observed supply capacity information is important in many real-life cases. We shall introduce some cases in a supply chain in which capacities can only be partially observed.

Suppliers' capacity uncertainties: Unexpected machine breakdowns and the resulting unexpected repairs lead to lowered production capacity. Moreover, uncertain repairing time affects the availability of certain facilities: the repair is not planned, or even planned, repair time is uncertain. These stochastic events lead to an uncertain supply capacity. Further, product quality can also be uncertain. When product quality is low or the rate of defects is high, the actual available products are unknown. Since these uncertainties cannot be predicted, the supply capacity is not observed directly.

Buyers' competition: A supplier serves multiple customers. When capacity is tight, the supplier reserves its capacity for individual customers. It is known in practice as the capacity allocation, or rationing. The most popular capacity allocation rules include the proportional capacity allocation policy (Lee, Padmanabhan and Whang [13]) and the turn-and-earn policy (Cachon and Lariviere [6]). Orders from buyers vary from time to time and are unknown to the other buyers. Hence, the reserved capacity for an individual buyer is uncertain, even when the supply allocation rules are known. As a result, a buyer doesn't know the capacity volume that is reserved for it.

Multiple sources of supply: The buyer also orders from multiple suppliers. Each supplier has its own reserved capacity for the buyer. Hence, the buyer can only estimate the capacity reserved for it, based on observed signals, such as the received ordering quantity. As a result, the available capacity is partially observed.

B. Summary of the Paper

This paper studies the optimal inventory-replenishment policy at the beginning of each period that can meet the stochastic demand with partially observed supply-capacity information. We allow the buyer to fulfill the out-of-capacity orders by paying an extra cost from other sources, such as a spot market. The demand is realized at the end of each period. Leftover inventories are salvaged while the unfulfilled demand incurs a penalty cost.

We consider a multiperiod problem in which the capacity observation of the current period will influence the capacity distribution and the value function of the next period. The available supply capacity may change from period to period, evolving as a Markovian process. The capacity for the current period is NOT fully observed by the buyer at the time when the buyer places its orders. When the order quantity is greater than the supply capacity, the supplier provides its maximum capacity to the buyer. Only at this time is the capacity fully observed by the buyer. The partially observed capacity information limits the buyer's capability to forecast

This work is supported in part by RGC Competitive Earmarked Research Grants CUHK4239/03E, CUHK4167/04E, and CUHK4177/06E.

Haifeng Wang is with the Center for Intelligent and Networked Systems, Department of Automation, Tsinghua University, Beijing 100084, China (email:wanghf99@mails.tsinghua.edu.cn).

Houmin Yan is with the Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, Shatin, N.T., Hong Kong. He is also a non-resident member of Center for Intelligent and Networked Systems, Department of Automation, Tsinghua University, Beijing 100084, China (email:yan@se.cuhk.edu.hk).

its needs and inventory optimization. In this paper, we prove the existence of an optimal ordering policy for a general Markovian capacity process.

C. Literature Review and Its Relation to Our Model

There is extensive literature on inventory decisions with given capacity constraints, including studies by Federgruen and Zipkin [7], Gallego and Scheller-Wolf [8], Gallego and Toktay [9], and Gallien and Wein [10]. In addition, there are papers dealing with uncertain capacity constraints. Khang and Fujiwara [12] address a discrete-time inventory model in which the maximum amount of supplies from which instantaneous replenishment orders can be placed is a random variable. They also characterize optimal ordering policies. Wang and Gerchak [17] extend the model to include variable capacity and random yields in periodic inventory systems.

The research on inventory decisions with partially observed information is quite recent. Bensoussan, Cakanyildirim and Sethi [4] study a multiperiod newsvendor model with partially observed demands. By assuming that the leftover inventories are salvaged and unsatisfied demands are lost in each period, they decouple the periods from the view point of the Beyesian demand update, and prove the existence of an optimal ordering policy. Aviv and Pazgal [1] introduce the partially observed Markovian decision process. In their setting, to maximize its revenue, a buyer faces a dynamic pricing problem of selling a given stock during a finite selling season. Bensoussan, Cakanyildirim and Sethi [2] and [3] study inventory systems with partially observed inventory levels. Lu, Song and Zhu [14] and [15] study the newsvendor model with censored demand data.

