The Effectiveness of a Simple Policy for Coordinating Inventory Control and Pricing Strategies

by

Zhibo Sun

B.B.A., National University of Singapore (2009)
B.Comp., National University of Singapore (2009)

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2010

Author ....
School of Engineering
July 30, 2010

Certified by ....
David Simchi-Levi
Professor of Civil and Environmental Engineering and Engineering Systems
Thesis Supervisor

Accepted by ....
Karen Willcox
Associate Professor of Aeronautics and Astronautics
Co-director, Computation for Design and Optimization Program

The Effectiveness of a Simple Policy for Coordinating Inventory Control and Pricing Strategies

by

Zhibo Sun

Submitted to the School of Engineering
on July 30, 2010, in partial fulfillment of the
requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract

We investigate the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy in a single product, periodic review, finite horizon model with stochastic multiplicative demand and fixed ordering cost, in which an (s, S, A, p) policy is optimal. An extensive numerical study shows that empirically an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy. We also formulate two alternative benchmark policies and find that the (s, S, p) policy is superior in terms of profit. In addition, we propose an efficient algorithm with simulated annealing and modified binary search to determine the (s, S, p) policy for the model.

Thesis Supervisor: David Simchi-Levi
Title: Professor of Civil and Environmental Engineering and Engineering Systems

Acknowledgments

This thesis would not have been completed without the support and assistance from the following people.

First and foremost, I would like to express my sincere gratitude to my thesis advisor, Professor David Simchi-Levi, for his constructive guidance and continuous encouragement throughout my research. To me, Professor Simchi-Levi is not only an insightful advisor, but also a great mentor. I really enjoyed working with him.

Second, I want to thank my friends, especially Laura, from CDO. They all helped me in one way or another and made my time at MIT both fruitful and fun.

Next, special appreciation goes to the Singapore-MIT Alliance, which provided me with financial aid and let me concentrate on my study.

Last yet most importantly, I am indebted to my mother and also my loving girlfriend, Naijun, for their love and inspiration. They have always been there for me and I will always be there for you.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1 Introduction **13**

1.1 Background and Motivation 13

1.2 Literature Review 14

1.3 Overview 17

2 The Model **19**

2.1 Notations and Assumptions 19

2.2 Symmetric k-convexity and Optimal Solution 21

3 Numerical Study **23**

3.1 Methodology 23

3.2 Sensitivity Analysis 24

3.2.1 Baseline Parameters 24

3.2.2 Initial Inventory Level 25

3.2.3 Total Number of Periods 27

3.2.4 Ordering Costs 28

3.2.5 Inventory Holding Cost 30

3.2.6 Price Range 31

3.2.7 Demand Function Parameters 33

3.2.8 Coefficient of Variation of Randomness 35

3.3 Overall Observation 36

4 Comparisons with Alternative Policies **39**

4.1 Comparison with a Deterministic Policy 39
4.2 Comparison with a Fixed Price Policy 40

5 Proposed Algorithm for (s, S, p) Policy **43**

5.1 Proposed Algorithm 43
5.2 Empirical Performance 45

6 Concluding Remarks **47**

6.1 Conclusion 47
6.2 Limitations and Extensions 48

List of Figures

2-1 Illustration of Symmetric k-Concavity and (s, S, A, p) Policy 22

3-1 Effectiveness against Initial Inventory Level 26
3-2 Effectiveness with a Combination of Two Conditions 27
3-3 Effectiveness against Total Number of Periods 28
3-4 Effectiveness against Fixed Ordering Cost with $c = 3$ 29
3-5 Effectiveness against Fixed Ordering Cost with $c = 2$ 29
3-6 Effectiveness against Variable Ordering Cost 30
3-7 Effectiveness against Percentage of Inventory Holding Cost 31
3-8 Movements of Reorder Point and Order Up-to-Level against Percentage of Inventory Holding Cost 32
3-9 Effectiveness against Price Floor 33
3-10 Effectiveness against Price Ceiling 33
3-11 Effectiveness against Parameter a 34
3-12 Effectiveness against Parameter b 35
3-13 Effectiveness against Coefficient of Variation of Randomness with Fixed Number of Elements 36
3-14 Effectiveness against Coefficient of Variation of Randomness with Different Number of Elements 37
3-15 Example of g Function Values against Inventory-Demand Pairs 38
3-16 Zoomed in Image at the Global Optimum of Figure 3-15 38

5-1 Effectiveness of Proposed Algorithm 46

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

4.1 Comparison with a Deterministic Policy 40
4.2 Comparison with a Fixed Price Policy 41
4.3 Comparison with a Fixed Price Policy with Low Ordering Cost 41

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Background and Motivation

Retail and manufacturing companies always strive to improve their operational profitability. Traditional inventory management models such as the Economic Order Quantity (EOQ) model and the newsvendor model assume exogenously determined demands and focus on effective inventory replenishment strategies to better meet customer demand. Recently, revenue management techniques have been applied in practice. More and more companies currently explore dynamic pricing strategies to adjust demands in order to boost their operation profits and bottom lines. For example, with the help of dynamic pricing, Dell Computers, Ford Motor and American Airlines all experienced major growth in revenue. For details, the interested reader is referred to Gallego and van Ryzin (1994), Agrawal and Kambil (2000), Leibs (2000) and Cook (2000).

In recent years, integrating inventory control and pricing strategies has gained increasing attention (See Elmaghraby and Keskinocak (2003) and Chan et al. (2004) for a review and classification of the recent literature). In particular, Chen and Simchi-Levi (2004a) investigate a single product, periodic review, finite horizon model with stochastic demand and fixed ordering cost. They prove that an (s, S, p) policy is optimal with additive demand. In this policy, the inventory control strategy is the well-known (s, S) policy: An order is placed to raise the inventory level to the order-up-to level, S_t, if the initial inventory level in period t is below the reorder point s_t. Otherwise, no order is placed. The optimal price p_t in period t depends on the initial inventory level at the beginning of period t. However, Chen and Simchi-Levi (2004a) show that this (s, S, p) policy may not necessarily be optimal with general, non-additive demand, e.g. multiplicative demand. They develop and employ the concept of symmetric k-convexity to discover an optimal (s, S, A, p) policy for this scenario. This policy includes a special set $A_t \in [s_t, \frac{s_t+S_t}{2}]$ for period t. When the initial inventory level in period t, x_t, is less than s_t or $x_t \in A_t$, an order is placed to raise the inventory level from x_t to S_t. Otherwise, no order is placed. Similar to an (s, S, p) policy, the optimal price of an (s, S, A, p) policy in each period depends on its initial inventory level.

Although an (s, S, A, p) policy is optimal for the general demand scenario, it is a complicated and less intuitive policy for many operations managers. The set A_t for period t is not easy to determine, may not be connected and does not always exist. Thus, if we can identify a simpler policy that yields a profit close to that of an (s, S, A, p) policy, it would become much more popular in practice. The (s, S, p) policy is naturally a good candidate for two reasons. First, Chen and Simchi-Levi(2004b, 2006) further prove that an (s, S, p) policy is optimal for a single product, periodic review, infinite horizon model and single product, continuous review models. It means that an (s, S, p) policy is the optimal policy in all the cases with the exception of the single product, periodic review, finite horizon model with general demand. Second, this policy is a more intuitive policy. It is easier for operations managers to understand and execute. The goal of this thesis is therefore to test numerically the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy, in scenarios with the latter as the optimal policy.

1.2 Literature Review

Research on coordinating price and inventory control strategies started with Whitin (1955), who investigates the newsvendor problem with price depending linearly on demand. Mills (1959), Lau and Lau (1988) and Polatoglu (1991) all examine single product, single period, additive demand models and they show that the optimal price with stochastic demand is no greater than that with deterministic demand. In contrast, Karlin and Carr (1962) address a single product, single period, multiplicative demand model and they find the optimal price with stochastic demand is no smaller than that with deterministic demand. Pertruzzi and Dada (1999) introduce a base price and demonstrate that the optimal price is equivalent to a base price plus a price premium.

Following Wagner and Whitin (1958), who first extend the single product, single period models to multi-period models with deterministic demand, Zabel (1972) and Thowsen (1975) add stochastic components to the demand processes and show the existence of a unique optimal solution under certain assumptions of the demand processes. Federgruen and Heching (1999) consider only variable ordering cost, linear demand model, and assumes backlogging of excess demand. They find that a base stock list price policy is optimal. In this policy, if the initial inventory level is below the base stock level, an order is placed to raise the inventory level to the base stock level. Otherwise, no order is placed and a discount price is offered. Thomas (1974) adds fixed ordering cost and analyzes a single product periodic review, finite horizon model. He conjectures that an (s, S, p) policy is optimal if all prices in an interval are under consideration. He also provides a counterexample to show that this policy may not be optimal if prices are chosen from a discrete set. Polatoglu and Sahin (2000) consider a similar model but allow lost sales. Although they point out sufficient conditions for the optimality of an (s, S, p) policy, they do not specify what kind of demand processes satisfy those conditions. Chen and Simchi-Levi (2004a) continue the research by Thomas (1974) and prove that when the demand process is additive, an (s, S, p) policy is optimal. When the demand process is general, however, an (s, S, A, p) policy is shown to be optimal by applying a new concept called symmetric k-convexity.

Amihud and Mendelson (1983) address an infinite horizon model and demonstrate the phenomenon of price smoothing in reaction to changes in inventory level. Chen and Simchi-Levi (2004b) use approaches similar to Iglehart (1963) and Veinott (1966) and characterize a stationary (s, S, p) policy as the optimal policy for the infinite horizon model, under both discounted and average profit criteria with general demand. Unlike Chen and Simchi-Levi (2004b), who consider a periodic review model, Feng and Chen (2002) develop a single product, continuous review, infinite horizon model with Poisson demand process and integral demand size. They show that an (s, S, p) policy is optimal for this model. Chao and Zhou (2006) further identify a close form solution and design an efficient algorithm to compute the optimal policy. Chen and Simchi-Levi (2006) generalize Feng and Chen (2002)'s result and demonstrate that there exists a stationary (s, S, p) policy that is optimal for the single product, continuous review, infinite horizon model.

