Probability in the Engineering and Informational Sciences, **20**, 2006, 287–306. Printed in the U.S.A.

OPTIMIZATION AND OPTIMALITY OF (s,S) STOCHASTIC INVENTORY SYSTEMS WITH NON-QUASICONVEX COSTS

FRANK Y. CHEN AND Y. FENG
Department of Systems Engineering & Engineering Management
The Chinese University of Hong Kong
Shatin, Hong Kong, China
E-mail: yhchen@se.cuhk.edu.hk; yyfeng@se.cuhk.edu.hk

This article considers the optimization and optimality of single-item/location, infinite-horizon, (s,S) inventory models. Departing from the conventional approach, we do not assume the loss function describing holding and shortage costs per period to be quasiconvex. As the existing optimization algorithms have been established on the condition of quasiconvexity, our goal in this article is to develop a computational procedure for obtaining optimal (s,S) policies for models with general loss functions. Our algorithm is based on the parametric method commonly used in fractional programming and is intuitive, exact, and efficient. Moreover, this method allows us to extend the optimality of (s,S) policies to a broader class of loss functions that can be non-quasiconvex.

1. INTRODUCTION

The (s,S) type of replenishment policy, also known as the Min-Max control rule, is widely implemented in single-item/location inventory systems. An (s,S) policy instructs that whenever the inventory position (on-hand stock plus outstanding orders minus any backorders, at review) drops to s, which is the reorder level, or below, an order is placed to raise it back up to S, which is the order-up-to level.

In the evaluation of any inventory models, including (s,S) systems, generally two cost components are considered: the ordering cost and the holding-shortage cost, where the second component is often referred to as the loss function. The

© 2006 Cambridge University Press 0269-9648/06 $16.00

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

optimization of an (s, S) system means finding a pair of s and S parameters to minimize the long-run average cost (per period) or discounted total cost. An assumption that is commonly made about (s, S) inventory models is that the expected loss function is convex or quasiconvex in the inventory position. (Quasiconvexity means the existence of an inventory position above which the function is nondecreasing and below which it is nonincreasing.) Most of the research that has been conducted on optimization is grounded on such a condition.

However, this condition is unwarranted in general, and when it is violated, the existing approaches are inoperative. Violation likely occurs when the shortage cost is nonlinear in the inventory positions; for example, when each unit that is backordered incurs a fixed cost, regardless of how long the unit remains in a backlog, or when a fixed cost is incurred whenever there is any backlog on the books in the period, regardless of the volume of the backlog that accumulates. These two types of cost have long been recognized as the shadow costs of fill-rate and ready-rate service constraint, respectively. The fill rate is defined as the fraction of demand that is fulfilled on time, and the ready rate measures the portion of periods during which no shortage occurs. In Section 2.3 we provide an example in which the loss function is non-quasiconvex (see Fig. 1).

This article aims to address models with a non-quasiconvex loss function. We have two objectives: first, to identify the conditions under which the (s, S) policy remains globally optimal even in a model with a non-quasiconvex loss function, and, second, to develop a computational procedure for optimizing (s, S) policies in such a generalized model. To fulfill the first objective, we extend the optimality result for the (s, S) policy by identifying a class of non-quasiconvex loss functions that can retain the optimality of (s, S) policies among all policies. Even though the (s, S)-type policy might be suboptimal for models with general loss functions, it is practically adequate because of its simplicity (see Zipkin [21] and Ballou [1]). Therefore, the optimization algorithm alone is of significance. To fulfill the second objective, we propose a descent-search algorithm based on the fractional programming approach that transforms the original model into a parametric problem through an auxiliary function. Emulating a cutting method, the algorithm based on the fractional programming approach is intuitive and easy to initialize and finds the optimal solution efficiently.

The auxiliary function measures the difference between the cost from an (s, S) policy and an artificial, or dummy, cost. The algorithm carries out a search in a three-dimensional space (γ, s, S), where γ denotes the dummy cost. It commences with an easily obtainable value of γ, which is also an upper bound of the optimal long-run average cost $C_{\min}$. At each dummy cost level γ, the auxiliary function is evaluated at a few reorder points and various order-up-to levels. When the sign of the auxiliary function becomes negative, the dummy cost is updated. As the dummy cost is lowered to $C_{\min}$, both s and S reach their optimal values, and the search ends.

There is a rich body of literature on single-item (s, S) inventory control systems. A significant result from inventory theory is the optimality of (s, S) policies in various models. Using the concept of K-convexity, Scarf [13] provided the first

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

optimality result in a finite horizon setting with setup costs. Iglehart [9] extended this result to an infinite horizon setting. These two seminal articles assumed that the loss function is convex in inventory positions. Veinott [17] adopted a different approach and obtained the optimality result for both finite and infinite horizon problems under a different set of assumptions: The loss function is quasiconvex, but demand over time is stochastically nondecreasing. This article thus generalizes the optimality results that are established by Iglehart [9] and Veinott [17] for the infinite horizon setting. Many authors have extended the previous results to different problem settings (e.g., Song and Zipkin [16] and Sethi and Cheng [14] in a Markovian demand environment, and Li, Xu, and Hayya [10] in a random supply context). Since the study of Veinott and Wagner [18], the optimization of infinite horizon, stationary (s, S) models has aroused much interest in the literature. Zheng and Federgruen [20] (ZF hereafter) have developed the best solution procedure so far for cases with quasiconvex loss functions. More recently, Feng and Xiao [7] (FX hereafter) proposed an alternative to ZF's. The method in the present article represents a further generalization of these two algorithms.

The rest of the article is organized as follows. Section 2 specifies the inventory model and shows an example in which a loss function is non-quasiconvex. Section 3 explores the properties of the method and proposes the algorithm. We address the optimality of (s, S) policies with non-quasiconvex loss functions in Section 4. Section 5 reports numerical examples, and Section 6 summarizes the article. A list of symbols is included in Section 2 for ease of reading.

2. THE INVENTORY MODEL AND LOSS FUNCTIONS

2.1. Notation

Symbols that are used throughout the sections are summarized here, whereas those that are used briefly will be defined locally.

x, y = inventory positions at the beginning of any period, before or after possible ordering

$G(y)$ = the loss function, which is the sum of the expected holding-shortage costs, which are conditional on the current inventory position y

K = the setup cost

$\Phi(\cdot)$ = probability distribution of demand in a period, with density $\phi_j = P[D = j]$, and the nth convolution of $\Phi(\cdot)$ (or ϕ_j) = $\Phi^n(\cdot)$ (or ϕ_j^n)

$m(j)$ = the renewal density, or the expected duration for which the cumulative demand is j: $m(0) = 1/(1 - \phi_0), \ldots, m(j) = \sum_{i=0}^{j} \phi_i \, m(j - i)$,

$M(x) = \sum_{i=0}^{x-1} m(i)$, where $x = 1, 2, \ldots$ and $M(0) = 0$

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

$C(s,S)$ = the long-run average cost of a given (s,S) policy

$C_{\min}$ = the minimum of $C(s, S)$ among all of the (s, S) policies

y_0 = a global minimizer of $G(y)$

γ = the dummy cost, and $\gamma^* = C_{\min}$

$l_\gamma(s,S)$ = an auxiliary function that relates to γ and the (s, S) policy

(s_γ, S_γ) = the optimal (s, S) policy that is associated with γ, which minimizes l_γ (s, S)

