

TECHNICAL NOTE

(s, S) Optimality in Joint Inventory-Pricing Control: An Alternate Approach

Woonghee Tim Huh

Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027,
huh@ieor.columbia.edu

Ganesh Janakiraman

Stern School of Business, New York University, New York, New York 10012, gjanakir@stern.nyu.edu

We study a stationary, single-stage inventory system, under periodic review, with fixed ordering costs and multiple sales levers (such as pricing, advertising, etc.). We show the optimality of (s, S) -type policies in these settings under both the backordering and lost-sales assumptions. Our analysis is constructive and is based on a condition that we identify as being key to proving the (s, S) structure. This condition is entirely based on the single-period profit function and the demand model. Our optimality results complement the existing results in this area.

Subject classifications: inventory/production: uncertainty, stochastic; marketing: pricing, advertising.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received May 2006; revisions received October 2006, March 2007; accepted April 2007. Published online in *Articles in Advance* February 20, 2008.

1. Introduction

In this paper, we study the optimal control problem for a periodically reviewed single-stage inventory system with fixed ordering costs and stochastic, but controllable, demand. In each period, a manager makes inventory decisions as well as decisions that influence demand—for example, the price choice or the advertisement budget. We refer to these decision variables as *sales levers*.

For a single-stage system with inventory and pricing control and a fixed ordering cost, the optimality of (s, S) -type inventory policies¹ has been established in recent papers such as Chen and Simchi-Levi (2004a, b) and Chen et al. (2006). They generalize the classical approach of Scarf (1960), who showed (s, S) optimality when demands are exogenous. The only sales lever in these papers is the price. The proofs in these papers rely heavily on induction arguments for finite-horizon results and rather involved convergence arguments for infinite-horizon results. A majority of their results require joint-concavity of the single-period expected profit function with respect to inventory and price. We extend the optimality of the (s, S) -type structure for *stationary systems* by allowing a multidimensional sales lever, and a less restrictive single-period expected profit function (permitting, for example, quasi-concavity). Moreover, we provide constructive proofs, primarily using sample-path arguments, which are completely different from the earlier proofs in the joint inventory-pricing literature; this is our main contribution. Our ideas are a generalization of

some arguments used by Veinott (1966), who provided an alternate proof for (s, S) optimality when demands are exogenous.

In the process of developing this proof technique, we identify a sufficient condition to guarantee the optimality of (s, S) -type policies (Condition 1). This condition can be verified by simply studying a single-period problem; meanwhile, the existing results in the literature are proved by studying the value function from the dynamic program, which is typically more difficult to analyze. This is our second contribution. The sufficiency of Condition 1 for (s, S) optimality has been used to prove new results in other settings—namely, Song et al. (2006) for lost-sales models with multiplicative demands, and Huh and Janakiraman (2008) in the context of selling through auction channels.

1.1. Brief Literature Review

With complete backlogging, no fixed ordering costs, and the concave demand model,² Federgruen and Heching (1999) show the optimality of the base-stock list-price policy for the nonstationary finite-horizon model as well as the stationary infinite-horizon model. (A base-stock list-price policy is defined as follows: If the starting inventory level is less than some level y_t^* , then order up to y_t^* , and charge a fixed list price in period t ; otherwise, do not order and offer a price discount.) They assume that the single-period expected profit function is jointly concave in inventory (after ordering) and price.³

With complete backlogging, positive fixed cost, and the additive demand model, Chen and Simchi-Levi (2004a) show that the (s, S, p) policy is optimal in the finite horizon. (An (s, S, p) policy is defined as follows: Order nothing if inventory exceeds s ; order up to S otherwise. The price chosen depends on y , the inventory level after ordering, through a specified function $p(y)$.) They also present an example indicating that the (s, S, p) policy may not be optimal in the finite-horizon problem when demand is not additive. With the linear demand model, Chen and Simchi-Levi (2004b) show the optimality of the (s, S, p) policy in infinite-horizon models, both with the discounted-profit and the average-profit criteria. They require joint concavity of the single-period expected-profit function for their results on the finite-horizon problem as well as the infinite-horizon discounted-profit problem. They develop the notion of *symmetric K -concavity*, a generalization of K -concavity, and show that the profit function in the finite-horizon dynamic program possesses this property. This is the key step in their proof, which uses inductive arguments and is quite involved.

Feng and Chen (2004) use fractional programming to establish (s, S, p) -optimality for the average-profit criterion, and provide an algorithm for computing the optimal parameters. The continuous-review extensions have been studied by Feng and Chen (2003) and Chen and Simchi-Levi (2006).

With the lost-sales assumption, Chen et al. (2006) study a periodic-review finite-horizon problem with the additive demand model. They introduce some restrictions on the function relating expected demand and price as well as additional restrictions on the distribution of ϵ . With these assumptions, they demonstrate the optimality of the (s, S, p) policy.

Subsequent to the initial version of the current paper, a number of papers on the optimality of (s, S) -type policies have been written. Yin and Rajaram (2007) extend the results of Chen and Simchi-Levi (2004a, b) to Markovian environments. Song et al. (2006) prove (s, S, p) optimality with lost sales and the multiplicative demand model. Chao and Zhou (2006) provide algorithmic results for models with a Poisson demand process.