We can see that the recent studies have shown that an (s, S, p) policy is optimal in most cases with the exception of the single product, periodic review, finite horizon model with stochastic general demand and fixed ordering cost, in which an (s, S, A, p) policy is optimal. An interesting question that naturally arises is how good an (s, S, p) policy is in a model with an (s, S, A, p) policy as the optimal policy. Although Chen and Simchi-Levi (2004a) prove the optimality of an (s, S, A, p) policy in such a model, they never specify how to compute the special set A. In fact, the majority of the research in coordinating pricing and inventory control strategies is theoretical and only a few, such as Federgruen and Heching (1999) and Chao and Zhou (2006), have done a numerical study. However, their numerical studies assume additive demand processes and, as a result, the optimal policy is an (s, S, p) policy. We know of no prior numerical studies focusing with multiplicative demand process, probably due to its more complicated optimal policy. In this thesis, we fill in the gap by conducting a comprehensive numerical study with stochastic multiplicative demand. We not only implement the dynamic programming process proposed by Chen and Simchi-Levi (2004a) to determine an optimal (s, S, A, p) policy, but also compare its profit with that from a corresponding (s, S, p) policy. We would like to show numerically that an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy and thus becomes a good substitute for the latter.

1.3 Overview

This thesis is organized as follows. In Section 2, we describe the model used in this thesis and review the main assumptions. In Section 3, we investigate the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy with an extensive numerical study. We would like to perform sensitivity analysis on the input parameters of the model and examine how they affect the effectiveness of an (s, S, p) policy. In Section 4, we formulate several alternative benchmark polices and compare them with an (s, S, p) policy to check whether the latter policy is superior. In Section 5, we propose a feasible and fairly efficient algorithm to compute an (s, S, p) policy. Finally, we provide the concluding remarks in Section 6.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

The Model

2.1 Notations and Assumptions

In this thesis, we follow the model described in Chen and Simchi-Levi (2004a) with stochastic multiplicative demand. We consider a firm that has to make simultaneous ordering and pricing decisions over a finite time horizon of T periods with time independent demand process. We assume lead time is zero and orders arrive immediately. For each period t, where $t = 1, 2, ..., T$, we define

w_t = actual demand observed in period t,

p_t = selling price in period t,

$\underline{p}_t, \overline{p}_t$ are price floor and price ceiling of p_t, respectively.

We assume $w_t := \alpha_t D_t(p_t) + \beta_t$, where α_t, β_t are two random variables with $E(\alpha_t) = 1$ and $E(\beta_t) = 0$. Clearly, this assumption can be made without loss of generality. If $\alpha_t = 1$ and $w_t = D_t(p_t) + \beta_t$, it is referred to as the additive case. If $\beta_t = 0$ and $w_t = \alpha_t D_t(p_t)$, it is referred to as the multiplicative case. Moreover, $D_t(p_t) = b_t - a_t p_t$ is referred to as the additive demand function whereas $D_t(p_t) = a p_t^{-b_t}$ is referred to as the multiplicative demand function. According to Petruzzi and Dada (1999), both are common demand functions in economics literatures. In this thesis, we focus on multiplicative demand process, which means $w_t = \alpha_t(a_t p_t^{-b_t})$. When the demand process is multiplicative, the expected revenue $R_t(d_t) = d_t p_t = d_t D_t^{-1}(d_t)$ is clearly a concave function of the expected demand d_t, where $D_t^{-1}(d_t)$ represents the inverse function of D_t and is continuous and strictly decreasing.

Let x_t and y_t be the inventory levels at the beginning of period t, just before and after placing an order, respectively. In other word, $y_t - x_t$ is the amount of products ordered in period t. The ordering cost consists of both a fixed ordering cost and a variable ordering cost. For every period $t \in T$, it is calculated as $k\delta(y_t - x_t) + c_t(y_t - x_t)$, where the fixed ordering cost k is time independent and the binary function $\delta(\mu)$ is defined as

$$\delta(\mu) := \begin{cases} 1 & \text{if } \mu > 0, \\ 0 & \text{otherwise.} \end{cases}$$

We also assume excess demand is backlogged. Thus, x_{t+1}, which represents the inventory level carried over from period t to period $t+1$, may be either positive or negative. It is easy to see that $x_{t+1} = y_t - w_t$. Inventory holding cost is incurred when $x_{t+1} > 0$ whereas backorder cost is incurred when $x_{t+1} < 0$. Let the corresponding cost function be $h_t(y_t, w_t) = c_{ht}(y_t - w_t)^+ + c_{bt}(y_t - w_t)^-$, where c_{ht} and c_{bt} represent the unit inventory holding cost and the unit backorder cost for period t.

Denote $G_t(y_t, p_t) = E[h_t(y_t - w_t)]$. Similar to Federgruen and Heching (1999), we assume $h_t(y_t, w_t)$ is a convex function with $\lim_{y \to \infty} G_t(y_t, p_t) = \lim_{y \to -\infty}[c_t y + G_t(y_t, p_t)] = \lim_{y \to \infty}[(c_t - c_{t+1})y + G_t(y_t, p_t)] = \infty$. We further assume that $0 \leq G_t(y_t, p_t) = O(|y|^\rho)$ for some integer ρ and $E[w_t]^\rho = E[\alpha_t D_t(p_t) + \beta_t]^\rho < \infty$ for all $p \in [\underline{p}_t, \overline{p}_t]$.

Since $w_t := \alpha_t D_t(p_t) + \beta_t$, there is a one to one relationship between the selling price p_t and the expected demand d_t. In particular, the expected demand floor $\underline{d}_t$ follows $\underline{d}_t = D_t(\overline{p}_t)$ whereas the expected demand ceiling $\overline{d}_t$ follows $\overline{d}_t = D_t(\underline{p}_t)$. We would like to maximize the total expected profit over the entire horizon by selecting an inventory level y_t and an expected demand d_t for each period t.

Let $v_t(x_t)$ be the profit-to-go function at the beginning of period t with an initial inventory level x_t. $v_t(x_t)$ can be computed recursively through a dynamic programming process with $v_{T+1}(x_{t+1}) = 0$.

For $t = 1, 2, ...T$, we have $v_t(x_t) = c_t x_t + max(-k\delta(y_t - x_t) + g_t(y_t, d_t(y_t))$, where $d_t(y_t)$ is the expected demand corresponding to the best selling price for a given inventory level y_t, i.e. $d_t = \arg_{\overline{d_t} \geq d_t \geq \underline{d}_t} \max(g_t(y_t, d_t))$.

The function $g_t(y_t, d_t)$ satisfies $g_t(y_t, d_t) = R_t(d_t) - c_t y + E[-h_t(y_t - \alpha_t d_t - \beta_t) + v_{t+1}(y_t - \alpha_t d_t - \beta_t)]$.

2.2 Symmetric k-convexity and Optimal Solution

Scarf (1960) uses a model similar to Chen and Simchi-Levi (2004a) and assumes stochastic demand. He shows that an (s, S) policy is optimal with the help of a concept called k-convexity. Chen and Simchi-Levi use an equivalent definition of k-convexity discovered by Porteus (1971) to show that the $g(y_t, d_t(y_t))$ function, i.e. the g function in short and hence the profit-to-go function are both k-concave and prove the optimality of an (s, S, p) policy, when the demand is stochastic and additive.

Definition 2.1. A real-valued function f is called k-convex for $k \geq 0$, if for any $x_0 \leq x_1$ and $\lambda \in [0, 1]$,

$$f((1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + \lambda k$$

A function f is called k-concave if $-f$ is k-convex.

For stochastic multiplicative demand process, however, Chen and Simchi-Levi (2004a) demonstrate that the g function is not always k-concave and an (s, S, p) policy may not necessarily be optimal. In order to identify the optimal policy, they employ a new concept called symmetric k-convexity.

Definition 2.2. A real-valued function f is called symmetric k-convex for $k \geq 0$, if for any x_0, x_1 and $\lambda \in [0, 1]$,

$$f((1-\lambda)x_0 + \lambda x_1) \leq (1-\lambda)f(x_0) + \lambda f(x_1) + max\{\lambda, 1-\lambda\}k.$$

A function f is called symmetric k-concave if $-f$ is symmetric k-convex.

In fact, k-convexity is a special case of symmetric k-convexity. For more details about properties of k-convexity and symmetric k-convexity, the interested reader is referred to Bertsekas (1995) and Chen and Simchi-Levi (2004a).

Chen and Simchi-Levi (2004a) show that the g function and hence the profit-to-go function, i.e. the v function, in the stochastic multiplicative demand model are both symmetric k-concave. They further show that for $t = T, T-1, ..., 1$, there exist a reorder point s_t and an order up-to-level S_t with $s_t \leq S_t$ and a special set $A_t \in [s_t, \frac{s_t+S_t}{2}]$ such that when the initial inventory level $x_t < s_t$ or $x_t \in A_t$, it is optimal to order $S_t - x_t$ and set the expected demand level $d_t = d_t(S_t)$; otherwise it is optimal to order nothing and set $d_t = d_t(x_t)$. As illustrated in Figure 2-1, when $x_t < s_t$ or $x_t \in A_t$, $g(S_t, d_t(S_t)) - k > g(x_t, d_t(x_t))$, and hence it is optimal to place an order to raise the inventory level to S_t. Otherwise, $g(S_t, d_t(S_t)) - k \leq g(x_t, d_t(x_t))$ and no order should be placed. This optimal policy is called an (s, S, A, p) policy.

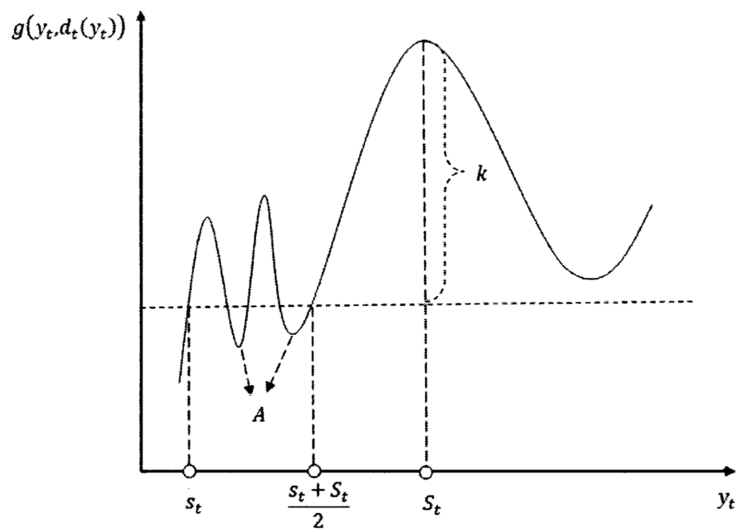

Figure 2-1: Illustration of Symmetric k-Concavity and (s, S, A, p) Policy

Chapter 3

Numerical Study

3.1 Methodology

Empirical research on (s, S, A, p) has been a challenging issue for two reasons. First, we do not know, for a given set of parameters, whether the special set A exists and we are not sure how many elements the set A contains, imposing difficulties in determining the set A. Second, solving the model numerically has to discretize the inventory and pricing variables to a certain level. However, Chen and Simchi-Levi (2004b) and Simchi-Levi et al. (2005) show that with discrete prices, an (s, S, p) policy and an (s, S, A, p) policy may not be optimal for the additive demand process and multiplicative demand process models, respectively. Nevertheless, Chen and Simchi-Levi (2004b) propose that it is feasible to discretize the inventory variables and then choose the best pricing variables accordingly from a continuous set, in order to solve a discretized model.

