Essays on Stochastic Inventory Systems

A DISSERTATION
SUBMITTED TO THE GRADUATE SCHOOL OF THE
UNIVERSITY OF MINNESOTA
BY

Rui Chen

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Advisor: Saif Benjaafar

JULY, 2015

Acknowledgements

First of all, I would like to express my sincere gratitude to my advisor, Professor Saif Benjaafar, for his excellent guidance, caring, patience, and for providing me with an excellent atmosphere for doing research. He gave so much of his time and effort to my research work and taught me so much about the tasks and meaning of scholarship. His advice on both research and my career has been priceless and helped me to grow as a scientific researcher. I have been able to gain invaluable experiences and accomplish many things that would not have been possible without his support, encouragement and patience.

I would also like to thank Professor William L. Cooper, who has taken time to comment on my work, check the proofs, and give me research guidance and career advice. I am inspired by the rigorous scholarship he has shown to me. Moreover, I would like to thank other members of my dissertation committee: Professor John Gunnar Carlsson and Professor Tony Cui, for letting my defense be an enjoyable moment, and for their brilliant comments and suggestions.

I would like to express my appreciation to Professor Mohsen Elhafsi and Professor Yimin Yu for many discussions that helped to complete this work. They provided opportunities to work on different research topics and helped me broaden my knowledge in the area of inventory management. I am grateful to their generous suggestions on research and career development. Special thanks must also go to my fellow student and coauthor Rowan Wang, who as a good friend, was always willing to help and give his best suggestions. It would have been a lonely lab in Minnesota and I would have had a lonely life in Singapore without him.

Along the way many friends in the Department of Industrial and Systems Engineering have helped me with support and inspiration. I thank each one of them. Many thanks must be given to Yibin Chen, Chenghao Du and Yan Liu, for “carrying” and “supporting” my leisure life, and for the great time we had together.

Special thanks to my parents for all of the sacrifices that they have made on my behalf. Words cannot express how grateful I am. Finally, I owe my present wonderful life to my wife and my son. They patiently and graciously endured countless hours when I occupied myself with my research instead of them. Without their constant support, encouragement, and understanding, it would not have been possible for me to complete this dissertation.

The research in this thesis was, in part funded, by National Science Foundation, the MIT-SUTD International Design Center, the Qatar National Research Foundation, and the University of Minnesota Center for Transportation Studies.

Dedication

To my parents for the great support and continuous care.

To my wife Shan Gao and my son Matthew Chen for the unyielding love, support, and encouragement.

Abstract

This thesis consists of three essays in stochastic inventory systems. The first essay is on the impact of input price variability and correlation on stochastic inventory systems. For a general class of such systems, we show that the expected cost function is concave in the input price. From this, it follows that higher input price variability in the sense of the convex order always leads to lower expected cost. We show that this is true under a wide range of assumptions for price evolution, including cases with i.i.d. prices and cases where prices are correlated and evolve according to an AR(1) process, a geometric Brownian motion, or a Markovian martingale. In addition, the result holds in cases where there is just a single period. We also examine the impact of price correlation over time and across inputs, and we find that expected cost is increasing in price correlation over time and decreasing in price correlation across components. We present results of a numerical study that provide insights on how various parameters influence the effects of price variability and correlation.

The second essay is on the optimal control of inventory systems with stochastic and independent leadtimes. We show that a fixed base-stock policy is sub-optimal and can perform poorly. For the case of exponentially distributed leadtimes, we show that the optimal policy is state-dependent and specified in terms of an inventory-dependent threshold function. Moreover, we show that this threshold function is non-increasing in the inventory level and characterized by at most m parameters. That is, once the threshold function starts to decrease it continues to

decrease with a rate that is at least one. Taking advantage of this structure, we develop an efficient algorithm for computing these parameters. In characterizing the structure of the optimal policy, we rely on an application of the Banach fixed point theorem. We compare the performance of the optimal policy to that of simpler heuristics. We also extend our analysis to systems with lost sales and systems with order cancellations.

The third essay is on the optimal policies for inventory systems with concave ordering costs. By extending the Scarf (1959) model to systems with piecewise linear concave ordering costs, we characterize the structure of optimal policies for periodic review inventory systems with concave ordering costs and general demand distributions. We show that, except for a bounded region, the generalized (s, S) policy is optimal. We do so by (a) introducing a conditional monotonicity property for the optimal order-up-to levels and (b) applying the notion of c-convexity. We also provide conditions under which the generalized (s, S) policy is optimal for all regions of the state space.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis consists of five chapters. In Chapter 2, 3 and 4, we present three completed research projects. In Chapter 5, we describes other ongoing research projects and future research directions. Chapters 2, 3 and 4 are self-contained, independent, and deal with separate topics. The following paragraphs are a brief summary of Chapters 2, 3, 4 and 5.

In Chapter 2, we explore the impact of input price variability in the context of an inventory system with stochastic demand and stochastic input prices. For a general class of such systems, we show that the expected cost function is concave in the input price. This implies that higher input price variability always leads to lower expected cost. We show that this is true under a wide range of assumptions for price evolution, including cases with i.i.d. prices and cases where prices are correlated and evolve according to an AR(1) process or a geometric Brownian motion. More significantly, we show that the result is true when prices evolve according to a Markovian martingale so that the expected price in the next period is equal to the realized price in the current period. This is perhaps surprising because one may attribute the results to a period-over-period effect whereby more (less) is ordered in one period because prices are expected to be lower (higher) in the next period. Although this temporal effect can be important, the result holds even if this temporal effect is absent and the problem is one of a single period. We also examine the impact of price correlation over time and across inputs. We find that expected cost is increasing in price correlation over time and decreasing in price correlation across components. This chapter is based on the paper "On the Impact of Input Price Variability and Correlation in Stochastic Inventory Systems", coauthored with Professor Saif Benjaafar and Professor William L. Cooper (see Chen et al. (2015a)).

In Chapter 3, we consider a continuous review inventory system with stochastic and independent leadtimes. Because orders may not be delivered in the same sequence in which they have been placed, characterizing the optimal policy is difficult and much of the available literature assumes a fixed base-stock policy. As we show, in this paper such policies are sub-optimal and can perform poorly. In this paper, we consider the case of exponentially distributed leadtimes and show that the optimal policy is not a fixed base-stock policy. Instead, the policy is state-dependent and specified in terms of an inventory-dependent threshold function. Moreover, we show that this threshold function is non-increasing in the inventory level and characterized by at most m parameters. That is, once the threshold function starts to decrease it continues to decrease with a rate that is at least one. Taking advantage of this structure, we develop an efficient algorithm for computing these parameters. In characterizing the structure of the optimal policy, we rely on an application of the Banach fixed point theorem. We compare the performance of the optimal policy to that of simpler heuristics. We also extend our analysis to systems with lost sales and systems with order cancellations. This chapter is based on the paper "Optimal Control of an Inventory System with Stochastic and Independent Leadtimes", coauthored with Professor Saif Benjaafar and Professor Mohsen Elhafsi (see Benjaafar et al. (2015a)).

In Chapter 4, we characterize the structure of optimal policies for periodic review inventory systems with concave ordering costs and general demand distributions. By extending the Scarf (1959) model to systems with piecewise linear concave ordering costs, we show that, except for a bounded region, the generalized (s, S) policy is optimal. We do so by (a) introducing a conditional monotonicity property for the optimal order-up-to levels and (b) applying the notion of c-convexity. We also provide conditions under which the generalized (s, S) policy is optimal for all regions of the state space. This chapter is based on the paper "Optimal Policies for Inventory Systems with Concave Ordering Costs", coauthored with Professor Yimin Yu and Professor Saif Benjaafar (see Yu et al. (2015)).

In Chapter 5, we provide conclusions and future research directions on the work presented in Chapters 2 and 3. We also briefly discuss other research projects, which includes (1) managing stochastic inventory systems with scarce resources, and (2) stochastic inventory systems with discount-driven backorders.

Chapter 2

On the Impact of Input Price Variability and Correlation in Stochastic Inventory Systems

2.1 Introduction

Stochastic input prices are common in practice. The prices of raw materials, precious metals, grain commodities, and electronic components, among many others, can fluctuate considerably over short periods. Such fluctuations may result from variations in supply and demand, changes in market conditions, or the introduction of new technology. Firms in some industries face input price variability because of their reliance on spot markets for procurement and, in the case of firms with global supply chains, because of exchange rate fluctuation.

The presence of stochastic input prices raises several important questions.

First, how does the presence of variability in input prices affect input ordering decisions and the nature of the optimal ordering policy? Second, how does price variability affect performance, and particularly cost? Does higher price variability increase or decrease overall costs? How does price correlation, over time or across inputs, interact with price variability and what is the net effect on cost? Is the effect of price variability more pronounced with higher correlation?

There is literature dealing with inventory systems with stochastic input prices; see Zhang (2012) for a comprehensive review. In a periodic review inventory system, Kalymon (1971) considers a single-item model with setup costs in which future input prices are determined by a Markovian stochastic process, and establishes that the optimal policy is a price-dependent (s, S) policy. Golabi (1985) considers a problem with an independent price process, negligible setup cost, and deterministic demand. He shows that the optimal policy is to always purchase a quantity that covers demands for the next several periods, and that this number of periods is decreasing in the current price. Gavirneni (2004) develops an efficient recursive procedure to calculate the base stock level when there are no setup costs and shows that myopic solutions are very effective under a non-speculative assumption. For continuous review inventory systems, Song and Zipkin (1993) characterize the optimal policy and develop algorithms for settings with Markov modulated purchasing price and Markov modulated demand. Yang and Xia (2009) consider a problem in which the input price follows a discrete-state Markov process and demand is a compound Poisson process. They show that the optimal policy is of the order-up-to type and identify conditions under which the order-up-to levels are decreasing in price. Berling and Martínez-de-Albéniz (2011) study a problem in which the price evolution is a continuous stochastic process and demand is Poisson. They characterize the optimal base-stock level using a series of threshold prices. Nie et al. (2014) consider a firm buying raw material from the spot market and selling a final product by submitting bids. They show that the optimal procurement policy is a price-dependent base-stock policy and the optimal bidding price decreases in the inventory level.

In the finance literature, Gibson and Schwartz (1990), Schwartz and Smith (2000), and Casassus and Collin-Dufresne (2005) develop multi-factor models to describe the dynamics of commodity prices. They test these models using empirical data and discuss implications for option valuation and investment decisions. There is a growing body of operations management literature concerned with traded commodities. This literature characterizes optimal operating policies for traded commodities regarding how much to buy, produce, and sell; see for example, Martínez-de-Albéniz and Simón (2009), Secomandi (2010), Devalkar et al. (2011), Goel and Gutierrez (2006, 2011), and Guo et al. (2011)

Another stream of literature studies the impact of a spot market on supply chain operations. Yi and Scheller-Wolf (2003), Boyabatlı et al. (2011), Inderfurth and Kelle (2011), Chen et al. (2013), and Secomandi and Kekre (2014) consider models in which a firm can procure a resource through long-term contracts or from the spot market. They characterize the optimal procurement policy under different assumptions. Park et al. (2012) study the inventory sharing problem for two firms where the firms can procure the commodity and sell excess inventory through either the spot or forward market. They show that inventory sharing is always beneficial. Haksöz and Seshadri (2007) provide a comprehensive review on the use of spot market operations to manage procurement in supply chains.

Much of this literature is concerned with describing the structure of the optimal ordering policy or with identifying other effective heuristics. There is only limited literature that studies the impact of input price variability in the context of inventory systems. Ho et al. (1998) analyze the impact of price format (the average price and the variance of the price) on the shopping frequency and purchasing behavior of a rational shopper using an economic order quantity (EOQ) model. They show that the optimal long run average cost is decreasing with the price variance. Berling and Rosling (2005) study how financial risks influence the optimal value of the order quantity and the reorder level in an inventory system with setup costs. They show that the systematic risk of demand has a negligible effect, but the systematic risk of the purchase price has a significant effect. Plambeck and Taylor (2013) study a problem where the firm is a price taker for both input and output products. They show that input price variability reduces the value of improving input efficiency (output produced per unit input) but increases that of capacity efficiency (the rate at which a production facility can convert input into output). Output price variability increases the value of capacity efficiency, but it increases the value of input efficiency only under certain conditions.

The papers by Janakiraman and Seshadri (2011) and Boyabatlı et al. (2011) are the most relevant to our study. Janakiraman and Seshadri (2011) examine a family of dynamic programs with stochastic cost parameters in which the vector of cost parameters evolves as a stochastic process. They show that if the single period cost is concave with respect to this vector, then the optimal cost is bounded above by the optimal cost for the dynamic program in which these stochastic cost parameters are replaced by their expectations in each period. However, the approach they employ cannot be used to compare two dynamic programs each with stochastic cost parameters. Boyabatlı et al. (2011) study optimal procurement, processing, and production policies for a meat-processing company which sources input through long-term contracts and from a spot market. They assume that the spot price follows a normal distribution and show that the optimal expected profit of the firm increases in the spot price variability under certain conditions.

In the economics literature, there is a stream of research that examines a firm's behavior when price or cost fluctuates. Sandmo (1971) and Batra and Ullah (1974) study the optimal output and input decisions for a competitive firm under price uncertainty and risk aversion. Anderson and Danthine (1981, 1983), Meyer (1987) and Kamara (1993) study how firms can use futures to hedge or speculate against price uncertainty. This literature relies on aggregate models of demand and supply and does not model operational decisions.

In this paper, we show that for a wide range of inventory problems and assumptions, higher input price variability (as measured by convex ordering of prices) leads to lower expected inventory costs over the planning horizon, where inventory costs include ordering, inventory holding, and shortage costs. One may initially attribute this phenomenon to the fact that higher variability affords more frequent opportunities to place large (small) orders in periods in which prices are anticipated to be higher (lower) in subsequent periods. Although we do observe such a period-over-period effect, the main result also holds when the input prices evolve as a martingale, where the price in a current period is equal to the conditional expected price in future periods. In addition, the result holds in systems with a single period where ordering decisions cannot be postponed to the future. We show that the benefit of input price variability can be traced to the concavity of the cost function with respect to the input price. This concavity in price is a consequence of the ability of the system manager to adjust the order quantity as prices change, leading to a cost that is lower than that which would be incurred if the order quantity were left unchanged.

We also examine the impact of correlation of prices over time. For certain types of input price sequences, we show that the expected cost decreases with increases in input price correlation. We also consider inventory systems with multiple inputs and allow for correlation among the prices of different inputs. For such systems, we use the notion of supermodular ordering to show that the expected total cost is decreasing in the correlation in input prices. Finally, we present numerical results illustrating how the benefit of input price variability is affected by various parameters. These results suggest, for instance, that the magnitude of the benefit of price variability is increasing in the length of the planning horizon and the correlation of prices of different inputs, and decreasing in the holding and backorder costs and the correlation in prices over time. The numerical results also suggest that the impact of price correlation over time and across components is more significant when the price variability is higher.

The rest of this chapter is organized as follows. In section 2.2, we describe and formulate the single-component inventory model and describe the structure of the optimal policy. In section 2.3, we analyze the impact of input price variability. In section 2.4, we study the impact of correlation of the input prices over time. In section 2.5, we consider inventory systems with multiple inputs and for such systems we study the impact of input price variability and the impact of correlation across component prices. In section 2.6, we provide numerical results and explore some of the implications of the results.

2.2 Problem Formulation

We consider a multi-period stochastic inventory control problem for a single product over a finite planning horizon consisting of $T \geq 1$ discrete time periods. Time $t = 1$ is the first period and time $t = T$ is the last period. Demand for the product occurs each period. We assume that demand forms an i.i.d. sequence of random variables with common distribution function $\Phi(\cdot)$ and density function $\phi(\cdot)$. We assume that one unit of the product is needed to fulfill one unit of demand. In each period, the ordering price, to which we also refer as the *input price*, is stochastic as well and is realized at the beginning of the period, before the realization of demand. An ordering decision (whether or not and how much to order) is made at the beginning of each period before the realization of demand but after the realization of the input price. There is no leadtime (the extension to positive leadtime is straightforward), and therefore quantities ordered in a period, if any, can be used to fulfill demands in that same period. Each unit of positive leftover inventory at the end of a period incurs a holding cost of h. Unfulfilled demand is backlogged and a backorder cost of b per unit backlogged per period is incurred. The one-period discount factor is denoted by $\beta \in (0, 1]$.