1.2. Problem Definition

Consider the following periodic-review system with a planning horizon of T periods (T can be finite or infinite) and a discount factor $\gamma \in (0, 1)$. All parameters of the system are assumed to be stationary. Periods are indexed forward. At the beginning of period t ($t \leq T$), we have x units of inventory. At this instant, an order can be placed to raise the inventory to some level y instantaneously (that is, there is no lead time). There is a fixed cost or a set-up cost, K , associated with ordering any strictly positive quantity. In addition, there is a set of levers to control the demand process. Examples of these levers are prices, salesforce incentives, and advertisements. We model the sales lever control

by a vector \mathbf{d} within some compact convex subset $\mathcal{D} \subseteq \mathfrak{R}^n$; the first component could denote the price discount and the second component could denote the advertisement expense and so on.⁴

After y and \mathbf{d} are chosen, the demand in period t is realized. It is a nonnegative random variable $D(\mathbf{d}, \epsilon)$, where ϵ is an exogenous random variable. (That is, the distribution of demand depends on the sales lever control, and ϵ is the source of randomness.) The net inventory at the end of the period is $y - D(\mathbf{d}, \epsilon)$, and a holding or a shortage cost is charged based on this quantity. We let $\pi(y, \mathbf{d})$ denote the expected profit in this period, excluding the set-up cost (that is, the total expected profit is $\pi(y, \mathbf{d}) - K \cdot 1[y > x]$); π includes sales revenue, the cost of choosing the sales lever, and the holding and shortage costs. The inventory level at the beginning of the next period is given by $\psi(y - D(\mathbf{d}, \epsilon))$, where either (a) $\psi(x) = x$ if excess demand is backordered, or (b) $\psi(x) = (x)^+$ if excess demand is lost.

The objective is to find a pair (y, \mathbf{d}) for every x and t that maximizes the expected discounted profits in periods $t, t + 1, \dots, T$.

A purchase cost, linear in the order size $y - x$, could also be present. However, a simple assumption about salvaging inventory left at the end of T periods can be used to transform the system into one in which this proportional cost is zero and the other cost parameters are suitably modified. Consequently, we will not consider this linear cost in our analysis. Our proofs for the finite-horizon results depend on this assumption. Although it is useful for notational simplicity, our infinite-horizon results and proofs do not require this assumption. Veinott (1966) makes a similar observation (see p. 1072).

We make the following assumption throughout the paper.

ASSUMPTION 1. (a) $\pi(y, \mathbf{d})$ is continuous, and

$$\max_{\mathbf{d}} \pi(y, \mathbf{d}) \rightarrow -\infty \quad \text{as } y \rightarrow \infty.$$

Furthermore, $\max_{(y, \mathbf{d})} \pi(y, \mathbf{d})$ exists.

(b) The demand model is stationary, that is, the sequence of ϵ s in time periods $\{1, 2, \dots, T\}$ is independent and identically distributed.

(c) $D(\mathbf{d}, \epsilon)$ is continuous in (\mathbf{d}, ϵ) . Moreover, for every ϵ , $D(\mathbf{d}, \epsilon)$ is componentwise monotone in \mathbf{d} .⁵

Let (y^*, \mathbf{d}^*) be a maximizer of $\pi(y, \mathbf{d})$, and let π^* be the maximum value. We remark that (y^*, \mathbf{d}^*) would be the solution chosen in a single-period problem in which there is no fixed cost and the starting inventory is lower than y^* .

The dynamic programming formulation for the finite-horizon problem is given by

$$U_t(x) = \max_{\{y: y \geq x\}} [V_t(y) - K \cdot 1[y > x]],$$

where

$$V_t(y) = \max_{\{\mathbf{d}: \mathbf{d} \in \mathcal{D}\}} [W_t(y, \mathbf{d})] \quad \text{and}$$

$$W_t(y, \mathbf{d}) = \begin{cases} \pi(y, \mathbf{d}) + \gamma \cdot E_\epsilon U_{t+1}(\psi(y - D(\mathbf{d}, \epsilon))) & \text{if } t < T, \\ \pi(y, \mathbf{d}) & \text{if } t = T. \end{cases}$$

(The subscript t in the above formulation denotes the period index.) It can be shown using standard arguments that the above maxima exist.⁶

An optimal policy for this finite-horizon problem specifies a feasible pair (y, \mathbf{d}) , i.e., $y \geq x$ and $\mathbf{d} \in \mathcal{D}$, for every x and t , that maximizes $W_t(y, \mathbf{d}) - K \cdot 1[y > x]$. The infinite-horizon optimal policy specifies a feasible pair (y, \mathbf{d}) , for every x , that maximizes $W(y, \mathbf{d}) - K \cdot 1[y > x]$, where $W(\cdot, \cdot)$ is the pointwise limit of $W_t(\cdot, \cdot)$ as T approaches infinity.⁷ In this paper, we study both finite-horizon problems and infinite-horizon discounted-profit problems.

The following definitions are based on the dynamic programming formulation and become useful later. Let $V(y) := \max\{W(y, \mathbf{d}) : \mathbf{d} \in \mathcal{D}\}$ and $Q(y) := \max\{\pi(y, \mathbf{d}) : \mathbf{d} \in \mathcal{D}\}$.

