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

Woonghee Tim Huh[†] Columbia University
Ganesh Janakiraman[‡] New York University

May 10, 2004; April 30, 2005; May 15, 2006; October 17, 2006

Abstract

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 which 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. We also provide improved bounds on the optimal (s, S) parameters.

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 policies[1] has been established in recent papers such as Chen and

*An earlier version of this paper was entitled "Optimality Results in Inventory-Pricing Control: An Alternate Approach"

[†]Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA. huh@ieor.columbia.edu.

[‡]Stern School of Business, New York University, New York, NY 10012, USA. gjanakir@stern.nyu.edu.

[1]We use the terminology "(s, S)-type policies", which are the obvious generalization of (s, S, p) policies, when we discuss our results.

Simchi-Levi (2004a,b). 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 and dynamic programming formulations, which are pervasive in the area of stochastic inventory theory. 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 multi-dimensional sales lever, and a less-restrictive single-period expected profit function (permitting, for example, quasi-concavity).

Our proofs, completely different from the earlier proofs in the joint inventory-pricing literature, are constructive and based on a condition which we identify as being key to proving the (s, S) structure. We believe that our proofs are important in their own right. 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.

1.1 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 forwards. 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. This ordering opportunity is a control on the supply or inflow process. Similarly, we have a set of levers to control the demand or outflow process. Examples of these levers are prices, sales-force incentives and advertisements. We model the sales lever control by a vector $\mathbf{d}$ within some compact convex subset $\mathcal{D} \subseteq \Re^n$; the first component could denote the price discount and the second component could denote the advertisement expense and so on.[2]

After y and $\mathbf{d}$ are chosen, the demand in period t is realized next. It is a 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

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

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 ψ is one of the following two operators: (a) an identity map if excess demand is completely backordered, or, (b) the positive part operator 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, $t + 1$, ..., 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 page 1072).

We make the following assumption throughout the paper.

Assumption 1.

(a) Either $\psi(x) = x$ (complete backlogging), or $\psi(x) = (x)^+$ (lost sales).

(b) $\pi(y, \mathbf{d})$ is continuous, and $\max_{\mathbf{d}} \pi(y, \mathbf{d}) \to -\infty$ as $y \to \infty$. Furthermore, $\max_{(y,\mathbf{d})} \pi(y, \mathbf{d})$ exists.

(c) The demand model is stationary, that is, the sequence of ϵ's in time periods $\{1, 2, \ldots, T\}$ is independent and identically distributed.

(d) $D(\mathbf{d}, \epsilon)$ is continuous in $(\mathbf{d}, \epsilon)$. Moreover, for every ϵ, $D(\mathbf{d}, \epsilon)$ is component-wise monotone in $\mathbf{d}$.[3]

Let $(y^*, \mathbf{d}^*)$ be the 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^*.

[3]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.

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

$$\begin{aligned} V_t(y) \;&=\; \max_{\{\mathbf{d}: \mathbf{d} \in \mathcal{D}\}} [W_t(y, \mathbf{d})], \quad \text{and} \\ W_t(y, \mathbf{d}) \;&=\; \begin{cases} \pi(y, \mathbf{d}) \;+\; \gamma \cdot E_\epsilon U_{t+1}(\psi(y - D(\mathbf{d}, \epsilon))), & \text{if } t < T; \\ \pi(y, \mathbf{d}), & \text{if } t = T. \end{cases} \end{aligned}$$

(The subscript t in the above formulation denotes the period index.) It can be shown using standard arguments that the above maxima exist.[4]

An optimal policy for this finite-horizon problem specifies a feasible pair $(y, \mathbf{d})$, 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 point-wise limit of $W_1(\cdot, \cdot)$ as T approaches infinity.[5] In this paper, we study finite-horizon problems and infinite-horizon discounted-profit problems.

The following definitions are based on the dynamic programming formulation and become useful later:

- Let y_t° be the smallest maximizer of $V_t(y)$, and y° be the smallest maximizer of $V(y)$, where $V(y) := \max\{W(y, \mathbf{d}) : \mathbf{d} \in \mathcal{D}\}$;
- $Q(y) := \max\{\pi(y, \mathbf{d}) \mid \mathbf{d} \in \mathcal{D}\}$.

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 $\mathbf{d}$ since $\mathcal{D} \subseteq \Re^1$.)

[4] Available upon request from the authors.

[5] Theorem 4.2.3 and Lemma 4.2.8 of Hernandez-Lerma and Lassere (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.

- **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 $\alpha \in \Re^n$ and $\beta \in \Re$. 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.

Among these models, the strongest structural results hold under the Additive Demand Model. For example, Chen and Simchi-Levi (2004a) show the optimality of (s, S, p) policies for finite horizon problems with backordering and the Additive Demand Model; however, they show that the result does *not* always hold under the Linear Demand Model. Similarly, Chen et al. (2006) require the Additive Demand Model in order to prove the optimality of (s, S, p) policies with the lost sales assumption. Under the backordering assumption, the infinite horizon results in Chen and Simchi-Levi (2004b) are proved using the Linear Demand Model. Notice 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. Federgruen and Heching (1999) introduced the Concave Demand Model when there is no fixed cost for ordering.

1.2 Literature Review and Summary of Contributions

The primary focus of this paper is to provide an alternate, constructive proof technique for important results in joint inventory-pricing control. In this section, we first review early classical work on inventory models with fixed ordering costs *without* any pricing or sales-lever decisions. Then, we review recent papers on joint inventory-pricing control, followed by an explanation of our contribution to the literature.

1.2.1 Classical Results without Pricing Decisions

Scarf (1960) and Veinott (1966), both seminal papers in Inventory Theory, establish the optimality of (s, S) policies for inventory systems with fixed ordering costs. (No sales lever decisions are considered.)

Scarf shows that the cost function in the finite horizon dynamic program possesses a property, he named as K-convexity. A key step in the proof is demonstrating *by induction* that this property holds for an $(n+1)$-period problem if it holds for the n period problem. The proof, although originally presented for stationary models, can be extended to non-stationary environments as noted by Veinott (1966). Scarf's main assumption is that the single-period expected holding and shortage cost function is convex with respect to the inventory level after ordering.