We assume that the sequence of ordering prices $\{X_t : t = 1, \ldots, T\}$ follows a Markov chain, where the ordering price X_{t+1} in period $t+1$ depends on the ordering price X_t in period t and another random variable ϵ_t. Specifically, we assume that

$$X_{t+1} = f_t(\epsilon_t) X_t + g_t(\epsilon_t), \qquad t = 1, \ldots, T-1, \tag{2.1}$$

where $\{\epsilon_t : t = 1, \ldots, T-1\}$ is a sequence of independent random variables. We denote the distribution function of ϵ_t as $\Psi_t(\cdot)$. We assume that the sequence $\{\epsilon_t\}$, the initial input price X_1, and the sequence of demands are mutually independent.

This assumption about price evolution is quite general. For instance, the case of i.i.d. ordering prices can be obtained by taking $f_t(\epsilon) = 0$, $g_t(\epsilon) = \epsilon$, and $\{\epsilon_t\}$ i.i.d. with the same distribution as X_1. Other special cases of (2.1) include a discrete-time analog of geometric Brownian motion as well as auto-regressive processes of order 1 (AR(1) processes). To obtain geometric Brownian motion with drift μ and volatility σ, we take $\{\epsilon_t\}$ to be i.i.d. normal random variables with mean μ and variance σ^2, $f_t(\epsilon) = e^{\epsilon}$, and $g_t(\epsilon) = 0$, in which case equation (2.1) becomes $X_{t+1} = X_t e^{\epsilon_t}$. To obtain an AR(1) process, we take $f_t(\epsilon) = \rho_t$, $g(\epsilon) = \epsilon + c_t$, and $E\epsilon_t = 0$, in which case (2.1) becomes $X_{t+1} = c_t + \rho_t X_t + \epsilon_t$. Moreover, through appropriate choices of $\{\epsilon_t\}$, $f_t(\cdot)$, and $g_t(\cdot)$, we can make the sequence of prices $\{X_t\}$ a martingale, supermartingale, or submartingale. We will discuss all these examples later. From here on, for notational simplicity, we only consider the case where $f_t(\cdot) = f(\cdot)$ and $g_t(\cdot) = g(\cdot)$. Our results also apply to cases where $f_t(\cdot)$ and $g_t(\cdot)$ or the holding and backorder cost parameters are time heterogeneous.

In view of the preceding assumptions, the problem can be viewed as a Markov decision process where the state of the system at the beginning of each period is a pair (s, x) that represents the net inventory level s and the ordering price x. In each period, the action, i.e., the decision to be made, is the order-up-to level $y \in [s, \infty)$. If in a particular period, net inventory is s, order-up-to level y is chosen, and demand is ξ, then the order quantity is $y - s$ and the net inventory level in the subsequent period is $y - \xi$.

The expected one-period holding and shortage costs can be expressed as a function of the action y as follows:

$$L(y) = \int_0^y h(y-\xi)\phi(\xi)d\xi + \int_y^\infty b(\xi - y)\phi(\xi)d\xi.$$

The objective is to determine in each period the optimal order-up-to level for each price such that the expected total discounted cost over the planning horizon is minimized. For $t = 1, \ldots, T$, let $v_t(s, x)$ be the optimal expected total cost from period t onward when the net inventory at the beginning of period t is s and the ordering price in period t is x. The optimality equations are given by

$$\begin{aligned} v_t(s,x) &= \min_{y \geq s} \left\{ x(y-s) + L(y) + \beta \int_\xi \int_\epsilon v_{t+1}(y-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi \right\} \\ &= \min_{y \geq s}\{w_t(y,x)\} - xs, \end{aligned} \tag{2.2}$$

where

$$\begin{aligned} w_t(y,x) &= xy + L(y) + \beta \int_\xi \int_\epsilon v_{t+1}(y-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi, \\ &= xy + L(y) + \beta \int_\xi E[v_{t+1}(y-\xi, X_{t+1})|X_t = x]\phi(\xi)d\xi \end{aligned} \tag{2.3}$$

and

$$v_{T+1}(s,x) = 0.$$

We let $y_t^*(s,x)$ denote a minimizer of (2.2). Then an optimal policy uses order-up-to level $y_t^*(s,x)$ if the state is (s,x) in period t, and the optimal order quantity is $y_t^*(s,x) - s$. The optimal expected total cost for the entire planning horizon (computed before learning the first ordering price) with starting inventory s is given by $V_1(s) = Ev_1(s, X_1)$.

In preparation for our analysis of the impact of input price variability, we next describe the form of the optimal policy for this inventory system. We begin with the following lemma.

Lemma 1. *The function $w_t(y,x)$ is convex in y for all x and $t = 1, \ldots, T$.*

The proof of Lemma 1 (and all other proofs not provided in the paper) can be found in the appendix. Let $y_t^\circ(x)$ denote a minimizer of $w_t(y,x)$ over $y \in (-\infty, \infty)$. An optimal policy is described in the following proposition, which follows immediately from Lemma 1.

Proposition 1. *There exists an optimal ordering policy for the multi-period inventory system with stochastic input prices that is a state-dependent base stock policy with base stock levels $y_t^\circ(x)$. That is, $y_t^*(s,x) = \max\{s, y_t^\circ(x)\}$ and the optimal order quantity in state (s,x) at time t is $\max\{0, y_t^\circ(x) - s\}$.*

The optimal base stock level $y_t^\circ(x)$ need not be decreasing in the realized price x. For example, consider a case where $T = 2$, $b = 0.5, h = 0.5$, $D_1 = D_2 = 10$ and $X_2 = 2X_1 - 5$, and suppose that the marginal distribution for the ordering price in period 1 is $P(X_1 = 4) = P(X_1 = 6) = 0.5$ and thus the marginal distribution of the ordering price in period 2 is $P(X_1 = 3) = P(X_1 = 7) = 0.5$. In this case, it is easy to check that it is optimal to order nothing if the realized ordering price in period 1 is 4 ($y_1^\circ(4) = 0$) and to order up to 20 if the realized ordering price in period 1 is 6 ($y_1^\circ(6) = 20$). Therefore, the optimal base stock is increasing with respect to the realized price. This is due to the strong positive correlation in the ordering price across periods. In the following proposition, we provide a sufficient condition under which this phenomenon does not occur and the base stock level is decreasing in the realized price.

Proposition 2. *If $E|f(\epsilon_t)| \leq 1$ for $t = 1, \ldots, T$, then $y_t^\circ(x)$ is decreasing in x for $t = 1, \ldots, T$.*

Examples that satisfy the condition $E|f(\epsilon_t)| \leq 1$ for $t = 1, \ldots, T$ include the case of i.i.d. input prices and the case where the input prices evolve according to an AR(1) process. In the first case $f(\epsilon) = 0$, and in the second case $f(\epsilon) = \rho \in [-1, 1]$. If the condition in the proposition is not satisfied, for example, if $Ef(\epsilon_t) > 1$, then it is possible that a high (low) price in one period would lead to a even higher (lower) expected price in the next period. In this case, it may be optimal to order more (less) when the price is high (low). Or, if $Ef(\epsilon_t) < -1$, then a low price in one period (say, now) would lead to a high expected price in the next period and an even lower expected price after two periods. In that case, one may wish to order more now in anticipation of a high price in the next period but to order less now in anticipation of an even lower price after two periods. It is possible that the second of these two effects is stronger. Therefore, it is possible that it is optimal to order more as price increases.

2.3 Impact of Price Variability

In this section, we discuss the impact of input price variability on the expected total cost and show that higher variability yields lower expected total cost. In our analysis, we use the tool of convex ordering to compare different levels of price variability. A random variable X is said to be smaller than $\widehat{X}$ in the *convex order* (written $X \leq_{cx} \widehat{X}$) if $Eu(X) \leq Eu(\widehat{X})$ for all convex functions $u(\cdot)$ such that the expectations exist. The concept of convex order is reviewed in, for example, Müller and Stoyan (2002) and Shaked and Shanthikumar (2007). If $X \leq_{cx} \widehat{X}$, then it is well known that $EX = E\widehat{X}$ and $\mathrm{Var}(X) \leq \mathrm{Var}(\widehat{X})$. For random variables drawn from various common distributions, convex ordering is equivalent to having ordered variances and identical means. For example, if we compare two normal random variables with the same mean, then the one with the smaller variance is smaller in the convex order. The same holds true for uniform, gamma and lognormal random variables as well. Below, we will frequently make use of the fact that if $u(\cdot)$ is concave and $X \leq_{cx} \widehat{X}$ then $Eu(X) \geq Eu(\widehat{X})$.

The next lemma establishes the concavity of the cost function $v_t(s, x)$ with respect to the ordering price x.

Lemma 2. *$v_t(s, x)$ is concave in x for all s and $t = 1, \ldots, T + 1$.*

To study the impact of ordering price variability, we compare two different inventory systems with ordering price sequences $\{X_t\}$ and $\{\widehat{X}_t\}$ and noise sequences $\{\epsilon_t\}$ and $\{\widehat{\epsilon}_t\}$ satisfying $X_{t+1} = f(\epsilon_t)X_t + g(\epsilon_t)$ and $\widehat{X}_{t+1} = f(\widehat{\epsilon}_t)\widehat{X}_t + g(\widehat{\epsilon}_t)$ respectively. All other parameters of the two systems are the same. We assume that each of the two systems individually satisfies the assumptions after (2.1) in Section 2.2. Let $\widehat{v}_t(s, x)$ be the optimal expected total cost-to-go in period

t when the inventory is s and the realization of price is x for the system with ordering prices $\{\widehat{X}_t\}$.

In preparation for our next result, let $X_{t+1}(x) = f(\epsilon_t)x + g(\epsilon_t)$ be a random variable that follows the conditional distribution of X_{t+1} given $X_t = x$. Likewise, let $\widehat{X}_{t+1}(x) = f(\widehat{\epsilon}_t)x + g(\widehat{\epsilon}_t)$ be a random variable that follows the conditional distribution of $\widehat{X}_{t+1}$ given $\widehat{X}_t = x$. With this notational device, $w_t(y, x)$ in (2.3) can be written as

$$w_t(y,x) = xy + L(y) + \beta \int_{\xi} E[v_{t+1}(y-\xi, X_{t+1}(x))]\phi(\xi)d\xi. \tag{2.4}$$

The following theorem describes the impact of price variability on the optimal expected total cost.

Theorem 1. *Consider $k \in \{1, \ldots, T-1\}$ and suppose $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ for all x and $t = k, \ldots, T-1$. Then $v_t(s,x) \geq \widehat{v}_t(s,x)$ for all (s,x) and $t = k, \ldots, T+1$ and $E[v_t(s, X_t)|X_{t-1} = x] \geq E[\widehat{v}_t(s, \widehat{X}_t)|\widehat{X}_{t-1} = x]$ for all x and $t = k+1, \ldots, T$.*

Proof. For a given $k = 1, \ldots, T-1$, we first prove that $v_t(s,x) \geq \widehat{v}_t(s,x)$ for all (s,x) and $t = k, \ldots, T+1$ by induction on t. We have $v_{T+1}(s,x) = 0 = \widehat{v}_{T+1}(s,x)$. Suppose $v_{t+1}(s,x) \geq \widehat{v}_{t+1}(s,x)$ for all (s,x). Then by (2.2) and (2.4), we have

$$\begin{aligned}
v_t(s,x) &= \min_{y \geq s} \left\{ x(y-s) + L(y) + \beta \int_{\xi} E[v_{t+1}(y-\xi, X_{t+1}(x))]\phi(\xi)d\xi \right\} \\
&\geq \min_{y \geq s} \left\{ x(y-s) + L(y) + \beta \int_{\xi} E[v_{t+1}(y-\xi, \widehat{X}_{t+1}(x))]\phi(\xi)d\xi \right\} \\
&\geq \min_{y \geq s} \left\{ x(y-s) + L(y) + \beta \int_{\xi} E[\widehat{v}_{t+1}(y-\xi, \widehat{X}_{t+1}(x))]\phi(\xi)d\xi \right\}
\end{aligned}$$

$$= \widehat{v}_t(s, x).$$

The first inequality above follows from the assumption that $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ and the fact that $v_{t+1}(s, x)$ is a concave function of x as shown in Lemma 2. The second inequality above follows from the inductive hypothesis. Thus, $v_t(s, x) \geq \widehat{v}_t(s, x)$ for all s, x, and $t = k, \ldots, T+1$. For $t = k+1, \ldots, T$ we have

$$\begin{aligned} E[v_t(s, X_t)|X_{t-1} = x] = Ev_t(s, X_t(x)) &\geq Ev_t(s, \widehat{X}_t(x)) \\ &\geq E\widehat{v}_t(s, \widehat{X}_t(x)) = E[\widehat{v}_t(s, \widehat{X}_t)|\widehat{X}_{t-1} = x], \end{aligned}$$

where the first inequality uses Lemma 2. □

The following corollary is an immediate consequence of Theorem 1, because for a given $k = 1, \ldots, T-1$, if $X_k \leq_{cx} \widehat{X}_k$, then $Ev_k(s, X_k) \geq E\widehat{v}_k(s, X_k) \geq E\widehat{v}_k(s, \widehat{X}_k)$, where the second inequality is due to the fact that $\widehat{v}_k(s, x)$ is concave in x.

Corollary 1. *Consider $k \in \{1, \ldots, T-1\}$ and suppose $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ for all x, and $t = k, \ldots, T-1$. If $X_k \leq_{cx} \widehat{X}_k$, then $Ev_k(s, X_k) \geq E\widehat{v}_k(s, \widehat{X}_k)$. In particular, if $X_1 \leq_{cx} \widehat{X}_1$, then $Ev_1(s, X_1) \geq E\widehat{v}_1(s, \widehat{X}_1)$.*

In view of the assumption of Markovian ordering prices, Theorem 1 indicates that given a history of the price realizations, the optimal expected total cost-to-go is decreasing with respect to the conditional variability of subsequent ordering prices. Corollary 1 shows that if no information is known about past prices, the unconditional optimal expected total cost-to-go is decreasing with respect to the unconditional variability of the current ordering price. In both cases, the more variable the price is, the lower the optimal expected total cost is. Our result implies that a risk neutral decision maker has a preference for suppliers with high price variability over suppliers with low price variability or suppliers with fixed prices. This contrasts with the effect of demand variability, where in many inventory systems, greater variability in demand leads to higher expected cost (or lower expected profit).

In the following proposition, we provide conditions under which the assumptions of Theorem 1 and Corollary 1 hold.

Proposition 3. *Suppose $X_1 \leq_{cx} \widehat{X}_1$. Then, the following statements hold.*

(a) If $\epsilon_t = \widehat{\epsilon}_t$ for $t = 1, \ldots, T$, then $X_t(x) \leq_{cx} \widehat{X}_t(x)$ for all x and $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$.

(b) If $\epsilon_t \leq_{cx} \widehat{\epsilon}_t$ and $f(\cdot)$ and $g(\cdot)$ are convex functions such that $Ef(\epsilon_t) = Ef(\widehat{\epsilon}_t)$ and $Eg(\epsilon_t) = Eg(\widehat{\epsilon}_t)$ for $t = 1, \ldots, T$, then $X_t(x) \leq_{cx} \widehat{X}_t(x)$ for $x \geq 0$ and $t = 2, \ldots, T$. Moreover, if X_t or $\widehat{X}_t$ is nonnegative a.s. for $t = 1, \ldots, T$, then $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$.

(c) If $\epsilon_t \leq_{cx} \widehat{\epsilon}_t$ and $g(\cdot)$ is a convex function such that $Eg(\epsilon_t) = Eg(\widehat{\epsilon}_t)$ for $t = 1, \ldots, T$, and $f(\cdot)$ is an affine function, then $X_t(x) \leq_{cx} \widehat{X}_t(x)$ for all x and $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$.