Next, we list the three kinds of demand models we use to capture the dependence of $D(\mathbf{d}, \epsilon)$ on \mathbf{d} .

- **Additive Demand Model.** For any $d \in \mathcal{D}$, where \mathcal{D} is a closed real interval, $D(d, \epsilon)$ is nonnegative for almost every ϵ , and can be expressed as $D(d, \epsilon) = d + \epsilon$. (In this model, we use d instead of \mathbf{d} because $\mathcal{D} \subseteq \mathbb{R}^1$.) Without loss of generality, we assume that $E[\epsilon] = 0$.

- **Linear Demand Model.** For any $\mathbf{d} \in \mathcal{D}$, $D(\mathbf{d}, \epsilon)$ is nonnegative for almost every ϵ , and can be expressed as $D(\mathbf{d}, \epsilon) = \alpha \cdot \mathbf{d} + \beta$, where $\epsilon = (\alpha, \beta)$, and $\alpha \in \mathbb{R}^n$ and $\beta \in \mathbb{R}$. Here, we assume that \mathcal{D} is a convex compact set. When $\beta = 0$, this is commonly known as the multiplicative demand model in the literature.

- **Concave Demand Model.** For any $\mathbf{d} \in \mathcal{D}$, $D(\mathbf{d}, \epsilon)$ is nonnegative, monotonic, and concave in \mathbf{d} for almost every ϵ . Again, we assume that \mathcal{D} is a convex compact set.

Note that the linear demand model is a generalization of the additive demand model, and the concave demand model is a generalization of the linear demand model.

We present our analysis in the following sequence. In §2, we present Conditions 1 and 2, sufficient conditions for (s, S) optimality. In §3, we prove the optimality of (s, S) -type policies when Condition 1 is satisfied. The validity of Condition 1 is shown in §4 for the backordering case, and the validity of Condition 2 is shown in §5 for the lost-sales case.

2. Single-Period Conditions for (s, S) Optimality

In this section, we present two sets of conditions on the expected single-period profit function, π , that lie at the core of our proofs. Although the conditions appear technical, we show in §§4 and 5 that common modeling assumptions found in the literature satisfy these conditions.

CONDITION 1. (a) $Q(y) = \max_{\mathbf{d} \in \mathcal{D}} \pi(y, \mathbf{d})$ is quasi-concave,⁸ and

(b) for any y^1 and y^2 satisfying $y^* \leq y^1 < y^2$ and \mathbf{d}^2 , there exists

$$\mathbf{d}^1 \in \{\mathbf{d} \mid \pi(y^1, \mathbf{d}) \geq \pi(y^2, \mathbf{d}^2)\} \quad (1)$$

such that for any ϵ ,

$$\psi(y^1 - D(\mathbf{d}^1, \epsilon)) \leq \max\{\psi(y^2 - D(\mathbf{d}^2, \epsilon)), y^*\}. \quad (2)$$

Recall that (y^*, \mathbf{d}^*) maximizes $\pi(y, \mathbf{d})$. Thus, y^* maximizes $Q(y)$. It follows from part (a) that the set in (1) is nonempty. An intuitive explanation of the condition follows. Part (a) indicates that the closer the starting inventory level (after ordering) is to y^* , the greater the single-period profit the system can generate. As a result, in a single-period problem without the fixed cost, it is optimal to order up to y^* if $y < y^*$ and to order nothing if $y \geq y^*$. By part (b), if $y^* \leq y^1 < y^2$, the system starting from y^1 is capable of ensuring that it will be at a “better” inventory level in the immediate following period than the system starting from y^2 ; in the next period, the starting inventory of the y^1 -system is closer to y^* than that of the y^2 -system, or ordering up to y^* is possible because the starting inventory of the y^1 -system is below y^* . We show that Condition 1 guarantees the optimality of (s, S) -type policies in infinite-horizon models.

THEOREM 1. *Suppose that Condition 1 holds. In the infinite-horizon discounted-profit model, there exist s and S such that if the inventory level, x , in the current period is at least s , it is optimal to not order, and otherwise to order up to S . That is, an (s, S) policy is optimal.⁹*

The proof of this result is presented in §3.2. It turns out that Condition 1 does not guarantee (s, S) optimality in finite-horizon models. That result requires a stronger condition presented below.

CONDITION 2. (a) Same as Condition 1(a).

(b) Same as Condition 1(b).

(c) If $y^* \geq y^1 > y^2$, there exists \mathbf{d}^1 satisfying (1) such that for any ϵ ,

$$\psi(y^1 - D(\mathbf{d}^1, \epsilon)) \geq \psi(y^2 - D(\mathbf{d}^2, \epsilon)). \quad (3)$$

We now state the (s, S) optimality result for finite-horizon models. The proof is presented in §3.3.