Veinott presents another proof for the (s, S) optimality result under different assumptions. He generalizes the assumption on the expected holding and shortage cost function to include quasi-convex functions, but requires the sequence of minimizers of these functions to be weakly increasing. His proof does not use induction based on the dynamic program. While his proof is algebraic, he sketches another proof that is constructive. Our work is inspired by this constructive approach.

1.2.2 Joint Inventory-Pricing Control

We now discuss the literature on joint inventory-pricing control. (We refer the reader to these papers for an exhaustive list of references.)

With complete backlogging, no fixed ordering costs, and the Concave Demand Model, Federgruen and Heching (1999) show the optimality of the base-stock list-price policy for the non-stationary 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 $\pi(y, \mathbf{d})$ is jointly concave in inventory (after ordering) and price.[6]

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 $\pi(y, \mathbf{d})$ for their results on the finite-horizon problem as well as the infinite-horizon discounted-profit problem. They

[6] 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.

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 (2005) 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.

1.2.3 Contributions

The papers mentioned in Section 1.2.2 use several inductive arguments for finite-horizon results and rather involved convergence arguments to establish the infinite-horizon results. The contrast between the simplicity of the structure in the optimal policies ((s, S, p) policies and base-stock list-price policies) and the complexity of the optimality proofs is quite striking. As mentioned earlier, the main contribution of this paper is to provide constructive proofs for these types of results.

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 (2004) in the context of selling through auction channels.

Finally, our results complement the existing literature in some ways. In particular, we relax the usual assumptions on the demand model and profit functions; for example, we replace the joint concavity of $\pi(y, \mathbf{d})$ with joint quasi-concavity. Instead, we require the stationarity assumption. Table 1 summarizes the recent work on single-stage systems with sales levers along with our results. In addition, we improve existing bounds on the optimal s and S values.

For joint inventory-pricing problems, there has been a remarkably long time gap between Thomas (1974)'s conjecture of the optimality of (s, S)-type policies and Chen and Simchi-Levi's proof. Even through their proof is inspired by Scarf's work, their work has rightfully gained a lot of attention because the generalization is significant both in terms of the result and the methodology. Our work is a similar advancement of the approach initiated by Veinott. More specifically, proving optimal policy structure with inventory and price as decision variables is much more involved than merely appending the price variable in the proofs of classical papers in which prices are exogenous.

1.3 Organization

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

2 A Single-Period Condition

In this section, we present a set of assumptions on the expected single-period profit π that lies at the core of our proofs. Although the following condition appears technical, we show in Sections 4 and 5 that common modeling assumptions found in the literature satisfy this condition.

Condition 1.

(a) $Q(y) = \max_{\mathbf{d} \in \mathcal{D}} \pi(y, \mathbf{d})$ *is quasi-concave*[7]*, and*

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

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

[7]A function $f : \Re^n \to \Re$ is *quasi-convex* if the level set $\{w : f(w) \leq l\}$ of f is convex for any $l \in \Re$. A convex function is quasi-convex. We say f is quasi-concave if $-f$ is quasi-convex.

Fixed Cost	Models	Earlier Results	Our Results [a]
Zero	Backorder Models (Finite & Infinite Horizon)	Concave Demand Model, jointly concave π: Base-stock list-price policy is optimal.[b]	Concave Demand Model: Base-stock list-price policy is optimal with weaker assumptions, e.g., quasi-concave π.
Positive	Backorder Models (Finite Horizon)	Additive Demand Model, jointly concave π: (s,S,p) is optimal. Linear Demand Model, jointly concave π: (s,S,p) may not be optimal, but (s,S,A,p) is optimal.[c]	Additive Demand Model: (s,S,p) is optimal with weaker assumptions, e.g., $\pi(r+d,d) = \phi^R(d) - \phi^H(r)$ where ϕ^R is quasi-concave and ϕ^H is quasi-convex.
	Backorder Models (Infinite Horizon)	Linear Demand Model, jointly concave π: (s,S,p) is optimal for the discounted-profit and average-profit problems.[d] "Strongly stochastic concave" demand model, jointly concave π: (s,S,p) is optimal for the average-profit problem.[e]	Concave Demand Model: (s,S,p) is optimal for the discounted-profit problems with weaker assumptions, e.g., quasi-concave π.
	Lost Sales Models (Finite Horizon)	Additive Demand Model: (s,S,p) is optimal with technical assumptions.[f]	Additive Demand Model: (s,S,p) is optimal with the same assumptions.
	Lost Sales Models (Infinite Horizon)	none	Additive Demand Model: (s,S,p) is optimal for the discounted-profit problems with the same assumptions as the finite-horizon case.

[a] All our results here assume stationarity of all parameters.
[b] Federgruen and Heching (1999). See Footnote 6.
[c] Chen and Simchi-Levi (2004a).
[d] Chen and Simchi-Levi (2004b).
[e] Feng and Chen (2004).
[f] Chen et al. (2006). They assume stationarity.

Table 1: Summary of (s,S)-Type Results under Periodic Review with Stationary Parameters

such that for any ϵ,

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

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.

The above assumption ensures that the problem has sufficient structure. 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 since the starting inventory of the y^1-system is below y^*.

3 (s, S) Optimality and Condition 1

This section contains the proof of the optimality of (s, S)-type policies under Condition 1. Preliminary results are established in Section 3.1. Infinite horizon and finite horizon optimality results are presented in Section 3.2 and Section 3.3, respectively. We provide improved bounds on s and S values in Section 3.4.

3.1 Preliminary Results

In this section, we establish that in period t, if the starting inventory level before ordering is below y_t°, *and* an order is placed, then it is optimal to order up to y_t° (Proposition 3.1). Furthermore, if the starting inventory level is above y_t°, then it is optimal not to order (Corollary 3.3).

Throughout the paper, we use x_t to denote the starting inventory level before ordering, whereas y_t is the inventory level after ordering.