Our approach in the numerical study follows the suggestion made by Chen and Simchi-Levi (2004b) and, for each period, we discretize the inventory variables to integer level and then choose the expected demand that associates with the best price. The reason for the choice of integer level discretization is threefold. First, integer level discretization approximates the solution of the continuous model very well. As we discretize the inventory variables further, we obtain more accurate results, which are closer to the solution of the continuous model. After comparing the results from different levels of discretization, we find that the solution and the associated expected profit of the integer level discretized model are very close to that from the further discretized models. In fact, in many cases, if we round the solutions of the further discretized models, we can get the solution of the integer level discretized model. Moreover, the effectiveness of an $(s.S, p)$ policy relative to an (s, S, A, p) policy from the integer level discretized model tends to be lower than those from the further discretized models, probably due to the cumulative rounding error of the discretization. In some sense, the integer level discretized model provides a lower bound for the effectiveness, which is useful for our investigation in this numerical study. Second, we can compute the optimal solution much faster with integer level discretization. For the same set of input parameters, the integer level discretized model may require less than 10 seconds to solve whereas the 0.1 discretized model takes several minutes to solve, and the time increases exponentially as we discretize the model further. Last and importantly, inventory level is always an integer in real-life practice. It therefore makes sense to use the integer level discretized model in our numerical study. In addition, since the special set A_t for any period t is hard to determine, we use an enumeration method to find the elements in the set A_t and make sure the computed (s, S, A, p) policy is accurate. In section 5, we propose a more efficient algorithm to determine the (s, S, p) policy instead of the (s, S, A, p) policy for the original continuous model.

3.2 Sensitivity Analysis

3.2.1 Baseline Parameters

In this section, we perform an extensive sensitivity analysis on all the parameters of our integer level discretized model. Our objective is to test the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy, where the latter is optimal, and understand how changes in different parameters affect the effectiveness.

We have tested a number of different combinations of input parameters with our model and, surprisingly, in most cases the special set A does not exist. As a result, an (s, S, p) policy is equivalent to an (s, S, A, p) policy. When there is a difference between the two policies, an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy and the effectiveness is usually above 98%.

After careful consideration, we select our baseline parameters as follows. The demand function parameters are $a = 500$ and $b = 1.5$, i.e. $D_t(p_t) = \alpha_t(500p_t^{-1.5})$. The randomness parameter α follows uniform distribution around 1 and $E(\alpha) = 1$, as specified in our assumptions. For simplicity, we use $\alpha \in \{\frac{2}{3}, 1, \frac{4}{3}\}$ in most analyses. The total number of time periods is $T = 3$. The range of price is $p \in [5, 20]$. The variable ordering cost and the fixed ordering cost are $c = 3$ and $k = 1$, respectively. The unit inventory holding cost is 10% of the average price and the unit backorder cost is 120% of the price ceiling. Moreover, the initial inventory level at the beginning of the first period x_1 is chosen such that $x_1 = 0$ if A does not exist; otherwise $x_1 \in A$ and the associated effectiveness is the lowest for all $x \in A$. We then generate the profit of an (s, S, p) policy and that of an (s, S, A, p) policy and compute the effectiveness by taking the ratio of the two profits. For the baseline parameters, we have $S_1 = 21$, $s_1 = 17$ and $x_1 = 18$. The effectiveness when $x_1 = 18$ is 99.8158%.

We choose this set of baseline parameters, because its associated special set A usually exists, when we conduct sensitivity analysis on a parameter. To make our numerical study more robust, we perform sensitivity analysis on all input parameters in the following subsections and examine how the effectiveness of an (s, S, p) policy changes, as we alter a particular input parameter.

3.2.2 Initial Inventory Level

We first test the influence of the initial inventory level on the effectiveness of an (s, S, p) policy. From the baseline parameters, we vary the initial inventory level at the beginning of the first period from 0 to 25 and plot the corresponding effectiveness in Figure 3-1. From the graph, we can see that the special set A_1 for the first period only contains one element at $x_1 = 18$ and the associated effectiveness is 99.8185%. At all the other initial inventory levels, the effectiveness remains 100%.

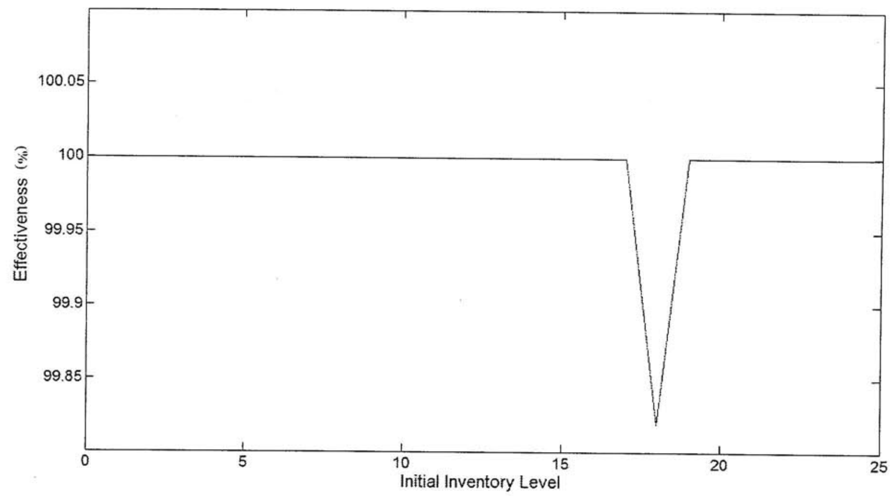

Figure 3-1: Effectiveness of an (s, S, p) Policy against Initial Inventory Level

We find that the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy may be less than 100%, when the initial inventory x_t at the beginning of period t falls in the special set A_t, provided A_t exists. There are generally three ways to make a difference between the two policies:

1. The special set A_1 for period $t = 1$ exists and $x_1 \in A_1$.

2. The special set $A_{t+1}(t > 1)$ exists and the initial inventory x_t $(x_t > S_t)$ in period t minus the actual demand w_t falls in A_{t+1}, i.e. $x_t - w_t \in A_{t+1}$.

3. The special set $A_{t+1}(t > 1)$ exists and an order is placed to increase the initial inventory $x_t(x_t < s_t)$ in period t to the inventory up to level S_t, which after meeting the actual demand w_t, falls in A_{t+1}, i.e. $S_t - w_t \in A_{t+1}$.

In general, the effectiveness reduces, when the input parameters satisfy a combination of two or three of the abovementioned conditions. Figure 3-2 shows an example of the three conditions, when we change the randomness parameters α in the set of the baseline parameters to be $\alpha \in \{\frac{1}{3}, \frac{4}{3}, \frac{4}{3}\}$ for the first period and $\alpha \in \{0.3, 1.1, 1.6\}$ for the next two periods. In this example, the special sets for the first two periods A_1 and A_2 both exist. At an initial inventory level of $x_1 = 22$, condition 2 is satisfied, as $x_1 > S_1$ and one instance of $x_1 - w_1 \in A_2$. The associated effectiveness is 99.9434%. At an initial inventory level of $x_1 = 19$, $x_1 \in A_1$ and one instance of $S_1 - w_1 \in A_2$. It satisfies both condition 1 and condition 3. The associated effectiveness hence becomes the lowest at $x_1 = 19$, which equals 99.6853%. Nevertheless, we can see that the (s, S, p) policy in this case is still highly effective.

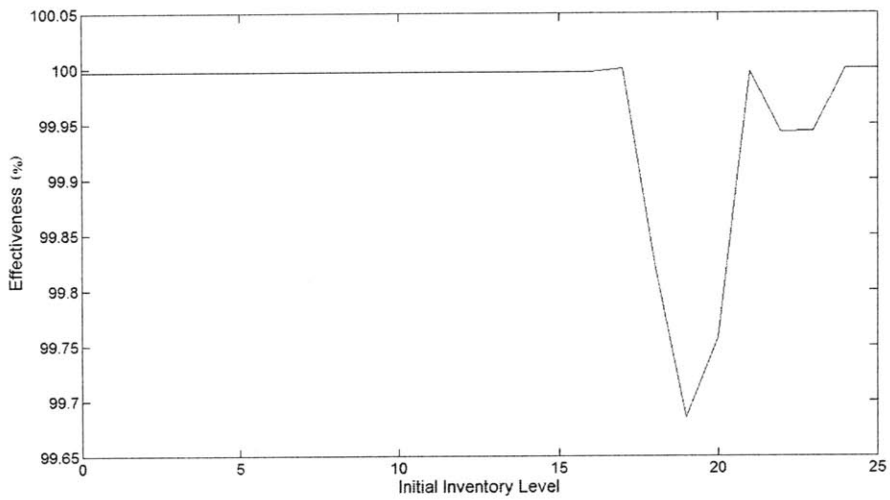

Figure 3-2: Effectiveness of an (s, S, p) Policy with a Combination of Two Conditions

3.2.3 Total Number of Periods

We then check the effect of the total number of periods on the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy. From the baseline parameters, we change the total number of periods, T, from 1 to 20. The graph of the effectiveness against the total number of periods is shown in Figure 3-3.

From the graph, we have three observations. First, when there is only one period, i.e. $T = 1$, the effectiveness is always 100%. This result is in accordance with a study by Simchi-Levi et al. (2005), which states for a single period, single product, periodic review, stochastic demand model, an (s, S, p) policy is always optimal. Thus, the special set A does not exist and an (s, S, p) policy is equivalent to an (s, S, A, p) policy. Second, after testing with different sets of parameters, we find the effectiveness tends to be the lowest when the special set A exists for the first time, as T increases. In this case, A first appears when $T = 2$ and the associated effectiveness reaches the lowest point at 99.7394%. Last, an (s, S, p) policy becomes more effective as T increases, provided A exists. If A does not exist, the effectiveness is 100%. Otherwise if A exists, the effectiveness approaches 100%, as $T \to \infty$. This result coincides with that from Chen and Simchi-Levi (2004b), proving that an (s, S, p) policy is optimal for an infinite horizon model.