THEOREM 2. *Suppose that Condition 2 holds. Then, in a finite-horizon model, an (s_t, S_t) policy is optimal in period t .*

3. Proofs of (s, S) Optimality

3.1. Preliminary Results

Let y_t^* be the smallest maximizer of $V_t(y)$. We first establish that in period t , if the starting inventory level before

ordering is below y_t^o , and an order is placed, then it is optimal to order up to y_t^o (Proposition 1). Furthermore, if the starting inventory level is above y_t^o , then it is optimal not to order (Corollary 1). We use x_t to denote the starting inventory level before ordering, whereas y_t is the inventory level after ordering. The proposition below follows directly from the definition of y_t^o .

PROPOSITION 1. *Suppose that $x_t \leq y_t^o$. Then, $y_t = y_t^o$ maximizes $\{V_t(y_t) : y_t \geq x_t\}$. That is, if the fixed cost K is waived in period t only, then it is optimal to order $y_t^o - x_t$ units.*

PROPOSITION 2. *Suppose that Condition 1 holds. For any y_t^1 and y_t^2 satisfying $y^* \leq y_t^1 < y_t^2$ or $y_t^2 < y_t^1 \leq y^*$, we have*

$$V_t(y_t^1) + \gamma K \geq V_t(y_t^2).$$

PROOF. We prove the result for the finite-horizon case. The infinite-horizon discounted-profit case follows directly. If $t = T$, Condition 1(a) implies $V_t(y_t^1) \geq V_t(y_t^2)$, and the required result holds. We proceed by assuming $t < T$.

We compare two systems starting with y_t^1 and y_t^2 , and use the superscripts 1 and 2 to denote each of them. (In this proof and all subsequent proofs, whenever we compare two systems, we assume without loss of generality that they experience the same sample paths of ϵ_s .) Suppose that the y^2 -system follows the optimal decision to attain $V_t(y_t^2)$. Let \mathbf{d}_t^2 be the sales lever decision of the y^2 -system.

We claim that there exists \mathbf{d}_t^1 such that $\pi(y_t^1, \mathbf{d}_t^1) \geq \pi(y_t^2, \mathbf{d}_t^2)$ and $x_{t+1}^1 \leq \max\{x_{t+1}^2, y^*\} \leq \max\{y_{t+1}^2, y^*\}$.

If $y_t^2 < y_t^1 \leq y^*$, Condition 1(a) implies the existence of \mathbf{d}_t^1 such that $\pi(y_t^1, \mathbf{d}_t^1) \geq \pi(y_t^2, \mathbf{d}_t^2)$, and we know $x_{t+1}^1 \leq y^*$ because $y_t^1 \leq y^*$. On the other hand, when $y^* \leq y_t^1 < y_t^2$, Condition 1(b) is applicable. In either case, the claim is true, and either (a) $x_{t+1}^1 \leq y_{t+1}^2$ or (b) $y_{t+1}^2 < x_{t+1}^1 \leq y^*$ holds.

• **Case (a):** $x_{t+1}^1 \leq y_{t+1}^2$. In the next period $t + 1$, set the ordering quantity of the y^1 -system to $y_{t+1}^2 - x_{t+1}^1$. Thus, $y_{t+1}^1 = y_{t+1}^2$. From period $t + 2$ onwards, let the y^1 -system mimic the y^2 -system.

• **Case (b):** $y_{t+1}^2 < x_{t+1}^1 \leq y^*$. In period $t + 1$, the y^1 -system does not order, i.e., set $y_{t+1}^1 := x_{t+1}^1$. By Condition 1(a), we choose \mathbf{d}_{t+1}^1 such that $\pi(y_{t+1}^1, \mathbf{d}_{t+1}^1) \geq \pi(y_{t+1}^2, \mathbf{d}_{t+1}^2)$. We continue choosing the sales lever in the y^1 -system in this way until we come across the first period in which Case (a) is encountered, or the end of the horizon is reached.

Therefore, the y^1 -system generates as much profit π as the y^2 -system in each period after period t . Furthermore, the y^1 -system does not place an order in periods in which the y^2 -system does not order, with possibly one exception (where the first case is applied). Thus, the ordering cost of the y^1 -system is at most γK more than the y^2 -system. Hence, the discounted profit of the y^1 -system is at worst γK less than the discounted profit in the y^2 -system. \square

A corollary of this proposition is the optimality of not placing any order when the starting inventory level x_t is at least y^* or y_t^o .

COROLLARY 1. *Under Condition 1, we have the following. If $x_t \geq \min\{y^*, y_t^o\}$, then it is optimal not to order in period t , i.e., $y_t = x_t$.*

PROOF. Suppose that the starting inventory before ordering is x_t , and we order up to y_t , where $y_t > x_t$. The ordering cost K is incurred in period t . One of the following cases occurs.

• **Case $x_t \geq y^*$.** Because $y_t > x_t \geq y^*$, we apply Proposition 2 to get

$$V_t(x_t) \geq V_t(y_t) - \gamma K \geq V_t(y_t) - K.$$

• **Case $y_t^o < x_t < y^*$.** By Proposition 2 and the choice of y_t^o , we have

$$V_t(x_t) \geq V_t(y_t^o) - \gamma K \geq V_t(y_t) - \gamma K > V_t(y_t) - K.$$

It follows that when the starting inventory level is greater than $\min\{y^*, y_t^o\}$, ordering a positive quantity does not increase the discounted profit. \square

REMARK. In online Appendix A, we discuss the implication of Corollary 1 to the zero fixed-cost case, i.e., the optimality of myopic base-stock policies. An electronic companion to this paper is available as part of the online version that can be found at <http://or.pubs.informs.org/ecompanion.html>.