Proposition 3.1. *Suppose $x_t \leq y_t^\circ$. Then, $y_t = y_t^\circ$ maximizes $\{V_t(y_t) \mid y_t \geq x_t\}$. That is, if the fixed cost K is waived in period t only, then it is optimal to order $y_t^\circ - x_t$ units.*

Proof. It follows from the optimality of y_t° and the zero variable ordering cost assumption. $\square$

Proposition 3.2. *Suppose Condition 1 holds. For any y_t^1 and y_t^2 satisfying $y^* \le y_t^1 < y_t^2$ or $y_t^2 < y_t^1 \le y^*$, we have*

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

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 $V_t(y_t^1) \ge V_t(y_t^2)$, 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 superscript 1 and 2 to denote each of them. Suppose 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 $\pi(y_t^1, \mathbf{d}_t^1) \ge \pi(y_t^2, \mathbf{d}_t^2)$ and $x_{t+1}^1 \le \max\{x_{t+1}^2,\ y^*\} \le \max\{y_{t+1}^2,\ y^*\}$.

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

- *Case (a): $x_{t+1}^1 \le y_{t+1}^2$:* 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): $y_{t+1}^2 < x_{t+1}^1 \le y^*$:* 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 $\pi(y_{t+1}^1, \mathbf{d}_{t+1}^1) \ge \pi(y_{t+1}^2, \mathbf{d}_{t+1}^2)$. 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.

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°.

Corollary 3.3. *Under Condition 1, we have the following. If* $x_t \geq \min\{y^*,\ y_t^\circ\}$*, 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* $x_t \geq y^*$*:* Since $y_t > x_t \geq y^*$, we apply Proposition 3.2 to get

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

- *Case* $y_t^\circ < x_t < y^*$*:* By Proposition 3.2 and the choice of y_t°, we have

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

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 by more than the fixed cost K of ordering. □

Remark. In Appendix A.1, we discuss the implication of Corollary 3.3 to the zero fixed cost case, i.e., the optimality of myopic base-stock policies.

3.2 Infinite-Horizon Discounted Profit Model with Positive K

In this section, we establish the optimality of (s, S)-type policies in the infinite-horizon discounted-profit model, in which $T = \infty$ and $\gamma < 1$.

Theorem 3.4. *Suppose Condition 1 holds. In the infinite-horizon discounted-profit model, there exist* s *and* S *such that if* $x_t \geq s$*, it is optimal to not order, and otherwise to order up to* S*. That is, an* (s, S) *policy is optimal.*[8]

Proof. 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.

[8]Since 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.

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 the current period t is 0, and drop the subscript t when $t = 0$. By Corollary 3.3, we proceed by assuming $x^1 < y_0^\circ = y^\circ$, since it is not optimal to order when the inventory level is above y°.

If we order when the starting inventory level is either x^1 or x^2, then the order-up-to level is y° by Proposition 3.1, and the maximum profit from period t onward is $v^\circ := V(y^\circ) - K$. Suppose 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 $Q(x^1) + \gamma v^\circ$. Since it is optimal to order in the current period, we must have $v^\circ \geq Q(x^1) + \gamma v^\circ$. Thus,

$$(1-\gamma)v^\circ \ \geq \ Q(x^1) \ \geq \ Q(x) \ \text{ for all } x \leq x^1, \tag{3}$$

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

Now, suppose the starting inventory level is x^2. Let J denote the period in which the next order is placed according to the optimal solution. Thus, $x^2 = x_0^2 \geq x_1^2 \geq x_2^2 \geq \cdots \geq x_{J-1}^2$. From $x^2 < x^1$ and (3), an upper bound on the present value of the maximum profit between periods 0 and period $J - 1$ is given by

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

In period J, an order is placed. Thus, the present value of the maximum expected profit from J onwards is $\gamma^J v^\circ$. In summary, the present value of the expected profit in all periods is bounded above by $(1-\gamma^J)v^\circ + \gamma^J v^\circ = v^\circ$, 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 3.3 implies $s \leq y^\circ$. Thus, by Proposition 3.1, the optimal order-up-to level is $S = y^\circ$ for all $x \leq s$. $\square$

3.3 Finite-Horizon Model with Positive K

In this section, we prove the optimality of an (s_t, S_t)-type policy, an (s, S) policy with time-dependent parameters, for the finite horizon problem. For proving this result, we need a stronger version of Condition 1. We will use this version in Sections 4.2 and 5.

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,

$$\psi(y^1 - D(\mathbf{d}^1, \epsilon)) \;\geq\; \psi(y^2 - D(\mathbf{d}^2, \epsilon)) \;. \tag{4}$$

Note that Condition 2 implies Condition 1.

Proposition 3.5. *Suppose Condition 2 holds. Then, $V_t(y_t)$ is nondecreasing in the interval $(-\infty, \min\{y^*, y_t^\circ\}]$ for each t.*

Proof. Suppose y_t^1 and y_t^2 satisfy $y_t^2 < y_t^1 \leq \min(y^*, y_t^\circ)$. We want to show $V_t(y_t^1) \geq V_t(y_t^2)$. Suppose 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 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, 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. This concludes the proof of the proposition. □

Theorem 3.6. *Suppose Condition 2 holds. Then, in a finite-horizon model, an (s_t, S_t) policy is optimal in period t.*

Proof. Proposition 3.1 and Corollary 3.3 state that it is optimal to order only if $x_t < \min(y^*, y_t^\circ)$, and the order-up-to level is $S_t = y_t^\circ$. 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 $x_t < \min(y^*, y_t^\circ)$, it is optimal to order if $V_t(x_t) \leq V_t(y_t^\circ) - K$, and not to order if $V_t(x_t) \geq V_t(y_t^\circ) - K$. Consequently, the monotonicity of V_t in the interval $(-\infty, \min\{y^*, y_t^\circ\}]$ (from Proposition 3.5) establishes the existence of s_t, which is any solution to to $V_t(s_t) = V_t(S_t) - K$. $\square$

Remark: Note that Theorem 3.6 depends on Condition 2 only through Proposition 3.5. Thus, this theorem still holds if Condition 2 is replaced with Condition 1 and the monotonicity of V_t in the interval $(-\infty, \min\{y^*, y_t^\circ\}]$.

3.4 Bounds on s and S

In Sections 3.2 and 3.3, we have shown the optimality of (s, S)-type policies. In this section, we establish bounds on the values of s and S in the case of the infinite-horizon discounted profit case, and on s_t and S_t in the case of the finite-horizon case. Our bounds improve those existing in the literature. Chen and Simchi-Levi (2004b) have first established bounds on s and S (or s_t and S_t also) for systems with backordering. Their bounds hold even for finite-horizon systems with the Linear Demand model, for which the optimal policy is not always an (s, S) policy; however, the derivation of our bounds (in Theorem 3.7) relies on the optimality of (s, S) policies. On the positive side, we note that our bounds are tighter than their bounds. In particular, as $K \to 0^+$, our bounds converge to y^* whereas their bounds do not. Furthermore, our bounds are also applicable to the lost sales case.