(s^*, S^*) = an optimal (s, S) policy (among the (s, S) class) (i.e., $(s_{\gamma^*}, S_{\gamma^*})$)

s_γ^i = any number that satisfies $\{y : G(y+1) \le \gamma < G(y)\}$

$P(\gamma)$ = the collection of s_γ^i, $i = 1,2,\ldots$, arranged in ascending order

S^γ = $\max\{y : G(y) \le \gamma\}$

2.2. A Periodic Review Inventory System

We consider an infinite horizon, periodic-review inventory system that faces discrete-valued independent and identically distributed (i.i.d.) demands. Demand that is not met immediately is fully backlogged and will be fulfilled on a first-come–first-served basis. The system replenishes inventory according to an (s,S) policy with a constant lead time of L periods. (The random lead time will be introduced after the basic model has been formulated.) The sequence of events in each period is as follows: The inventory is reviewed, an order that was placed L periods ago is received, an ordering decision is made, and the demand in the period is then subtracted out of the inventory. These events are supposed to happen at the beginning of each period.

It is well known (e.g., Veinott and Wagner [18], and ZF [20]) that

$$C(s,S) = \frac{K + \sum_{l=0}^{S-s-1} m(y)G(S-l)}{M(S-s)}.$$

For ease of analysis, we impose a minor condition on $G(y)$, which is also assumed in most of the previous work:

Assumption 1: $\lim_{y\to\pm\infty} G(y) > K + G(y_0)$.

It should be noted that when L is a random variable, $G(y)$ should be understood as the expected loss function over demand and the lead time

$$G(y) = E_L G(y|L).$$

Under the conditions that are explicated by Zipkin [21], the distribution of demand over the deterministic lead time L is just replaced by the distribution of total demand over the stochastic lead time. The conditions are met when the system faces an

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

exogenous, sequential supply source in which the lead times are i.i.d. over time and no order-crossing occurs.

We also assume that demand takes a discrete value, but that continuous distributions can be discretized with an arbitrary precision.

2.3. An Example of Non-quasiconvex Loss Function

In a recent article, Rosling [11] considered loss functions with nonlinear shortage costs. His basic model is a periodic-review model that is almost identical to ours; in his model, two nonlinear shortage (or backorder) costs are added to the traditional linear shortage cost. Here, we only give their mathematical expressions and refer the reader to Rosling [11] for a further explanation.

Suppose that the inventory position after possible ordering is y; then

$$G(y) = pB_p(y) + \pi B_\pi(y) + bB_b(y) + H(y), \tag{1}$$

where $p, \pi, b \geq 0$; $B_p(y) = \sum_{\xi=y}^{\infty}[1 - \Phi^{L+1}(\xi)]$, which is the shortage cost that accumulates at a constant rate proportional to the number of backorders on the books; $B_\pi(y) = \sum_{\xi=y}^{\infty}[\Phi^{L}(\xi) - \Phi^{L+1}(\xi)]$, which is the cost that occurs only for the incremental backorders that arise in a period (Hadley and Whitin [8], Chen and Zheng [3], and Bashyam and Fu [2]); $B_b(y) = 1 - \Phi^{L+1}(\xi)$, which is the cost that is incurred whenever there are any backorders in a period (Silver, Pyke, and Peterson [15]); and $H(y)$ is the expected holding cost (per period). (For brevity, we simply denote $G(y|L)$ by $G(y)$.) Setting $\pi = b = 0$, we have a loss function with a traditional linear backorder cost. Generally, b and π are viewed as Lagrangian multipliers, or shadow costs, of the ready and fill rates, respectively.

Assuming that the holding cost incurs at a constant rate and that the lead time for replenishment orders is constant, Rosling [11] showed that when $p, \pi, b \geq 0$, $G(y)$ is quasiconvex if $\Phi(\xi)$ is log-concave. This condition is met by many demand distributions, including the Normal, Poisson, Negative Binomial, gamma family, Weibull, and beta with parameters (r, q) such that both r and $q \geq 1$. However, there are significant exceptions. Take an example that is given in Eppen and Martin [5]: The demand of each time period is normally distributed with mean 20 units and variance 10 $(\text{units})^2$. The lead time (in periods) can take on the values 2 and 4 with probabilities of 0.8 and 0.2, respectively. In such a case, the (expected) loss function fails to be quasiconvex when the backlogging cost is nonlinear. Figure 1 illustrates a case in which $p = \$2/\text{unit/period}$, $\pi = 0$, $b = \$200/\text{period}$, and the holding cost rate is $\$1/\text{unit/period}$. (The line that is associated with $C_{\min}$ will be explained later.)

Although the loss function can be non-quasiconvex in the inventory position, $G(y)$ generally increases monotonically after y becomes sufficiently large, and it is nonincreasing when $y < 0$. For example, a test can be carried to see whether $G(y)$ will become monotonic. Suppose that $H(y)$ is convex. If $H(y_1 + 1) - H(y_1) \geq pB_p(y_1) + \pi B_\pi(y_1) + bB_b(y_1)$ (see (1)), then $G(y)$ is increasing for all $y > y_1$.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

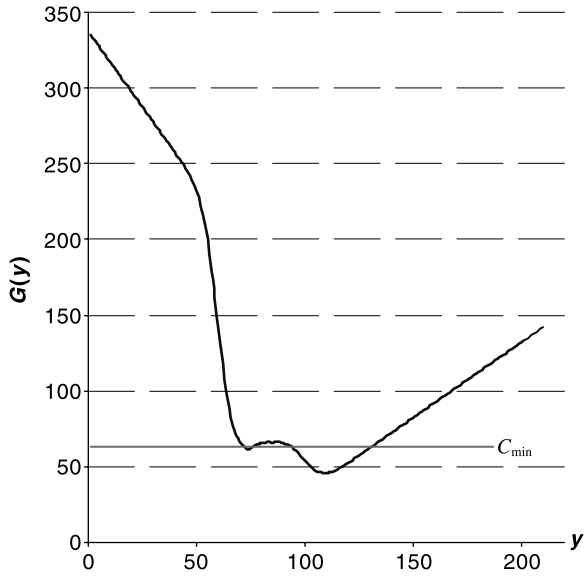


FIGURE 1. Non-quasiconvexity and random lead time.

3. FINDING AN OPTIMAL (*s, S*) POLICY

When the loss function $G(y)$ is quasiconvex, one can apply the algorithm developed by ZF [20] or by FX [7] to search for an optimal (s, S) policy. These algorithms are exact and efficient. Here, our focus is on a model with a non-quasiconvex loss function, and we will discuss whether it is possible to extend ZF's algorithm to the current setting later.

We will first reveal the properties of an optimal (s, S) policy and then derive an algorithm for its computation.