3.2. Infinite-Horizon Models: Proof of Theorem 1

We first prove the following claim: For $x^2 < x^1 \leq y^*$, if it is optimal to place an order when the beginning inventory level is x^1 , then it is also optimal to order when the starting inventory level is x^2 .

Let y^o be the smallest maximizer of $V(y)$. For the infinite-horizon discounted-profit model, the profit function and the optimal policy are both stationary. (See the comments following Assumption 1.) Thus, we assume that the current period t is 0, and drop the subscript t when $t = 0$. By Corollary 1, we proceed by assuming $x^1 < y_0^o = y^o$ because it is not optimal to order when the inventory level is above y^o .

If we order when the starting inventory level is either x^1 or x^2 , then the order-up-to level is y^o by Proposition 1, and the maximum profit from period t onward is $v^o := V(y^o) - K$. Suppose that the starting inventory level is x^1 . By deferring order placement to the next period, we can obtain a present value of profit equal to $Q(x^1) + \gamma v^o$. Because it is optimal to order in the current period, we must have $v^o \geq Q(x^1) + \gamma v^o$. Thus,

$$(1 - \gamma)v^o \geq Q(x^1) \geq Q(x) \quad \text{for all } x \leq x^1, \quad (4)$$

where the second inequality follows from Condition 1(a).

Now suppose that the starting inventory level is x^2 . Let J denote the period in which the next order is placed according to the optimal solution, or $J = \infty$ if the x^2 -system never

orders. Thus, $x^2 = x_0^2 \geq x_1^2 \geq x_2^2 \geq \dots \geq x_{J-1}^2$. From $x^2 < x^1$ and (4), an upper bound on the present value of the maximum profit between periods 0 and period $J - 1$ is given by

$$\begin{aligned} (1 + \gamma + \dots + \gamma^{J-1})Q(x^2) &\leq (1 + \gamma + \dots + \gamma^{J-1})Q(x^1) \\ &\leq (1 + \gamma + \dots + \gamma^{J-1}) \cdot (1 - \gamma)v^0 \\ &= (1 - \gamma^J)v^0. \end{aligned}$$

In period J , an order is placed. Thus, the present value of the maximum expected profit from J onwards is $\gamma^J v^0$. In summary, the present value of the expected profit in all periods is bounded above by $(1 - \gamma^J)v^0 + \gamma^J v^0 = v^0$, which is attained if we order in the current period. Thus, we complete the proof of the claim.

Let $s = \max\{x: \text{it is optimal to order when the inventory level is } x\}$. Corollary 1 implies $s \leq y^0$. Thus, by Proposition 1, the optimal order-up-to level is $S = y^0$ for all $x \leq s$. This completes the proof of Theorem 1.

3.3. Finite-Horizon Models: Proof of Theorem 2

PROPOSITION 3. *Suppose that Condition 2 holds. Then, $V_t(y_t)$ is nondecreasing in the interval $(-\infty, \min\{y^*, y_t^0\})$ for each t .*

PROOF. Suppose that y_t^1 and y_t^2 satisfy $y_t^2 < y_t^1 \leq \min\{y^*, y_t^0\}$. We want to show $V_t(y_t^1) \geq V_t(y_t^2)$. Suppose that the y^2 -system with the starting inventory level y_t^2 follows an optimal policy and attains $V_t(y_t^2)$. Below, we construct a policy for the y^1 -system such that

- the expected profit of the y^1 -system before accounting for the ordering cost in every period is at least that of the corresponding quantity in the y^2 -system, and
- the y^1 -system places an order in a period only if the y^2 -system places an order in that period.

Let \mathbf{d}_t^2 be the optimal sales lever of the y^2 -system in period t . By Condition 2(c), there exists \mathbf{d}_t^1 such that $\pi(y_t^1, \mathbf{d}_t^1) \geq \pi(y_t^2, \mathbf{d}_t^2)$ and $x_{t+1}^1 \geq x_{t+1}^2$ for any realization of ϵ in the demand model. Thus, if $x_{t+1}^1 < y_{t+1}^2$ occurs, an order must have been placed in the y^2 -system. In that case, let the y^1 -system place an order such that $y_{t+1}^1 = y_{t+1}^2$, and mimic the y^2 -system for the remaining periods. Otherwise, we repeat the above process until both the systems have the same ending inventory level, or the end of the horizon is reached. This concludes the proof of the proposition. \square

Proposition 1 and Corollary 1 state that it is optimal to order only if $x_t < \min\{y^*, y_t^0\}$, and the order-up-to level is $S_t = y_t^0$. Thus, if an order is placed in period t , the optimal profit from the remaining periods is independent of the starting inventory level x_t . For $x_t < \min\{y^*, y_t^0\}$, it is optimal to order if $V_t(x_t) \leq V_t(y_t^0) - K$, and not to order if $V_t(x_t) \geq V_t(y_t^0) - K$. Consequently, the monotonicity of V_t in the interval $(-\infty, \min\{y^*, y_t^0\})$ (from Proposition 3) establishes the existence of s_t , which is any solution to $V_t(s_t) = V_t(S_t) - K$. This completes the proof of Theorem 2.