Let

$$\begin{aligned}
\overline{m} &= \inf \ \{ \ y \mid Q(y^*) - Q(y') > \gamma K \text{ for each } y' \geq y \ \} \\
\underline{m} &= \sup \ \{ \ y \mid Q(y^*) - Q(y') > K \text{ for each } y' \leq y \} \\
\underline{M} &= \sup \ \{ \ y \mid Q(y^*) - Q(y') > (1+\gamma) K \text{ for each } y' \leq y \ \} \ .
\end{aligned}$$

Clearly, $\underline{M} \leq \underline{m} \leq y^* \leq \overline{m}$.

Theorem 3.7. *The following statements are true:*

(a) In Theorems 3.4, there exist s and S satisfying $\underline{m} \ \leq \ s \ \leq \ S \ \leq \ \overline{m}$, and $s \leq y^*$.

(b) In Theorems 3.6, there exist s_t and S_t satisfying $\underline{M} \ \leq \ s_t \ \leq \ y^* \ \leq \ S_t \ \leq \ \overline{m}$.

Proof. These bounds are derived using the constructive proof technique we employed in proving the sufficiency of Condition 1 for (s, S) optimality earlier in this section. See Appendix A.2 for the proof. $\square$

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 in Section 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 Condition 1 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 the quasi-concavity of the expected single-period profit function π, together, imply Condition 1. These modeling assumptions are more general than those used in the literature. Federgruen and Heching (1999) assume both the concavity of $D(p, \epsilon)$ in p, and the joint concavity of the single-period profit $\pi(y, p)$. For the infinite-horizon discounted-profit criterion, Chen and Simchi-Levi (2004b) use a Linear Demand Model and assume that the single-period profit function is concave. Thus, the following proposition implies that these models also satisfy Condition 1.

Note that if f is a quasi-concave function and w^* is its maximizer, then $f(w^* + \lambda v)$ is nonincreasing in $\lambda \geq 0$ for any v.

Proposition 4.1. *With the Concave Demand Model, the joint quasi-concavity of π implies Condition 1.*

Proof. Since π is jointly quasi-concave, $Q(y) = \max_{\mathbf{d}} \pi(y, \mathbf{d})$ is quasi-concave, and part (a) of Condition 1 is satisfied.

Suppose 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 $\mathbf{d}^2$: $\mathbf{d}^1 = \lambda \mathbf{d}^* + (1 - \lambda)\mathbf{d}^2$. Since π is quasi-concave,

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

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 $\psi(y^1 - D(\mathbf{d}^1,\epsilon)) \leq \max\{y^*,\ \psi(y^2 - D(\mathbf{d}^2,\epsilon))\}$, satisfying part (b) of Condition 1. □

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

Theorem 4.2. *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 4.1, we know that Condition 1 is satisfied if π is jointly quasi-concave. The result now follows from Theorem 3.4. □

The above theorem shows the optimality of an (s,S)-type policy in the infinite-horizon discounted-profit model when excess demand is backordered. 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. While they use the Linear Demand Model, we use the more general Concave Demand Model. In addition, our proof holds with multi-dimensional 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: Another Sufficient Condition for Condition 1

We now show that when the Additive Demand Model is used, a separability-like condition on π is sufficient to guarantee Condition 1. 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. Notice 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 4.3. *Under the Additive Demand Model, Assumption 2 implies Condition 1 and Condition 2.*

Proof. Suppose 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, $d^1 \in \mathcal{D}$ since $\mathcal{D}$ 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, $\phi^H(r^1) = \phi^H(r^2)$. Since $r^1 \leq y^1 - d^*$, we get

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

  Since $d^1 := y^1 - r^1$, 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 $\pi(y^1, d^1) = \phi(r^1, d^1) \geq \phi(r^2, d^2) = \pi(y^2, d^2)$. 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 4.4. *Consider the Additive Demand Model. Suppose Assumption 2 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. Recall that Assumption 2 satisfies Condition 1 and Condition 2 by Proposition 4.3. Thus, the infinite horizon optimality result follows from Theorem 3.4. The finite horizon result follows from Theorem 3.6. □

Chen and Simchi-Levi (2004a) also show the finite horizon optimality result. Our result is more general in that, for instance, the single-period profit function does not even need to be quasi-concave, whereas they assume concavity of this profit function. (Also see comments following Assumption 2.) On the other hand, their results hold in non-stationary 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 will demonstrate that their assumptions are sufficient to prove the result for both the finite and the infinite horizon discounted profit models using our proof technique. We apply results from Section 3.

They use the Additive Demand Model in which the sales lever is the per-unit selling price p per unit[9], 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.

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

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 $\mathcal{P}$, and*

(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 5.1. *In a finite-horizon model, Assumption 3 implies Condition 1 and Condition 2.*

Proof. See Appendix A.3. □

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 5.2. *Suppose 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 5.1 that Condition 1 is satisfied; the result now follows from Theorem 3.4. Moreover, for the finite-horizon case, Proposition 5.1 implies Condition 2 holds. Therefore, the finite horizon result now follows from Theorem 3.6. □

The appendix of this paper can be found online. In addition to the proofs omitted here, it also contains a section on stochastically increasing, additive demands.

Acknowledgement

We sincerely thank John Muckstadt of Cornell University for several motivating discussions with us when we started this work. We also thank Robin Roundy for helpful comments. Constructive comments by anonymous referees and the Associate Editor have helped us in improving the exposition of this paper.

References

Chao, X., and S. X. Zhou. 2006. Joint inventory-and-pricing strategy for a stochastic continuous-review system. *IIE Transactions* 38 (5): 401–408.

Chen, X., and D. Simchi-Levi. 2004a. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case. *Operations Research* 52 (6): 887–896.

Chen, X., and D. Simchi-Levi. 2004b. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Ininite Horizon Case. *Mathematics of Operations Research* 29 (3): 698–723.