3.1. Auxiliary Function Method and Preliminary Results

To elaborate our computational procedure, we introduce a *dummy cost* γ, which is assumed to be in the range $\gamma \in A = \{\gamma : G(y_0) \leq \gamma \leq K + G(y_0)\}$. As will be shown, the assumption about γ can be trivially met (see Section 3.4). We then define an auxiliary function that relates to this γ as follows:

$$l_\gamma(s, S) = K + \sum_{j=0}^{S-s-1} m(j)[G(S-j) - \gamma]. \tag{2}$$

In fractional programming, the problem of minimizing $l_\gamma(s, S)$ over $s < S$ is called a parametric problem. The use of such an approach in inventory theory was first proposed by Feng and Sun [6].

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

One might think of the dummy cost γ as a budget level per period. Because $C(s,S)\cdot M(S-s) = K + \sum_{y=0}^{S-s-1} m(y)G(S-y)$ is the expected total cost per ordering cycle, the difference

$$l_\gamma(s,S) = (C(s,S) - \gamma)M(S-s)$$

gives the cumulative deficit in operating the system per renewal cycle; that is, if the budgeted cost exceeds the expected operating cost, then there will be a surplus, and if the expected operating cost exceeds the budgeted cost, then there will be a fiscal shortfall. The objective is therefore converted to one that minimizes $l_\gamma(s,S)$.

Our interest in the auxiliary function is due to the fundamental properties of fractional programming regarding the value γ^* and the function w_γ, which are defined respectively as $\gamma^* = \{\gamma \in A : \min_{s\le S} l_\gamma(s,S) = 0\}$ and $w_\gamma = \min_{s<S} l_\gamma(s,S)$. In what follows, we denote by s_γ and S_γ the pair that minimize $l_\gamma(s,S)$ under a given γ. We also term (s_γ, S_γ) optimal with respect to the auxiliary function; that is,

$$l_\gamma(s_\gamma, S_\gamma) = \min_{s<S} l_\gamma(s,S) = w_\gamma.$$

LEMMA 1:

(a) The minimum achievable cost $C_{\min} = \gamma^*$, and the function $w(\gamma)$ is strictly decreasing in γ.

(b) For a given $\gamma \ge G(y_0)$, if a pair of s and S exists such that $l_\gamma(s,S) < 0$, then γ is an upper bound of $C_{\min}$.

(c) Defining

$$\gamma_1 = \gamma + \frac{l_\gamma(s,S)}{M(S-s)},$$

$C_{\min} \le \gamma_1 < \gamma$, $\gamma_1 = C(s,S)$, and $l_{\gamma_1}(s,S) = 0$. In particular, $C_{\min} = \gamma^* = C(s_{\gamma^*}, S_{\gamma^*})$.

In view of Lemma 1, the minimization of $C(s,S)$ is equivalent to finding a $\gamma \ge G(y_0)$, such that $\min_{s<S} l_\gamma(s,S) = 0$. Specifically, we begin with an initial dummy cost $\gamma_0 \ge C_{\min}$ and then iteratively lower it until the condition $\min_{s<S} l_\gamma(s,S) = 0$ is satisfied.

For later use, we rearrange the auxiliary function

$$\begin{aligned} l_\gamma(s,S) &= K + \sum_{y=1}^{S-s-1} \sum_{j=0}^{y} \phi_j m(y-j)[G(S-y) - \gamma] + m(0)[G(S) - \gamma] \\ &= K\left(1 - \sum_{j=0}^{S-s-1} \phi_j\right) + \sum_{j=0}^{S-s-1} \phi_j l_\gamma(s, S-j) + G(S) - \gamma. \end{aligned} \tag{3}$$

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

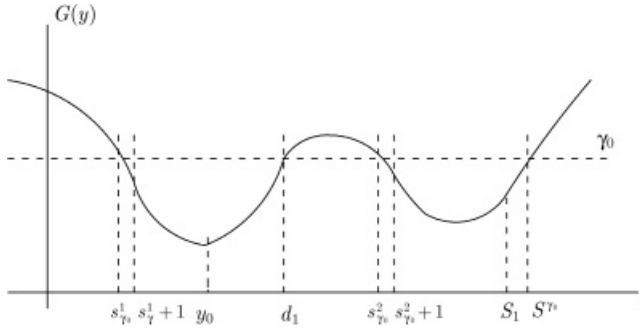


FIGURE 2. An illustration of candidates for s_{γ_0} and ranges for S_{γ_0}.

As y_0 is a global minimizer of $G(y)$ given a $\gamma \geq G(y_0)$, it follows that

$$l_\gamma(s_\gamma, S_\gamma) \leq l_\gamma(y_0 - 1, y_0) = K + m(0)[G(y_0) - \gamma] \leq K. \tag{4}$$

Let $P(\gamma) = \{y : G(y+1) \leq \gamma < G(y)\}$. In Figure 2 (where the curve is smoothed and $\gamma = \gamma_0$), there are two elements in $P(\gamma_0)$: $s^1_{\gamma_0}$ and $s^2_{\gamma_0}$. Lemma 2 shows that $P(\gamma)$ is the set that contains all potential candidates for s_γ. If we denote these candidates as $s^1_\gamma, s^2_\gamma, \ldots$, in ascending order of value, then $P(\gamma) = \{s^1_\gamma, s^2_\gamma, \ldots\}$.

LEMMA 2:

(a) For any given S such that $G(S) < \gamma$, the optimal reorder level s_γ satisfies

$$G(s_\gamma + 1) \leq \gamma < G(s_\gamma).$$

(b) For a given γ and $s_\gamma \in P(\gamma)$, the optimal order-up-to level $S_\gamma \leq S^\gamma$, where $S^\gamma = \max\{y : G(y) \leq \gamma\}$.

PROOF: Part (a) follows from the fact that

$$l_\gamma(s, S) - l_\gamma(s+1, S) = [G(s+1) - \gamma]m(S - s - 1).$$

For part (b), we assume to the contrary that $G(S_\gamma) > \gamma$. Then, following (3),

$$l_\gamma(s_\gamma, S_\gamma) > K\left(1 - \sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j\right) + \sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j l_\gamma(s_\gamma, S_\gamma),$$

which implies that

$$\left(1 - \sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j\right) l_\gamma(s_\gamma, S_\gamma) > K\left(1 - \sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j\right).$$

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

If $\sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j = 1$, then we have $0 > 0$, which is an impossible event. Thus, $\sum_{j=0}^{S_\gamma - s_\gamma - 1} \phi_j < 1$, leading to

$$l_\gamma(s_\gamma, S_\gamma) > K,$$

which is a contradiction of (4). Thus, the lemma follows. ■

Lemma 2 implies that $s_\gamma \in P(\gamma)$. Because of Assumption 1 and the assumption that $G(y_0) \le \gamma \le K + G(y_0)$, the set $P(\gamma)$ is clearly nonempty and finite. It can easily be seen that the lower bound of s^* is $s^1_{\gamma_0}$. Furthermore, not every value of $y \le S^\gamma$ can be an effective order-up-to level. For example, in Figure 2, it can be seen that for an (s^i_γ, S) to be optimal, $y_0 \le S \le d_1$ or $S > s^2_\gamma$, where $d_1 = \min\{y > s^1_\gamma : G(y) \le \gamma\}$ and $i = 1$ or 2. There are more unattractive candidates for S_γ that can be eliminated. However, for notational simplicity, we will compromise some efficiency by ignoring these properties in the sequel algorithm design.