(d) If $f(\epsilon_t) \leq_{cx} f(\widehat{\epsilon}_t)$ for $t = 1, \ldots, T$ and $g(\cdot)$ is a constant, then $X_t(x) \leq_{cx} \widehat{X}_t(x)$ for all x and $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$.

Proof. (a) Suppose $u(\cdot)$ is an arbitrary convex function. Then

$$Eu(X_{t+1}(x)) = Eu(f(\epsilon_t)x + g(\epsilon_t)) = Eu(f(\widehat{\epsilon}_t)x + g(\widehat{\epsilon}_t)) = Eu(\widehat{X}_{t+1}(x)).$$

Therefore, $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ for all x and $t = 2, \ldots, T$.

To show that $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$, we only need to show that if $X_t \leq_{cx} \widehat{X}_t$, then $X_{t+1} \leq_{cx} \widehat{X}_{t+1}$. Suppose $X_t \leq_{cx} \widehat{X}_t$. Let $u(\cdot)$ be an arbitrary convex function and let $\kappa(x) = Eu(X_{t+1}(x))$ and $\widehat{\kappa}(x) = Eu(\widehat{X}_{t+1}(x))$. Then $\kappa(x) \leq \widehat{\kappa}(x)$ for all x because $u(\cdot)$ is convex and we have already shown that $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ for all x. Moreover, $\kappa(x) = Eu(f(\epsilon_t)x + g(\epsilon_t))$ and hence κ is convex in x. Therefore, we have

$$Eu(X_{t+1}) = E\kappa(X_t) \leq E\kappa(\widehat{X}_t) \leq E\widehat{\kappa}(\widehat{X}_t) = Eu(\widehat{X}_{t+1}).$$

Thus, we have $X_{t+1} \leq_{cx} \widehat{X}_{t+1}$.

(b) We will use the fact that $X \leq_{cx} Y$ is equivalent to $EX = EY$ and $Eu(X) \leq Eu(Y)$ for all increasing convex functions $u(\cdot)$. See, for example, Theorem 1.5.3 of Müller and Stoyan (2002). We have

$$EX_{t+1}(x) = E[f(\epsilon_t)x + g(\epsilon_t)] = E[f(\widehat{\epsilon}_t)x + g(\widehat{\epsilon}_t)] = E\widehat{X}_{t+1}(x).$$

Suppose now that $u(\cdot)$ is an arbitrary increasing convex function. The function $\eta(\epsilon) = f(\epsilon)x + g(\epsilon)$ is convex in ϵ for $x \geq 0$. Therefore, $\widetilde{u}(\epsilon) = u(\eta(\epsilon)) = u(f(\epsilon)x + g(\epsilon))$ is a convex function of ϵ for $x \geq 0$. Hence, for $x \geq 0$ we have

$$Eu(X_{t+1}(x)) = E\widetilde{u}(\epsilon_t) \leq E\widetilde{u}(\widehat{\epsilon}_t) = Eu(\widehat{X}_{t+1}(x)).$$

Thus, $EX_{t+1}(x) = E\widehat{X}_{t+1}(x)$ and $Eu(X_{t+1}(x)) \leq Eu(\widehat{X}_{t+1}(x))$ for any increasing and convex function $u(\cdot)$ when $x \geq 0$. This implies that $X_{t+1}(x) \leq_{cx} \widehat{X}_{t+1}(x)$ for $x \geq 0$.

Next we show that if X_t or $\widehat{X}_t$ is nonnegative a.s. for $t = 1, \ldots, T$, then $X_t \leq_{cx} \widehat{X}_t$ for $t = 2, \ldots, T$. Suppose $X_t \leq_{cx} \widehat{X}_t$. Let $u(\cdot)$ be an arbitrary convex function and let $\kappa(x) = Eu(X_{t+1}(x))$ and $\widehat{\kappa}(x) = Eu(\widehat{X}_{t+1}(x))$. Then

$\kappa(x) \le \widehat{\kappa}(x)$ for $x \ge 0$ because $u(\cdot)$ is convex and because we have already shown that $X_{t+1}(x) \le_{cx} \widehat{X}_{t+1}(x)$ for $x \ge 0$. Moreover, $\kappa(x) = Eu(f(\epsilon_t)x + g(\epsilon_t))$ and $\widehat{\kappa}(x) = Eu(f(\widehat{\epsilon}_t)x + g(\widehat{\epsilon}_t))$ are convex in x. Therefore, if X_t is nonnegative a.s., we have

$$Eu(X_{t+1}) = E\kappa(X_t) \le E\widehat{\kappa}(X_t) \le E\widehat{\kappa}(\widehat{X}_t) = Eu(\widehat{X}_{t+1}).$$

If $\widehat{X}_t$ is nonnegative a.s., we have

$$Eu(X_{t+1}) = E\kappa(X_t) \le E\kappa(\widehat{X}_t) \le E\widehat{\kappa}(\widehat{X}_t) = Eu(\widehat{X}_{t+1}).$$

Thus, we have $X_{t+1} \le_{cx} \widehat{X}_{t+1}$.

The proofs of (c) and (d) are similar to the proof of (b) and are omitted. □

Property (a) of the above lemma implies that for systems with the same sequence $\{\epsilon_t\}$, higher variability of the price in the first period will lead to higher variability of prices in all subsequent periods. In property (b) and property (c), $Ef(\epsilon_t) = Ef(\widehat{\epsilon}_t)$ and $Eg(\epsilon_t) = Eg(\widehat{\epsilon}_t)$ together imply that $E[X_{t+1}|X_t = x] = E[\widehat{X}_{t+1}|\widehat{X}_t = x]$, which is a necessary condition for $X_{t+1}(x) \le_{cx} \widehat{X}_{t+1}(x)$. Properties (b), (c), and (d) state that under some conditions, if the random influence ϵ_t or $f(\epsilon_t)$ on the prices becomes more variable in some period t, then the prices will also become more variable in all subsequent periods. We next provide a few specific examples for which the preceding results tell us that the greater input price variability leads to lower expected costs.

- $\{X_t\}$ and $\{\widehat{X}_t\}$ are both i.i.d. price sequences: Suppose that $X_1 \le_{cx} \widehat{X}_1$. To place this setting in our framework, we may take $f(\epsilon) = 0$, $g(\epsilon) = \epsilon$, and $\{\epsilon_t\}$ i.i.d. [respectively, $\{\widehat{\epsilon}_t\}$ i.i.d.] with the same distribution as X_1 [resp., $\widehat{X}_1$]. In this case, the conditions in (b) and (c) of Proposition 3 hold and hence we have that the system with more-variable input prices $\{\widehat{X}_t\}$ has lower expected costs as indicated by Theorem 1 and Corollary 1.

- $\{X_t\}$ and $\{\widehat{X}_t\}$ are both stationary AR(1) price sequences: Let $\mu = c/(1-\rho)$ for constants c and $\rho \in (-1, 1)$. Suppose that $X_{t+1} = \rho X_t + \epsilon_t + c$ where $\{\epsilon_t\}$ are i.i.d. $N(0, \sigma^2)$ and $X_1 \sim N(\mu, \sigma^2/(1-\rho^2))$ and $\widehat{X}_{t+1} = \rho\widehat{X}_t + \widehat{\epsilon}_t + c$ where $\{\widehat{\epsilon}_t\}$ are i.i.d. $N(0, \widehat{\sigma}^2)$ and $\widehat{X}_1 \sim N(\mu, \widehat{\sigma}^2/(1-\rho^2))$. Suppose that $\sigma \leq \widehat{\sigma}$. For normal random variables, $X \leq_{cx} Y$ is equivalent to $EX = EY$ and $\mathrm{Var}(X) \leq \mathrm{Var}(Y)$. Therefore, $\epsilon_t \leq_{cx} \widehat{\epsilon}_t$ and $X_t \leq_{cx} \widehat{X}_t$ for $t = 1, \ldots, T$. To place this setting in our framework, we may take $f(\epsilon) = \rho$ and $g(\epsilon) = \epsilon + c$. The conditions in (c) of Proposition 3 hold and hence the system with more-variable input prices again has lower expected costs.

- $\{X_t\}$ and $\{\widehat{X}_t\}$ are both (discrete-time) geometric Brownian motions: Suppose that $X_1 = \widehat{X}_1 = x$, $X_{t+1} = X_t e^{\epsilon_t}$, and $\widehat{X}_{t+1} = \widehat{X}_t e^{\widehat{\epsilon}_t}$. Suppose that $\{\epsilon_t\}$ are i.i.d. $N(\mu, \sigma^2)$ and $\{\widehat{\epsilon}_t\}$ are i.i.d. $N(\widehat{\mu}, \widehat{\sigma}^2)$ where $2\mu + \sigma^2 = 2\widehat{\mu} + \widehat{\sigma}^2$ and $\sigma \leq \widehat{\sigma}$. Take $f(\epsilon) = e^\epsilon$ and $g(\epsilon) = 0$ to place this within our framework. Hence, $\{f(\epsilon_t)\}$ and $\{f(\widehat{\epsilon}_t)\}$ are i.i.d. lognormal random variables for which $Ef(\epsilon_t) = Ef(\widehat{\epsilon}_t)$ and $\mathrm{Var}(f(\epsilon_t)) \leq \mathrm{Var}(f(\widehat{\epsilon}_t))$. Moreover, $f(\epsilon_t) \leq_{cx} f(\widehat{\epsilon}_t)$; see page 63 of Müller and Stoyan (2002). Application of part (d) of Proposition 3 allows us to conclude that the system with more-variable input prices $\{\widehat{X}_t\}$ has lower expected costs.

- $\{X_t\}$ and $\{\widehat{X}_t\}$ are both Markovian martingales: Suppose that $X_1 = \widehat{X}_1 = x$, $X_{t+1} = X_t + \epsilon_t$, and $\widehat{X}_{t+1} = \widehat{X}_t + \widehat{\epsilon}_t$ where $\{\epsilon_t\}$ and $\{\widehat{\epsilon}_t\}$ are sequences of independent random variables with $E\epsilon_t = E\widehat{\epsilon}_t = 0$ and $\epsilon_t \leq_{cx} \widehat{\epsilon}_t$. This example fits our framework with $f(\epsilon) = 1$ and $g(\epsilon) = \epsilon$, and part (c) of Proposition 3 allows us to apply Theorem 1 and Corollary 1 to conclude that the system with more-variable input prices has lower expected costs.

One may at first be tempted to attribute the lower costs associated with higher input price variability solely to the more frequent opportunities afforded by higher variability to place large (small) orders in periods in which prices are anticipated to be higher (lower) in subsequent periods. As we note in Section 2.6, this period-over-period effect is indeed important (there, for example, we observe that the relative reduction in cost due to higher variability is increasing in the length of the planning horizon and is decreasing in the correlation in prices over time). However, higher price variability yields lower expected total cost even when the input prices form a martingale (wherein the price in a current period is equal to the conditional expected price in future periods) and also when the problem has only one period.

The effect of variability can be traced to the concavity of the expected cost as a function of the input price. This concavity arises from the ability to adjust order quantities based on price realization. The order quantity in each period is determined by trading off input price, inventory holding cost, backorder cost, and expectations about future prices. The firm can benefit from lower prices by ordering more and, therefore, reducing backorder costs. Higher prices are of course harmful, but the effect is mitigated by the ability of the firm to order less and instead incur higher backorder costs. If input prices are sufficiently high, the firm stops ordering and instead incurs the backorder cost. Beyond a certain threshold, expected total cost becomes invariant to price. The above effects are easiest to see in the context of a single period problem, which we explore next.

The Single Period Case. Consider a single period version of the problem where there is only a single opportunity to order after price is revealed but before demand is realized. If demand falls below the order quantity, an overage cost per unit is incurred while if demand exceeds the order quantity, a shortage cost is incurred. To be consistent with the multi-period problem, let h denote the unit overage cost and b the unit shortage cost. Given the realized price, this is of course an instance of the classic newsvendor problem.

Let $X = X_1$ denote the random input price. Demand is denoted by D with distribution function $\Phi(\cdot)$ and density function $\phi(\cdot)$. Given price realization x, the expected total cost is

$$v(x) = \min_{y \geq 0}[xy + L(y)] = \min_{y \geq 0} w(x, y),$$

where $w(x, y) = xy + L(y)$ is the cost when price is x and the ordering quantity is y. The optimal order quantity is

$$y^*(x) = \begin{cases} \Phi^{-1}(\frac{b-x}{b+h}) & \text{if } x \leq b, \\ 0 & \text{if } x > b. \end{cases}$$

Substituting into the expression for the expected total cost leads to

$$v(x) = \begin{cases} b\int_{\Phi^{-1}(\frac{b-x}{b+h})}^{\infty} \xi\phi(\xi)d\xi - h\int_0^{\Phi^{-1}(\frac{b-x}{b+h})} \xi\phi(\xi)d\xi & \text{if } x \leq b, \\ bE[D] & \text{if } x > b, \end{cases}$$

from which we can easily show that $v(x)$ is a concave function in x. In turn, this leads to the result that, for the single period case, higher price variability leads

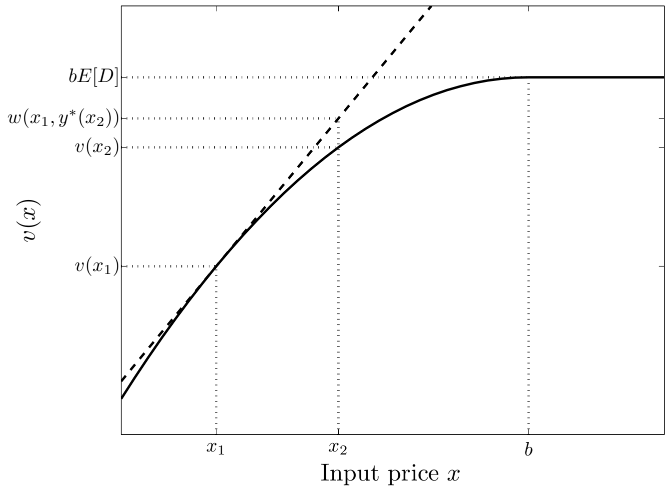

Figure 2.1: Cost as a function of the ordering price for stochastic demand

to lower expected total cost (if $X \leq_{cx} \widehat{X}$, then $Ev(X) \geq Ev(\widehat{X})$). This is true regardless of the distribution of demand.

The concavity of the expected cost can be explained as follows. At a given price $x_1 \leq b$, the optimal order quantity is $y^*(x_1) = \Phi^{-1}(\frac{b-x_1}{b+h})$ and the associated expected cost is $x_1 y^*(x_1) + L(y^*(x_1))$. If the input price increases (decreases) from x_1 to x_2 and the order quantity is not adjusted, the expected cost would increase (decrease) linearly with rate $y^*(x_1)$ to $w(x_2, y^*(x_1)) = x_2 y^*(x_1) + L(y^*(x_1))$. However, if the order quantity is adjusted and chosen optimally, then the order quantity $y^*(x_1)$ would be lower (higher) and the optimal cost $x_2 y^*(x_2) + L(y^*(x_2))$ would be lower than that if the order quantity is not adjusted. As a consequence, the optimal expected total cost is concave in the input price. This is illustrated in Figure 2.1.

Note that a special case is when demand is deterministic and assumes a single value $D = d$. The optimal cost function in that case is linear with slope d for

$x \leq b$ and equal to bd for $x > b$, i.e.,

$$v(x) = \begin{cases} xd & \text{if } x \leq b, \\ bd & \text{if } x > b. \end{cases}$$

If the input price is either $\mu + \alpha$ or $\mu - \alpha$ with equal probability (in which case α is the standard deviation of the input price) and $\mu \leq b$, then the optimal cost is

$$V(\alpha) = \begin{cases} \frac{1}{2}(\mu + \alpha)d + \frac{1}{2}(\mu - \alpha)d = \mu d & \text{if } \alpha \leq b - \mu, \\ \frac{1}{2}bd + \frac{1}{2}(\mu - \alpha)d = \frac{1}{2}(b + \mu - \alpha)d & \text{if } \alpha > b - \mu. \end{cases}$$

Clearly, $V(\alpha)$ is decreasing in α.