Chen, X., and D. Simchi-Levi. 2006. Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Continuous Review Model. *Operations Research Letters* 34:323–332.

Chen, Y., S. Ray, and Y. Song. 2006. Optimal Pricing and Inventory Control Policy in Periodic-Review Systems with Fixed Ordering Cost and Lost Sales. *Naval Research Logistics* 53 (2): 117–136.

Federgruen, A., and A. Heching. 1999. Combined Pricing and Inventory Control Under Uncertainty. *Operations Research* 47 (3): 454–475.

Feng, Y., and 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 Univ. of Hong Kong, Shatin, N.T., Hong Kong.

Feng, Y., and 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., and J. Lassere. 1996. *Discrete-time markov control processes: Basic optimality criteria*. Springer-Verlag.

Huh, W., and G. Janakiraman. 2004. Optimality of (s, S) Replenishment Policies for Inventory Systems with Heterogeneous Sales Channels. Working paper, Department of Industrial Engineering and Operations Research, Columbia University, New York, U.S.A.

Porteus, E. L. 2002. *Foundations of stochastic inventory theory.* Stanford University Press.

Scarf, H. 1960. The Optimality of (s, S) Policies for the Dynamic Inventory Problem. In *Proceedings of the First Stanford Symposium on Mathematical Methods in the Social Sciences*. Stanford, CA: Stanford University Press.

Song, Y., S. Ray, and 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, Canada.

Thomas, L. 1974. Price and production decisions with random demand. *Operations Research* 26:513–518.

Veinott, A. 1966. On the Optimality of (s,S) Inventory Policies: New Conditions and A New Proof. *SIAM Journal on Applied Mathematics* 14 (5): 1067–1083.

Yin, R., and K. Rajaram. 2005. Joint Pricing and Inventory Control with A Markovian Demand. Working paper, UCLA Anderson School of Management, Los Angeles, USA.

A Appendix

A.1 A Special Case: Zero Fixed Cost

In this section, we prove the optimality of base-stock list-price policies in the absence of the fixed ordering cost. This result is similar to Federgruen and Heching (1999).

Theorem A.1. *If Condition 1 holds and $K = 0$, then, the optimal ordering policy for the finite horizon problem or the infinite horizon discounted profit problem has the following structure. If $x_t < y^*$, then order up to y^*, and set the sales lever to $\mathbf{d}^*$. Otherwise, do not order.*

Proof. By Theorem 6.1 of Porteus (2002), it is easy to verify that $(y^*, \mathbf{d}^*)$ is the optimal action pair if $x_t < y^*$. If $x_t \geq y^*$, it is optimal not to order by Corollary 3.3. $\square$

Remark. We make the following two remarks on Theorem A.1: (a) If $x_t \leq y^*$, then the myopic policy is optimal; (b) Theorem A.1 does not specify the optimal sales lever $\mathbf{d}_t$ in period t for $x_t > y^*$. This computation can be efficiently performed as follows.

Suppose the horizon T is finite or the profit is discounted, i.e., $\gamma < 1$. Federgruen and Heching (1999) give the following dynamic programming formulation. Recall $U_t(x_t)$ is the optimal value function as a function of the beginning inventory x_t in period t, i.e., $U_t(x_t) = \max_{y_t \geq x_t} V_t(y_t)$. Then,

$$U_t(x_t) = \max_{y_t \geq x_t,\ \mathbf{d}_t} \pi(y_t, \mathbf{d}_t) + \gamma E[U_{t+1}(\psi(y_t - D_t(\mathbf{d}_t, \epsilon)))]$$

for $t \leq T$, and $U_{T+1}(\cdot) = 0$. This dynamic program is well-defined when T is finite.

Once the starting inventory level x_t falls below y^*, the expected profit in each subsequent period is $Q(y^*)$. Thus, the expected profit (discounted to period t) from period t to T is

$$Q(y^*)(1 + \gamma + \gamma^2 + \cdots + \gamma^{T-t}) \quad = \quad Q(y^*) \cdot (1 - \gamma^{T-t+1})/(1 - \gamma) \ .$$

As a result, we simplify the dynamic programming recursion as follows: for any $t \leq T$,

$$U_t(x_t) \quad = \quad \begin{cases} \max_{\mathbf{d}_t} \pi(x_t, \mathbf{d}_t) + \gamma E[U_{t+1}(\psi(x_t - D_t(\mathbf{d}_t, \epsilon)))], & \text{for } x_t > y^*, \\ Q(y^*) \cdot (1 - \gamma^{T-t+1})/(1 - \gamma), & \text{for } x_t \leq y^*, \end{cases}$$

and $U_{T+1}(\cdot) = 0$. It is well-defined when T is finite. For the infinite-horizon discounted-profit problem, we let $T \to \infty$ and substitute U_t with its pointwise limiting function.

We point out the following important differences between the results in this section and the corresponding results in Federgruen and Heching (1999) under the backordering assumption. They require the concavity of the expected single-period profit whereas we require Condition 1, which, for example, is satisfied if π is quasi-concave (see Section 4.1). Their sales lever is single dimensional whereas ours is multi-dimensional. However, our simplification relies upon the stationarity of all the parameters of the system, whereas their finite-horizon results do not. For the lost sales case, Chen et al. (2006) prove this result under some assumptions which imply Condition 1, as we have shown in Section 5. It should be noted that subject to the validity of Condition 1, our result holds for both the backordering and the lost sales cases.

A.2 Proof of Theorem 3.7

In this section, we prove Theorem 3.7 which contains the following results: (i) the order-up-to levels (S and S_t) in Theorems 3.4 and 3.6 are bounded above by $\overline{m}$; and (ii) the reorder points (s and s_t) in Theorems 3.4 and 3.6 are bounded below by $\underline{m}$ and $\underline{M}$, respectively.

By the quasi-concavity of Q (Condition 1), we have $\underline{M} \leq \underline{m} \leq y^* \leq \overline{m}$. First, we show that $\overline{m}$ is an upper bound on the value of S or S_t. By Proposition 3.1, the order-up-to level in Theorem 3.4 (Theorem 3.6) is y_t° (y°), which maximizes $V_t(\cdot)$ ($V(\cdot)$).