Lemma 2 represents the major extension of the work of FX. In the rest of the section, we discuss a method for the upgrading of order-up-to levels and when to end the search. Interestingly, the properties that are revealed in these discussions are direct extensions of those of FX and in fact these discussions do not engage the loss function.

We close this subsection with a brief comparison of the average-cost approach of ZF and our auxiliary-function approach. By Eq. (6) in ZF and by the proof of Lemma 1 in Zheng [19], one can derive a lower bound for s^* and an upper bound for S^*, respectively. If their algorithm starts with an initial policy that gives γ_0 and ends with the lowest average cost $C_{\min}$ being identified, then these two bounds coincide with our $s^1_{\gamma_0}$ and $S^{C_{\min}}$. However, a full enumeration within these bounds is required in their approach, whereas our approach allows a more efficient search. The difference arises from the fact that they directly use the cost function $C(s,S)$ to design their algorithm, whereas we use the auxiliary function $l_\gamma(s,S)$. Although both functions vary with the policy (s,S), $l_\gamma(s,S)$ has a simpler form.

3.2. Updating Order-up-to Levels

In this subsection, we discuss a method for updating the optimal order-up-to level in conjunction with the upgrading of the dummy cost. We use Figure 2 to demonstrate how the dummy cost will be upgraded in our algorithm. Under $\gamma = \gamma_0$, suppose that for all $S < S_1$, $l_\gamma(s^i_\gamma, S) \ge 0$, $i = 1$ or 2, but $l_\gamma(s^1_\gamma, S_1) < 0$ (or $l_\gamma(s^2_\gamma, S_1) < 0$). We then immediately upgrade the dummy cost from γ to $\gamma_1 = C(s^1_\gamma, S_1)$ (or $C(s^2_\gamma, S_1)$). In the sequel, when we say that the *algorithm updates the dummy cost* γ *at pair* (s_1, S_1), we mean that under γ, for all the potential (s,S) policies with $s < S < S_1$, $l_\gamma(s,S) \ge 0$, but by raising S to S_1, a pair (s_1, S_1) emerges such that $l_\gamma(s_1, S_1) < 0$, where $s_1 < S_1, s_1 \in P(\gamma)$ and $S_1 \le S^\gamma$. This is how and when the dummy cost γ is upgraded in the proposed algorithm: As soon as $l_\gamma(s,S)$ turns negative, the dummy cost is lowered from γ to γ_1. The following discussion is based on such an upgrading scheme.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

Now suppose that the dummy cost γ is upgraded to γ_1 by a given policy (s_1, S_1), where $s_1 \in P(\gamma)$, and a new set $P(\gamma_1)$ is formed as $s^1_{\gamma_1} < s^2_{\gamma_1} < \cdots$. Do we need to evaluate all $(s^i_{\gamma_1}, S_1 - j)$ policies, where $j = 1, 2, \ldots$? The answer is no, which is deduced from the following lemma.

LEMMA 3: *Suppose that γ is upgraded to $\gamma_1 < \gamma$ by a policy (s_1, S_1), where $s_1 \in P(\gamma)$, $S_1 < S^{\gamma}$, and*

$$l_{\gamma}(s_1, S_1) < 0.$$

Then for any $s^i_{\gamma_1} \in P(\gamma_1)$ and $y < S_1$, we have

$$l_{\gamma_1}(s^i_{\gamma_1}, y) \geq 0.$$

PROOF: Apparently, given $s_1 < S_1$, $l_{\gamma}(s_1, S_1)$, as a function of $\gamma > 0$, strictly decreases in γ. As a result, for $s^i_{\gamma_1} \in P(\gamma_1)$ and $y < S_1$,

$$l_{\gamma_1}(s^i_{\gamma_1}, y) > l_{\gamma}(s^i_{\gamma_1}, y) \geq \min_{s^j_{\gamma} \in P(\gamma),\, y < S_1} l_{\gamma}(s^j_{\gamma}, y) \geq 0.$$

The last inequality results from the updating scheme of the dummy cost that was described earlier. ■

3.3. Identifying $C_{\min}$ and an Optimal (s, S) Policy

Once γ is upgraded to $\gamma^* = C_{\min}$, the algorithm needs to be terminated and an optimal (s, S) policy identified. This subsection discusses when the algorithm should terminate.

LEMMA 4: *Suppose that the algorithm upgrades the dummy cost γ at (s_1, S_1) to $\gamma_1 = C(s_1, S_1)$ where $s_1 \in P(\gamma)$. In addition, suppose that for all $s \in P(\gamma_1)$ and $S_1 \leq y \leq S^{\gamma_1}$,*

$$l_{\gamma_1}(s, y) \geq 0.$$

Then, $\gamma^* = C_{\min} = \gamma_1$, $s^* = s_{\gamma^*} = s_1$, and $S^* = S_{\gamma^*} = S_1$.

PROOF: In view of Lemma 3, for a case in which $s \in P(\gamma_1)$ and $y < S_1$, $l_{\gamma_1}(s, y) \geq 0$. With these assumptions, we have

$$\min_{s \in P(\gamma_1)} \min_{s < y \leq S^{\gamma_1}} l_{\gamma_1}(s, y) \geq 0$$

and, thus, $\gamma_1 \leq C_{\min}$. However, as $\gamma_1 = C(s_1, S_1) \geq C_{\min} = \gamma^*$, we have $\gamma_1 = \gamma^*$. Furthermore, as $l_{\gamma^*}(s_1, S_1) = 0$, we can conclude that $s^* = s_{\gamma^*} = s_1$ and $S^* = S_{\gamma^*} = S_1$. ■

This lemma indicates that in the final γ upgrading attempt, the search will continue until S^{γ^*}, which is the largest order-up-to level to be evaluated, is reached.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

3.4. The Algorithm

We now present the descent-search algorithm, which starts with an initial dummy cost $\gamma_0 = C(s,S) \leq K + G(y_0)$. This kind of dummy cost can be easily obtained. For instance, when $s = y_0 - 1$ and $S = y_0$,

$$C(y_0 - 1, y_0) = \frac{K + m(0)G(y_0)}{m(0)} = \frac{K}{m(0)} + G(y_0) \leq K + G(y_0)$$

can be selected as γ_0. It is generally easy to locate $s^1_{\gamma_0}$; for example, we can start with $y = 0$. If $G(0) > \gamma_0$, then $s^1_{\gamma_0} \geq 0$, and if $G(0) \leq \gamma_0$, then $s^1_{\gamma_0} < 0$. This is because $G(y)$ is decreasing when $y < 0$. We assume that a number $\bar{y}$ is also available so that for $y \geq \bar{y}$ (see the remark at the end of Section 2.3), $G(y)$ is increasing. In finding $\bar{y}$, a global minimizer of $G(y)$, y_0, also becomes known. In the algorithm, we do not describe how S^γ is computed, because such a procedure is straightforward. The algorithm will be called the *descent-search algorithm*, because the dummy cost is lowered after each iteration.