To the best of our knowledge, the problem of inventory decisions with a partially observed Markovian capacity remains open. We model an inventory system with a partially observed Markovian supply capacity. Similar to Bensoussan, Cakanyildirim and Sethi [4], we use the *unnormalized probability* to linearize the highly nonlinear dynamic programming equation. We solve the inventory-management problem with the partial supply-capacity information, while Bensoussan, Cakanyildirim and Sethi [4] solve the inventory-management problem with the partial demand information. The methodology in Bensoussan, Cakanyildirim and Sethi [4] is extended in our model to prove the existence of an optimal inventory policy.

D. Plan of the Paper

In the next section, we formulate the problem. In our model, the evolution of a partially observed Markovian capacity is characterized, and the unnormalized probability is used to linearize the dynamic programming equation. In Section 3, we prove the existence of an optimal ordering policy. Finally, we conclude the paper in Section 4.

II. Problem Formulation and Dynamic Programming Equations

A. Sequence of Events and Cost Structure

The sequence of events is as follows: at the beginning of each period, the buyer informs the supplier of its intended order quantity q_t; then, the supplier checks the available capacity reserved for this buyer and tries its best to satisfy the requested purchase. Denote the reserved supply capacity as Q_t. Assume that $\{Q_t\}$ is a Markovian process with a transition probability $p(Q_{t+1} \mid Q_t)$, which is exogenously decided as an input. Note that many commonly used forecasting methods satisfy the Markovian property, such as the last-value forecasting method, averaging forecasting method, and exponential-smoothing forecasting method (see Hillier and Lieberman [11]). The buyer gets 1) $a_t = q_t \wedge Q_t$ at the normal unit purchasing cost c and 2) $q_t - a_t$, if the amount of unsatisfied orders is greater than 0, at a higher unit purchasing cost $c + e$, where $e \geq 0$ indicates that the extra order can be more expensive. Note that the extra order may come from the same supplier by asking for an expensive and emergent order, or come from some other supply resource, such as a spot market. At the end of this selling season, all the unsatisfied demand is lost. The shortage cost at the unit cost of b and the salvaged value at the unit revenue of h are recorded.

The following is the list of notations used in this paper:

q_t :	order quantity from the buyer;
Q_t :	available capacity of the supplier;
α :	discount factor;
D_t :	demand of period t;
$p(Q_{t+1} \mid Q_t)$:	transition probability of a Markovian process $\{Q_t\}$;
$Pr\{A\}$:	probability of the event A;
c :	normal unit purchasing cost;
a_t :	amount of order charged at normal purchasing cost c;
e :	additional unit purchasing cost for the extra unit when $q_t > Q_t$;
h :	unit salvage value;
b :	unit shortage cost;
$F(\cdot)$:	cumulative demand distribution function;
$\pi_t(\cdot)$:	probability density function of the supply capacity Q_t;
$V_t(\cdot)$:	value function.

B. Dynamic Forecast Updates for the Supply Capacity

In each period, the buyer observes a_t, which is the amount of ordering that is charged at the normal purchasing cost c. Then, the two scenarios are as follows.

1) When the order quantity is greater than the supply capacity, i.e. $q_t \geq Q_t$, the supplier cannot satisfy the request. Then the supplier exhausts all of its capacity to produce products, that is, $a_t = Q_t$. In this scenario, the supply capacity is revealed through the observed signal a_t. The density probability function of the capacity for the next period can be obtained as follows:

$$\pi_{t+1}(Q_{t+1}) = p(Q_{t+1} \mid a_t). \tag{1}$$

2) On the other hand, the order quantity is less than the supply capacity, and the supplier satisfies the request completely, that is, $q_t \leq Q_t$ and $a_t = q_t$. In this

scenario, the buyer knows that the capacity is greater than its order quantity, that is, $Q_t \geq a_t$, but does not know the exact value of the current-period supply capacity. We denote this case as the partially observed supply capacity. Then