Proposition A.2. *Under Condition 1, $y_t^\circ \leq \overline{m}$ for each t, and $y^\circ \leq \overline{m}$.*

Proof. Suppose, by way of contradiction, $y_t^\circ > \overline{m}$ for some t. We compare two systems starting with $y_t^1 = y^*$ and $y_t^2 = y_t^\circ$, respectively. Thus, $y_t^1 < y_t^2$.

Suppose the y^2-system follows the optimal decision to attain $V_t(y_t^2)$, and let $\mathbf{d}_t^2$ be its sales lever decision in period t. Let $\mathbf{d}_t^1$ be such that $Q(y_t^1) = \pi(y_t^1, \mathbf{d}_t^1)$. Then, from $y_t^\circ > \overline{m}$, the expected profits of two systems satisfy

$$\pi(y_t^1, \mathbf{d}_t^1) \;=\; Q(y^*) \;>\; Q(y_t^2) + \gamma K \;\geq\; \pi(y_t^2, \mathbf{d}_t^2) + \gamma K \;. \tag{6}$$

Consider the next period $t+1$. We have either (i) $x_{t+1}^1 < y_{t+1}^2$ or (ii) $y_{t+1}^2 \leq x_{t+1}^1$. In the first case, the y^1-system can incur the fixed cost of K, and order up to y_{t+1}^2, i.e., $y_{t+1}^1 = y_{t+1}^2$. Thus, in period $t+1$,

$$U_{t+1}(x_{t+1}^1) \;\;\geq\;\; V_{t+1}(y_{t+1}^2) - K \;. \tag{7}$$

In the second case, let the y^1-system place no order, i.e., $y^1_{t+1} = x^1_{t+1}$. It follows $y^2_{t+1} \le y^1_{t+1} \le y^*$. By Proposition 3.2, we obtain

$$V_{t+1}(y^1_{t+1}) \quad \ge \quad V_{t+1}(y^2_{t+1}) - K \ . \tag{8}$$

In either case, combining (7) or (8) with inequality (6), we obtain $V_t(y^1_t) > V_t(y^2_t)$, contradicting the choice of $y^2_t = y^\circ_t$.

The same proof shows the result for y° by replacing U_{t+1} and V_{t+1} with U and V, respectively. □

We now provide lower bounds for the value of the reorder point. In Proposition A.3, we show $\underline{M}$ is a lower bound on s or s_t.

Proposition A.3. *Reorder points s in Theorem 3.4 and s_t in Theorem 3.6 are both bounded below by $\underline{M}$.*

Proof. It suffices to show that for any $y^2_t < \underline{M}$,

$$V_t(y^*) \quad > \quad V_t(y^2_t) + K \tag{9}$$

holds. Then, when the starting inventory is $x^2_t = y^2_t$, it is better to order up to y^* than not to order. Let $\mathbf{d}^2_t$ be any sales lever decision of the y^2-system.

Let $y^1_t = y^*$, and let $\mathbf{d}^1_t$ be such that $Q(y^1_t) = \pi(y^1_t, \mathbf{d}^1_t)$. Then, $y^2_t < y^1_t = y^*$ implies $x^2_{t+1} \le x^1_{t+1} \le y^*$. For any y^2_{t+1} decision of the y^2-system, let $y^1_{t+1} = \max\{y^2_{t+1}, x^1_{t+1}\}$. Then, it follows

$$\mathbf{1}[y^1_{t+1} > x^1_{t+1}] \quad \le \quad \mathbf{1}[y^2_{t+1} > x^2_{t+1}] \ ,$$

and either $y^1_{t+1} = y^2_{t+1}$ or $y^2_{t+1} \le y^1_{t+1} \le y^*$ holds. Thus, by Proposition 3.2,

$$V_{t+1}(y^1_{t+1}) \quad \ge \quad V_{t+1}(y^2_{t+1}) - \gamma K \ .$$

Furthermore, from $y^2_t < \underline{M}$, we have

$$\pi(y^1_t, \mathbf{d}^1_t) \ = \ Q(y^*) \ > \ Q(y^2_t) + (1+\gamma)K \ \ge \ \pi(y^2_t, \mathbf{d}^2_t) + (1+\gamma)K \ ,$$

which is an analogous statement to (6). From these three results, we obtain

$$\begin{aligned} V_t(y^*) \quad &\ge \quad \pi(y^1_t, \mathbf{d}^1_t) - \gamma K \cdot \mathbf{1}[y^1_{t+1} > x^1_{t+1}] + \gamma V_{t+1}(y^1_{t+1}) \\ &\ge \quad \pi(y^2_t, \mathbf{d}^2_t) + (1+\gamma)K - \gamma K \cdot \mathbf{1}[y^2_{t+1} > x^2_{t+1}] + \gamma V_{t+1}(y^2_{t+1}) - \gamma^2 K \\ &> \quad \pi(y^2_t, \mathbf{d}^2_t) - \gamma K \cdot \mathbf{1}[y^2_{t+1} > x^2_{t+1}] + \gamma V_{t+1}(y^2_{t+1}) + K \end{aligned}$$

for any choice of y_t^2, $\mathbf{d}_t^2$ and y_{t+1}^2 in the y^2-system. Thus, we obtain (9). $\square$

We recall that Theorem 3.4 considers the infinite-horizon discounted-profit case. In this case, we are able to prove a stronger lower bound on s. The following result is used in establishing Proposition A.5.

Proposition A.4. *In Theorem 3.4, $S \geq \underline{m}$ holds.*

Proof. Suppose, by the way of contradiction, $S = y^\circ$ is strictly less than $\underline{m}$. We compare two systems starting with $y_1^1 = y^*$ and $y_1^2 = y^\circ < \underline{m}$. Let y^1-system follow order-up-to y^* and choose $\mathbf{d}^*$ in each period. The expected single-period profit of the y^1-system is $Q(y^*) - K$ in each period, except the first period in which it is $Q(y^*)$. Let y^2-system follow the optimal decision to attain $V(y^\circ)$. By Theorem 3.4, y_t^2 never exceeds $\underline{m}$ in each period t, and the y^2-system's expected single-period profit $\pi(y_t^2, \mathbf{d}_t^2)$ is bounded from above by $Q(y^*) - K$. Thus, we obtain $V(y^\circ) < V(y^*)$, a contradiction. $\square$

Proposition A.5. *The reorder point s in Theorem 3.4 is bounded below by $\underline{m}$.*

Proof. By Proposition A.4, $\underline{m}$ is a lower bound on y°, the optimal order-up-to level. It suffices to show that for each t, $y_t^2 < \underline{m}$ implies

$$V(y^\circ) \quad > \quad V(y_t^2) + K \ .$$

Without loss of generality, we proceed by assuming $t = 1$.