Descent-search algorithm

- **Step 0 (Initialization)**
  Let $\gamma_1 = \gamma_0, \gamma^* = \gamma_1, I = \emptyset$, and $C_{\min} = \gamma_1$. Compute $s^1_{\gamma_1}$ and add $s^1_{\gamma_1}$ into set I. Let $s = s^1_{\gamma_1}$, and $S = s + 1$. $s^* = s$, and $S^* = S$. Go to Step 1.1.
- **Step 1**
  1.1. For all $s \in I$, compute $l_{\gamma_1}(s,S)$ and $\min_{s\in I} l_{\gamma_1}(s,S)$. If $\min_{s\in I} l_{\gamma_1}(s,S) < 0$ for $s = s' \in I$, then go to Step 2.1. Otherwise (i.e., $\min_{s\in I} l_{\gamma_1}(s,S) \geq 0$), go to Step 1.2.
  1.2. Let $S := S + 1$. If $S > S^{\gamma_1}$, then go to Step 3. If $G(S-1) > \gamma_1 \geq G(S)$, then add $S - 1$ into I. Go to Step 1.1.
- **Step 2 (Updating the dummy cost)**
  2.1. Update γ_1 by computing

$$\gamma_1 := \frac{l_{\gamma_1}(s',S)}{M(S - \bar{s})} + \gamma_1;$$

  let $s^* = \bar{s}$ and $S^* = S$. Go to Step 2.2.
  2.2. Let $S := S^*$ and compute $s^1_{\gamma_1}$. Set $I = \emptyset$ and add $s^1_{\gamma_1}$ into I. Go to Step 1.1.
- **Step 3 (Termination)**
  Stop with $C_{\min} = \gamma^*$ and the optimizing policy (s^*, S^*).

The algorithm works as follows. It starts with γ_0. The set of potential reorder levels associated with the dummy cost γ_0 consists of $s^i_{\gamma_0}$, $i = 1,2,\ldots,\tau$. The set I contains the levels in $P(\gamma_0)$ that are currently used for evaluation. At the very beginning, the set I contains just one reorder level $s = s^1_{\gamma_0}$, and the algorithm begins with the computation of $l_{\gamma_0}(s^1_{\gamma_0}, s^1_{\gamma_0} + 1)$.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

Step 1.1 evaluates all $l_{\gamma_1}(s,S)$, with S being fixed but s being taken from I. The minimum of $l_{\gamma_1}(s,S)$ and $l_{\gamma_1}(s',S)$ is then obtained. Once this minimum becomes negative, the algorithm immediately upgrades γ_1 in Step 2 by the (s',S) policy.

If the minimum of $l_{\gamma_1}(\ )$ is nonnegative, then the algorithm proceeds to Step 1.2. The current order-up-to level S is replaced by $S+1$; meanwhile, a test is carried out to see whether $S > S^{\gamma}$. If this is the case, then the algorithm jumps to Step 3 in terms of Lemma 2b. Otherwise, another test is undertaken to see whether $G(S-1) > \gamma \geq G(S)$ and $S-1$ can be added to I (by Lemma 2a). Afterward, the algorithm goes back to Step 1.1.

In Step 2.1, the dummy cost γ_1 is upgraded to a "new" γ_1 by (s',S) according to Lemma 1b. The "default" optimal policy (s^*,S^*) and dummy cost γ^* are also updated at this step. In view of Lemma 3, the search for a "new" order-up-to level begins with S^* in Step 2.2. With the new dummy cost, the algorithm regenerates the first candidate for the optimal reorder level and adds it into set I.

Finally, the computation comes to an end in Step 3, in which the optimal (s,S) policy and $C_{\min}$ are identified. The reason why it is possible to have $C_{\min}$ in Step 3 is given by Lemma 4.

We use Figures 2 and 3 to illustrate how the algorithm is executed. At γ_0, the algorithm first searches for the smallest candidate for s_{γ_0} (i.e., $s^1_{\gamma_0}$). It then sets $S_0 = s^1_{\gamma_0} + 1$ to start the process that will upgrade γ_0. As S_0 is increased to S_0', $s^2_{\gamma_0}(= S_0' - 1)$ is encountered, and both $(s^1_{\gamma_0}, S_0')$ and $(s^2_{\gamma_0}, S_0')$ are evaluated. As S is further increased to S_1, $l_{\gamma_0}(s^i_{\gamma_0}, S_1)$ turns to a negative value ($i = 1$ or 2) and γ_0 is lowered to γ_1. Under γ_1, two candidates of s_{γ_1} are identified: $s^1_{\gamma_1}$ and $s^2_{\gamma_1}$. The

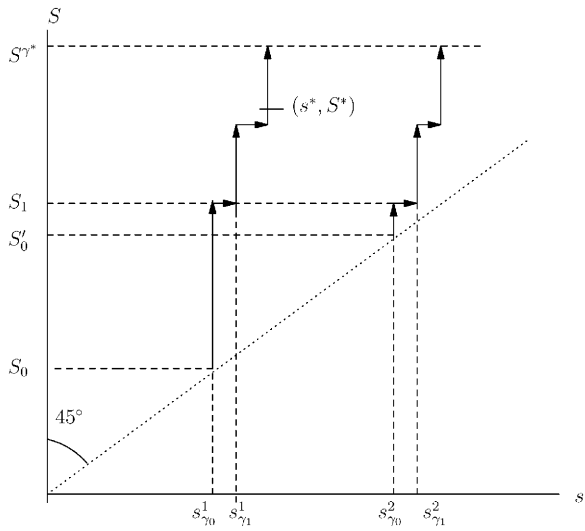


FIGURE 3. A typical execution of the algorithm.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

next around of dummy-cost upgrading then begins with the evaluation of both $(s^i_{\gamma_1}, S_1)$, $i = 1$ and 2. (For the reason for skipping the order-up-to level below S_1, refer to the beginning of Section 3.2.) This process continues until γ_1 is upgraded, and eventually the optimal policy (s^*, S^*) is confirmed after S^{γ^*} has been reached.

It is well recognized that an evaluation of $C(s^1_{\gamma_1}, S)$, and hence $l_{\gamma_1}(s^1_{\gamma_1}, S)$, is a major operation. We thus call each computation of $l_{\gamma_1}(s^1_{\gamma_1}, S)$ a *critical* evaluation. Each time S advances by a unit, $l_{\gamma_1}(s^1_{\gamma_1}, S)$ has to be "completely" recalculated, with virtually nothing from the previous evaluation of $l_{\gamma_1}(s^1_{\gamma_1}, S - j)$, $j \geq 1$, being usable. Therefore, the complexity of the algorithm is mainly determined by the number of critical evaluations (NCE) of $l_{\gamma_1}(s^1_{\gamma_1}, S)$ pertaining to different S. Because the NCE is determined by the cost structure and initial dummy cost γ_0, the complexity of the algorithm is dependent on the holding-shortage cost and the initial dummy cost. Given an initial dummy cost of γ_0, the NCE can be greater than $(S^{\gamma^*} - s^1_{\gamma_0})$, because if the dummy cost is upgraded at (s', S) from γ to γ_1, then S will probably be evaluated twice. Nevertheless, the NCE is bounded by $2(S^{\gamma^*} - s^1_{\gamma_0})$.