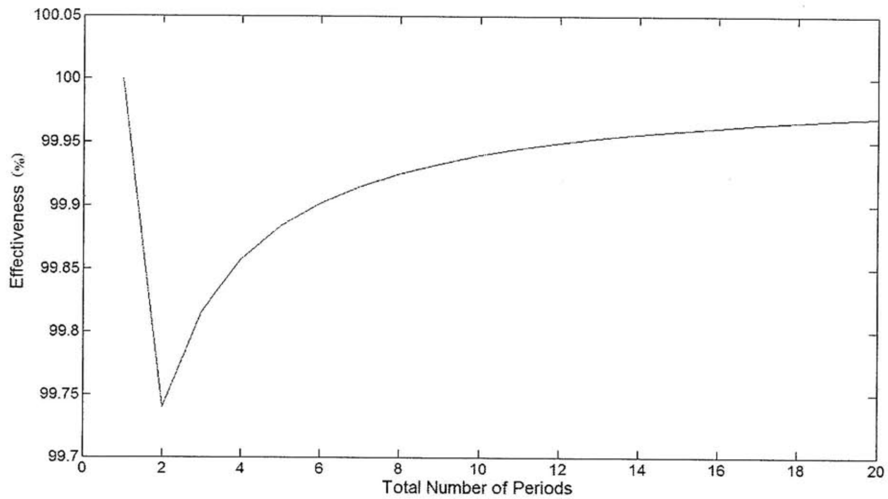

Figure 3-3: Effectiveness of an (s, S, p) Policy against Total Number of Periods

3.2.4 Ordering Costs

For the sensitivity analysis of the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy on the ordering costs, we first look at the fixed ordering cost and then proceed to the variable ordering cost. From the baseline parameters, we increase the fixed ordering cost from 0 to 25 with an incremental interval equal to 1. Figure 3-4 shows the graph of the effectiveness of an (s, S, p) policy against the fixed ordering cost. We also change the variable ordering cost to 2 and then perform the same analysis. The result is plotted in Figure 3-5.

The results show that, in general, it is more likely for an (s, S, p) policy to be equivalent to an (s, S, A, p) policy, when the fixed ordering cost becomes very large. However, we have found some counterexamples, which demonstrate that it is still possible for the special set A to exist, when the ratio of the fixed ordering cost to the price is very high. Moreover, when the fixed ordering cost is 0, the effectiveness is always 100% and $S_t = s_t$ for every period $t \in T$. This result is in line with Federgruen and Heching (1999), who prove a base stock list price policy is optimal, i.e. $S_t = s_t$, when there is no fixed ordering cost. Another interesting observation

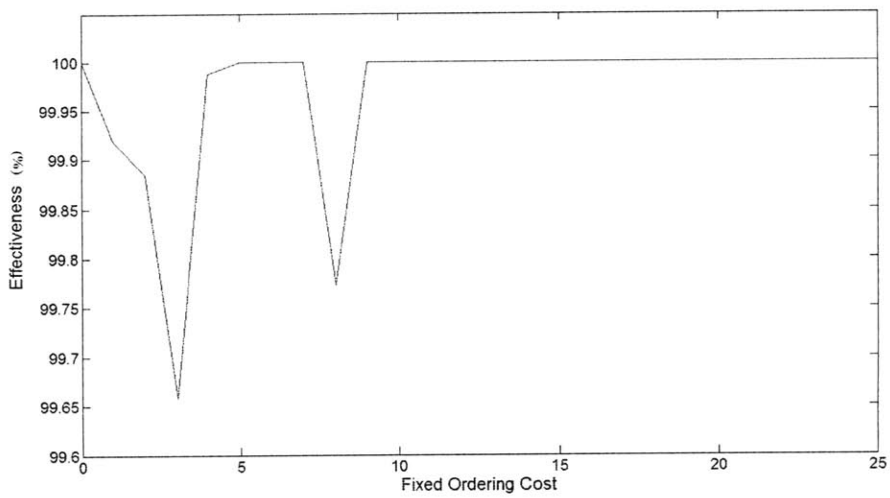

Figure 3-4: Effectiveness of an (s, S, p) Policy against Fixed Ordering Cost with $c = 3$

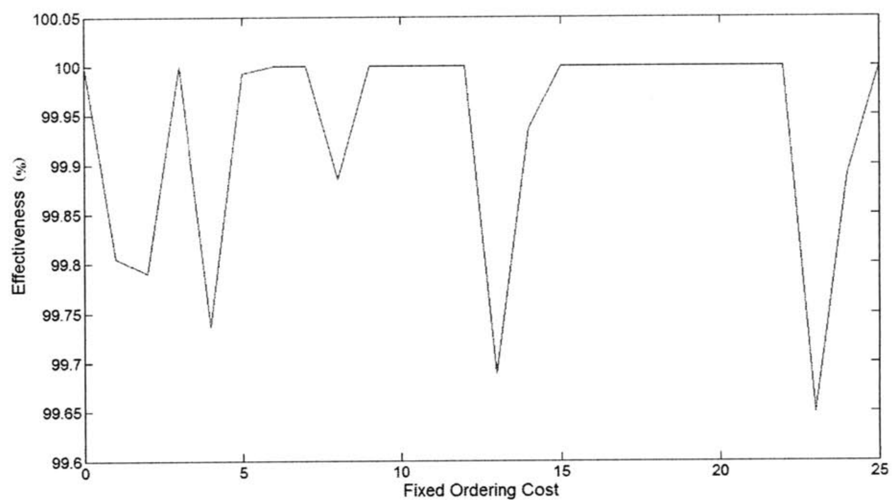

Figure 3-5: Effectiveness of an (s, S, p) Policy against Fixed Ordering Cost with $c = 2$

is that as the fixed ordering cost increases, the order up-to-level S_t may increase or remain unchanged whereas the reorder point s_t may decrease or remain unchanged for period t. Whenever a new set of (s_t, S_t) appears and A_t also exists, the effectiveness tends to be the lowest. In addition, for the same set of (s_t, S_t), the effectiveness tends to be non-decreasing, as the fixed ordering cost increases.

For the sensitivity analysis on the variable ordering cost, we increase the variable ordering cost from 0 to 5.5 (the ratio of the variable ordering cost to the price floor changes from 0% to 110%) with an incremental interval equal to 0.5. The graph of the effectiveness against the variable ordering cost is shown in Figure 3-6. The graph shows that when the variable ordering cost is close to 0 or very high (greater than the price floor in this case), the effectiveness becomes 100% and there is no difference between an (s, S, p) policy and an (s, S, A, p) policy. In particular, when there is no variable ordering cost, the sets of (s_t, S_t) are the same for all period t and an (s, S, p) policy is always optimal. Furthermore, unlike the fixed ordering cost, as the variable ordering cost increases, both the reorder point s_t and the order up-to-level S_t are non-increasing for every period $t \in T$.

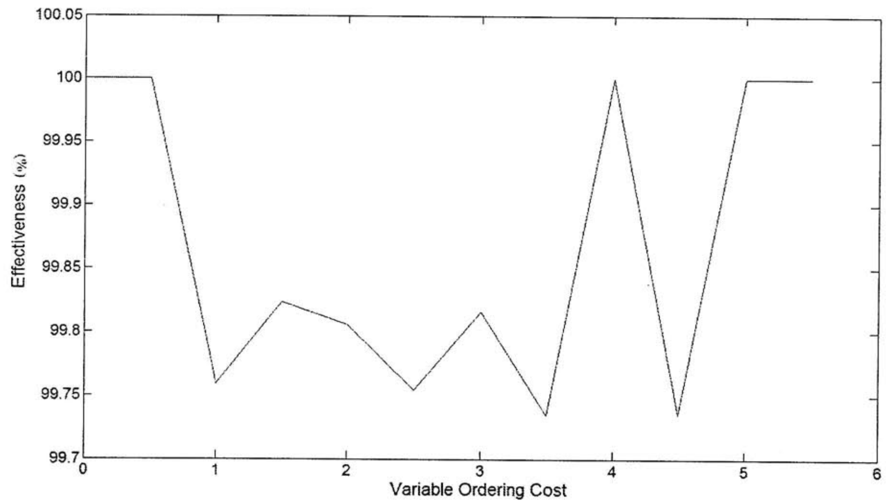

Figure 3-6: Effectiveness of an (s, S, p) Policy against Variable Ordering Cost

3.2.5 Inventory Holding Cost

Next, we perform sensitivity analysis on the inventory holding cost to examine its impact on the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy. In the set of baseline parameters, we specify the unit inventory holding cost as 10% of the average price. In this sensitivity analysis, we change the unit inventory holding cost from 0% to 80% of the average price with an incremental interval equal to 5% and compute the corresponding effectiveness. The graph of the effectiveness against the percentage of inventory holding cost is shown in Figure 3-7.

From the graph, we can see that when the inventory holding cost is close to 0 or very high (greater than 50% in this case), the effectiveness of an (s, S, p) policy is 100%. We notice that for this set of parameters, when the inventory holding cost is about 45% of the average price, the effectiveness can be as low as 98.2017%, though it is still considered to be highly effective. Similar to the variable ordering cost, as the inventory holding cost increases, both the reorder point s_t and the order up-to-level S_t are non-increasing with s_t decreasing first for every period $t \in T$. Figure 3-8 illustrates how the reorder point s_1 and the order up-to-level S_1 of the first period change, as the inventory holding cost increases. We also observe that, for the same S_t, the associated effectiveness is non-increasing, as the inventory holding cost increases. In addition, we find the effectiveness of an (s, S, p) policy is insensitive to change in the backorder cost, because backorder is very costly and undesirable.