We conclude this section by noting that the benefit of input price variability is also present in other inventory systems, including systems with an infinite planning horizon, systems with lost sales instead of backorders, systems with a fixed ordering cost, and systems with fixed leadtimes. For the sake of brevity, we omit the details.

2.4 Impact of Price Correlation over Time

In this section, we study the impact of price correlation over time on the optimal expected total cost. To do so, we compare two different inventory systems that are identical except that they have different stationary AR(1) ordering price sequences $\{X_t\}$ and $\{\widehat{X}_t\}$ such that $X_1, \widehat{X}_1 \sim N(\mu, \sigma^2)$,

$$X_{t+1} = (1 - \rho)\mu + \rho X_t + \sqrt{1 - \rho^2}\epsilon_t, \tag{2.5}$$

$$\widehat{X}_{t+1} = (1 - \widehat{\rho})\mu + \widehat{\rho}\widehat{X}_t + \sqrt{1 - \widehat{\rho}^2}\epsilon_t \tag{2.6}$$

for $t = 1, \dots, T-1$, and $\{\epsilon_t\}$ are i.i.d. normal random variables with mean 0 and variance σ^2 that are independent of X_1 and $\widehat{X}_1$. It is easy to check that $X_t, \widehat{X}_t \sim N(\mu, \sigma^2)$ and $\mathrm{Corr}(X_t, X_{t+j}) = \rho^j$, $\mathrm{Corr}(\widehat{X}_t, \widehat{X}_{t+j}) = \widehat{\rho}^j$ for $j \geq 0$. It will be helpful to view the two price sequences as random vectors, which we denote by $\mathbf{X} = (X_1, \dots, X_T)$ and $\widehat{\mathbf{X}} = (\widehat{X}_1, \dots, \widehat{X}_T)$. As before, we suppose that the assumptions in Section 2.2 hold for each of the two systems viewed in isolation.

Let $V_1(s) = Ev_1(s, X_1)$ and $\widehat{V}_1(s) = E\widehat{v}_1(s, \widehat{X}_1)$ be the optimal expected total costs for the two systems, $y_t^*(s, x)$ and $\widehat{y}_t^*(s, x)$ be the optimal order-up-to levels for the two systems in period t when the inventory is s and the ordering price is x, and $y_t^\circ(x)$ and $\widehat{y}_t^\circ(x)$ be the base stock levels for the two systems in period t when the ordering price is x. Recall from Proposition 1 that $y_t^*(s, x) = \max\{s, y_t^\circ(x)\}$ and $\widehat{y}_t^*(s, x) = \max\{s, \widehat{y}_t^\circ(x)\}$. Below, we compare $V_1(s)$ and $\widehat{V}_1(s)$

In the following developments, we will use the tool of supermodular ordering of random vectors. The supermodular order is reviewed in, e.g., Müller and Stoyan (2002) and Shaked and Shanthikumar (2007). A function $u(\cdot)$ on $\mathbb{R}^T$ is said to be supermodular if $u(\mathbf{x}+\varepsilon\mathbf{e}^i+\delta\mathbf{e}^j)-u(\mathbf{x}+\varepsilon\mathbf{e}^i)-u(\mathbf{x}+\delta\mathbf{e}^j)+u(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbb{R}^T$, all $i, j = 1, \dots, T$ with $i < j$ and all $\varepsilon, \delta > 0$. A function $u(\cdot)$ is submodular if $-u(\cdot)$ is supermodular. If $u(\cdot)$ is twice differentiable then $u(\cdot)$ is supermodular if and only if $\frac{\partial^2 u}{\partial x^i \partial x^j}(\mathbf{x}) \geq 0$ for all $\mathbf{x}$ and all i, j with $i < j$. A random vector $\mathbf{X} = (X_1, \dots, X_T)$ is said to be smaller than a random vector $\widehat{\mathbf{X}} = (\widehat{X}_1, \dots, \widehat{X}_T)$ in the supermodular order, written $\mathbf{X} \leq_{sm} \widehat{\mathbf{X}}$, if $Eu(\mathbf{X}) \leq Eu(\widehat{\mathbf{X}})$ for all supermodular functions $u(\cdot)$ such that the expectations exist. The condition that $\mathbf{X} \leq_{sm} \widehat{\mathbf{X}}$ can be interpreted to mean that the entries of $\widehat{\mathbf{X}}$ have greater positive dependence than do the entries of $\mathbf{X}$; see, Müller and Stoyan (2002) or Shaked and Shanthikumar (2007). If $\mathbf{X}$

and $\widehat{\mathbf{X}}$ are normally distributed random vectors, then $\mathbf{X} \leq_{sm} \widehat{\mathbf{X}}$ is equivalent to $\mathbf{X}$ and $\widehat{\mathbf{X}}$ having the same marginal distributions and $\mathrm{Corr}(X_i, X_j) \leq \mathrm{Corr}(\widehat{X}_i, \widehat{X}_j)$ for all $i \neq j$ (see Theorem 3.13.5 of Müller and Stoyan 2002). Therefore, $\mathbf{X}$ and $\widehat{\mathbf{X}}$ in (2.5)–(2.6) satisfy $\mathbf{X} \leq_{sm} \widehat{\mathbf{X}}$ if $0 \leq \rho \leq \widehat{\rho}$.

In preparation for the proof of the main result of this section, for $k = 1, \ldots, T$, consider $\mathbf{X}_k = (X_{1,k}, \ldots, X_{T,k})$, where $X_{1,k} \sim N(\mu, \sigma^2)$,

$$X_{i+1,k} = \begin{cases} (1-\widehat{\rho})\mu + \widehat{\rho} X_{i,k} + \sqrt{1-\widehat{\rho}^2}\epsilon_i & \text{for } i = 1, \ldots, k-1, \\ (1-\rho)\mu + \rho X_{i,k} + \sqrt{1-\rho^2}\epsilon_i & \text{for } i = k, \ldots, n-1, \end{cases}$$

and $\{\epsilon_t\}$ are i.i.d. normal random variables with mean 0 and variance σ^2 that are independent of $X_{1,k}$. Note that $\mathbf{X} = \mathbf{X}_1$ and $\widehat{\mathbf{X}} = \mathbf{X}_T$. Note also that $\mathbf{X}_k$ is a non-stationary AR(1) process with $X_{i,k} \sim N(\mu, \sigma^2)$ and

$$\mathrm{Corr}(X_{i,k}, X_{j,k}) = \begin{cases} \widehat{\rho}^{j-i} & \text{for } i < j \leq k, \\ \rho^{j-i} & \text{for } k \leq i < j, \\ \widehat{\rho}^{k-i}\rho^{j-k} & \text{for } i < k < j. \end{cases}$$

It is easy to check that if $0 \leq \rho \leq \widehat{\rho}$, then $\mathrm{Corr}(X_{i,k}, X_{j,k}) \leq \mathrm{Corr}(X_{i,k+1}, X_{j,k+1})$ for all $1 \leq i < j \leq T$, from which we immediately obtain the following lemma.

Lemma 3. *Suppose that $0 \leq \rho \leq \widehat{\rho}$. Then $\mathbf{X}_k \leq_{sm} \mathbf{X}_{k+1}$ and $(X_{t,k}, X_{t+1,k}) \leq_{sm} (X_{t,k+1}, X_{t+1,k+1})$ for all $t = 1, \ldots, T-1$ and $k = 1, \ldots, T-1$.*

Let $v_{t,k}(s, x)$ be the optimal expected total cost from time t onward when the inventory is s and the ordering price is x for the system with input price sequence $\mathbf{X}_k$. Then $v_{t,k}(s,x) = \min_{y \geq s} w_{t,k}(y, x) - xs$ where

$$w_{t,k}(y,x) = xy + L(y) + \beta \int_{\xi} E[v_{t+1,k}(y-\xi, X_{t+1,k}) | X_{t,k} = x] \phi(\xi) d\xi.$$

These expressions are simply equations (2.2) and (2.3) for a system with input prices $\mathbf{X}_k$. Below we also use the notation $y^*_{t,k}(s,x)$ and $y^\circ_{t,k}(x)$ for optimal order up-to-levels and base-stock levels in this system. (Note that to place the prices $\mathbf{X}_k$ into the form (2.1), we must allow the functions $f_t(\cdot)$ and $g_t(\cdot)$ to depend upon t. As we noted in Section 2.2, our results still hold for such non-homogeneous cases.) We have $v_t(s,x) = v_{t,1}(s,x)$, $\widehat{v}_t(s,x) = v_{t,T}(s,x)$, $V_1(s) = Ev_{1,1}(s, X_{1,1})$, and $\widehat{V}_1(s) = Ev_{1,T}(s, X_{1,T})$.

Lemma 4. *Suppose that $\rho, \widehat{\rho} \geq 0$. Then $v_{t,k}(s,x)$ is submodular in (s,x) for $t = 1, \ldots, T+1$ and $k = 1, \ldots, T$.*

The main result of this section is the following theorem, which describes the impact of price correlation over time on the optimal expected cost and indicates that the optimal expected cost is increasing in that correlation if the correlation is positive.

Theorem 2. *If $0 \leq \rho \leq \widehat{\rho}$, then $V_1(s) \leq \widehat{V}_1(s)$ for all s.*

Proof. We will show that for all decreasing functions $u(\cdot)$, we have

$$Ev_{t,k}(u(X_{t,k}), X_{t,k}) \leq Ev_{t,k+1}(u(X_{t,k+1}), X_{t,k+1}), \tag{2.7}$$

for $k = 1, \ldots, T-1$ and $t = 1, \ldots, T$. From this the theorem follows, because for a given inventory level s, we may take $u(x) = s$ to obtain $V_1(s) = Ev_{1,1}(s, X_{1,1}) \leq Ev_{1,T}(s, X_{1,T}) = \widehat{V}_1(s)$. We establish (2.7) by considering the cases $t \geq k+1$, $t = k$, and $t \leq k-1$ separately.

Fix k. Given a decreasing function $u(\cdot)$, define

$$\theta_{t,k}(x) = x[y^*_{t,k}(u(x), x) - u(x)] + L(y^*_{t,k}(u(x), x)).$$

To avoid a proliferation of subscripts in the remainder of the proof, let $\overline{X}_t = X_{t,k}$ and $\underline{X}_t = X_{t,k+1}$ for $t = 1, \ldots, T$. This allows us to express $v_{t,k}(u(x), x)$ as

$$v_{t,k}(u(x), x) = \theta_{t,k}(x) + \beta \int_{\xi} E[v_{t+1,k}(y^*_{t,k}(u(x), x) - \xi, \overline{X}_{t+1}) | \overline{X}_t = x] \phi(\xi) d\xi.$$

In a given period t, the expected cost from time t onward depends upon that period's realized price x and starting inventory level s, as well as the conditional distribution of future prices in periods $t+1, \ldots, T$ given the price x in period t. Likewise, the optimal base stock level in a given period t depends only on the realized price x in that period and the conditional distribution of future prices. It does not depend on the prices or distributions of the prices in the past. Therefore, we have $y^\circ_{t,k}(x) = y^\circ_{t,k+1}(x)$ and $v_{t,k}(s, x) = v_{t,k+1}(s, x)$ for $t \geq k+1$. Hence, (2.7) holds for $t \geq k+1$ and for any decreasing function $u(\cdot)$ because $\overline{X}_t$ and $\underline{X}_t$ have the same distribution (both are $N(\mu, \sigma^2)$) for $t \geq k+1$.

When $t = k$, we have

$$\begin{aligned} Ev_{t,k}(u(\overline{X}_t), \overline{X}_t) &\leq E[w_{t,k}(y^*_{t,k+1}(u(\overline{X}_t), \overline{X}_t), \overline{X}_t) - \overline{X}_t u(\overline{X}_t)] \\ &= E[\theta_{t,k+1}(\overline{X}_t)] + \beta \int_{\xi} E[v_{t+1,k}(y^*_{t,k+1}(u(\overline{X}_t), \overline{X}_t) - \xi, \overline{X}_{t+1})] \phi(\xi) d\xi. \end{aligned} \tag{2.8}$$

In the preceding, we can replace $\overline{X}_t$ by $\underline{X}_t$ in the argument of $E(\theta_{t,k+1}(\cdot))$ because both have the same distribution. For the second term in (2.8), note that $v_{t+1,k}(s, x) = v_{t+1,k+1}(s, x)$ because $t+1 \geq k+1$. Moreover, $v_{t+1,k}(s, x)$ is submodular in (s, x) by Lemma 4, and $y^*_{t,k+1}(u(x), x) = \max\{u(x), y^\circ_{t,k+1}(x)\}$ is decreasing in x by Proposition 2. Consequently, $v_{t+1,k}(y^*_{t,k+1}(u(x_t), x_t) - \xi, x_{t+1})$ is supermodular in (x_t, x_{t+1}). By Lemma 3 we also have $(\overline{X}_t, \overline{X}_{t+1}) \leq_{sm} (\underline{X}_t, \underline{X}_{t+1})$.

Therefore,

$$\begin{aligned}E[v_{t+1,k}(y^*_{t,k+1}(u(\overline{X}_t),\overline{X}_t)-\xi,\overline{X}_{t+1})] &\leq E[v_{t+1,k}(y^*_{t,k+1}(u(\underline{X}_t),\underline{X}_t)-\xi,\underline{X}_{t+1})]\\ &= E[v_{t+1,k+1}(y^*_{t,k+1}(u(\underline{X}_t),\underline{X}_t)-\xi,\underline{X}_{t+1})].\end{aligned}$$

As a consequence, by (2.8) we have

$$\begin{aligned}&Ev_{t,k}(u(\overline{X}_t),\overline{X}_t)\\ &\qquad\leq E[\theta_{t,k+1}(\underline{X}_t)]+\beta\int_\xi E[v_{t+1,k+1}(y^*_{t,k+1}(u(\underline{X}_t),\underline{X}_t)-\xi,\underline{X}_{t+1})]\phi(\xi)d\xi\\ &\qquad= Ev_{t,k+1}(u(\underline{X}_t),\underline{X}_t),\end{aligned}$$

and so (2.7) holds for $t=k$.