$$\begin{aligned}
\pi_{t+1}(Q_{t+1}) &= Pr\{Q_{t+1} \mid Q_t \geq a_t\} \\
&= \frac{Pr\{Q_{t+1}, Q_t \geq a_t\}}{Pr\{Q_t \geq a_t\}} \\
&= \frac{\int_{a_t}^{+\infty} \pi_t(\xi) p(Q_{t+1} \mid \xi) d\xi}{\int_{a_t}^{+\infty} \pi_t(\xi) d\xi}, \quad (2)
\end{aligned}$$

where the second equality is due to Bayesian equations, and the third equality is due to the conditional probability equation, which is $Pr\{Q_{t+1}, Q_t \geq a_t\} = \int_{a_t}^{+\infty} Pr\{Q_{t+1} | Q_t = \xi\} \pi_t(\xi) d\xi$.

Hence, the supply capacity Q_t is partially observed. This approach is similar to that of Bensoussan, Cakanyildirim and Sethi [4] in which the partially observed demand signal is characterized.

C. Dynamic Programming Equations

With the sequence of events described above, it is possible for us to derive one-period cost function as $L(D_t, q_t) + e(q_t - Q_t)^+$, where

$$\begin{aligned}
L(x, y) &= \begin{cases} cy - h(y - x), & \text{if } x \leq y \\ cy + b(x - y), & \text{if } y \leq x, \end{cases} \\
&= \begin{cases} cy - h(y - x), & \text{if } x \leq y \\ bx + (c - b)y, & \text{if } y \leq x. \end{cases} \quad (3)
\end{aligned}$$

To eliminate trivial cases, it is reasonable to assume that $0 \leq h < c < b$ and $0 < e$. Then it is straightforward to see that

$$L(x, y) \leq \begin{cases} cy, & \text{if } x \leq y \\ bx, & \text{if } y \leq x \end{cases} \quad \text{if } y \geq 0. \quad (4)$$

Recall that $\pi_t(\cdot)$ is the probability density function of the supply capacity Q_t. Then the value function can be expressed as follows:

$$\begin{aligned}
& V_t(\pi_t(x)) \\
= & \inf_{q_t \geq 0} \{ \int_0^{q_t} E_{D_t}[L(D_t, q_t) + e(q_t - Q_t) \\
& + \alpha V_{t+1}(p(x \mid Q_t))] \pi_t(Q_t) dQ_t \\
& + \int_{q_t}^{+\infty} E_{D_t}[L(D_t, q_t) + \\
& \alpha V_{t+1}(\frac{\int_{q_t}^{+\infty} \pi_t(\xi) p(x \mid \xi) d\xi}{\int_{q_t}^{+\infty} \pi_t(\xi) d\xi})] \pi_t(Q_t) dQ_t \} \\
= & \inf_{q_t \geq 0} \{ E_{D_t}[L(D_t, q_t)] + \int_0^{q_t} e(q - Q_t) \pi_t(Q_t) dQ_t \\
& + \alpha \int_0^{q_t} V_{t+1}(p(x \mid Q_t)) \pi_t(Q_t) dQ_t \\
& + \alpha V_{t+1}(\frac{\int_{q_t}^{+\infty} \pi_t(\xi) p(x \mid \xi) d\xi}{\int_{q_t}^{+\infty} \pi_t(\xi) d\xi}) \int_{q_t}^{+\infty} \pi_t(Q_t) dQ_t \},
\end{aligned} \quad (5)$$

where the first term on the right-hand side is the cost when the supply capacity Q_t is smaller than the buyer's intended order quantity, and the second term of $+\int_{q_t}^{+\infty} E_{D_t}[L(D_t, q_t) + \alpha V_{t+1}(\frac{\int_{q_t}^{+\infty} \pi_t(\xi) p(\cdot|\xi) d\xi}{\int_{q_t}^{+\infty} \pi_t(\xi) d\xi})] \pi_t(Q_t) dQ_t$ is the cost when the supply capacity Q_t is larger than the buyer's intended order quantity.