Let the y^2-system follow the optimal policy to achieve $V(y_1^2)$. Let $\tau = \min\{t \mid y_t^2 > x_t^2\}$ be the first period in which the y^2-system places an order (and incurs the fixed cost of K). By Theorem 3.4, $y_\tau^2 = y^\circ$. For each $t < \tau$, we have $x_t^2 \leq y_1^2 < \underline{m}$, and thus the expected single-period profit in period t is at most $Q(y_1^2) < Q(y^*) - K$. It follows

$$V(y_1^2) \quad < \quad \sum_{t=1}^{\tau-1} \gamma^{t-1}[Q(y^*) - K] \ + \ \gamma^{\tau-1}[V(y^\circ) - K] \ . \tag{10}$$

We make the following claim:

$$\frac{Q(y^*) - K}{1 - \gamma} \quad \leq \quad V(y^\circ) - K.$$

Consider a system which starts at y^* and orders up to y^* in every period. The infinite-horizon discounted profit of this system is given by

$$Q(y^*) + \sum_{t=2}^{\infty} \gamma^{t-1}(Q(y^*) - K) \;=\; K + \sum_{t=2}^{\infty} \gamma^{t-1}(Q(y^*) - K) \;=\; K + \frac{Q(y^*) - K}{1-\gamma} \; .$$

However, this cost is bounded above by $V(y^\circ)$ by the choice of y°. Thus we complete the proof of the above claim.

Now from (10) and the claim, we obtain

$$\begin{aligned} V(y_1^2) &\le \sum_{t=1}^{\tau-1} \gamma^{t-1}(1-\gamma)[V(y^\circ) - K] \;+\; \gamma^{\tau-1}[V(y^\circ) - K] \\ &= V(y^\circ) - K \; . \end{aligned}$$

This completes the proof. □

Now, $s \le y^*$ and $s_t \le y^*$ follows from Corollary 3.3. It remains to prove the following result.

Proposition A.6. *In Theorem 3.6, $S_t \ge y^*$ holds.*

Proof. It suffices to show that under Condition 2,

$$V_t(y^2) \;\; \le \;\; V_t(y^1) \quad \text{if} \quad y^2 \le y^* \text{ and } y^1 = y^*.$$

In each period $t' \ge t$, let $\mathbf{d}_{t'}^1$ of the y^1-system be as given by Condition 2 for any $\mathbf{d}_{t'}^2 \in \mathcal{D}$ of the y^2-system. Furthermore, whenever the y^2-system orders, let the y^1-system order up to $\max\{y^*, y_{t'}^2\}$ where $y_{t'}^2$ is the order-up-to level of the y^2-system. Then, in each period t', $y^2 \le y^1 \le y^*$ or $y^2 = y^1$. Thus, it can be shown that $\pi(y_{t'}^1, \mathbf{d}_{t'}^1) \ge \pi(y_{t'}^2, \mathbf{d}_{t'}^2)$, and that the ordering cost is no higher in the y^1-system than in the y^2-system. □

A.3 The Chen et al. (2006) Model: Proof of Proposition 5.1

We introduce the following notation and definitions used in Chen et al. (2006). Let $z := y - d(p)$ be the "riskless" leftover inventory at the end of a period. Let $G(z, p)$ be the expected single-period profit without considering the fixed ordering cost, i.e., $G(z, p) = \pi(z + d(p), p)$.[10] Let $P(y)$ be the

[10] In this section, p corresponds to the sales lever $\mathbf{d}$ used in earlier sections.

optimal price when the starting inventory level after ordering is y, i.e., $P(y) := \arg\max_p \pi(y,p)$. Thus, $\pi(y, P(y)) = Q(y)$. Let $Z(y) := y - d(P(y))$. It follows that

$$G(Z(y), P(y)) \;=\; Q(y) \;=\; \pi(y, P(y)). \tag{11}$$

Also, let $p(z) := \arg\max_p G(z,p)$. We denote the maximizer of $G(z, p(z))$ by Z. Then, the single-period optimal stocking quantity satisfies $y^* = Z + d(p(Z))$.

Chen et al. (2006) give the following list of properties based on Assumption 3. They are related to the expected single-period profit function.

Fact 1 *(Equation (6)).*

$$\frac{\partial^2 G(z,p)}{\partial p^2} < 0.$$

Fact 2 *(Lemma 1 and Its Corollary).* $p(z)$ is continuous on $[0, +\infty)$. Furthermore, $p(z)$ is increasing on $[0, \infty)$.

Fact 3 *(Lemma 3).* $Z(y)$ is nonnegative and increasing on $[0, +\infty)$.

Fact 4 *(Theorem 1).* $P(y) \le p(z)$ for $y > y^*$, and $P(y) \ge p(z)$ for $y < y^*$ where $y = z + d(p(z))$. Furthermore, $G(Z(y), P(y))$ is unimodal on $[0, +\infty)$ and y^* is its maximizer.

Proof of Proposition 5.1

Fact 4 shows the quasi-concavity of $Q(y) = G(Z(y), P(y))$, implying Condition 1 (a).

<u>Case $y^* \le y^1 < y^2$.</u> Suppose that a pair of y^1 and y^2 satisfy $y^* \le y^1 < y^2$, and p^2 is given. Let $z^2 := y^2 - d(p^2)$. By (11), Fact 4, and the definition of $P(y)$,

$$\pi(y^1, P(y^1)) \;=\; G(Z(y^1), P(y^1)) \;\ge\; G(Z(y^2), P(y^2)) \;=\; \pi(y^2, P(y^2)) \;\ge\; \pi(y^2, p^2).$$

Thus, if $Z(y^1) \le z^2$ holds, $p^1 = P(y^1)$ satisfies Condition 2 (b), i.e., (2). As a result, we proceed by supposing

$$Z(y^1) \;>\; z^2. \tag{12}$$

We consider the following two cases: $z^2 + d(p(z^2)) \le y^1$ and $z^2 + d(p(z^2)) > y^1$.