Consider some $t \leq k$. Suppose inductively that $Ev_{t,k}(u(\overline{X}_t),\overline{X}_t) \leq Ev_{t,k+1}(u(\underline{X}_t),\underline{X}_t)$ for all decreasing functions $u(\cdot)$. Consider an arbitrary decreasing function $u(\cdot)$. We have

$$\begin{aligned}Ev_{t-1,k}(u(\overline{X}_{t-1}),\overline{X}_{t-1}) &\leq E[w_{t-1,k}(y^*_{t-1,k+1}(u(\overline{X}_{t-1}),\overline{X}_{t-1}),\overline{X}_{t-1})-\overline{X}_{t-1}u(\overline{X}_{t-1})]\\ &= E[\theta_{t-1,k+1}(\overline{X}_{t-1})]+\beta\int_\xi Ev_{t,k}(y^*_{t-1,k+1}(u(\overline{X}_{t-1}),\overline{X}_{t-1})-\xi,\overline{X}_t)\phi(\xi)d\xi.\end{aligned}$$

Observe that $(\overline{X}_{t-1},\overline{X}_t)$ are normal random variables each with mean μ and variance σ^2 and with correlation ρ. Recall that $\overline{X}_t = (1-\rho)\mu+\rho\overline{X}_{t-1}+\sqrt{1-\rho^2}\epsilon_{t-1}$. Then $\overline{X}_{t-1}$ can be written as $\overline{X}_{t-1}=\pi(\overline{X}_t,\tilde{\epsilon}_{t,k})$, where $\pi(x,\epsilon)=(1-\rho)\mu+\rho x+\sqrt{1-\rho^2}\epsilon$ and $\tilde{\epsilon}_{t,k}$ is normally distributed with mean 0 and variance σ^2 and is independent of $\overline{X}_t$. Similarly, we have $\underline{X}_{t-1}=\pi(\underline{X}_t,\tilde{\epsilon}_{t,k+1})$, where $\tilde{\epsilon}_{t,k+1}$ is normally distributed with mean 0 and variance σ^2 and is independent of $\underline{X}_t$. Note that $\tilde{\epsilon}_{t,k}$ and $\tilde{\epsilon}_{t,k+1}$ have the same distribution (which is the distribution of

ϵ_t). Let $\eta(x,\epsilon) = y^*_{t-1,k+1}(u(\pi(x,\epsilon)),\pi(x,\epsilon))$. Since $\rho \geq 0$, we have that $\pi(x,\epsilon)$ is an increasing function of x. Therefore, $\eta(x,\epsilon)$ is a decreasing function of x by Proposition 2. By the inductive assumption, we have

$$Ev_{t,k}(\eta(\overline{X}_t,\epsilon)) - \xi, \overline{X}_t) \leq Ev_{t,k+1}(\eta(\underline{X}_t,\epsilon)) - \xi, \underline{X}_t)$$

for any realization ϵ. As a consequence,

$$\begin{aligned}
&Ev_{t-1,k}(u(\overline{X}_{t-1}),\overline{X}_{t-1})\\
&\quad\leq E[\theta_{t-1,k+1}(\overline{X}_{t-1})]) + \beta \int_\xi \int_\epsilon Ev_{t,k}(\eta(\overline{X}_t,\epsilon) - \xi, \overline{X}_t)\Psi(d\epsilon)\phi(\xi)d\xi\\
&\quad\leq E[\theta_{t-1,k+1}(\underline{X}_{t-1})]) + \beta \int_\xi \int_\epsilon Ev_{t,k+1}(\eta(\underline{X}_t,\epsilon) - \xi, \underline{X}_t)\Psi(d\epsilon)\phi(\xi)d\xi\\
&\quad= E[\theta_{t-1,k+1}(\underline{X}_{t-1})]) + \beta \int_\xi E[v_{t,k+1}(y^*_{t-1,k+1}(u(\underline{X}_{t-1}),\underline{X}_{t-1}) - \xi, \underline{X}_t)]\phi(\xi)d\xi\\
&\quad= Ev_{t-1,k+1}(u(\underline{X}_{t-1}),\underline{X}_{t-1}).
\end{aligned}$$

Therefore (2.7) holds for all $t \leq k$ by induction on t. $\square$

The preceding theorem shows that if the input prices follow a stationary AR(1) process, then greater positive price correlation over time yields larger expected total cost. This result should be intuitive. With high positive correlation in prices, an unusually high price is often followed by another unusually high price, and therefore delaying purchase will likely not avoid high costs. Therefore, high correlation in prices over time leads to high expected total cost. On the other hand, with low positive correlation, if the price is unusually high in one period then the probability that the price in the next period will continue to be high is comparatively small, and hence purchases can be delayed in expectation of a price decrease.

2.5 Inventory Systems with Multiple Inputs

In this section, we extend our analysis to systems with multiple input components, where one unit of each of n input components is needed to satisfy one unit of demand. The ordering prices of these n components are stochastic (deterministic prices can be treated as a special case). The holding cost of component $i = 1, \ldots, n$ is h^i.

As in the single component model, we assume that the price X_{t+1}^i of component i in period $t+1$ is dependent on the price X_t^i of component i in period t:

$$X_{t+1}^i = f^i(\epsilon_t^i)X_t^i + g^i(\epsilon_t^i), \qquad\qquad t = 1, \ldots, T-1,$$

where $\{\boldsymbol{\epsilon}_t = (\epsilon_t^1, \ldots, \epsilon_t^n)' : t = 1, \ldots, T-1\}$ is a sequence of independent random vectors. Let $\mathbf{X}_t = (X_t^1, \ldots, X_t^n)'$ for $t = 1, \ldots, T$. We assume that $\{\boldsymbol{\epsilon}_t\}$, $\mathbf{X}_1$, and the sequence of demands are independent. The prices of different components in the same period may be correlated. Other assumptions are the same as those of the single component model.

The problem can be viewed as a Markov decision process where the state of the system at the beginning of each period is $(\mathbf{s}, \mathbf{x})$ where $\mathbf{s} = (s^1, \ldots, s^n)'$ is the vector of net inventory levels and $\mathbf{x} = (x^1, \ldots, x^n)'$ is the vector of input prices. In each period, the action, i.e., the decision to be made, is the vector of order-up-to net inventory levels $\mathbf{y} = (y^1, \ldots, y^n)'$ where $y_i \in [s_i, \infty)$ for $i = 1, \ldots, n$. If, in a particular period, we bring the net inventory up to $\mathbf{y}$, and the realized demand is ξ, then the net inventory level in the subsequent period is $\mathbf{y} - \xi$.

For a given state $(\mathbf{s}, \mathbf{x})$ at the beginning of period t, let $s^k = \min\{s^1, \ldots, s^n\}$. To compute the ordering cost and the one-period holding and shortage cost, we consider two cases: (i) $s^k \geq 0$ and (ii) $s^k < 0$. Let $\widehat{y} = \min\{y^1, \ldots, y^n\}$. In case (i), we have no backorders, and the inventory level of component i is s^i for $i = 1, \ldots, n$. If we decide to bring the net inventory level up to $\mathbf{y}$, then the ordering cost is $\sum_{i=1}^{n} x^i(y^i - s^i)$. Note that in this period we can satisfy at most $\widehat{y}$ units of demand. If demand D is less than or equal to $\widehat{y}$, the holding cost for component i is $h^i(y^i - D)$. If demand D is larger than $\widehat{y}$, the holding cost for component i is $h^i(y^i - \widehat{y})$ and the backorder cost is $b(D - \widehat{y})$. Therefore, the one-period holding and shortage cost is

$$\begin{aligned} L(\mathbf{y}) &= \sum_{i=1}^{n} h^i \left(\int_0^{\widehat{y}} (y^i - \xi)\phi(\xi)d\xi + \int_{\widehat{y}}^{\infty} (y^i - \widehat{y})\phi(\xi)d\xi \right) + b\int_{\widehat{y}}^{\infty} (\xi - \widehat{y})\phi(\xi)d\xi \\ &= \sum_{i=1}^{n} h^i E(\widehat{y} - D)^+ + bE(D - \widehat{y})^+ + \sum_{i=1}^{n} h^i(y^i - \widehat{y}). \end{aligned} \tag{2.9}$$

In case (ii), we have $-s^k$ units of backorders, and the inventory level of component i is $s^i - s^k$ for $i = 1, \ldots, n$. If we decide to bring the net inventory level up to $\mathbf{y}$, (or equivalently, we decide to bring the inventory level up to $\mathbf{y} - s^k$), then the ordering cost is

$$\sum_{i}^{n} x^i[(y^i - s^k) - (s^i - s^k)] = \sum_{i=1}^{n} x^i(y^i - s^i).$$

After bringing the inventory levels to $\mathbf{y} - s^k$, we can satisfy at most $\min\{y^1 - s^k, y^2 - s^k, \ldots, y^n - s^k\} = \widehat{y} - s^k$ units of backorders and new demand. Backorders and new demand combined equal $D - s_k$. So, by the same argument that gave us (2.9), the one-period holding and shortage cost is

$$\begin{aligned} \tilde{L}(\mathbf{y}, \mathbf{s}) = \sum_{i=1}^{n} & h^i E[(\widehat{y} - s^k) - (D - s^k)]^+ \\ & + bE[(D - s^k) - (\widehat{y} - s^k)]^+ + \sum_{i=1}^{n} h^i[(y^i - s^k) - (\widehat{y} - s^k)] \end{aligned}$$

$$= \sum_{i=1}^{n} h^i E(\widehat{y} - D)^+ + bE(D - \widehat{y})^+ + \sum_{i=1}^{n} h^i (y^i - \widehat{y})$$

$$= L(\mathbf{y}).$$

For both cases, the ordering cost is $\sum_{i=1}^{n} x^i(y^i - s^i)$ and the one-period holding and shortage cost is

$$L(\mathbf{y}) = \sum_{i=1}^{n} h^i \left(\int_0^{\widehat{y}} (y^i - \xi)\phi(\xi)d\xi + \int_{\widehat{y}}^{\infty} (y^i - \widehat{y})\phi(\xi)d\xi \right) + b \int_{\widehat{y}}^{\infty} (\xi - \widehat{y})\phi(\xi)d\xi,$$

where $\widehat{y} = \min\{y^1, \ldots, y^n\}$.

Let $\mathbf{g}(\boldsymbol{\epsilon}_t) = (g^1(\epsilon_t^1), \ldots, g^n(\epsilon_t^n))'$ and let $\Lambda(\boldsymbol{\epsilon}_t)$ be the $n \times n$ matrix with diagonal entries $f^1(\epsilon_t^1), \ldots, f^n(\epsilon_t^n)$ and other entries 0. Then we have $\mathbf{X}_{t+1} = \Lambda(\boldsymbol{\epsilon}_t)\mathbf{X_t} + \mathbf{g}(\boldsymbol{\epsilon}_t)$. The optimality equations are

$$\begin{aligned} v_t(\mathbf{s}, \mathbf{x}) &= \min_{\mathbf{y} \geq \mathbf{s}} \left\{ \mathbf{x}'(\mathbf{y} - \mathbf{s}) + L(\mathbf{y}) + \beta \int_{\xi} E[v_{t+1}(\mathbf{y} - \xi, \mathbf{X_{t+1}}) | \mathbf{X} = \mathbf{x}]\phi(\xi)d\xi \right\} \\ &= \min_{\mathbf{y} \geq \mathbf{s}} \left\{ \mathbf{x}'\mathbf{y} + L(\mathbf{y}) + \beta \int_{\xi} Ev_{t+1}(\mathbf{y} - \xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t))\phi(\xi)d\xi \right\} - \mathbf{x}'\mathbf{s} \\ &= \min_{\mathbf{y} \geq \mathbf{s}} w_t(\mathbf{y}, \mathbf{x}) - \mathbf{x}'\mathbf{s} \end{aligned}$$

and $v_{T+1}(\mathbf{s}, \mathbf{x}) = 0$, where

$$w_t(\mathbf{y}, \mathbf{x}) = \mathbf{x}'\mathbf{y} + L(\mathbf{y}) + \beta \int_{\xi} Ev_{t+1}(\mathbf{y} - \xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t))\phi(\xi)d\xi.$$

Lemma 5. *$L(\mathbf{y})$ is a convex and submodular function of $\mathbf{y}$ and $v_t(\mathbf{s}, \mathbf{x})$ is a convex and submodular function of $\mathbf{s}$ for all $\mathbf{x}$ and $t = 1, \ldots, T + 1$.*

The following theorem follows directly from Lemma 5.

Theorem 3. *The optimal policy is a state-dependent base stock policy for each component. For component i, there exists a base stock level $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$ where $\mathbf{s}^{-i} = (s^1, \ldots, s^{i-1}, s^{i+1}, \ldots, s^n)$ such that if the starting net inventory s^i in period t is less than $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$, then we order up to $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$; otherwise, we do not order. That is, the optimal order-up-to level for component i in state $(\mathbf{s}, \mathbf{x})$ is $\max\{s^i, y_t^i(\mathbf{s}^{-i}, \mathbf{x})\}$. In addition, the base stock level $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$ is increasing in each s_j for $j \neq i$.*

The structure of the optimal policy is illustrated in Figure 2.2 for a system with two components, where Figures 2.2(a) and 2.2(b) illustrate the policy in period 1 for two different realized prices. When the starting inventory for the two components is in region I, we order both components; in region II, we order only component 2; in region III, we order only component 1; and in region IV, we order nothing. The figure provides some insights into the effect of the price of component 1 (the price of component 2 is fixed in this example). First, notice that a decrease in the price of component 1 leads to higher order up to levels for both components 1 and 2. Second, notice that the optimal policy may not always seek to balance the inventory of both components. For example, when the starting inventory is in region I and the price is high, it is optimal to balance the inventory of the two components. However, when the price is low, it is optimal to bring the inventory of component 1 to a higher level than that of component 2 to take advantage of the lower price of component 1 (more of component 2 can always be ordered in future periods at the same price).

We provide conditions under which the base stock levels are decreasing with respect to the realized price of each component in the following proposition.

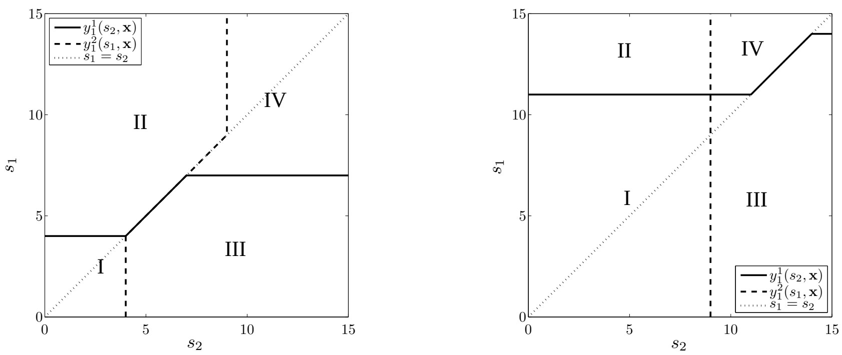

(a) $X_1^1 = 70$ (b) $X_1^1 = 30$

Figure 2.2: Structure of the optimal policy for different realizations of input prices. In this example, the price of component 1 is stochastic and the price of component 2 is fixed. Demand is uniformly distributed on [1,15], $T = 12$, $P(X_t^1 = 30) = P(X_t^1 = 70) = 0.5$ for all t, $X_t^2 = 40$ for all t, $\beta = 0.99$, $b = 50$, $h_1 = 20$, and $h_2 = 40$.

Proposition 4. *If* $0 \leq f^i(\epsilon) \leq 1$ *for all* ϵ *and* $i = 1, \ldots, n$, *then* $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$ *is decreasing in each* x_j *for* $j = 1, \ldots, n$.

The base stock level $y_t^i(\mathbf{s}^{-i}, \mathbf{x})$ need not be decreasing in x_j if the condition in the above proposition is not satisfied. If $f^i(\epsilon) > 1$ for some i, it is possible that a high (low) price of component i in one period would lead to an even higher (lower) expected price of component i in the next period, and it may be optimal to order more (less) of component i when the price of component i is high (low). If $f^i(\epsilon) < 0$ for some i, an increase in the price of component i would lead to a decrease in the expected price of component i in the next period and possibly an increase in the order up to level for component i in the next period. To keep up with a higher order up to level of component i in the next period, it may be optimal to order more of other components. Therefore, in this case, the order up to level for the other components may be increasing in the price of component i.

Impact of Price Variability. With regard to the impact of price variability on the optimal expected total cost, we have similar results as in the single component case. Consider two different inventory systems with input price sequences $\{\mathbf{X}_t\}$ and $\{\widehat{\mathbf{X}}_t\}$ satisfying $X^i_{t+1} = f^i(\epsilon^i_t)X^i_t + g^i(\epsilon^i_t)$ and $\widehat{X}^i_{t+1} = f^i(\widehat{\epsilon}^i_t)\widehat{X}^i_t + g^i(\widehat{\epsilon}^i_t)$. All other parameters of the two systems are the same. Let $v_t(\mathbf{s},\mathbf{x})$ and $\widehat{v}_t(\mathbf{s},\mathbf{x})$ be the optimal total cost-to-go in period t when the net inventory levels are $\mathbf{s}$ and the input prices are $\mathbf{x}$ in period t for the two systems.

The following theorem shows that higher variability in the input prices yields lower optimal expected total cost. Here we use the notion of convex orders of random vectors. A random vector $\mathbf{X}$ is said to be smaller than $\widehat{\mathbf{X}}$ in the *convex order* (written $\mathbf{X} \leq_{cx} \widehat{\mathbf{X}}$) if $Eu(\mathbf{X}) \leq Eu(\widehat{\mathbf{X}})$ for all convex functions $u(\cdot)$ such that the expectations exist. The convex order of random vectors is reviewed in, for example, Müller and Stoyan (2002) and Shaked and Shanthikumar (2007). If $\mathbf{X} = (X^1, \ldots, X^n)$ and $\widehat{\mathbf{X}} = (\widehat{X}^1, \ldots, \widehat{X}^n)$ each have independent components, then $\mathbf{X} \leq_{cx} \widehat{\mathbf{X}}$ is equivalent to $X^i \leq_{cx} \widehat{X}^i$ for all $i = 1, \ldots, n$. If $\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\widehat{\mathbf{X}} \sim N(\widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\Sigma}})$, then $\mathbf{X} \leq_{cx} \widehat{\mathbf{X}}$ if and only if $\boldsymbol{\mu} = \widehat{\boldsymbol{\mu}}$ and $\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}$ is positive semidefinite.