The buyer chooses a set of order quantities $\{q_t\}$ to minimize its total cost. Ordering too much results in a high purchasing cost; on the other hand, ordering too less results in a high penalty cost and failures to observe the current-period supply capacity.

D. Unnormalized Probability

It is difficulty to directly solve the dynamic programming Equation (5) owing to its complex fractional parts. Both the denominator and the numerator include the decision variable q_t. In this part, we use a similar technique of a variable substitution as in Bensoussan, Cakanyildirim and Sethi [4] so that Equation (5) can be expressed in a simpler and solvable form.

From Equations (1) and (2), we know

$$\begin{aligned}
\pi_{t+1}(x) &= I_{a_t = q_t} \frac{\int_{q_t}^{+\infty} \pi_t(\xi) p(x \mid \xi) d\xi}{\int_{q_t}^{+\infty} \pi_t(\xi) d\xi} \\
&\quad + I_{a_t < q_t} p(x | a_t), \ t \geq 1,
\end{aligned}$$

where the indicator function $I_A = 1$ when event A occurs; otherwise, $I_A = 0$. Define recursively that

$$\begin{aligned}
\rho_{t+1}(x) &:= I_{a_t = q_t} \int_{q_t}^{+\infty} p(x|\xi) \rho_t(\xi) d\xi \\
&\quad + I_{a_t < q_t} p(x | a_t), \ t \geq 1; \\
\rho_1(x) &:= \pi_1(x).
\end{aligned}$$

Also let

$$\lambda_t := \int \rho_t(x) dx.$$

Then,

$$\lambda_{t+1} = I_{a_t = q_t} \int_{q_t}^{+\infty} \rho_t(\xi) d\xi + I_{a_t < q_t}, \ t \geq 1; \lambda_1 = 1.$$

We can check recursively that the following equation is true (it holds when $t = 1$; and by supposing the correctness for t, check that it holds for $t + 1$):

$$\rho_t(x) = \pi_t(x) \lambda_t.$$

Hence, we have

$$\pi_t(x) = \frac{\rho_t(x)}{\lambda_t} = \frac{\rho_t(x)}{\int \rho_t(x) dx}.$$

In what follows, we consider the infinite-horizon case. To simplify the notation, the sub-script is omitted.

We define $W(\rho)$ as

$$W(\rho) := \int \rho(x) dx \cdot V(\frac{\rho}{\int \rho(x) dx}). \quad (6)$$

It is obvious that $W(\cdot)$ can be obtained by $V(\cdot)$. By the definition of $W(\rho)$, we have

$$\begin{aligned} & W(\rho(\cdot)) \\ = & \int \rho(x)dx \cdot \inf_{q\geq 0} \{ \int L(D,q)f(D)dD \\ & + \int_0^q e(q-Q)\frac{\rho(Q)}{\int \rho(x)dx}dQ \\ & +\alpha \int_0^q V(p(\cdot \mid Q))\frac{\rho(Q)}{\int \rho(x)dx}dQ \\ & +\alpha V(\frac{\int_q^{+\infty} \rho(\xi)p(\cdot \mid \xi)d\xi}{\int_q^{+\infty}\rho(\xi)d\xi}) \int_q^{+\infty} \frac{\rho(Q)}{\int \rho(x)dx}dQ \Big\} \\ = & \inf_{q\geq 0} \Big\{ \int L(D,q)f(D)dD \int \rho(x)dx \\ & + \int_0^q e(q-Q)\rho(Q)dQ \\ & +\alpha \int_0^q W(p(\cdot \mid Q))\rho(Q)dQ \\ & +\alpha\, W(\int_q^{+\infty} \rho(\xi)p(\cdot \mid \xi)d\xi) \Big\} \\ := & \inf_{q\geq 0} G(q). \end{aligned} \tag{7}$$

Note that Equation (7) is a Bellman equation in ρ.

III. The Existence of an Optimal Feedback Solution

In this section, we first show the existence and uniqueness of the solution W for Equation (7). Then we prove the existence of an optimal feedback control, q^*.