REMARK. Note that Theorem 2 depends on Condition 2(c) only through Proposition 3. Thus, this theorem still holds if Condition 2 is replaced with Condition 1 and the monotonicity of V_t in the interval $(-\infty, \min\{y^*, y_t^0\})$.

4. Application to the Backordering Model

In this section, we apply the results of the previous section to the case where excess demand is backordered. The specification of Condition 1 and Condition 2 in §2 is quite technical. However, we will now show that it is a generalization of common modeling assumptions found in the literature. We present two sets of sufficient conditions for Conditions 1 or 2 to hold, and formally establish the optimality of (s, S) -type policies for finite- and infinite-horizon models under these conditions.

4.1. Joint Quasi-Concavity of π : A Sufficient Condition for Condition 1

It is shown in the following proposition that the concave demand model and the quasi-concavity of the expected single-period profit function π , together, imply Condition 1. The models studied in Federgruen and Heching (1999) and Chen and Simchi-Levi (2004b) satisfy these assumptions.

PROPOSITION 4. *Assume that excess demand is backordered. With the concave demand model, the joint quasi-concavity of π implies Condition 1.*

PROOF. Because π is jointly quasi-concave, $Q(y) = \max_{\mathbf{d}} \pi(y, \mathbf{d})$ is quasi-concave (see p. 102 of Boyd and Vandenberghe 2004 for a proof), and part (a) of Condition 1 is satisfied.

Suppose that y^1 and y^2 satisfy $y^* \leq y^1 < y^2$. Let $\lambda := (y^2 - y^1)/(y^2 - y^*)$. For any \mathbf{d}^2 , let \mathbf{d}^1 be the following convex combination of \mathbf{d}^* and \mathbf{d}^2 : $\mathbf{d}^1 = \lambda \mathbf{d}^* + (1 - \lambda) \mathbf{d}^2$. Because π is quasi-concave and (y^1, \mathbf{d}^1) is a convex combination of (y^2, \mathbf{d}^2) and (y^*, \mathbf{d}^*) ,

$$\pi(y^1, \mathbf{d}^1) \geq \pi(y^2, \mathbf{d}^2).$$

The concavity of $D(\mathbf{d}, \epsilon)$ in \mathbf{d} implies

$$D(\mathbf{d}^1, \epsilon) \geq \lambda D(\mathbf{d}^*, \epsilon) + (1 - \lambda) D(\mathbf{d}^2, \epsilon).$$

Thus,

$$\begin{aligned} y^1 - D(\mathbf{d}^1, \epsilon) &\leq \lambda[y^* - D(\mathbf{d}^*, \epsilon)] + (1 - \lambda)[y^2 - D(\mathbf{d}^2, \epsilon)] \\ &\leq \max\{y^* - D(\mathbf{d}^*, \epsilon), y^2 - D(\mathbf{d}^2, \epsilon)\} \\ &\leq \max\{y^*, y^2 - D(\mathbf{d}^2, \epsilon)\}. \end{aligned}$$

It follows that

$$\psi(y^1 - D(\mathbf{d}^1, \epsilon)) \leq \max\{y^*, \psi(y^2 - D(\mathbf{d}^2, \epsilon))\},$$

satisfying part (b) of Condition 1. \square

We are now ready to show the infinite-horizon optimality result holds under the concave demand model if π is jointly quasi-concave.

THEOREM 3. *Assume that excess demand is backordered. Assume the concave demand model and the joint quasi-*

concavity of π . Then, for the infinite-horizon discounted-profit model, an (s, S) policy is optimal.

PROOF. By Proposition 4, we know that Condition 1 is satisfied if π is jointly quasi-concave. The result now follows from Theorem 1. \square

We point out the following differences between our results and those contained in Chen and Simchi-Levi (2004b). They assume the concavity of the expected single-period profit function, whereas quasi-concavity is found to be sufficient for our results. Whereas they use the linear demand model, we use the more general concave demand model. In addition, our proof holds with multidimensional sales levers. However, they are able to show the optimality of (s, S, p) policies even in the infinite-horizon average-profit case, which we have not shown.

4.2. Additive Demand Model and the “Separability” of π : A Sufficient Condition for Condition 2

We now show that when the additive demand model is used, a separability-like condition on π is sufficient to guarantee Condition 2. Here, the sales lever d is single dimensional, corresponding to the expected demand associated with the decision.

ASSUMPTION 2. *There exist quasi-concave ϕ^R and quasi-convex ϕ^H such that $\pi(r + d, d) = \phi^R(d) - \phi^H(r)$.*

This assumption holds, for example, under the additive demand model with backordering in which revenues are received based on total demand and not on demand satisfied. In this case, ϕ^R denotes the revenue function and ϕ^H denotes the holding and shortage cost function. Note that this assumption does not imply the joint quasi-concavity of π . Assumption 2 is more general than the additive demand model of Chen and Simchi-Levi (2004a), which assumes concave ϕ^R and convex ϕ^H .

Let $d^* := \arg \max \phi^R(d)$ and $r^* := \arg \min \phi^H(r)$. Clearly, $y^* = r^* + d^*$.