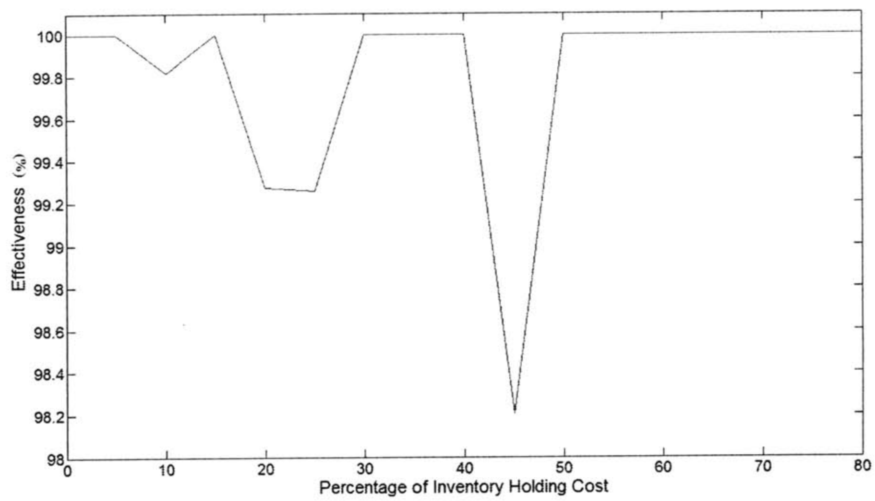

Figure 3-7: Effectiveness of an (s, S, p) Policy against Percentage of Inventory Holding Cost

3.2.6 Price Range

In the set of baseline parameters, the price ranges from 5 to 20. In this subsection, we conduct sensitivity analysis on both the price floor and the price ceiling in order to test their effects on the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy. It should be noted that since the unit inventory holding cost is set to be 10% of the average price, the inventory holding cost is also affected, as the price range changes.

In Figure 3-9, we plot the graph of the effectiveness against the price floor, which

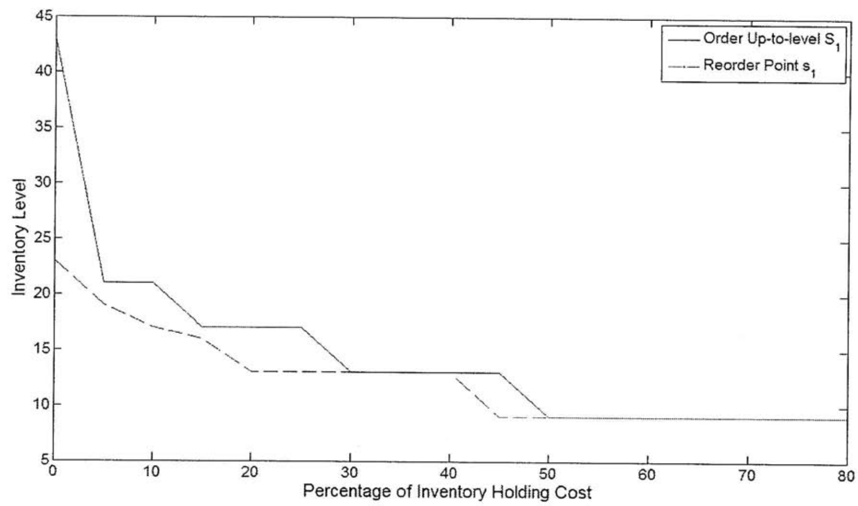

Figure 3-8: Movements of Reorder Point s_1 and Order Up-to-level S_1 against Percentage of Inventory Holding Cost

increases from 3 to 20 with an incremental interval of 1. Notice the range is chosen because we want to make sure the price floor is no less than the variable ordering cost, which equals 3, and no greater than the price ceiling, which equals 20. The graph shows that the effectiveness is quite steady around 99.8% until it jumps to 100%, when the price floor reaches 11. From 11 to 20, the special set A does not exist and hence the effectiveness remains 100%. When the price floor increases from 3 to 10, the sets of reorder point s_t and the order up-to-level S_t remain unchanged and the best prices corresponding to S_1 and S_2 are around 10, which are within the price range. However, the best prices of S_1 and S_2 in this case are bounded and equal to the price floor, when the price floor increases to 11 and above. As a result, an (s, S, p) policy is equivalent to an (s, S, A, p) policy, when the price floor increases from 11 to 20.

We also change the price ceiling from 5 to 40 with an incremental interval of 1 from the set of baseline parameters and plot the effectiveness against the price ceiling in Figure 3-10. The graph shows the effectiveness is quite steady around 99.8%, when the price ceiling ranges from 11 to 27. This is because the special set A exists and the set of best price remain unchanged, as the price ceiling varies within this range. When the price ceiling becomes greater than 27, the best prices increase, which is

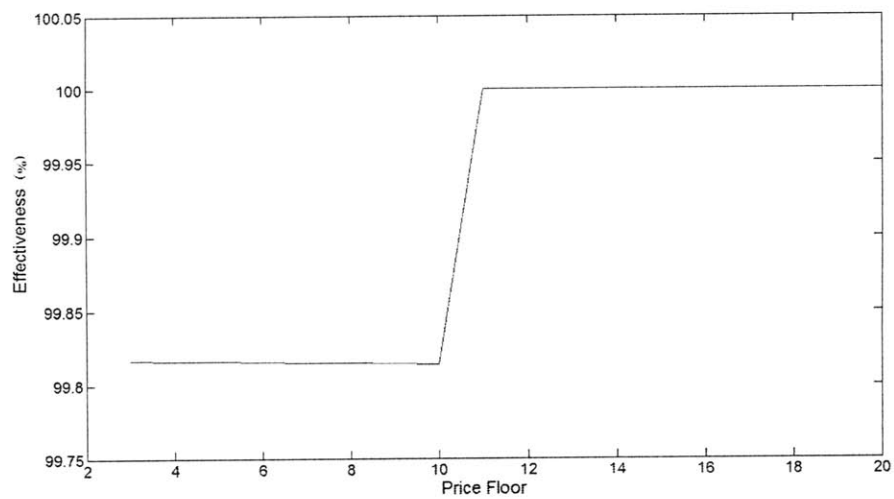

Figure 3-9: Effectiveness of an (s, S, p) Policy against Price Floor

partially due to the increase in the inventory holding cost, and the special set A no longer exists. As a result, the effectiveness becomes 100%. In particular, when the price floor equal to the price ceiling, the price and the expected demand become fixed for all the periods and the associated effectiveness is hence always equal to 100%.

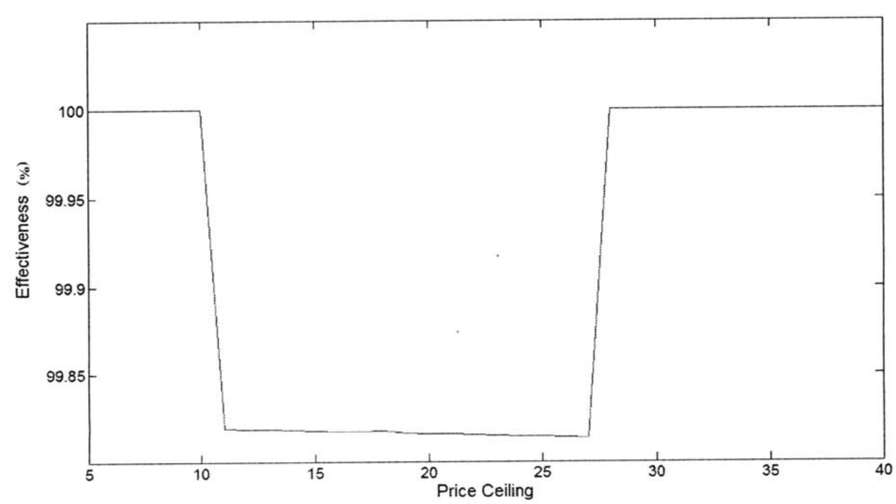

Figure 3-10: Effectiveness of an (s, S, p) Policy against Price Ceiling

3.2.7 Demand Function Parameters

After looking at the inventory, time, cost and price parameters, we turn our attention to the demand function parameters, namely a and b. From $D_t(p_t) = ap_t^{-b_t}$, we know demand is an increasing function of parameter a, but a decreasing function of parameter b. We first increase the parameter a from 100 to 2000 with an incremental interval of 50 and plot the graph of the effectiveness in Figure 3-11.

The graph shows that when the parameter a is very small, the effectiveness of an (s, S, p) policy is equal to 100%. This is because if a is too small, the demand range would be very narrow and the special set A does not exist. We observe that both the reorder point s_t and the order up-to-level S_t increase for period t, as the parameter a increases. For this set of parameters, we can see that the effectiveness fluctuate less and becomes more stable around 99.9%, as the parameter a increases. However, in general, we do not observe any trend of the effectiveness on the parameter a, when it is large enough for the special set A to exist.

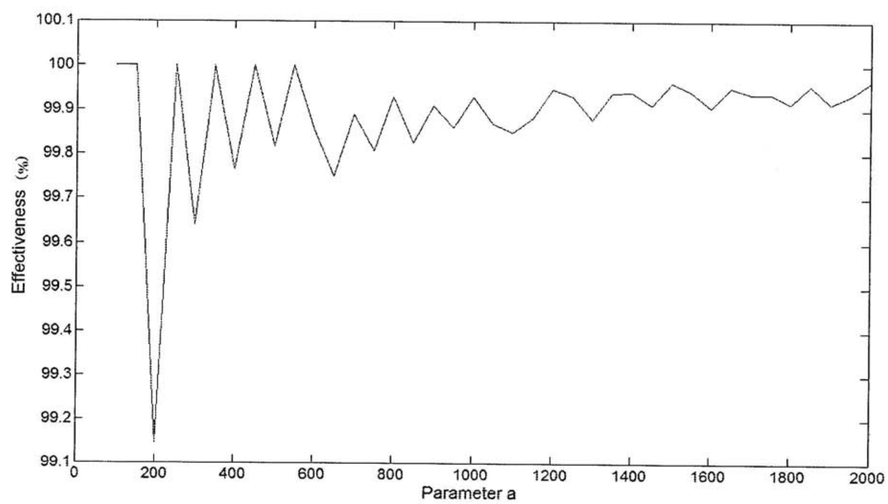

Figure 3-11: Effectiveness of an (s, S, p) Policy against Parameter a

We next increase the parameter b from 1.0 to 3.0 with an incremental interval of 0.1. The graph of the effectiveness against the parameter b is shown in Figure 3-12. We can see from the graph, when b is very close to 1, the effectiveness is 100% and $S_t = s_t$ for every period t, which indicates a baseline list price policy is optimal. When b is very large (greater than 2.7 in this case), the effectiveness is also 100%, because the demand range becomes very narrow and the special set A does not exist. Unlike the parameter a, both the reorder point s_t and the order up-to-level S_t decrease, as the parameter b increases. For this set of parameters, the effectiveness is as low as

98.1635% when $b = 2.5$. This is because the special sets $A_1$and A_2 both exist and the associated initial inventory level satisfy the conditions 1 and 3 abovementioned in section 3.2.2. In general, there is no noticeable trend of the effectiveness on the parameter b.