A. Existence of a Unique Solution W

In this subsection we start with a definition of a function space, $\mathcal{B}$. Then we prove that if there is a solution to Equation (7), it must exist in the above-defined function space. As a result, the existence and uniqueness of the solution follows. Note that similar definitions of function spaces $\mathcal{H}$ and $\mathcal{B}$ are used in Bensoussan, Cakanyildirim and Sethi [4].

Define a function space $\mathcal{H}$ of function ρ

$$\mathcal{H} := \{\rho \in L^1(\mathcal{R}^+) : \int_0^\infty x|\rho(x)|dx < \infty\}, \tag{8}$$

where $L^1(\mathcal{R}^+)$ is the space of integrable functions whose domain is the set of nonnegative real numbers, and

$$\mathcal{H}^+ := \{\rho \in \mathcal{H} \mid \rho \geq 0\}, \tag{9}$$

with the norm

$$||\rho|| = \int_0^{+\infty} |\rho(x)|dx + \int_0^\infty x|\rho(x)|dx. \tag{10}$$

Also define the following space $\mathcal{B}$ of function ϕ

$$\mathcal{B} = \left\{\phi(\rho) : \mathcal{H}^+ \to \mathcal{R} \mid \sup_{x>0} \frac{|\phi(\rho)|}{||\rho||} < \infty \right\}, \tag{11}$$

with the norm

$$||\phi||_\mathcal{B} = \sup_{\rho\in\mathcal{H}^+} \frac{|\phi(\rho)|}{||\rho||}. \tag{12}$$

For technical convenience, we make the following assumption:

Assumption 3.1: Assume, for any $\rho \in \mathcal{H}^+$,

$$\int x \int p(x|\xi)\rho(\xi)d\xi dx \leq c_0 \int \xi\rho(\xi)d\xi, \quad \text{with } \alpha c_0 < 1. \tag{13}$$

This assumption is necessary to complete proofs of the following lemmas and the theorem. It is satisfied by a specific probability-transition function $p(\cdot \mid \cdot)$, such that

$$\alpha \int x \int p(x|\xi)\rho(\xi)d\xi dx \leq \alpha c_0 \int \xi\rho(\xi)d\xi < \int \xi\rho(\xi)d\xi,$$

that is,

$$\begin{aligned} & \text{the } \textit{discounted} \text{ expected demand of the} \\ & \text{next period forecasted by } p(\cdot|\cdot) \\ < \; & \text{the mean of the current-period demand.} \end{aligned} \tag{14}$$

As the discount factor is always strictly smaller than 1, this assumption holds when the expected demand of the next-period forecasted by $p(\cdot|\cdot)$ is equal to the mean of the current-period demand.

We need the following lemma. Its proof appears in Wang and Yan [16].

Lemma 3.1: If equation (7) has a solution W, the solution is in $\mathcal{B}$.

Now define the mapping $T(W)$ as,

$$\begin{aligned} T(W) \;\; := \;\; & \min_{q\geq 0} \{ \int L(D,q)f(D)dD \int \rho(x)dx \\ & + \int_0^q e(q-Q)\rho(Q)dQ \\ & +\alpha \int_0^q W(p(\cdot \mid Q))\rho(Q)dQ + \\ & \alpha W(\int_q^{+\infty} \rho(\xi)p(\cdot \mid \xi)d\xi) \} . \end{aligned} \tag{15}$$

Then we can obtain the following lemma. See Wang and Yan [16] for its proof.

Lemma 3.2: $||T(W) - T(\tilde{W})||_\mathcal{B} \leq \alpha \max\{1, c_0\}||W - \tilde{W}||_\mathcal{B}$.

Theorem 3.1: There exists a unique solution W for the dynamic programming Equation (7).