Theorem 4. *Suppose* $\mathbf{X}_1 \leq_{cx} \widehat{\mathbf{X}}_1$. *If*

(a) $\boldsymbol{\epsilon}_t = \widehat{\boldsymbol{\epsilon}}_t$ *for* $t = 1, \ldots, T$, *or*

(b) $\boldsymbol{\epsilon}_t \leq_{cx} \widehat{\boldsymbol{\epsilon}}_t$ *for* $t = 1, \ldots, T$ *and* $f^i(\cdot)$ *and* $g^i(\cdot)$ *are affine functions for* $i = 1, \ldots, n$, *or*

(c) $(f^1(\epsilon^1_t), \ldots, f^n(\epsilon^n_t)) \leq_{cx} (f^1(\widehat{\epsilon}^1_t), \ldots, f^n(\widehat{\epsilon}^n_t))$ *for* $t = 1, \ldots, T$ *and* $g^i(\cdot)$ *is a constant for* $i = 1, \ldots, n$,

then

(1) $v_t(\mathbf{s},\mathbf{x}) \geq \widehat{v}_t(\mathbf{s},\mathbf{x})$ *for all* $\mathbf{s}$, $\mathbf{x}$ *and* $t = 1,\ldots,T$;

(2) $E[v_t(\mathbf{s},\mathbf{X}_t)|\mathbf{X}_{t-1} = \mathbf{x}] \geq E[\widehat{v}_t(\mathbf{s},\widehat{\mathbf{X}}_t)|\widehat{\mathbf{X}}_{t-1} = \mathbf{x}]$ *for all* $\mathbf{s}$, $\mathbf{x}$ *and* $t = 2,\ldots,T$; *and*

(3) $Ev_t(\mathbf{s},\mathbf{X}_t) \geq E\widehat{v}_t(\mathbf{s},\widehat{\mathbf{X}}_t)$ *for all* $\mathbf{s}$ *and* $t = 1,\ldots,T$.

Impact of Correlation across Component Prices. Next, we study the impact of correlation across component prices on the optimal expected total cost. We compare the expected costs of two different inventory systems where the correlations across component prices in one system are larger than those in the other system in every period. More precisely, we consider two systems such that for each time t, the random input price vectors $\mathbf{X}_t$ and $\widehat{\mathbf{X}}_t$ of the two systems are comparable in the supermodular order; i.e., $\mathbf{X}_t \leq_{sm} \widehat{\mathbf{X}}_t$. Recall that this implies that the price correlations are ordered as well; i.e., $\mathrm{Corr}(X_t^i, X_t^j) \leq \mathrm{Corr}(\widehat{X}_t^i, \widehat{X}_t^j)$ for all i, j.

Let $\mathbf{X}_{t+1}(\mathbf{x}) = \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t)$ be a random vector that follows the conditional distribution of $\mathbf{X}_{t+1}$ given $\mathbf{X}_t = \mathbf{x}$ and let $\widehat{\mathbf{X}}_{t+1}(\mathbf{x}) = \Lambda(\widehat{\boldsymbol{\epsilon}}_t)\mathbf{x} + \mathbf{g}(\widehat{\boldsymbol{\epsilon}}_t)$ be a random vector that follows the conditional distribution of $\widehat{\mathbf{X}}_{t+1}$ given $\widehat{\mathbf{X}}_t = \mathbf{x}$.

Lemma 6. *Consider two price sequences* $\{X_t^i\}$ *and* $\{\widehat{X}_t^i\}$, *where* $X_{t+1}^i = f^i(\epsilon_t^i)X_t^i + g^i(\epsilon_t^i)$ *and* $\widehat{X}_{t+1}^i = f^i(\widehat{\epsilon}_t^i)\widehat{X}_t^i + g^i(\widehat{\epsilon}_t^i)$. *If* $\mathbf{X}_1 \leq_{sm} \widehat{\mathbf{X}}_1$, $f^i(\epsilon^i)f^j(\epsilon^j) \geq 0$ *for all* ϵ^i, ϵ^j, $i \neq j$, *and either:*

(a) $\boldsymbol{\epsilon}_t = \widehat{\boldsymbol{\epsilon}}_t$ *for* $t = 1,\ldots,T$ *or*

(b) $\boldsymbol{\epsilon}_t \leq_{sm} \widehat{\boldsymbol{\epsilon}}_t$ *for* $t = 1, \ldots, T$, $f^i(\cdot)$ *is a constant for* $i = 1, \ldots, n$, *and* $g^i(\cdot)$ *is either increasing for all* $i = 1, \ldots, n$ *or decreasing for all* $i = 1, \ldots, n$,

then $\mathbf{X}_t \leq_{sm} \widehat{\mathbf{X}}_t$ *and* $\mathbf{X}_t(\mathbf{x}) \leq_{sm} \widehat{\mathbf{X}}_t(\mathbf{x})$ *for all* $\mathbf{x}$ *and* $t = 1, \ldots, T$.

One example of property (b) is the case where each component price evolves according to an AR(1) process; i.e., $X^i_{t+1} = \rho^i X^i_t + \epsilon^i_t + c^i$ where $\rho^i \rho^j \geq 0$ for all $i \neq j$.

Lemma 7. *If* $f^i(\epsilon^i) f^j(\epsilon^j) \geq 0$ *for all* ϵ^i, ϵ^j, $i \neq j$, *then* $v_t(\mathbf{s}, \mathbf{x})$ *is a submodular function of* $\mathbf{x}$ *for all* $\mathbf{s}$ *and* $t = 1, \ldots, T+1$.

From the definition of the supermodular order, we have the following theorem describing the impact of correlation over component prices on the optimal expected total cost.

Theorem 5. *Suppose the conditions in Lemma 6 hold. Then*

(1) $v_t(\mathbf{s}, \mathbf{x}) \leq \widehat{v}_t(\mathbf{s}, \mathbf{x})$ *for all* $\mathbf{s}$, $\mathbf{x}$ *and* $t = 1, \ldots, T$;

(2) $E[v_t(\mathbf{s}, \mathbf{X}_t) | \mathbf{X}_{t-1} = \mathbf{x}] \leq E[\widehat{v}_t(\mathbf{s}, \widehat{\mathbf{X}}_t) | \widehat{\mathbf{X}}_{t-1} = \mathbf{x}]$ *for all* $\mathbf{s}$, $\mathbf{x}$ *and* $t = 2, \ldots, T$; *and*

(3) $E v_t(\mathbf{s}, \mathbf{X}_t) \leq E \widehat{v}_t(\mathbf{s}, \widehat{\mathbf{X}}_t)$ *for all* $\mathbf{s}$ *and* $t = 1, \ldots, T$.

Proof. It is easy to check by backward induction that $v_t(\mathbf{s}, \mathbf{x}) \leq \widehat{v}_t(\mathbf{s}, \mathbf{x})$ for all $\mathbf{s}$ and $\mathbf{x}$. By Lemma 6, we have $\mathbf{X}_t(\mathbf{x}) \leq_{sm} \widehat{\mathbf{X}}_t(\mathbf{x})$ and $\mathbf{X}_t \leq_{sm} \widehat{\mathbf{X}}_t$ for $\mathbf{x}$ and $t = 1, \ldots, T$. Therefore,

$$E[v_t(\mathbf{s}, \mathbf{X}_t) | \mathbf{X}_{t-1} = \mathbf{x}] = E v_t(\mathbf{s}, \mathbf{X}_t(\mathbf{x}))$$

$$\le E\widehat{v}_t(\mathbf{s}, \mathbf{X}_t(\mathbf{x}))$$
$$\le E\widehat{v}_t(\mathbf{s}, \widehat{\mathbf{X}}_t(\mathbf{x}))$$
$$= E[\widehat{v}_t(\mathbf{s}, \widehat{\mathbf{X}}_t) | \widehat{\mathbf{X}}_{t-1} = \mathbf{x}],$$

where the second inequality follows from Lemma 7. Similarly, $Ev_t(\mathbf{s}, \mathbf{X}_t) \le E\widehat{v}_t(\mathbf{s}, \mathbf{X}_t) \le E\widehat{v}_t(\mathbf{s}, \widehat{\mathbf{X}}_t)$. $\square$

This theorem shows that expected total cost is decreasing in correlation across component prices, implying that higher correlation across component prices is beneficial. To provide some intuition as to why such higher correlation leads to lower costs, consider the single period case. In that case, it is optimal to always order the same quantity of each component (assuming equal starting inventory levels). Therefore, the problem reduces to one of a single component with unit price equal to the sum of the unit prices of all the components. If we let X be the random variable that describes this "equivalent" unit price and $X^1, \dots, X^n$ be the individual component prices, then $X = X^1 + X^2 + \cdots + X^n$ and $\mathrm{Var}(X) = \sum_{i=1}^{n} \mathrm{Var}(X^i) + \sum_{i \ne j} \mathrm{Cov}(X^i, X^j)$. As we can see, higher price correlation leads to higher price variance, which for several common distributions, also implies higher price variability as measured by the convex order. Therefore, in such cases, higher correlation would lead to lower cost. Correlation can also impact cost by affecting order up to levels. For example, in the settings described by Proposition 4, the order up to level of each component is decreasing in the price of all other components. Therefore, when the price of one component is low and it is desirable to order more, it is preferable that the prices of other components are also low. Otherwise, the opportunity to take advantage of price variability is diminished.

2.6 Numerical Results

In this section, we provide numerical results that illustrate how the relative benefits of price variability, price correlation over time, and price correlation over components are affected by various problem parameters. First, we compare the performance of systems with and without price variability and compare the performance of systems with and without correlation over periods for systems with a single component. Then for systems with multiple input components, we compare the performance of systems with and without correlation across component prices.

Let $v_t(s, x)$ be the optimal cost-to-go in period t when the beginning inventory is s and the realization of price is x in period t for the system with input price sequence $\{X_t\}$ and let $\bar{v}_t(s)$ be the optimal cost-to-go in period t when the beginning inventory is s for the system with fixed input price sequence $\{\mu_t\}$, where $\mu_t = EX_t$. The relative benefit of price variability, which we denote by δ_{v}, is defined as follows:

$$\delta_{\mathrm{v}} = \frac{\bar{v}_1(s) - Ev_1(s, X_1)}{\bar{v}_1(s)}.$$

In Figures 2.3(a)–2.3(d), we examine the relative benefit of price variability for different lengths of planning horizons, different holding costs, different backorder costs, and different levels of price correlation over time. In all the numerical examples, we set the initial inventory to be 0, the discount factor to be $\beta = 0.99$, and demand to be uniformly distributed on [1,30]. (Results are qualitatively the same for other common distributions we tested.) In Figures 2.3(a)–2.3(c), the input prices are i.i.d. across periods with $P(X_t = 80 + \alpha) = P(X_t = 80 - \alpha) = 0.5$, in which case the standard deviation of input prices is α. In Figure 2.3(d), the

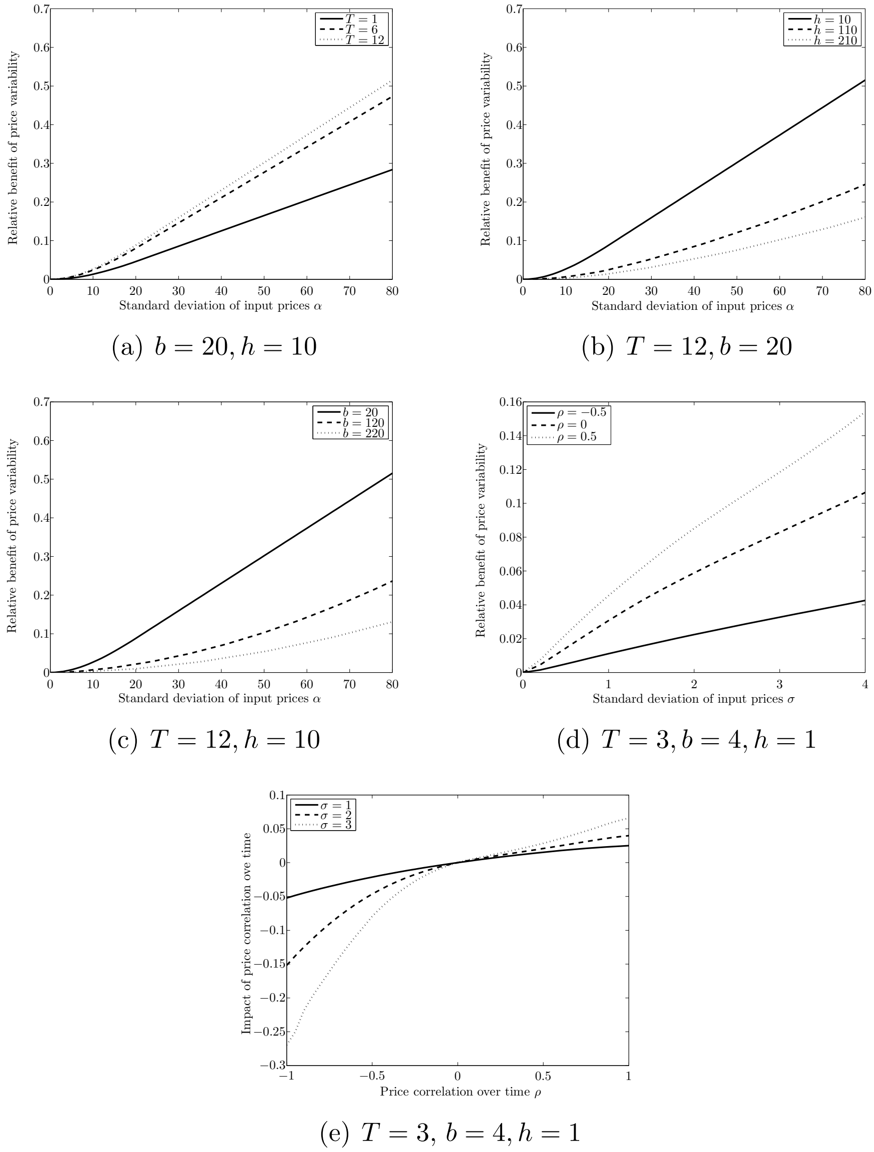

(a) $b = 20, h = 10$

(b) $T = 12, b = 20$

(c) $T = 12, h = 10$

(d) $T = 3, b = 4, h = 1$

(e) $T = 3$, $b = 4, h = 1$

Figure 2.3: Impact of price variability and price correlation over time in systems with a single input.

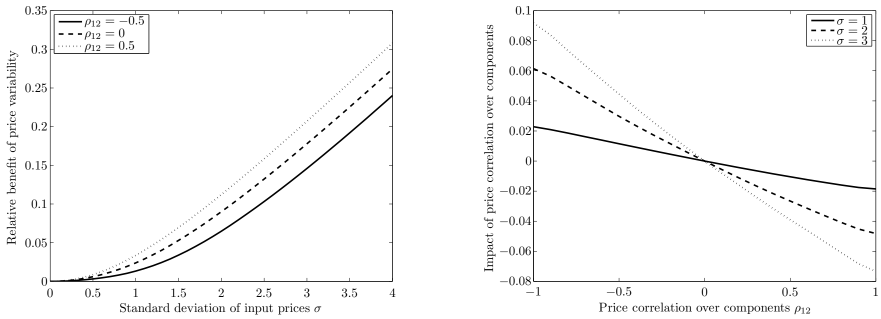

(a) $b = 2, h_1 = h_2 = 2$ (b) $b = 2, h_1 = h_2 = 2$

Figure 2.4: Impact of price variability and correlation across components in systems with multiple inputs.

input prices follow a stationary AR(1) process, namely, $X_{t+1} = (1-\rho)\mu + \rho X_t + \sqrt{1-\rho^2}\epsilon_t$, where $\mu = EX_1 = 6$, $\{\epsilon_t\}$ are normally distributed with mean 0 and variance σ^2, and $\sigma^2 = \text{Var}(X_1)$.