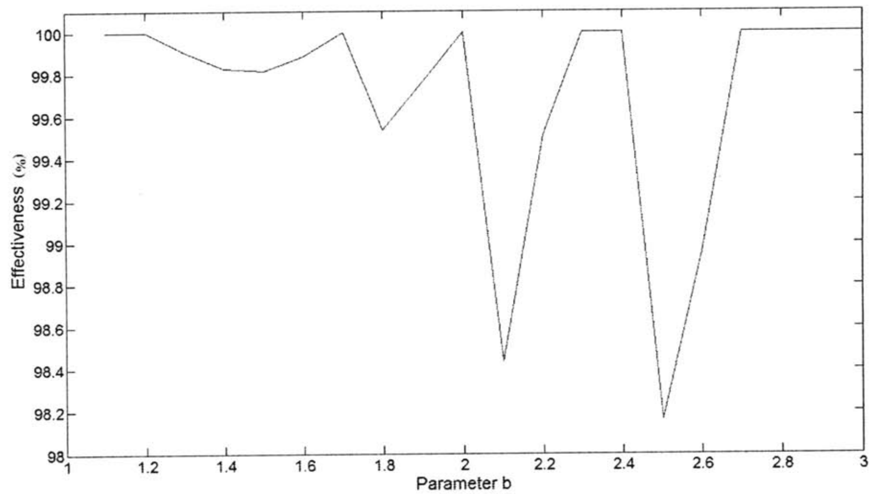

Figure 3-12: Effectiveness of an (s, S, p) Policy against Parameter b

Another interesting question to ask is whether the effectiveness would be lower if different demand functions are used in different periods. We find that the effectiveness behaves similarly as the uniform demand function case. For example, from the baseline parameters, if we change the demand parameters to be $(a_1, b_1) = (1500, 2.5)$ for the first period, $(a_2, b_2) = (1000, 2.0)$ for the second period and $(a_3, b_3) = (500, 1.5)$ for the last period, the lowest effectiveness equals 99.7589% with an initial inventory of 22, which is still highly effective.

3.2.8 Coefficient of Variation of Randomness

Finally, we perform sensitivity analysis on the coefficient of variation of the randomness parameter α. The coefficient of variation (CV) is a measure of statistical dispersion of randomness, which is defined as the ratio of the standard deviation to the mean, i.e. $cv = \frac{\sigma(\alpha)}{E(\alpha)}$. Since $E(\alpha) = 1$ in our stochastic multiplicative demand process, the coefficient of variation is effectively equal to the standard deviation, $\sigma(\alpha)$.

We perform two tests on the coefficient of variation of randomness. First, from the baseline parameters, we fix the number of elements in the set of α and increase the coefficient of variation from 0 to 0.95 with an incremental interval of 0.05. The graph of the effectiveness against the coefficient of variation of randomness is shown in Figure 3-13. Second, we increase the number of elements in the set of α from 1 to 10 with randomly generated elements and plot the graph of the effectiveness in Figure 3-14. Not surprisingly, the two graphs show that the effect of randomness on the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy is random. However, when the coefficient of variation is equal to 0, the effectiveness of an (s, S, p) policy is 100%. This is because an (s, S, p) policy is always optimal if there is no randomness in the demand process.

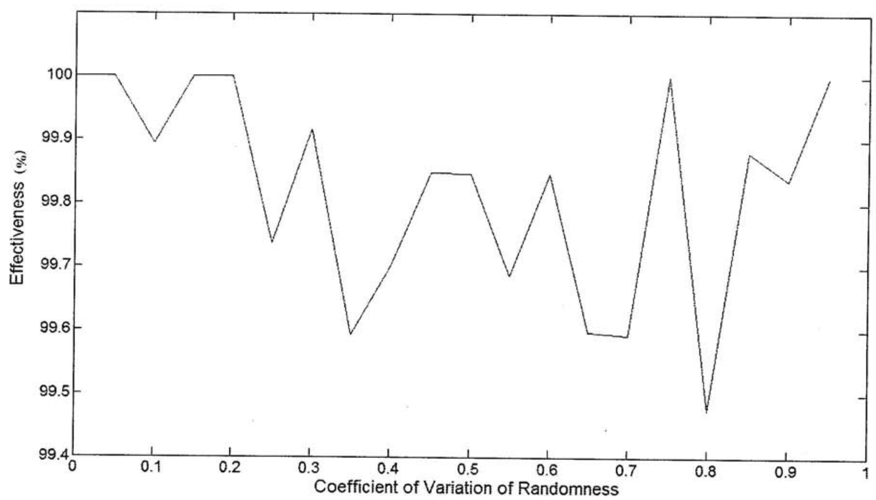

Figure 3-13: Effectiveness of an (s, S, p) Policy against Coefficient of Variation of Randomness with Fixed Number of Elements in the Set of α

3.3 Overall Observation

In general, the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) depends on the existence of the special set A_t and whether the initial inventory x_t of period t falls into A_t. If this condition does not hold for all periods in the horizon, then the effectiveness is always 100% and an (s, S, p) policy is equivalent to an (s, S, A, p) policy. Otherwise, the effectiveness is less than 100%. The existence of the special set A_t may be affected by the time parameters, cost parameters, price parameters and

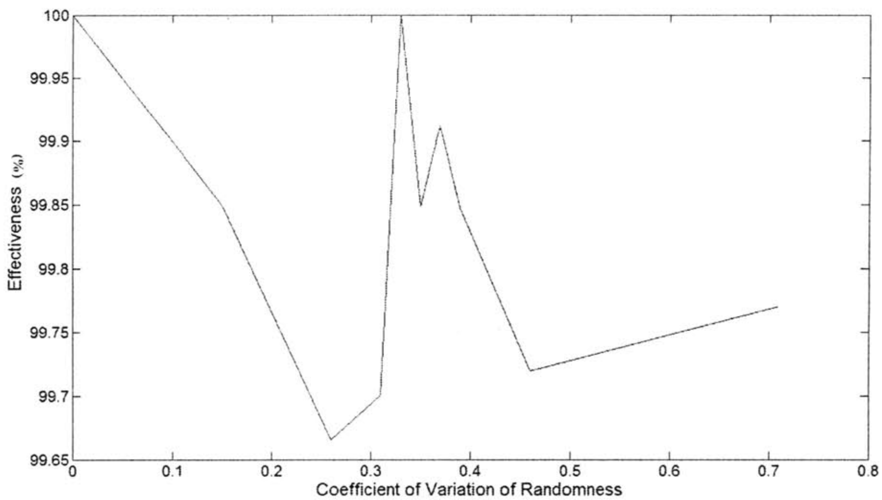

Figure 3-14: Effectiveness of an (s, S, p) Policy against Coefficient of Variation of Randomness with Different Number of Elements in the Set of α

the demand parameters, as we have discussed above.

We find, in most cases, the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy is very close to 100%. In Figure 3-15, we plot the three dimensional graph of the g function values against the inventory level after placing an order y and the expected demand d for the first period with the set of baseline parameters. The graph seems to be quite flat around its global optimal value. We zoom in at the global optimum, which is displayed in Figure 3-16. The image shows that the g function, which has been proven to be symmetric k-concave, is not concave and there are some local optima around the global optimum. However, the values of the local optima are close to that of the global optimum. It explains why an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy for the baseline parameters. In some extreme cases, when we use different sets of carefully selected input parameters in different periods, we can obtain an effectiveness as low as 95%, though it implies that an (s, S, p) policy is still very effective. Thus, we can conclude that empirically an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy in this stochastic multiplicative demand model.

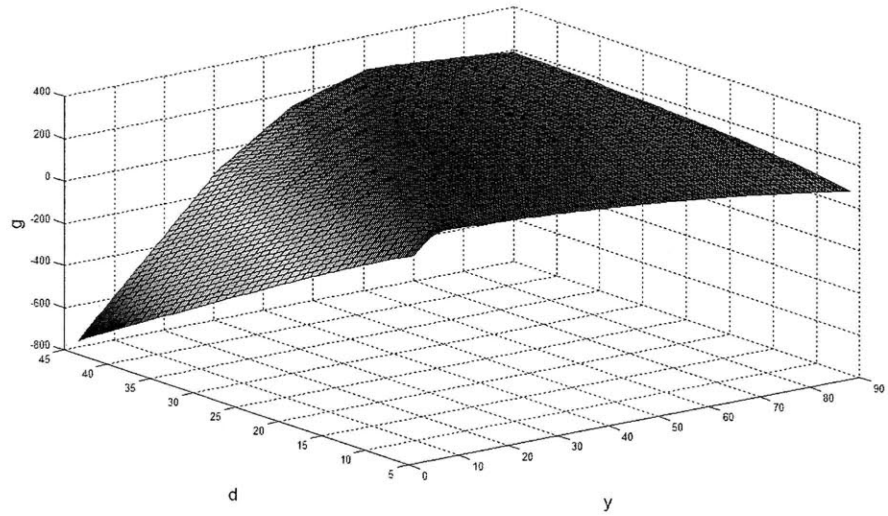

Figure 3-15: Example of g Function Values against Inventory-Demand Pairs

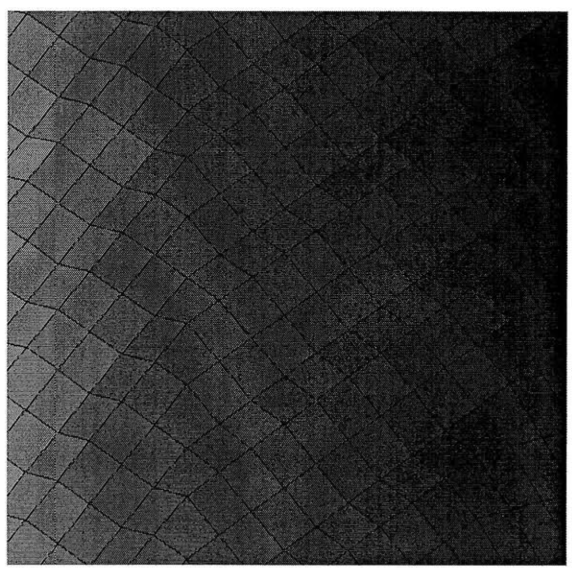

Figure 3-16: Zoomed in Image at the Global Optimum of Figure 3-15

Chapter 4

Comparisons with Alternative Policies

4.1 Comparison with a Deterministic Policy

We have seen that, as a relatively simple policy, an (s, S, p) policy closely approximates an (s, S, A, p) policy in our model. But what is its performance in comparison with some even simpler policies? If an (s, S, p) policy does not outperform an simpler alternative policy, we may want to use the alternative one to approximate an (s, S, A, p) policy, because it is simpler and takes less time to compute.