Proof: From Assumption 3.1, we know that $\alpha c_0 < 1$, such that $\alpha \max\{1, c_0\} < 1$. Hence, by Lemma 3.2, the mapping $T : \mathcal{B} \to B$ is a contraction mapping. Then by the *Contraction Mapping Theorem* [5], there exists a unique solution W such that $T(W) = W$. Then, we have proved the desired results. $\square$

B. The Existence of an Optimal Feedback Control

By Equation (3), we can easily obtain $L(x, y) \geq (c-h)y$. By substituting it into Equation (7), we can then prove that $W(\rho) \geq (c-h)q \int \rho(x)dx$. With Equations (52) and (56) in

Wang and Yan [16], we know that

$$
\begin{aligned}
&(c-h)q\int\rho(x)dx \le W(\rho)\\
\le\quad & \left(\frac{b\mu_D}{\mu_Q}+\alpha\max\{1,c_0\}||W||_{\mathcal{B}}\right)||\rho||\\
\le\quad & \left(\frac{b\mu_D}{\mu_Q}+\alpha\max\{1,c_0\}\frac{b\mu_D}{\mu_Q(1-\alpha\max\{1,c_0\})}\right)||\rho||.
\end{aligned}
$$

Note that $\int\rho(x)dx = \frac{\int x\rho(x)dx}{\int x\pi(x)dx} = \frac{\int x\rho(x)dx}{\mu_Q}$, then $\frac{||\rho||}{\int\rho(x)dx} = \frac{\int\rho(x)dx+\int x\rho(x)dx}{\int\rho(x)dx} = 1+\mu_Q$. Hence we obtain that

$$
0\le q\le\frac{b\mu_D(1+\mu_Q)}{\mu_Q(c-h)(1-\alpha\max\{1,c_0\})} := B. \tag{16}
$$

Now we define $W_B(\rho)$ as

$$
\begin{aligned}
& W_B(\rho)\\
= \quad & \min_{q\in[0,B]}\Big\{\int L(D,q)f(D)dD\int\rho(x)dx\\
& +\int_0^q e(q-Q)\rho(Q)dQ\\
& +\alpha\int_0^q W_B(p(\cdot\mid Q))\rho(Q)dQ\\
& +\alpha W_B(\int_q^{+\infty}\rho(\xi)p(\cdot\mid\xi)d\xi)\Big\}.
\end{aligned} \tag{17}
$$

Note that $W_B(\rho)$ can be shown to be a unique solution to Equation (17) as $W(\rho)$ is shown to be unique by Lemma 3.2. If $W_B(\rho)$ is continuous in ρ, we obtain the continuity of $G(q)$ (defined in Equation (7)) in q. After that, the existence of the minimizer q^* is established. See Wang and Yan [16] for proofs of the following lemma.

Lemma 3.3: For any $\rho,\tilde\rho\in\mathcal{H}^+$, we have

$$
|W_B(\rho)-W_B(\tilde\rho)| \quad\le\quad H_B||\rho-\tilde\rho||, \tag{18}
$$

where H_B is a constant that is independent of ρ.

Theorem 3.2: There exists an optimal feedback control, that is, the optimal order quantity.

Proof: Recall the definition of $G(q)$ in Equation (7), and we can write

$$
W_B(\rho)=\min_{q\in[0,B]}G(q), \tag{19}
$$

where

$$
\begin{aligned}
G(q)\quad = \quad & \int L(D,q)f(D)dD\int\rho(x)dx\\
& +\int_0^q e(q-Q)\rho(Q)dQ\\
& +\alpha\int_0^q W_B(p(\cdot\mid Q))\rho(Q)dQ\\
& +\alpha W_B(\int_q^{+\infty}\rho(\xi)p(\cdot\mid\xi)d\xi).
\end{aligned} \tag{20}
$$

The first, second, and third terms of Equation (20) are continuous in q. By Lemma 3.3, we know that $W(\rho)$ is continuous in ρ, which leads to the continuity of the fourth term of Equation (20) in q. Hence, $G(q)$ is continuous in q. Moreover, q belongs to a bounded and closed set $[0,B]$. Then, by *Weierstrass' Theorem* [5], there exists a minimum in the bounded and closed set to minimize such a continuous function. The minimizer q^* exists. □

IV. Conclusions and Future Research

We have studied a multiperiod newsvendor problem with partially observed supply-capacity information. The partial supply-capacity observations make dynamic programming a space of probability distributions. Un-normalized probability is used and the existent of a unique value function is proved. The existence of the optimal purchasing policy is also provided.