In Figure 2.3(e), we examine the relative impact of price correlation over time for different levels of price variability. Here, the sequence of input prices $\{X_t\}$ is a stationary AR(1) process with mean $\mu = 6$. Let $v_t(s, x, \rho)$ be the optimal cost-to-go in period t when the beginning inventory is s and the realization of price is x in period t for the above system. Note that $v_t(s, x, \rho)$ is increasing in $\rho \geq 0$ by Theorem 2. The relative impact of price correlation over time, which denoted by δ_{ct}, is defined as follows:

$$\delta_{\text{ct}} = \frac{Ev_1(s, X_1, \rho) - Ev_1(s, X_1, 0)}{Ev_1(s, X_1, 0)}.$$

For inventory systems with multiple inputs, we examine how the benefit of price variability is affected by price correlation across components and how the relative impact of price correlation across components is affected by other parameters. We consider a 2-component, 2-period problem as an example. The input prices of the components are i.i.d. and in each period the prices are normally distributed with mean $\boldsymbol{\mu}$ and variance matrix $\boldsymbol{\Sigma}$, where $\mu_1 = \mu_2 = 6$, $\sigma_{11} = \sigma_{22}$, and $\sigma_{12} = \sigma_{21} = \sigma_{11}\rho_{12}$. Let $v_t(\mathbf{s}, \mathbf{x}, \rho_{12})$ be the optimal cost-to-go in period t when the beginning inventory levels are $\mathbf{s}$, the realizations of prices are $\mathbf{x}$ and the price correlation between the two components is ρ_{12} in period t. Then the relative impact of price correlation across components, which is denoted by δ_{cc}, is defined as follows:

$$\delta_{\text{cc}} = \frac{Ev_1(\mathbf{s}, \mathbf{X}, \rho_{12}) - Ev_1(\mathbf{s}, \mathbf{X}, 0)}{Ev_1(\mathbf{s}, \mathbf{X}, 0)}.$$

Figure 2.4(a) shows the relative benefit of price variability for different levels of component price correlation and Figure 2.4(b) shows the relative impact of component price correlation for different levels of price variability.

Based on Figure 2.3 and Figure 2.4, we can make the following observations (we provide some intuition to explain these observations; however, we caution that, in general, the interactions between various factors can be quite complex).

Observation 1: *The relative benefit of price variability is increasing in the length of the planning horizon.* This is illustrated in Figure 2.3(a). When prices are high, a firm can order less and take advantage of the possibility of backordering and fulfilling demand in future periods. Similarly, when prices are low, a firm can order more and take advantage of the possibility of holding inventory and using this inventory to fulfill demand in future periods. The advantage derived from the flexibility of either backordering or carrying inventory across periods (to which we refer as the period-over-period effect) increases with the length of the planning horizon, as the opportunity to exercise this flexibility also increases.

Observation 2: *The relative benefit of price variability is decreasing in the* *holding and backorder costs.* This is illustrated in Figures 2.3(b) and 2.3(c). When either the holding or the backorder cost is high, taking advantage of the period-over-period effect (ordering more and holding inventory or ordering less and backordering) becomes less desirable. In turn, this diminishes the benefit that may be derived from higher price variability.

Observation 3: *The relative benefit of price variability is decreasing in the price correlation over time.* This is illustrated in Figure 2.3(d). The benefit derived from ordering more (less) in periods when prices are low (high) diminishes with correlation over time, as a low (high) price period tends to be followed by another low (high) price period.

Observation 4: *The relative benefit from lower correlation over time is increasing in price variability.* This is illustrated in Figure 2.3(e). Lower correlation over time provides an opportunity to take advantage of the period-over-period effect. This opportunity is enhanced when price variability is high. We can also see from Figure 2.3(e) that for the examples depicted there, systems with uncorrelated prices have greater expected cost than systems with negatively correlated prices. This suggests that we can perhaps relax the condition that correlations are positive in Theorem 2, at least in some cases.

Observation 5: *The relative benefit of price variability is increasing in the price correlation across components.* This is illustrated in Figure 2.4(a). In general, the interaction between price variability and price correlation across components is complex and depends on the correlations of prices of components over time. However, higher price correlation among components typically enables a firm to take better advantage of variability. For example, when the price of a component is low, we may prefer to buy more of that component. The value of doing so is greater when we also prefer to buy more of other components (recall that one unit of each component is needed to fulfill demand). Such opportunities will arise more frequently when prices across components are more correlated than when they are less so.

Observation 6: *The relative benefit of higher price correlation across components is increasing in price variability.* This is illustrated in Figure 2.4(b). Higher correlation typically implies that when it is preferable to order more (less) of one component it is also preferable to order more (less) of other components. This matching of inventory levels across components is more valuable when the price variability of components is high and the benefit from adjusting order quantities is greater.

2.7 Appendix: Proofs

Proof of Lemma 1: Observe first that $w_T(y,x) = xy + L(y)$ is convex in y for all x. Consider arbitrary $t \in \{1,\ldots,T-1\}$ and suppose inductively that $\frac{\partial^2 w_{t+1}}{\partial y^2}(y,x) \geq 0$. (It can be verified that $v_t(s,x)$ and $w_t(s,x)$ are continuously differentiable by the Envelope Theorem. At any point where $v_t(s,x)$ or $w_t(s,x)$ is not twice differentiable, it can be checked that the left limit of the derivative at this point is less than or equal to the right limit. A similar approach can be used whenever we use second derivatives in the proofs.) Let $y_t^\circ(x)$ denote a minimizer of $w_t(y,x)$ over $y \in (-\infty,\infty)$. Then

$$v_{t+1}(s,x) = \begin{cases} w_{t+1}(s,x) - xs & \text{if } s \geq y_{t+1}^\circ(x), \\ w_{t+1}(y_{t+1}^\circ(x),x) - xs & \text{otherwise,} \end{cases}$$

and

$$\frac{\partial^2 v_{t+1}}{\partial s^2}(s,x) = \begin{cases} \frac{\partial^2 w_{t+1}}{\partial y^2}(s,x) & \text{if } s \geq y_{t+1}^\circ(x), \\ 0 & \text{otherwise.} \end{cases}$$

Thus, $\frac{\partial^2 v_{t+1}}{\partial s^2}(s,x) \geq 0$ for all (s,x), and therefore

$$\frac{\partial^2 w_t}{\partial y^2}(y,x) = L''(y) + \beta \int_\xi \int_\epsilon \frac{\partial^2 v_{t+1}}{\partial s^2}(y-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi \geq 0.$$

By backward induction, we have $\frac{\partial^2 w_t}{\partial y^2}(y,x) \geq 0$ for all x and $t = 1,\ldots,T$. $\square$

Proof of Proposition 2: We first prove by backward induction that $\frac{\partial^2 v_t}{\partial s \partial x}(s,x) \in [-1,1]$ for all (s,x) and $t = 1,\ldots,T+1$. It is true when $t = T+1$ because $v_{T+1}(s,x) = 0$. Suppose $\frac{\partial^2 v_{t+1}}{\partial s \partial x}(s,x) \in [-1,1]$ for all (s,x). By (2.2) we have

$v_t(s,x) = w_t(y_t^*(s,x),x) - sx$. By Proposition 1 if $s \le y_t^\circ(x)$, then $v_t(s,x) = w_t(y_t^\circ(x),x) - sx$ and $\frac{\partial^2 v_t}{\partial s \partial x}(s,x) = -1$. If $s > y_t^\circ(x)$, then

$$v_t(s,x) = w_t(s,x) - sx = L(s) + \beta \int_\xi \int_\epsilon v_{t+1}(s-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi,$$

and

$$\begin{aligned}\left|\frac{\partial^2 v_t}{\partial s \partial x}(s,x)\right| &\le \beta \int_\xi \int_\epsilon \left| f(\epsilon)\frac{\partial^2 v_{t+1}}{\partial s \partial x}(s-\xi, f(\epsilon)x + g(\epsilon))\right| \Psi_t(d\epsilon)\phi(\xi)d\xi \\ &\le E|f(\epsilon_t)| \le 1.\end{aligned}$$

The second inequality above follows from the inductive hypothesis.

Since $y_t^\circ(x)$ is a minimizer of $w_t(y,x)$ over $y \in (-\infty,\infty)$ and $w_t(y,x)$ is convex in y, it follows that $\frac{\partial w_t}{\partial y}(y_t^\circ(x),x) = 0$. Note that

$$\frac{\partial^2 w_t}{\partial y \partial x}(y,x) = 1 + \beta \int_\xi \int_\epsilon \frac{\partial^2 v_{t+1}}{\partial s \partial x}(y-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi \ge 1-\beta \ge 0.$$

Hence, $\frac{\partial w_t}{\partial y}(y,x)$ is increasing in x. Therefore, for any $x' < x$, $\frac{\partial w_t}{\partial y}(y_t^\circ(x),x') \le 0$. By the definition of $y_t^\circ(x')$ and the convexity of $w_t(y,x')$ in y, we have $y_t^\circ(x') \ge y_t^\circ(x)$. Thus, $y_t^\circ(x)$ is decreasing in x. □

Proof of Lemma 2: We have $v_{T+1}(s,x) = 0$, which is a concave function of x. Suppose $\frac{\partial^2 v_{t+1}}{\partial x^2}(s,x) \le 0$ for all x. Using (3), we have

$$\frac{\partial^2 w_t}{\partial x^2}(y,x) = \beta \int_\xi \int_\epsilon f^2(\epsilon)\frac{\partial^2 v_{t+1}}{\partial x^2}(y-\xi, f(\epsilon)x + g(\epsilon))\Psi_t(d\epsilon)\phi(\xi)d\xi \le 0.$$

Thus, $w_t(y,x)$ is a concave function of x. Since concavity is preserved under minimization, $v_t(s,x) = \min_{y \ge s}\{w_t(y,x)\} - xs$ is also a concave function of x. By backward induction, $\frac{\partial^2 v_t}{\partial x^2}(s,x) \leq 0$ for all x and $t = 1, \ldots, T+1$ and $v_t(s,x)$ is convex in x. $\square$

Proof of Lemma 4: The proof is similar to that of Proposition 2. $\square$

Proof of Lemma 5: We first show that $L(\mathbf{y})$ is a convex and submodular function of $\mathbf{y}$. We have

$$\frac{\partial L}{\partial y^i}(\mathbf{y}) = \begin{cases} h^i - (\sum_{j=1}^{n} h^j + b) \int_{y^i}^{\infty} \phi(\xi) d\xi & \text{if } y^i \leq \min\{y^1, \ldots, y^{i-1}, y^{i+1}, \ldots, y^n\}, \\ h^i & \text{if } y^i > \min\{y^1, \ldots, y^{i-1}, y^{i+1}, \ldots, y^n\}, \end{cases}$$

which is decreasing in y^j for $i \leq j$. Therefore, $L(\mathbf{y})$ is a submodular function of $\mathbf{y}$. Note that $L(\mathbf{y})$ can be written in the following form

$$L(\mathbf{y}) = \widehat{L}(\widehat{y}) + \sum_{i=1}^{n} h^i (y^i - \widehat{y}),$$

where $\widehat{L}(y) = \sum_{i=1}^{n} h^i \int_0^y (y - \xi)\phi(\xi)d\xi + \int_y^{\infty} b(\xi - y)\phi(\xi)d\xi$. Let $\mathbf{e}^i$ be the ith unit vector of dimension n. Since $h^i \int_0^{y^i} \phi(\xi)d\xi - (\sum_{j \neq i} h^j + b) \int_{y^i}^{\infty} \phi(\xi)d\xi \leq h^i$, we have $\frac{\partial L}{\partial y^i}(\mathbf{y}) \leq h^i$ for all $\mathbf{y}$. Therefore,

$$L(\mathbf{y} + \mathbf{e}^i) - L(\mathbf{y}) \leq h^i. \tag{2.10}$$

Given $\mathbf{y}$, define $\bar{\mathbf{y}}^i = (\max\{y^1, y^i\}, \ldots, \max\{y^n, y^i\})$ for $i = 1, \ldots, n$. Using (2.10) we have

$$\begin{aligned} L(\mathbf{y}) &\geq L(\bar{\mathbf{y}}^i) + \sum_{j: y^j \leq y^i} h^j (y^j - y^i) \\ &= \widehat{L}(y^i) + \sum_{j: y^j > y^i} h^j (y^j - y^i) + \sum_{j: y^j \leq y^i} h^j (y^j - y^i) \end{aligned}$$

$$= \widehat{L}(y^i) + \sum_{j=1}^{n} h^j(y^j - y^i). \tag{2.11}$$

To show $L(\mathbf{y})$ is convex in $\mathbf{y}$, we need to show that $\lambda L(\mathbf{y}) + (1-\lambda)L(\widetilde{\mathbf{y}}) \geq L(\lambda\mathbf{y} + (1-\lambda)\widetilde{\mathbf{y}})$ for any $\mathbf{y}$, $\widetilde{\mathbf{y}}$, and $\lambda \in [0,1]$. Suppose $\lambda y^k + (1-\lambda)\widetilde{y}^k = \min_{i=1,\ldots,n}\{\lambda y^i + (1-\lambda)\widetilde{y}^i\}$. Then

$$\begin{aligned}
&\lambda L(\mathbf{y}) + (1-\lambda)L(\widetilde{\mathbf{y}}) \\
&\quad \geq \lambda[\widehat{L}(y^k) + \sum_{j=1}^{n} h^j(y^j - y^k)] + (1-\lambda)[\widehat{L}(\widetilde{y}^k) + \sum_{j=1}^{n} h^j(\widetilde{y}^j - \widetilde{y}^k)] \\
&\quad = \lambda\widehat{L}(y^k) + (1-\lambda)\widehat{L}(\widetilde{y}^k) + \sum_{j=1}^{n} h^j[(\lambda y^j + (1-\lambda)\widetilde{y}^j) - (\lambda y^k + (1-\lambda)\widetilde{y}^k)] \\
&\quad \geq \widehat{L}(\lambda y^k + (1-\lambda)\widetilde{y}^k) + \sum_{j=1}^{n} h^j[(\lambda y^j + (1-\lambda)\widetilde{y}^j) - (\lambda y^k + (1-\lambda)\widetilde{y}^k)] \\
&\quad = L(\lambda\mathbf{y} + (1-\lambda)\widetilde{\mathbf{y}}),
\end{aligned}$$

where the first inequality follows from (2.11) and the second from the convexity of $\widehat{L}(y)$.

Next we prove that $v_t(\mathbf{s},\mathbf{x})$ is a convex and submodular function of $\mathbf{s}$ for $t = 1,\ldots,T+1$. We prove this by backward induction. We have that $v_{T+1}(\mathbf{s},\mathbf{x}) = 0$ is convex and submodular in $\mathbf{s}$. Suppose that $v_{t+1}(\mathbf{s},\mathbf{x})$ is convex and submodular in $\mathbf{s}$. Then

$$w_t(\mathbf{y},\mathbf{x}) = \mathbf{x}'\mathbf{y} + L(\mathbf{y}) + \beta\int_{\xi} Ev_{t+1}(\mathbf{y}-\xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t))\phi(\xi)d\xi$$

is a convex and submodular function of $\mathbf{y}$ since $L(\mathbf{y})$ is convex and submodular in $\mathbf{y}$ and $Ev_{t+1}(\mathbf{y}-\xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x}+\mathbf{g}(\boldsymbol{\epsilon}_t))$ is convex and submodular in $\mathbf{y}$ by the inductive assumption. Submodularity is preserved under minimization on a lattice, so $v_t(\mathbf{s},\mathbf{x}) = \min_{\mathbf{y}\geq\mathbf{s}} w_t(\mathbf{y},\mathbf{x})$ is submodular in $\mathbf{s}$.