Case $z^2 + d(p(z^2)) \le y^1$: Choose p^3 such that $d(p^3) = y^1 - z^2$. (This is possible because $d(\cdot)$ is continuous and $y^1 - z^2$ is bounded below by $d(p(z^2))$ and bounded above by $y^2 - z^2$, which is the same as $d(p^2)$.) Since $z^2 + d(p^3) = y^1$ and $z^2 + d(p^2) = y^2$, it follows $y^1 - d(p^3) = z^2 = y^2 - d(p^2)$. We also obtain that

$$z^2 + d(p(z^2)) \;\le\; y^1 \;<\; y^2 \;,$$

which is equivalent to

$$d(p(z^2)) \;\le\; d(p^3) \;<\; d(p^2).$$

Therefore we obtain $p(z^2) \ge p^3 > p^2$. Now, the definition of $p(z^2) = \arg\max_p G(z^2, p)$ and Fact 1 imply that $G(z^2, p)$ is increasing in p for $p \le p(z^2)$. Thus,

$$\pi(y^1, p^3) \;=\; G(z^2, p^3) \;\ge\; G(z^2, p^2) \;=\; \pi(y^2, p^2).$$

Thus, Condition 2 (b) is satisfied with p^3.

Case $z^2 + d(p(z^2)) > y^1$: If $Z(y^1) \le y^*$, then Condition 2 (b) holds with $p^1 = P(y^1)$. We proceed by assuming otherwise, i.e., $Z(y^1) > y^*$.

Let $y' := z^2 + d(p(z^2)) > y^1$. Then, $y' > y^1 \ge Z(y^1) > y^*$. By Fact 4, we obtain $P(y') \le p(z^2)$, which implies

$$z^2 \;=\; y' - d(p(z^2)) \;\ge\; y' - d(P(y')) \;=\; Z(y') \;.$$

Thus, by (12), $Z(y^1) > z^2 \ge Z(y')$, which contradicts the monotonicity of Z (Fact 3) since $y^1 < y'$. Thus, this case cannot happen.

<u>CASE $y^* \ge y^1 > y^2$.</u> Consider a pair y^1 and y^2 such that $y^2 < y^1 \le y^*$. For any p^2, let $z^2 := y^2 - d(p^2)$. We will show that there exist p^1 such that (i) $z^1 := y^1 - d(p^1) \ge z^2$, and (ii) $G(z^1, p^1) \ge G(z^2, p^2)$.

Recall $Z(y^1) = y^1 - d(P(y^1))$. If $z^2 \le Z(y^1)$, then $p^1 = P(y^1)$ satisfies (i). Furthermore, the quasi-concavity of Q implies

$$\pi(y^1, P(y^1)) \;=\; Q(y^1) \;\ge\; Q(y^2) \;=\; \pi(y^2, P(y^2)) \;\ge\; \pi(y^2, p^2),$$

which shows (ii). Thus, we proceed by assuming

$$Z(y^1) < z^2. \tag{13}$$

Let $y' := z^2 + d(p(z^2))$.

Case $y' \geq y^1$: We choose p^1 such that $y^1 = z^2 + d(p^1)$. (Note that such a p^1 can be chosen because $d(p(\cdot))$ is a continuous function, and $y^1 - z^2$ is bounded below by $d(p^2) = y^2 - z^2$ and above by $d(P(y^1)) = y^1 - Z(y^1)$.) Then, $z^2 = y^1 - d(p^1)$ shows (i). Now, note that

$$y^2 = z^2 + d(p^2) \ < \ y^1 = z^2 + d(p^1) \ \leq \ y' = z^2 + d(p(z^2))$$

implies $p^2 > p^1 \geq p(z^2)$, and $p = p(z^2)$ maximizes $G(z^2, p)$. Thus, Fact 1 shows $\pi(y^1, p^1) = G(z^2, p^1) \geq G(z^2, p^2) = \pi(y^2, p^2)$, verifying (ii).

Case $y' < y^1$: Recall $z^2 + d(p(z^2)) = y'$. Also, by the monotonicity of Z (Fact 3) and (13), we have $Z(y') \leq Z(y^1) < z^2$. Thus,

$$d(p(z^2)) \ = \ y' - z^2 \ < \ y' - Z(y') \ = \ d(P(y')) \ ,$$

implying $p(z^2) > P(y')$. Then, by Fact 4, $y' > y^*$ must hold. However, it contradicts the assumption $y' < y^1 \leq y^*$. Thus, this case cannot happen.

A.4 The Case of Stochastically Increasing, Additive Demands

All the results we have shown in this paper so far assume that all cost parameters and the demand distributions are stationary. In this section, we extend our results to a special class of non-stationary problems. We use the Additive Demand Model in this section. Due to non-stationarity, we will use the time index as a subscript for all parameters.

Theorem A.7. *In the finite-horizon model, suppose the following conditions hold:*

(a) Demand is additive.

(b) $\{y_t^*\}$ is an increasing sequence.[11]

(c) $K_t \ \geq \ \gamma \cdot K_{t+1}$ holds for every t.

(d) Either (i) excess demand is backordered and Assumption 2 holds for every t; or, (ii) excess demand is lost and Assumption 3 holds for every t.[12]

[11] In models with stationary cost parameters, this assumption usually holds when $\{\epsilon_t\}$ is a sequence of stochastically increasing random variables.

[12] In (i), Assumption 2 now involves π_t, ϕ_t^R and ϕ_t^H. In (ii), Assumption 3 now involves $d_t(p)$, ϵ_t, B_t, f_t, F_t, r_t, π_t, h_t and b_t.

Then, an (s_t, S_t) policy is optimal in each period t.

Proof. The proof of this theorem is similar to the stationary case presented in the paper, and we provide a high-level sketch of this proof. In the backordering case, an argument similar to Proposition 4.3 shows modified versions of Condition 1 and Condition 2, where y^* is replaced with y_t^*. In the lost sales case, likewise, Proposition 5.1 can be adapted to show the same result. Then, the proof of Theorem 3.6 can be adjusted to show the required result using the modified versions of Condition 1 and Condition 2. $\square$