The following proposition shows that under the additive demand model, Assumption 2 implies Condition 2.

PROPOSITION 5. *Assume that excess demand is backordered. Under the additive demand model, Assumption 2 implies Conditions 1 and 2.*

PROOF. Suppose that y^1 and y^2 satisfy $y^* \leq y^1 < y^2$. For any fixed d^2 , set $r^2 := y^2 - d^2$. Set $r^1 := \min\{r^2, y^1 - d^*\}$ and $d^1 := y^1 - r^1$. (It can be shown that d^1 is sandwiched between d^* and d^2 . Therefore, $d^1 \in \mathcal{D}$ because \mathcal{D} is convex.) Clearly, by $r^1 \leq r^2$ and the additivity of demand, (2) is satisfied.

Now consider the following two cases.

- *Case $r^1 = r^2$.* Clearly, $\phi^H(r^1) = \phi^H(r^2)$. Because $r^1 \leq y^1 - d^*$, we get

$$d^* \leq y^1 - r^1 < y^2 - r^1 = y^2 - r^2 = d^2.$$

Because $d^1 := y^1 - r^1$, we have $d^* \leq d^1 < d^2$. It follows from the quasi-concavity of ϕ^R that $\phi^R(d^1) \geq \phi^R(d^2)$.

- *Case $r^1 = y^1 - d^*$.* It follows $d^1 = d^*$, implying $\phi^R(d^1) = \phi^R(d^*) \geq \phi^R(d^2)$. We have

$$r^2 \geq r^1 = y^1 - d^1 \geq y^* - d^1 = y^* - d^* = r^*.$$

By the quasi-convexity of ϕ^H , it follows that $\phi^H(r^2) \geq \phi^H(r^1)$.

Therefore, for every d^2 , there exists d^1 such that $\pi(y^1, d^1) \geq \pi(y^2, d^2)$. We thus verify Condition 2(b). Furthermore, we obtain that $Q(y)$ is nonincreasing for $y \geq y^*$.

If y^1 and y^2 satisfy $y^* \geq y^1 > y^2$, let $r^1 := \max\{r^2, y^1 - d^*\}$ instead. Clearly, $r^1 \geq r^2$. A similar argument shows that Condition 2(c) holds and $Q(y)$ is nondecreasing for $y \leq y^*$. Thus, it follows that Q is quasi-concave. \square

We now establish that Assumption 2 implies the optimality of (s, S) -type policies for finite- and infinite-horizon problems.

THEOREM 4. *Assume that excess demand is backordered. Consider the additive demand model. Suppose that Assumption 2 holds. In the finite-horizon model, an (s, S_t) policy is optimal in period t . For the infinite-horizon case, an (s, S) policy is optimal.*

PROOF. Recall that Assumption 2 satisfies Conditions 1 and 2 by Proposition 5. Thus, the infinite-horizon optimality result follows from Theorem 1. The finite-horizon result follows from Theorem 2. \square

Chen and Simchi-Levi (2004a) also show the finite-horizon optimality result. Our result is more general in that, for example, the single-period profit function does not even need to be quasi-concave, whereas they assume concavity of this profit function. (Also see the comments following Assumption 2.) On the other hand, their results hold in nonstationary systems also.

5. Application to the Lost-Sales Model

Chen et al. (2006) study a finite-horizon model in which any unsatisfied demand is lost. They introduce some technical assumptions that facilitate their dynamic programming, induction-based proof of the (s, S) optimality result. In this section, we show this result for both the finite- and the infinite-horizon discounted-profit models using our proof technique. We apply results from §3.

They use the additive demand model in which the sales lever is the per-unit selling price p per unit,¹⁰ i.e.,

$$D(p, \epsilon) = d(p) + \epsilon, \tag{5}$$

where $d(p)$ is the deterministic part of demand. Let f and F denote the probability density and cumulative distribution functions of ϵ , respectively. Let $\mathcal{P} := [0, P^u]$ be the domain of p . Assume that any $p \in \mathcal{P}$ satisfies $d(p) \geq 0$. They impose the following additional technical assumptions.

ASSUMPTION 3 (CHEN ET AL. 2006). Demand is additive and is given by (5), where $d(p)$ is a deterministic function, and ϵ is a nonnegative, continuous random variable defined on a closed interval $[0, B]$. The probability density of ϵ is strictly positive on $(0, B)$. We have

- (a) $d(p)$ is decreasing, concave, and $3d'' + pd''' \leq 0$ on \mathcal{P} ;
- (b) the failure rate function $r(u) := f(u)/(1 - F(u))$ of ϵ satisfies $r'(u) + 2[r(u)]^2 > 0$ for any $u \in (0, B)$;
- (c) the expected single-period profit function is given by

$$\pi(y, p) = p \cdot E[\min(y, D(p, \epsilon))] - h \cdot E[(y - D(p, \epsilon))^+] - b \cdot E[(D(p, \epsilon) - y)^+],$$

where h is a holding cost per unit and b is a penalty cost per unit.

This assumption is satisfied by a wide range of demand functions and distribution functions of ϵ ; see Chen et al. (2006) for a discussion. Our proof makes use of some intermediate results they derived using the technical assumption above.