In this section, we first compare the performance of an (s, S, p) policy with that of a deterministic policy. In this policy, we assume the demand process is always deterministic. If there is no randomness in the demand, an (s, S, p) policy is always optimal with the reorder point s_t equal to the order up-to-level S_t for every period t. We call this policy a deterministic policy. Of course, the deterministic policy is different from the original (s, S, p) policy, as they are based on different demands. We then compute the total profit of the deterministic policy under stochastic demand. In Table 4.1, we present the profits of the two policies and their computation time, respectively, with the set of baseline parameters.

The results show that, for the baseline parameters, the profit of the deterministic policy is only 69.8% of that of the (s, S, p) policy. In other words, the (s, S, p) policy

Table 4.1: Comparison between an (s, S, p) Policy and a Deterministic Policy

Policy	Total Profit ($)	Computation Time (s)
Deterministic	245.86	0.372
(s, S, p)	352.24	0.527

generates significantly more profit than the deterministic policy, though the latter policy is faster to compute. Indeed, after testing with different parameters, we find an (s, S, p) policy is strictly superior to a deterministic policy. This is because in a deterministic policy, the price and inventory up-to-level are set in such a way that the inventory level after an order is placed is equal to the expected demand. As the expected randomness is equal to 1, there must be cases under stochastic demand, in which the demand cannot be met with the inventory at stock and the additional demand has to be backordered, which is very costly. As a result, a deterministic policy performs worse than an (s, S, p) policy. In addition, we find the profit of a deterministic policy tends to be closer to that of an (s, S, p) policy when the coefficient of variation of the randomness in demand is smaller.

4.2 Comparison with a Fixed Price Policy

We next formulate another alternative policy and compare it with an (s, S, p) policy. In this alternative policy, we first ignore the randomness and choose a price that maximizes the expected revenue of each period. As there is a one-to-one relationship between the expected demand d_t and the price p_t, i.e. $d_t = ap_t^{-b_t}$, the expected revenue $R_t = d_t p_t = ap_t^{1-b_t}$. Since the demand function parameters $a_t > 0, b_t > 1$, clearly R_t is maximized when p_t is set to be the price floor, i.e. $p_t = \underline{p}_t$. For the baseline parameters, we thus let $p_t = 5$ for all the three periods. By fixing the price, we effectively fix the expected demand d_t in each period. We then consider the randomness in demand and the optimal policy in this case becomes the well-known (s, S) policy. We call this alternative policy a fixed price policy. Table 4.2 shows the profits of the two policies and their computation time, using the set of baseline parameters,.

Table 4.2: Comparison between an (s, S, p) Policy and a Fixed Price Policy

Policy	Total Profit ($)	Computation Time (s)
Fixed Price	220.35	0.034
(s, S, p)	352.24	0.527

We can see from the table that the fixed price policy generates only 62.6% of the profit of the (s, S, p) policy in this case. The reason is this policy aims to maximize the expected revenue instead of the expected profit and it hence performs poorly when the cost is high relative to the revenue. On the other hand, when the cost plays a less significant role in the profit, a fixed price policy can be very effective in terms of profit. Table 4.3 shows another comparison between the two policies, when we lower the variable ordering cost from 3 to 1 in the set of baseline parameters. This time the profit of the fixed price policy becomes 99.2% of that of the (s, S, p) policy. Furthermore, though an (s, S, p) policy generates greater profit than a fixed price policy, the latter policy requires much less computation time, as we no longer need to spend much time to determine the best price. In this case, the fixed price policy takes only about 6% of the computation time of the (s, S, p) policy. Thus, when the cost revenue ratio is low, a fixed price policy can become a good alternative to an (s, S, p) policy and it is a much more efficient policy.

Table 4.3: Comparison between an (s, S, p) Policy and a Fixed Price Policy with Low Ordering Cost

Policy	Total Profit ($)	Computation Time (s)
Fixed Price	478.35	0.036
(s, S, p)	482.05	0.527

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Proposed Algorithm for (s, S, p) Policy

5.1 Proposed Algorithm

We have shown in Chapter 3 that an (s, S, p) policy is empirically highly effective relative to an (s, S, A, p) policy under stochastic multiplicative demand, in which the latter policy is optimal. The question now is how to efficiently compute an (s, S, p) policy. Chen and Simchi-Levi (2004a) suggest a natural dynamic programming procedure. But solving this dynamic programming procedure imposes an empirical challenge. The enumeration method we have used previously is time consuming, when the total number of periods is large or the model is further discretized. In this section, we propose an fairly efficient algorithm using simulated annealing and modified binary search to solve the dynamic programming procedure in the original continuous model.

To solve the dynamic programming procedure we first need to find the global optimum of the g function, i.e. the best (y_t, d_t) pair for period t, which corresponds to the order up-to-level S_t and its associated best price p_t. We feel that gradient methods are not suitable for this problem, because in the dynamic programming procedure, the g function calls the v function iteratively and thus it is difficult to obtain an explicit expression for the gradient of the g function. After observing the plot in Figure 3-15, we can see that though the graph is not convex, it has a rather smooth shape with some, but not too many local optima. This inspires us to use an simulated annealing algorithm to determine the global optimum.

Simulated annealing is a popular generic probabilistic heuristic algorithm for optimization problems. From a starting point, the algorithm updates its current position in every iteration, by searching randomly in its neighborhood. If a neighbor point has a better value than the current point, the algorithm moves to this neighbor point at the end of the iteration. Otherwise, it still has a certain probability to move to this neighbor point. This probability depends on the cooling scheme of the algorithm, which typically starts from a very high energy state and gradually moves to a stable, low energy state. The algorithm terminates when the energy state is low enough. In general, simulated annealing is proven to be a more efficient algorithm than enumeration in solving optimizations with multiple local optima.

For the simulated annealing algorithm used to solve the original continuous problem, we use an initial temperature of 10000 and adopt a cooling scheme with $T_0^{n+1} = 0.95T_0^n$, where T_0 and n represent the temperature and iteration number, respectively. We believe this cooling scheme is effective enough for the algorithm to jump out of most local optima and return a near-optimal solution, if not the global optimum itself. For the starting point, the initial expected demand d_0 is set to be the mid-point in the demand range and the initial inventory up-to-level y_0 is set to be equal to d_0. The initial g function value at the starting point is hence equal to $g(y_0, d_0)$. The neighbor (y, d) pairs are randomly selected with the criteria $|y_{neighbor} - y_{current}| \leq 1$ and $|d_{neighbor} - d_{current}| \leq 1$ and the new (y, d) pair must be within the feasible set. After determining the best (y_t, d_t) pair for period t, we set the order up-to-level S_t to $S_t = y_t$.

Next, we need to locate the reorder point s_t for period t. Since the g function is symmetric k-concave, we have $g(s_t, d(s_t)) = g(S_t, d(S_t)) - k$, where $d(s_t)$ and $d(S_t)$ are the best expected demands associated with s_t and S_t, respectively. Let $g_s = g(s_t, d(s_t))$. If the g function is k-concave, as in the additive demand model, there is only one point smaller than S_t with a g function value of g_s and we can use binary search to find s_t easily. However, in our multiplicative demand model, the g function is symmetric k-concave and there are multiple points smaller than S_t, which have g function values equal to g_s. The reorder point s_t corresponds to the smallest y_t that has $g(y_t, d(y_t)) = g_s$.

We decide to use a modified binary search to find the reorder point s_t. Since $s_t \in (0, S_t)$, we let the upper bound $\overline{s}_t$ be S_t and the lower bound $\underline{s}_t$ be 0 at the beginning of the modified binary search. In each iteration, we take the mid-point $s_{mid} = \frac{\overline{s}_t + \underline{s}_t}{2}$ and check its associated g function value, $g(s_{mid}, d(s_{mid}))$. If $g(s_{mid}, d(s_{mid})) > g_s$, we reset the upper bound to $\overline{s}_t = s_{mid}$ and continue the modified binary search. Otherwise, if $g(s_{mid}, d(s_{mid})) \leq g_s$, we check whether there is any point $s' \in (0, s_{mid})$ with a g function value greater than or equal to g_s, using simulated annealing. If $g(s', d(s')) \geq g_s$, we reset the upper bound to $\overline{s}_t = s'$ and continue the modified binary search. Otherwise if the maximum g function value in the range of $(0, s_{mid})$ is smaller than g_s, we reset the lower bound $\underline{s}_t = s_{mid}$ and continue the modified binary search. The modified binary search terminates when $g(s_{mid}, d(s_{mid})) = g_s$ and $\underline{s}_t = s_{mid}$. We believe our algorithm of simulated annealing and modified binary search is more efficient in determining the (s, S, p) policy of the continuous model.

5.2 Empirical Performance

The empirical results and computation time of our proposed algorithm vary slightly, because each time the randomly generated neighbor pairs may be different, which affect the performance of the algorithm. It should be noted that our proposed algorithm only returns a good approximation of the true (s, S, p) policy, since simulated annealing is a heuristic method. However, as the g function has a flat shape at its optimum, our near-optimal solution works very well and generates a profit that is almost identical to that of the true (s, S, p) solution. In Figure 5-1, we compare the profit of our proposed algorithm with that of the enumeration method in 20 different instances and the results show that our proposed algorithm is highly effective. Moreover, for the set of baseline parameters, our proposed algorithm generally finishes within 10 minutes whereas it requires the further discretized model more than

5 hours to achieve the same level of accuracy. It shows that our proposed algorithm is indeed much more efficient than the enumeration method. In addition, it is also worth noting that the choice of the cooling scheme has a significant impact on the empirical performance of the algorithm. If a faster cooling scheme is adopted, the computation time will decrease. But the quality of the approximated (s, S, p) policy is likely to be compromised.