Our future research will focus on studying multi-period inventory decisions with a partially observed supply lead time. We also plan to consider the case in which leftover inventories can be carried over to the next period instead of lost sales. This case would be much more complicated as it would introduce the inventory level as an additional state variable.

Acknowledgment

We would like to thank two anonymous reviewers and an associate editor for their helpful comments. This work is supported in part by RGC Competitive Earmarked Research Grants CUHK4239/03E, CUHK4167/04E, and CUHK4177/06E.

References

[1] Y. Aviv and A. Pazgal, "A Partially Observed Markov Decision Process for Dynamic Pricing," *Management Science*, Vol. 51, No. 9, 2005, pp. 1400-1416.

[2] A. Bensoussan, M. Cakanyildirim, and S. P. Sethi, "On the Optimal Control of Partially Observed Inventory Systems," *To appear in the Comptes Rendus de l'Academié des Sciences*, 2005a.

[3] A. Bensoussan, M. Cakanyildirim, and S. P. Sethi, "Partially Observed Inventory Systems: The Case of Zero Balance Walk," *Working paper*, 2005b.

[4] A. Bensoussan, M. Cakanyildirim, and S. P. Sethi, "A Multiperiod Newsvendor Problem with Partially Observed Demand," *Working paper*, 2005c.

[5] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, Belmont, Massachusetts, 1995.

[6] G. P. Cachon and M. A. Lariviere, "Capacity Allocation Using Past Sales: When to Turn-and-Earn", *Management Science*, Vol. 45, No. 5, 1999, pp. 685-703.

[7] A. Federgruen and P. Zipkin, "An Inventory Model with Limited Production Capacity and Uncertain Demands I. The Average-Cost Criterion," *Mathematics of Operations Research*, Vol. 11, No. 2, 1986, pp. 193-215.

[8] G. Gallego and A. Scheller-Wolf, "Capacitated Inventory Problems with Fixed Order Costs: Some Optimal Policy Structure," *European Journal of Operational Research*, Vol. 126, No. 3, 2000, pp. 603-613.

[9] G. Gallego and L. B. Toktay, "All-or-Nothing Ordering under a Capacity Contraint," *Operations Research*, Vol. 52, No. 6, 2003, pp. 1001-1002.

[10] J. Gallien and L. M. Wein, "A Smart Market for Industrial Procurement with Capacity Constraints," *Management Science*, Vol. 51, No. 1, 2005, pp. 76-91.

[11] F. S. Hillier and G. J. Lieberman, "*Introduction to Operations Research*," Seventh Edition, McGraw-Hill, 2001.

[12] Do Ba Khang and O. Fujiwara, "Optimality of Myopic Ordering Policies for Inventory Model with Stochastic Supply," *Operations Research*, Vol. 48, No. 1, 2000, pp. 181-184.

[13] Hau L Lee, V. Padmanabhan, and S. Whang, "Information Distortion in a Supply Chain: the Bullwhip Effect," *Management Science*, Vol. 43, No. 4, 1997, pp. 546-558.

[14] X. Lu, J. S. Song, and K. Zhu, “The Censored Newsvendor and the Optimal Acquisition of Information,” *Operations Research, Forthcoming*.

[15] X. Lu, J. S. Song, and K. Zhu, “Analysis of Perishable-Inventory Systems with Censored Demand Data,” *Working paper*, 2006.

[16] H. Wang and H. Yan, “Optimal Inventory Decisions in the Multi-period News-vendor Problem with Partial-Observed Markovian Supply Capacities,” *Working paper*, the Chinese University of Hong Kong, 2006.

[17] Y. Wang and Y. Gerchak, “Periodic Review Production Models with Variable Capacity, Random Yield, and Uncertain Demand,” *Management Science*, Vol. 42, No. 1, 1996, pp. 130-137.