PROPOSITION 6. Assume that excess demand is lost. Assumption 3 implies Conditions 1 and 2.

PROOF. See online Appendix B. \square

We are now ready to demonstrate the application of our proof technique to the optimality of (s, S, p) policies for both the finite- and the infinite-horizon discounted-profit models with lost sales.

THEOREM 5. Assume that excess demand is lost. Suppose that Assumption 3 holds. In the finite-horizon model, an (s_t, S_t) policy is optimal in period t . For the infinite-horizon case, an (s, S) policy is optimal.

PROOF. For the infinite-horizon case, we know from Proposition 6 that Condition 1 is satisfied; the result now follows from Theorem 1. Moreover, for the finite-horizon case, Proposition 6 implies Condition 2 holds. Therefore, the finite-horizon result now follows from Theorem 2. In addition to the proofs omitted here, it also contains a section on stochastically increasing, additive demands. \square

6. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.pubs.informs.org/ecompanion.html>.

Endnotes

1. We use the terminology “ (s, S) -type policies,” which are the obvious generalization of (s, S, p) policies, when we discuss our results.
2. The three types of demand models discussed in this section—namely, concave, linear, and additive—will be defined precisely in §1.2.

3. As Chen and Simchi-Levi (2004a) point out, the only sufficient condition provided by Federgruen and Heching (1999) for this assumption is the linearity of the demand model.

4. We also allow the possibility of \mathcal{D} containing a single element; this represents traditional inventory control without any sales levers.

5. The monotonicity requirement of demand with respect to each sales lever does not cause any loss of generality in a practical sense. For example, price discounts and advertising budgets have a clear monotone effect on the demand.

6. The details are available upon request from the authors.
7. Theorem 4.2.3 and Lemma 4.2.8 of Hernandez-Lerma and Lasserre (1996) give conditions for the existence of an optimal stationary policy and for the convergence of the finite-horizon dynamic program to the infinite-horizon dynamic program. The models studied in this paper satisfy these conditions. The details are available upon request. Also, the functions W_t depend on T , but we have suppressed this dependence for the sake of conciseness.

8. A function $f: \mathfrak{R}^n \rightarrow \mathfrak{R}$ is *quasi-convex* if the level set $\{w: f(w) \leq l\}$ of f is convex for any $l \in \mathfrak{R}$. A convex function is quasi-convex. We say f is *quasi-concave* if $-f$ is quasi-convex. Quasi-concave functions are also referred to as unimodal functions.

9. Because we do not make any claims about the optimal sales lever to be chosen, we prefer referring to this as an (s, S) policy rather than an (s, S, \mathbf{d}) policy.

10. In this section, we use p in place of \mathbf{d} because in this model, the only sales lever is the price, p .

Acknowledgments

The authors sincerely thank John Muckstadt of Cornell University for several motivating discussions with them when they started this work. They also thank Robin Roundy for helpful comments. Constructive comments by anonymous referees and the associate editor helped improve the exposition of this paper. An earlier version of this paper was entitled “Optimality Results in Inventory-Pricing Control: An Alternate Approach.”

References

- Boyd, S., L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press, Cambridge, UK.
- Chao, X., S. X. Zhou. 2006. Joint inventory-and-pricing strategy for a stochastic continuous review system. *IIE Trans.* **38**(5) 401–408.
- Chen, X., D. Simchi-Levi. 2004a. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* **52**(6) 887–896.
- Chen, X., D. Simchi-Levi. 2004b. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The infinite horizon case. *Math. Oper. Res.* **29**(3) 698–723.
- Chen, X., D. Simchi-Levi. 2006. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The continuous review model. *Oper. Res. Lett.* **34** 323–332.

- Chen, Y., S. Ray, Y. Song. 2006. Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales. *Naval Res. Logist.* **53**(2) 117–136.
- Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3) 454–475.
- Feng, Y., F. Y. Chen. 2003. Joint pricing and inventory control with setup costs and demand uncertainty. Working paper, Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong.
- Feng, Y., F. Y. Chen. 2004. Optimality and optimization of a joint pricing and inventory-control policy for a periodic-review system. Working paper, Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong.
- Hernandez-Lerma, O., J. B. Lasserre. 1996. *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. Springer-Verlag, New York.
- Huh, W. T., G. Janakiraman. 2008. Inventory management with auctions and other sales channels: Optimality of (s, S) policies. *Management Sci.* **54**(1) 139–150.
- Scarf, H. 1960. The optimality of (s, S) policies for the dynamic inventory problem. *Proc. First Stanford Sympos. Math. Methods in the Soc. Sci.* Stanford University Press, Stanford, CA.
- Song, Y., S. Ray, T. Boyaci. 2006. Optimal dynamic joint pricing and inventory control for multiplicative demand with fixed order costs. Working paper, Desautels Faculty of Management, McGill University, Montreal, Quebec, Canada.
- Veinott, A. F. 1966. On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM J. Appl. Math.* **14**(5) 1067–1083.
- Yin, R., K. Rajaram. 2007. Joint pricing and inventory control with a Markovian demand model. *Eur. J. Oper. Res.* **182**(1) 113–126.