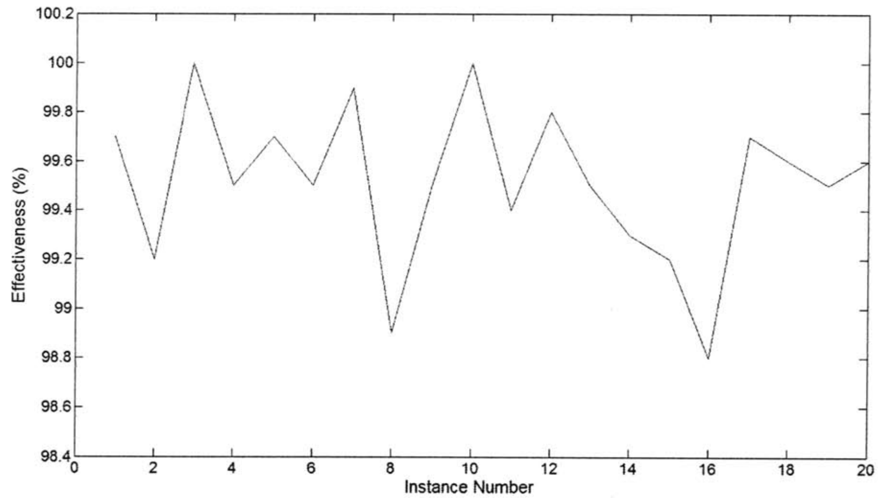

Figure 5-1: Effectiveness of Proposed Algorithm relative to Enumeration Method

Chapter 6

Concluding Remarks

6.1 Conclusion

In this thesis, we investigate the effectiveness of an (s, S, p) policy for coordinating inventory control and pricing strategies. We show empirically that an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy in the single product, periodic review, finite horizon model with stochastic multiplicative demand and fixed ordering cost, in which the latter policy is optimal. We find that the effectiveness depends on whether the initial inventory x_t of period t falls into A_t, if the special set $A_t \in [s_t, \frac{s_t+S_t}{2}]$ exists. If this condition is satisfied for at least one period in the time horizon, the effectiveness is less than 100%. Otherwise, the effectiveness equals 100% and an (s, S, p) policy is equivalent to an (s, S, A, p) policy. Our empirical results show that the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy is equal to 100% in most cases, and even when it is less than 100%, the effectiveness is generally above 95%. We also conduct sensitivity analysis on the input parameters of the model, which jointly affect the effectiveness of an (s, S, p) policy. In particular, we find the effectiveness increases as the total number of periods increases. When the fixed ordering cost is 0, the reorder point s_t is equal to the order up-to-level S_t for all periods, which indicates that a base stock list price policy is optimal. Moreover, we compare the performances of a deterministic policy and a fixed price policy with that of an (s, S, p) policy and find the latter policy is superior in terms of profit, though the fixed price policy can be an effective and more efficient policy, when the cost is insignificant in comparison with the revenue. In addition, we propose an algorithm with simulated annealing and modified binary search to determine the (s, S, p) policy for the continuous model.

The fact that an (s, S, p) policy is highly effective relative to an (s, S, A, p) policy has important practical meanings. Chen and Simchi-Levi (2004a, 2004b, 2006) prove that an (s, S, p) policy is optimal for all kinds of single product, stochastic demand models with fixed ordering cost, with the exception of the periodic review, finite horizon model with stochastic non-additive demand model, in which this policy is highly effective relative to the optimal (s, S, A, p) policy. Furthermore, an (s, S, p) policy is also more intuitive and simpler than an (s, S, A, p) policy. Thus, we can apply the (s, S, p) policy in industry practice to help supply chain managers coordinate their inventory control and pricing strategies and achieve superior profits. Our proposed algorithm can then be used to determine the (s, S, p) policy for the supply chain managers.

6.2 Limitations and Extensions

It is necessary to discuss some of the limitations and possible extensions of our work in this thesis. First, our tests on the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy are on the integer level discretized model instead of the original continuous model. As Chen and Simchi-Levi (2004b) point out, the structure of the optimal policy may be different after discretization of the original continuous model. Since every numerical study has to discretize the inventory variables and the pricing variables to a certain extent, the results from the numerical study serve only as an approximation. Second, although we have tried to conduct our numerical study with many different sets of input parameters, it is infeasible to cover all cases. Thus, it is entirely possible that, in some extreme cases, the effectiveness of an (s, S, p) policy is much lower. Third, in this thesis, we only examine the effectiveness of an (s, S, p) policy under stochastic multiplicative demand. Chen and Simchi-Levi (2004a) proves that an (s, S, A, p) policy is optimal under stochastic general demand.

Thus, in future studies, the effectiveness of an (s, S, p) policy should also be tested, when the demand is stochastic, non-additive and non-multiplicative. Nevertheless, we feel our conclusion that an (s, S, p) policy is highly effective relative to an (s, S, p) policy still holds in general. Perhaps the best way to overcome these problems is to investigate the effectiveness of an (s, S, p) policy relative to an (s, S, A, p) policy theoretically, though this goal could be difficult to achieve. In further works, it would also be interesting to find out how to determine the special set A efficiently in the original continuous model.

In addition, there are some limitations in the model formulated by Chen and Simchi-Levi (2004a), which is used in this thesis. The model ignores capacity constraint and lead time, which are two important factors in supply chain management. Indeed, there are examples showing that an (s, S, A, p) policy fails to be optimal in the model under capacity constraint. With lead time, the price when an order is placed may be different from the price when the product arrives. Thus, the structure of the optimal policy would be affected, when lead time is considered. Moreover, Chen and Simchi-Levi (2004a) also suggest that it may be more appropriate to use a non-decreasing function of price as the backorder cost. Further work should be done to incorporate these factors in the model and analyze its optimal policy.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] Agrawal, V., A. Kambil. Dynamic pricing strategies in electronic commerce. Working paper, Stern Business School, New York University, New York, 2000.

[2] Amihud, Y., H. Mendelson. Price smoothing and inventory. Review of Economics Studies, Vol. 50, 1983, pp. 87-98.

[3] Bertsekas, D. *Dynamic Programming and Optimal Control*, Vol. 1. Athena Scientific, Belmont, MA, 1995.

[4] Chan, L. M. A., Z. J. Max Shen, D. Simchi-Levi, J. Swann. Coordination of pricing and inventory decisions: A survey and classification. D. Simchi-Levi, S. D. Wu, Z. J. Max Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Chap. 9. Kluwer, Boston, MA, 2004.

[5] Chao, X., S. X. Zhou. Joint inventory-and-pricing strategy for a stochastic continuousreview system. *IIE Transactions*, 2006, pp. 401-408.

[6] Chen, X., D. Simchi-Levi. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: the finite horizon case, *Operations Research*, Vol. 52 (6), 2004a, pp 887-896.

[7] Chen, X., D. Simchi-Levi. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: the infinite horizon case, *Mathematics of Operations Research*, Vol. 29 (3), 2004b, pp. 698-723.

[8] Chen, X., D. Simchi-Levi. Coordinating inventory control and pricing strategies: The continuous review model, *Operations Research Letters*, Vol. 34, 2006, pp. 323-332.

[9] Cook, T. Creating competitive advantage in the airline industry. Seminar sponsored by the MIT Global Airline Industry Program and the MIT Operations Research Center, Massachusetts Institute of Technology, Boston, MA, 2000.

[10] Elmaghraby, W., P. Keskinocak. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, Vol. 49, 2003, pp. 1287-1309.

[11] Federgruen, A., A. Heching. Combined pricing and inventory control under uncertainty. *Operations Research*, Vol. 47(3), 1999, pp. 454-475.

[12] Feng, Y., F. Chen. Joint pricing and inventory control with setup costs and demand uncertainty. Working paper, Chinese University of Hong Kong, Hong Kong, 2002.

[13] Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, Vol. 40, 1994, pp. 999-1020.

[14] Iglehart, D. Optimality of (s, S) policies in the infinite-horizon dynamic inventory problem. *Management Science*, Vol. 9, 1963, pp. 259-267.

[15] Karlin, S., C. R. Carr. Prices and optimal inventory policy. *Studies in Applied Probability and Management Science*. K. Arrow, S. Karlin, H. Scarf, eds. Stanford University Press, Stanford, CA, 1962.

[16] Lau, A. H., H. Lau. The newsboy problem with pricedependent demand distribution. *IIE Transactions*, Vol. 20, 1988, pp. 168-175.

[17] Leibs, S. Ford heads the profits. *CFO The Magazine*, Vol. 16(9), 2000, pp. 33-35.

[18] Mills, E. S. Uncertainty and price theory. *Quarterly Journal of Economics*, Vol. 73, 1959, pp. 116-130.

[19] Petruzzi, N. C., M. Dada. Pricing and the newsvendor model: A review with extensions. *Operations Research*, Vol. 47, 1999, pp. 183-194.

[20] Polatoglu, L. H. Optimal order quantity and pricing decisions in single period inventory systems. *International Journal of Production Economics*, Vol. 23, 1991, pp. 175-185.

[21] Polatoglu, H., I. Sahin. Optimal procurement policies under pricedependent demand. *International Journal of Production Economics*, Vol. 65, 2000, pp. 141-171.

[22] Porteus, E. On the optimality of the generalized (s, S) policies. *Management Science*, Vol. 17, 1971, pp. 411-426.

[23] Scarf, H. The optimality of (s, S) policies for the dynamic inventory problem. *Proceedings of the First Stanford Symposium on Mathematical Methods in the Social Sciences.* Stanford University Press, Stanford, CA, 1960.

[24] Simchi-Levi, D., X. Chen, J. Bramel. Integration of Inventory and Pricing. Simchi-Levi, D., X. Chen, J. Bramel, eds. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*, Chap. 9. Springer, New York, NY, 2005.

[25] Thomas, L. J. Price and production decisions with random demand. *Operations Research*, Vol. 22, 1974, pp. 513-518.

[26] Thowsen, G. T. A dynamic, nonstationary inventory problem for a price/quantity setting firm. *Naval Research Logistics Quarerly*, Vol. 22, 1975, pp. 461-476.

[27] Veinott, A. F. On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal of Applied Mathematics*, Vol. 14(5), 1966, pp. 1067-1083.

[28] Wagner, H. M., T. M. Whitin. Dynamic problems in the theory of the firm. *Naval Research Logistics Quarerly.* Vol. 7, 1958, pp. 7-12.

[29] Whitin, T. M. 1955. Inventory control and price theory. *Management Science*, Vol. 2, 1955, pp. 61-80.

[30] Zabel, E. Multiperiod monopoly under uncertainty. *Journal of Economics Theory*, Vol. 5, 1972, pp. 524-536.