Although the main computational effort relates to the NCE, after each updating of γ, a set of new candidates for the optimal reorder level must be relocated. For this reason, we introduce another criterion: the number of dummy profit cost upgrades (NDC), which measures the number of times that γ is upgraded. Both the NCE and NDC will be obtained later in numerical examples.

We conclude this section with a brief discussion of the complexity of the algorithm. Note that the algorithm ends in a finite number of iterations, because each γ upgrading will push s and/or S forward by at least one. Also note that when the loss function is quasiconvex, the algorithm essentially collapses to that proposed by FX. Note that FX's algorithm has the same complexity as ZF's—about twice the effort that is required to evaluate a specific (s, S) policy. When the loss function is nonquasiconvex, additional effort is required to evaluate perhaps several (s^i_γ, S) policies simultaneously for each S. However, according to ZF, in calculating $C(s^1_\gamma, S)$, the $C(s^i_\gamma, S)$ are readily available, where $i = 2, \ldots$ and $s^i_\gamma > s^1_\gamma$. Thus, our algorithm has roughly the same complexity as the two other algorithms (see Chen and Feng [4] for more details).

4. OPTIMALITY OF (s, S) POLICIES: AN EXTENSION

This section extends the optimality analysis of (s, S) policies to a larger class of loss functions, taking a quasiconvex loss function as a special case.

Suppose that we have applied our algorithm to the loss function $G(y)$ and obtained the long-run average cost, $C_{\min}$, of the optimal (s, S) policy (i.e., $\gamma^* = C_{\min}$). Without causing any confusion and for notational convenience, we drop $*$ from γ^* in the rest of this section. We denote by (s^*, S^*) the optimal (s, S) policy, as in the previous sections. The (s, S) optimization analysis confines the optimal parameters (s^*, S^*) between s^1_γ, the smallest or the first candidate for the optimal reorder

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

level, and an upper bound, S^γ, of S^*, which is defined by $G(S^\gamma) \le \gamma < G(y)$ for all $y > S^\gamma$.

The class of loss functions that we are interested in is defined as follows.

DEFINITION 1. *A loss function $G(y)$ is called first reorder-point optimal if, under γ, the optimal (s^*, S^*) policy is such that $s^* = s^1_\gamma < S^*$, and for $s^* < x \le S^\gamma$,*

$$\gamma \ge \frac{\sum_{j=0}^{x-s^*-1} m(j)G(x-j)}{M(x-s^*)}. \tag{5}$$

In words, a first reorder-point optimal loss function requires that $l_\gamma(s^1_\gamma, S^*) = 0$ and $l_\gamma(s, S) \ge 0$ for all $s < S$. Note that the right-hand side of (5) is a weighted average of $G(s^*+1), G(s^*+2), \ldots, G(x)$, where $S^\gamma \ge x > s^* = s^1_\gamma$. This weighted average can be interpreted as the long-run average cost of the policy (s^*, x) as if there were no setup cost. Take Figure 2 as an illustration. Suppose that $\gamma = \gamma_0$ in the figure. If $G(y)$ is first reorder-point optimal, then it must follow that $s^* = s^1_{\gamma_0}$, and the right-hand side of (5) is bounded by γ_0 for all $x = s^1_{\gamma_0} + 1, \ldots, S^{\gamma_0}$.

It is obvious that a quasiconvex loss function is first reorder-point optimal. Another type of loss function that is first reorder-point optimal might be referred to as a *γ-uniformly-bounded loss function*.

DEFINITION 2: *A loss function $G(y)$ is called γ-uniformly-bounded if for any pair of integers $a < b$,* $\max\{G(a), G(b)\} \le \gamma$ *implies* $G(y) \le \gamma$ *for all* $a \le y \le b$.

Figure 4 depicts an example of a γ-uniformly-bounded function. If $G(y)$ is found to be γ-uniformly-bounded, then $s^* = s^1_\gamma$, and for all x, $s^* \le x \le S^\gamma$, $\gamma \ge G(x)$. As a result, (5) is met for all $x \in (s^*, S^\gamma]$, and the loss function is first reorder-point optimal.

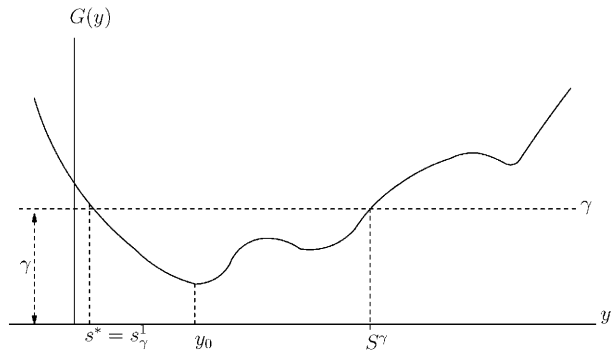


FIGURE 4. An illustration of a γ-uniformly-bounded function.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

To examine whether the loss function, $G(y)$, is first reorder-point optimal, we need to (1) apply our algorithm to obtain the (s^*, S^*) policy and its cost $\gamma = C_{\min}$, (2) check whether $s^* = s_\gamma^1$, and (3) confirm whether (5) holds for every $x \in (s^*, S^\gamma]$ should $s^* = s_\gamma^1$ be true.

First reorder-point optimal loss functions lead to the global optimality of the (s^*, S^*) policy, as is claimed by the ensuing theorem.

THEOREM 5: *Suppose that the loss function $G(y)$ is first reorder-point optimal. The (s^*, S^*) policy is then the optimal policy of all of the admissible ordering policies.*

PROOF: According to the basic theory of dynamic programming (see [12, Chap. V, Thm. 2.1]), the average cost $\gamma = C_{\min}$ is the minimum average that can be attained by any admissible policy if a bounded function $h(y)$ exists such that the following optimality equation holds:

$$h(x) = \inf_{y \geq x} \left\{ K \cdot \delta_{\{y > x\}} + G(y) - \gamma + \sum_{l=0}^{\infty} \phi_l h(y - l) \right\} \quad \text{for all integers } x, \qquad \textbf{(6)}$$

where a stationary policy that chooses actions to minimize the right-hand side of (6) is admissible, and $\delta_{\{\cdot\}}$ is an indicator function of the event described by $\{\cdot\}$: if the argument is true (i.e., $y > x$), then $\delta = 1$; otherwise, $y = 0$. This set of conditions is often called an optimality criterion. From Ross [12], if (6) holds, then the implied policy that minimizes the right-hand side of (6) is optimal among all admissible policies and γ is its long-run average cost. Consequently, the (s^*, S^*) policy is also optimal, because it also achieves γ. It should be noted that the implied policy and the (s^*, S^*) policy need not be the same.

It seems impossible to construct a bounded function $h(x)$ such that the optimal equation (6) holds with γ. For this reason, we adopt the following argument that was proposed by Zheng [19], who provided a simple proof for the optimality of (s, S) policies in models with quasiconvex loss functions. Suppose that inventory can be returned to the supplying source; that is, the firm can place orders of negative quantities. Returning of any quantity incurs the same setup cost as placing an order for a positive quantity, but there is no variable cost. Allowing the firm to reduce excess inventory might lower the long-run average cost. A model with such an additional option will be called a relaxed model. If we can verify that the minimum cost that the relaxed model can achieve is still γ, then the (s^*, S^*) policy is globally optimal in the original model.

With a return option, the *optimality equation* (6) has to be modified for the relaxed model:

$$h(x) = \inf_{-\infty < y < +\infty} \left\{ K \cdot \delta_{\{y \neq x\}} + G(y) - \gamma + \sum_{j=0}^{\infty} \phi_j h(y - j) \right\} \quad \text{for all integers } x. \qquad \textbf{(7)}$$

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

To ensure (7), we construct a function $h(x)$ as follows:

$$h(x) = \begin{cases} K & \text{if } x \le s^* \\ l_\gamma(s^*, x) & \text{if } s^* < x \le S^\gamma \\ \min\left(K, G(x) - \gamma + \sum_{j=0}^{\infty} \phi_j h(x-j)\right) & \text{if } x > S^\gamma, \end{cases} \tag{8}$$

where $l_\gamma(\ )$ is as previously defined.

First, we show that such a function $h(x)$ is bounded; in particular,

$$0 \le h(x) \le K$$

for all integers x. This is true by the definition of $h(x)$ when $x \le s^*$. For $x \in (s^*, S^\gamma]$, as it has already been shown in Section 4 that $l_\gamma(s^*, x) \ge l_\gamma(s^*, S^*) = 0$, $h(x) = l_\gamma(s^*, x) \ge 0$. To confirm that $h(x) \le K$, note that for every $x \in (s^*, S^\gamma]$,

$$\gamma \ge \frac{\sum_{j=0}^{x-s^*-1} m(j)G(x-j)}{M(x-s^*)} \Leftrightarrow K + \sum_{j=0}^{x-s^*-1} m(j)[G(x-j) - \gamma] \le K.$$

By (2), we thus have

$$h(x) = l_\gamma(s^*, x) \le K$$

for $s^* < x \le S^\gamma$.

When $x > S^\gamma$, because $h(x) \le K$ by construction, we only need to confirm that $h(x) \ge 0$. Applying induction, we begin with $x = S^\gamma + 1$. Note that either

$$h(S^\gamma + 1) = G(S^\gamma + 1) - \gamma + \sum_{j=0}^{\infty} \phi_j h(S^\gamma + 1 - j) > 0,$$

because $G(S^\gamma + 1) > \gamma$, or $h(S^\gamma + 1) = K$. In either case, $h(S^\gamma + 1) > 0$, which suffices for the induction. Thus, $K \ge h(x) > 0$ for $x > S^\gamma$.

We now verify the optimality equation (7). By defining a new function as

$$V(x) = G(x) - \gamma + \sum_{l=0}^{\infty} \phi_l h(x-l),$$

the optimality equation (7) can be rewritten as

$$h(x) = \inf_y \{K \cdot 1_{\{y \ne x\}} + V(y)\}.$$

Note that for $s^* < x \le S^\gamma$, $V(x) = l_\gamma(s^*, x)$. Because $G(x) > \gamma$ when $x \le s^*$ and $x > S^\gamma$, it is implied by the nonnegativity of $h(x)$ that $H(x) > 0$ for those x. When

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

$s^* < x \leq S^\gamma$, $V(x) = h(x) = l_\gamma(s^*, x) \geq 0$ and the minimum of $V(x)$ is achieved at $x = S^*$, where $V(S^*) = h(S^*) = 0$. Therefore, the criterion (7) is equivalent to

$$h(x) = \min(K + V(S^*), V(x)) = \min(K, V(x)). \tag{9}$$

An optimal action that minimizes the right-hand side of (9) is to decide between staying at the current inventory position x or adjusting x to S^*, depending on which yields a smaller value.

To show criterion (9), the following three ranges are considered: (1) $x \leq s^*$, (2) $S^* < x \leq S^\gamma$, and (3) $S^\gamma < x$. In range (1), $V(x) = G(x) - \gamma + K > K$, which indicates that an optimal policy is to set $y = S^* \geq x$ (i.e., to place an order of $(S^* - x)$ units). In range (3), $h(x) = \min(K, V(x))$ by definition, implying an optimal policy of either returning some inventories to lower x to S^* or doing nothing. Finally, in range (2), $h(x) = V(x) \leq K$, suggesting not to do anything. Hence, the optimality equation is met for all x, and the optimal policy at each x is admissible. This concludes that the minimum average cost is γ in the relaxed model and, consequently, that the (s^*, S^*) policy is also optimal in the relaxed and original models. ■

This argument is intuitive. Note that the optimal policy minimizing the right-hand side of (7) differs from the (s^*, S^*) policy only in a finite number of initial periods. This is because when the initial inventory position is above S^γ, the former policy might issue an order for a negative quantity, whereas the latter policy does nothing. In either case, the inventory position will fall to S^γ or below in a finite number of periods, after which the two policies work identically. In terms of long-run average cost, both policies obtain γ, because the cost difference in a finite number of initial periods does not affect the performance in the long run.

Theorem 5 broadens the optimality of the (s, S) ordering policies in the sense that the quasiconvex loss functions become just a subclass of the first reorder-point functions. We have thus extended the optimality result of Iglehart [9] and Veinott [17].

5. NUMERICAL EXAMPLES

The descent algorithm was used to solve several sets of problems to observe its operating characteristics. In the algorithm implementation, a module is included to test whether the first-reorder-point condition is met in each instance.

In all of the examples, the period demand follows the Normal distribution (μ, σ^2), and the lead time is either two periods (w.p. 0.8) or four periods (w.p. 0.2). We consider only one nonlinear shortage cost: in Eq. (1), $p > 0$ and $b > 0$, but $\pi = 0$. In addition, we take the linear holding cost with a rate $h > 0$.

The results for these problems are displayed in Table 1, in which the first block records the cost parameter values; the second block displays three bounds: s_0 and S^u, which represent the first candidate of reorder point that is associated with C_0 (the initial dummy cost), and the upper bound of S^* associated with $C_{\min}$; and the

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

Table 1. Results for an example set

No.	K	h	p	b	s_0	S^u	C_0	CoV	s^*	S^*	$C_{\min}$	Opt[a]	NCE[b]	NDC[c]
1	24	1	2	150	65	127	64.43	nq	67	116	59.87	Yes	64	44
2				200	68	132	69.66	nq	91	124	64.86	Unknown	87	41
3				250	94	133	70.72	nq	96	126	66.18	Yes	41	27
4			4	150	69	130	69.38	qu	75	123	63.78	Yes	67	37
5				200	90	132	70.46	qu	94	125	65.81	Yes	46	23
6				250	95	134	71.52	qu	97	127	67.01	Yes	41	19
7		3	2	150	56	84	149.81	qu	62	70	94.36	Yes	34	3
8				200	59	85	150.87	nq	63	70	105.74	Yes	30	3
9				250	61	85	151.94	nq	64	71	116.71	Yes	27	3
10			4	150	58	88	150.34	qu	63	70	107.41	Yes	34	3
11				200	61	90	151.63	qu	64	71	118.39	Yes	31	3
12				250	62	109	152.72	qu	65	72	129.16	Yes	50	3
13	44	1	2	150	62	134	88.82	nq	66	128	66.26	Yes	76	50
14				200	64	138	89.66	nq	68	129	71.08	Yes	76	37
15				250	66	143	90.72	nq	71	130	75.50	Yes	872	35
16			4	150	64	137	89.49	qu	68	126	70.18	Yes	77	29
17				200	67	141	90.77	qu	74	128	74.45	Yes	78	34
18				250	69	143	91.62	qu	92	128	76.79	Yes	86	47
19		3	2	150	56	104	167.81	qu	61	90	119.68	Yes	53	14
20				200	57	105	170.45	nq	62	90	121.36	Yes	51	16
21				250	60	110	171.13	nq	63	108	131.64	Yes	53	20
22			4	150	58	109	170.73	qu	62	92	128.52	Yes	55	20
23				200	59	110	171.86	qu	63	91	130.21	Yes	54	20
24				250	61	113	172.74	qu	64	109	137.94	Yes	55	28

Note: $(\mu, \sigma^2) = (20,10)$. Lead time: 2,w.p. 0.8; 4,w.p. 0.2.
[a]Global optimality.
[b]Number of calculations of $l_\gamma(\ )$.
[c]Number of γ upgrading.

quasiconvexity of the loss function is tested and indicated under "CoV," where *nq* stands for non-quasiconvex and *qu* stands for quasiconvex. The third block reports the optimal solution: the optimal policy (s^*, S^*) and the minimum cost $C_{\min}$. The column under "Opt" indicates whether the (s^*, S^*) policy is globally optimal. If the first-reorder-point condition is satisfied, then the answer is "Yes"; otherwise, the optimality is unknown. The column NCE reports the number of times that the auxiliary function is evaluated, and the NDC displays the number of dummy cost upgrades. As explained earlier, these two numbers measure the computational effort that is expended in obtaining the optimal policy parameters. The global minimizer of $G(y)$ is not reported, although in all examples, $C_0 = C(y_0 - 1, y_0)$.

In all instances except one, the (s^*, S^*) policy is tested to be globally optimal. In the only exception, example 2, the test is inconclusive, which is shown graphically in Figure 1. As can be seen from Table 1 and Figure 1, the optimal reorder point s^* does not take the first candidate for the reorder point. The achieved ready rates that were calculated but not reported here are between 0.64 and 0.97. Several different sets of problem data were also computed, for example, with different

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

demand and lead-time distributions, but the results are similar and are thus not reported here. In general, when the setup cost K or the holding cost h is relatively large, the optimality test of the (s, S) policy is more conclusive, because the average cost is high and, hence, the loss function is γ-uniformly-bounded. Although it is generally true that the smaller the nonlinear shortage cost p or π is, the more likely it is that the first-reorder-point condition is met, there are exceptions. Comparing Example 2 with Example 3 in Table 1, we see that as b decreases, the condition is actually violated.

6. CONCLUDING REMARKS

In this article we develop an intuitive and efficient algorithm for computing an optimal (s, S) policy. Unlike most of the existing methods, this algorithm can handle models with non-quasiconvex loss functions. Based on a parametric problem associated with the original optimization problem, the algorithm was devised using an auxiliary function with a dummy cost. It starts with an easily identifiable dummy cost and then upgrades the dummy cost reiteratively until the optimizing policy is identified. The computational effort is mainly expended on the evaluation of the value of the auxiliary function with each order-up-to level. Because it advances the order-up-to level unidirectionally from left to right, the algorithm is efficient. We also establish conditions under which an (s, S) policy is optimal among all policies in the model with a non-quasiconvex loss function. This result extends the standard model with a quasiconvex loss function and is also readily extendable to other models. However, we refer the reader to Zheng and Federgruen [20] for a detailed discussion of these models.

Acknowledgments

This research was partially supported by Hong Kong RGC/CUHK grant No. 2050310. The authors are grateful to the anonymous referee for constructive comments, which have improved the presentation of the paper. The computational assistance of Mr. Tong Wang is also appreciated.

References

1. Ballou, R.H. (1999). *Business logistics management: Planning, organizing, and controlling the supply chain*, 4th ed. Upper Saddle River, NJ: Prentice-Hall.
2. Bashyam, S. & Fu, M.C. (1998). Optimization of (s, S) inventory systems with random lead times and a service level constraint. *Management Science* 44: 243–256.
3. Chen, F. & Zheng, Y.-S. (1993). Inventory models with general backorder costs. *European Journal of Operational Research* 65: 175–186.
4. Chen, Y. & Feng, Y. (2004). Optimization and optimality of (s, S) stochastic inventory systems with non-quasiconvex costs. Unabridged version, Department of System Engineering & Engineering Management, The Chinese University of Hong Kong, Shatin, Hong Kong, China. Downloadable from www.se.cuhk.edu.hk/ỹhchen.
5. Eppen, G.D. & Martin, R.K. (1988). Determining safety stock in the presence of stochastic lead time and demand. *Management Science* 34: 1380–1390.
6. Feng, Y. & Sun, J. (2001). An algorithm for deriving (r, R, d, D) inventory replenishment policy. *Operations Research* 49: 1380–1390.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press

7. Feng, Y. & Xiao, B. (2000). New algorithms for computing optimal (s, S) policies in a stochastic single item/location inventory system. *IIE Transactions* 32: 1081–1090.
8. Hadley, G. & Whitin, T.M. (1963). *Analysis of inventory systems*. Upper Saddle River, NJ: Prentice-Hall.
9. Iglehart, D. (1963). Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science* 10: 259–267.
10. Li, Z., Xu, S.H., & Hayya, J. (2004). A periodic-review inventory system with supply interruptions. *Probability in the Engineering and Informational Sciences* 18: 33–53.
11. Rosling, K. (2002). Applicable cost rate functions for single-item inventory control. *Operations Research* 50: 1007–1017.
12. Ross, S. (1983). *Stochastic processes.* New York: Wiley.
13. Scarf, H. (1960). The optimailtiy of (s, S) policies in the dynamic inventory problem. In K. Arrow, S. Karlin, & P. Suppes (eds.), *Mathematical methods in the social sciences*. Stanford, CA: Stanford University Press, Chap. 22.
14. Sethi, S. & Cheng, F. (1997). Optimality of (s, S) policies in inventory models with Markovian demands. *Operations Research* 45: 931–939.
15. Silver, E.A., Pyke D., & Peterson, R.A. (1998). *Inventory management and production planning and scheduling*, 3rd ed. New York: Wiley.
16. Song, J.-S. & Zipkin, P. (1993). Inventory control in a fluctuating demand environment. *Operations Research* 41: 351–370.
17. Veinott, A., Jr. (1966). On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal of Applied Mathematics* 14: 1067–1083.
18. Veinott, A., Jr. & Wagner, H.M. (1965). Computing optimal (s, S) inventory policies. *Management Science* 11: 690–723.
19. Zheng, Y.-S. (1991). A simple proof for optimality of (s, S) policies in infinite-horizon inventory systems. *Journal of Applied Probability* 28: 802–810.
20. Zheng, Y.-S. & Federgruen, A. (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations Research* 39: 654–665.
21. Zipkin, P. (1999) *Foundations of inventory management*. Singapore: McGraw-Hill Higher Education, pp. 196–198.

https://doi.org/10.1017/S0269964806060189 Published online by Cambridge University Press