To show $v_t(\mathbf{s},\mathbf{x})$ is convex in $\mathbf{s}$, we need to show that $\lambda v_t(\mathbf{s},\mathbf{x})+(1-\lambda)v_t(\widetilde{\mathbf{s}},\mathbf{x}) \geq v_t(\lambda\mathbf{s}+(1-\lambda)\widetilde{\mathbf{s}},\mathbf{x})$ for any $\mathbf{s}$, $\widetilde{\mathbf{s}}$, and $\lambda \in [0,1]$. Let $\mathbf{y}_t^*(\mathbf{x}) = \arg\min_{\mathbf{y}\geq\mathbf{s}} w_t(\mathbf{y},\mathbf{x})$ and $\widetilde{\mathbf{y}}_t^*(\mathbf{x}) = \arg\min_{\mathbf{y}\geq\widetilde{\mathbf{s}}} w_t(\mathbf{y},\mathbf{x})$. Then we have

$$\begin{aligned}
\lambda v_t(\mathbf{s},\mathbf{x})+(1-\lambda)v_t(\widetilde{\mathbf{s}},\mathbf{x}) &= \lambda[w_t(\mathbf{y}_t^*(\mathbf{x}),\mathbf{x})-\mathbf{x}'\mathbf{s}]+(1-\lambda)[w_t(\widetilde{\mathbf{y}}_t^*(\mathbf{x}),\mathbf{x})-\mathbf{x}'\widetilde{\mathbf{s}}] \\
&\geq w_t(\lambda\mathbf{y}_t^*(\mathbf{x})+(1-\lambda)\widetilde{\mathbf{y}}_t^*(\mathbf{x}),\mathbf{x})-\mathbf{x}'[\lambda\mathbf{s}+(\mathbf{1}-\lambda)\widetilde{\mathbf{s}}] \\
&\geq \min_{\mathbf{y}\geq\lambda\mathbf{s}+(\mathbf{1}-\lambda)\widetilde{\mathbf{s}}} w_t(\mathbf{y},\mathbf{x})-\mathbf{x}'(\lambda\mathbf{s}+(\mathbf{1}-\lambda)\widetilde{\mathbf{s}}) \\
&= v_t(\lambda\mathbf{s}+(1-\lambda)\widetilde{\mathbf{s}},\mathbf{x}).
\end{aligned}$$

The first inequality is due to the convexity of $w_t(\mathbf{y},\mathbf{x})$ in $\mathbf{y}$ and the second inequality is due to the fact that $\lambda\mathbf{y}_t^*(\mathbf{x})+(1-\lambda)\widetilde{\mathbf{y}}_t^*(\mathbf{x}) \geq \lambda\mathbf{s}+(\mathbf{1}-\lambda)\widetilde{\mathbf{s}}$. $\square$

Proof of Proposition 4: We only need to show that $w_t(\mathbf{y},\mathbf{x})$ is supermodular in (y^i,x^j) for $t=1,\ldots,T$ and $i,j=1,\ldots,n$. The proof of this is similar to that of Proposition 2. $\square$

Proof of Theorem 4: We will show only that $v_t(\mathbf{s},\mathbf{x})$ is concave in $\mathbf{x}$ for $t = 1,\ldots,T+1$. The rest of the proof is similar to the proofs of Theorem 1, Corollary 1, and Lemma 3. We prove the concavity by induction.

For $t=T+1$, we have $v_{T+1}(\mathbf{s},\mathbf{x})=0$. Suppose $v_{t+1}(\mathbf{s},\mathbf{x})$ is concave in $\mathbf{x}$. Then

$$w_t(\mathbf{y},\mathbf{x}) = \mathbf{x}'\mathbf{y}+L(\mathbf{y})+\beta\int_{\xi} Ev_{t+1}(\mathbf{y}-\xi,\Lambda(\boldsymbol{\epsilon}_t)\mathbf{x}+\mathbf{g}(\boldsymbol{\epsilon}_t))\phi(\xi)d\xi$$

is a concave function of $\mathbf{x}$. Concavity is preserved under minimization, so $v_t(\mathbf{s},\mathbf{x})$ is concave in $\mathbf{x}$ for $t=1,\ldots,T+1$. $\square$

Proof of Lemma 6: We prove this by induction. Suppose $\mathbf{X}_t \leq_{sm} \widehat{\mathbf{X}}_t$. Let $u(\cdot)$ be an arbitrary supermodular function and fix $\mathbf{x}$. Then for case (a), we have

$$Eu(\mathbf{X}_{t+1}(\mathbf{x})) = Eu(\Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t)) = Eu(\Lambda(\widehat{\epsilon}_\mathbf{t})\mathbf{x} + \mathbf{g}(\widehat{\epsilon}_\mathbf{t})) = Eu(\widehat{\mathbf{X}}_{t+1}(\mathbf{x})).$$

For case (b), let $\tilde{u}(\boldsymbol{\epsilon}) = u(\Lambda(\boldsymbol{\epsilon})\mathbf{x}+\mathbf{g}(\boldsymbol{\epsilon}))$. Recall that each $f^i(\cdot)$ is assumed to be a constant, say a^i, and therefore $\Lambda(\boldsymbol{\epsilon}) = A$ where A is the matrix with a^i in the ith diagonal position for $i = 1, \ldots, n$ and zeros elsewhere. Hence $\tilde{u}(\boldsymbol{\epsilon}) = u(A\mathbf{x}+\mathbf{g}(\boldsymbol{\epsilon}))$.

We will now argue that $\tilde{u}(\boldsymbol{\epsilon})$ is supermodular. Suppose that $\varepsilon, \delta > 0$ and $i \neq j$. We have

$$\begin{aligned}
&\tilde{u}(\boldsymbol{\epsilon} + \varepsilon\mathbf{e}^i + \delta\mathbf{e}^j) - \tilde{u}(\boldsymbol{\epsilon} + \varepsilon\mathbf{e}^i) - \tilde{u}(\boldsymbol{\epsilon} + \delta\mathbf{e}^j) + \tilde{u}(\boldsymbol{\epsilon}) \\
&\qquad = u(A\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon} + \varepsilon\mathbf{e}^i + \delta\mathbf{e}^j)) - u(A\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon} + \varepsilon\mathbf{e}^i)) \\
&\qquad\quad - u(A\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon} + \delta\mathbf{e}^j)) + u(A\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon})) \\
&\qquad = u(\mathbf{z} + \tilde{\varepsilon}\mathbf{e}^i + \tilde{\delta}\mathbf{e}^j) - u(\mathbf{z} + \tilde{\varepsilon}\mathbf{e}^i) - u(\mathbf{z} + \tilde{\delta}\mathbf{e}^j) + u(\mathbf{z}) \\
&\qquad \geq 0
\end{aligned}$$

where we define $\mathbf{z} = A\mathbf{x}+\mathbf{g}(\boldsymbol{\epsilon})$, $\tilde{\varepsilon} = g^i(\epsilon^i+\varepsilon)-g^i(\epsilon^i)$, and $\tilde{\delta} = g^j(\epsilon^j+\delta)-g^j(\epsilon^j)$. The inequality above follows because $u(\cdot)$ is supermodular and because $\tilde{\epsilon}$ and $\tilde{\delta}$ have the same sign owing to the assumption that $g^i(\cdot)$ and $g^j(\cdot)$ are either both decreasing or both increasing. Hence we have established that $\tilde{u}(\boldsymbol{\epsilon})$ is supermodular.

As a consequence,

$$Eu(\mathbf{X}_{t+1}(\mathbf{x})) = E\tilde{u}(\boldsymbol{\epsilon}_t) \leq E\tilde{u}(\widehat{\epsilon}_\mathbf{t}) = Eu(\widehat{\mathbf{X}}_{t+1}(\mathbf{x})).$$

Thus, for both cases (a) and (b), we have $\mathbf{X}_{t+1}(\mathbf{x}) \leq_{sm} \widehat{\mathbf{X}}_{t+1}(\mathbf{x})$. Let $\eta(\mathbf{x}) = Eu(\mathbf{X}_{t+1}(\mathbf{x}))$ and $\widehat{\eta}(\mathbf{x}) = Eu(\widehat{\mathbf{X}}_{t+1}(\mathbf{x}))$. Then, we have $\eta(\mathbf{x}) \leq \widehat{\eta}(\mathbf{x})$ for all $\mathbf{x}$. In addition, it can be verified that $\widehat{\eta}(\mathbf{x})$ is a supermodular function of $\mathbf{x}$ (here we use $f^i(\epsilon^i)f^j(\epsilon^j) \geq 0$). Hence,

$$Eu(\mathbf{X}_{t+1}) = E\eta(\mathbf{X}_t) \leq E\widehat{\eta}(\mathbf{X}_t) \leq E\widehat{\eta}(\widehat{\mathbf{X}}_t) = Eu(\widehat{\mathbf{X}}_{\mathbf{t+1}}).$$

The second inequality holds because $\mathbf{X}_t \leq_{sm} \widehat{\mathbf{X}}_t$. This completes the induction and the proof. □

Proof of Lemma 7: We prove the result by backward induction. We have that $v_{T+1}(\mathbf{s},\mathbf{x}) = 0$ is trivially submodular in $\mathbf{x}$ for all $\mathbf{s}$. Suppose $v_{t+1}(\mathbf{s},\mathbf{x})$ is submodular in $\mathbf{x}$ for all $\mathbf{s}$ and consider i,j with $i \neq j$. We will establish that $\frac{\partial^2 v_t}{\partial x^i \partial x^j}(\mathbf{s},\mathbf{x}) \leq 0$.

In period t, we need to solve the optimization problem

$$\min \; w_t(\mathbf{y},\mathbf{x}) = \mathbf{x}'\mathbf{y} + L(\mathbf{y}) + \beta \int\limits_{\xi} Ev_{t+1}(\mathbf{y}-\xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t))\phi(\xi)d\xi$$

$$\text{s.t. } \; \mathbf{y} \geq \mathbf{s}.$$

The optimal solution $\mathbf{y}_t^*(\mathbf{s},\mathbf{x}) = (y_t^{*1}(\mathbf{s},\mathbf{x}),\ldots,y_t^{*n}(\mathbf{s},\mathbf{x}))$ satisfies the KKT conditions:

$$\begin{aligned} \frac{\partial w_t}{\partial y^k}(\mathbf{y},\mathbf{x}) - \lambda^k = 0 \quad & \text{for } k = 1,\ldots,n, \\ \lambda^k(y^k - s^k) = 0 \quad & \text{for } k = 1,\ldots,n, \\ \lambda^k \geq 0 \quad & \text{for } k = 1,\ldots,n. \end{aligned}$$

By Theorem 3, for each $k = 1,\ldots,n$, we know that $y_t^{*k}(\mathbf{s},\mathbf{x}) = \max\{s^k, y_t^k(\mathbf{s}^{-k},\mathbf{x})\}$, where $y_t^k(\mathbf{s}^{-k},\mathbf{x})$ is the optimal base-stock level. If $s^k < y_t^k(\mathbf{s}^{-k},\mathbf{x})$, then $y_t^{*k}(\mathbf{s},\mathbf{x}) = y_t^k(\mathbf{s}^{-k},\mathbf{x}) > s^k$. Thus, $\lambda^k = 0$ and $\frac{\partial w_t}{\partial y^k}(\mathbf{y}_t^*(\mathbf{s},\mathbf{x}),\mathbf{x}) = 0$

by the KKT conditions. On the other hand, if $s^k \geq y_t^k(\mathbf{s}^{-k}, \mathbf{x})$, then $y_t^{*k}(\mathbf{s}, \mathbf{x}) = s^k$ and therefore $\frac{\partial y_t^{*k}}{\partial x^i}(\mathbf{s}, \mathbf{x}) = 0$ for any i.

Since $v_t(\mathbf{s}, \mathbf{x}) = w_t(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) - \mathbf{s}'\mathbf{x}$, we have

$$\begin{aligned}\frac{\partial v_t}{\partial x^i}(\mathbf{s}, \mathbf{x}) =& \sum_{k=1}^{n} \frac{\partial w_t}{\partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) \frac{\partial y_t^{*k}}{\partial x^i}(\mathbf{s}, \mathbf{x}) + \frac{\partial w_t}{\partial x^i}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) - s^i \\ =& \frac{\partial w_t}{\partial x^i}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) - s^i.\end{aligned}$$

Letting $I = \{k : s^k < y_t^k(\mathbf{s}^{-k}, \mathbf{x})\}$, we obtain

$$\frac{\partial^2 v_t}{\partial x^i \partial x^j}(\mathbf{s}, \mathbf{x}) = \sum_{k \in I} \frac{\partial^2 w_t}{\partial x^i \partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) \frac{\partial y_t^k}{\partial x^i}(\mathbf{s}^{-i}, \mathbf{x}) + \frac{\partial^2 w_t}{\partial x^i \partial x^j}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}). \quad (2.12)$$

For the second term on the right side of (2.12), by the inductive hypothesis we have

$$\begin{aligned}&\frac{\partial^2 w_t}{\partial x^i \partial x^j}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) \\ &\quad = \beta \int_{\xi} E\left[f^i(\epsilon_t^i) f^j(\epsilon_t^j) \frac{\partial^2 v_{t+1}}{\partial x_i \partial x_j}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}) - \xi, \Lambda(\boldsymbol{\epsilon}_t)\mathbf{x} + \mathbf{g}(\boldsymbol{\epsilon}_t)) \right] \phi(\xi) d\xi \\ &\quad \leq 0.\end{aligned}$$

If we can establish that the first term on the right side of (2.12) is non-positive as well, then we will be done with the proof. To this end, note that $\frac{\partial w_t}{\partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) = 0$ for $k \in I$. Differentiating with respect to x^i yields

$$\sum_{\ell \in I} \frac{\partial^2 w_t}{\partial y^\ell \partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) \frac{\partial y_t^\ell}{\partial x^i}(\mathbf{s}^{-i}, \mathbf{x}) + \frac{\partial^2 w_t}{\partial x^i \partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) = 0.$$

Summing the preceding over $k \in I$ we get

$$\sum_{k \in I} \frac{\partial^2 w_t}{\partial x^i \partial y^k}(\mathbf{y}_t^*(\mathbf{s}, \mathbf{x}), \mathbf{x}) \frac{\partial y_t^k}{\partial x^i}(\mathbf{s}^{-i}, \mathbf{x})$$

$$= -\sum_{k\in I}\sum_{\ell\in I}\frac{\partial^2 w_t}{\partial y^\ell \partial y^k}(\mathbf{y}_t^*(\mathbf{s},\mathbf{x}),\mathbf{x})\frac{\partial y_t^\ell}{\partial x^i}(\mathbf{s}^{-i},\mathbf{x})\frac{\partial y_t^k}{\partial x^i}(\mathbf{s}^{-i},\mathbf{x})$$

$$= -\mathbf{z}'U\mathbf{z},$$

where $\mathbf{z}$ is the $|I|$-vector with entries $\frac{\partial y_t^\ell}{\partial x^i}(\mathbf{s}^{-i},\mathbf{x})$ for $\ell \in I$, and U is the $|I| \times |I|$ matrix with entries $\frac{\partial^2 w_t}{\partial y^\ell \partial y^k}(\mathbf{y}_t^*(\mathbf{s},\mathbf{x}),\mathbf{x})$ for $\ell, k \in I$. The matrix U is positive semidefinite because $w_t(\mathbf{y},\mathbf{x})$ is convex in $\mathbf{y}$. As a consequence, $-\mathbf{z}'U\mathbf{z} \leq 0$. This establishes that the first term on the right side of (2.12) is non-positive and therefore completes the proof. $\square$

Chapter 3

Optimal Control of an Inventory System with Stochastic and Independent Leadtimes

3.1 Introduction

Inventory systems with stochastic and independent leadtimes are notoriously difficult to analyze. This difficulty arises in part because of the possibility of order *crossovers* (that is, units ordered are not necessarily delivered in the same sequence in which they have been placed). In settings where the leadtimes of units ordered at the same time are also independent (the case we consider in this paper), this difficulty is compounded because the time until the next replenishment can be affected by the size of the order (larger orders can speed up delivery times). To our knowledge, there are no known results that characterize the structure of the optimal policy for systems with independent leadtimes, including for the system we consider in this paper. For example, Zipkin (2000, p. 410) notes:

"There is no general optimality theory for such systems, to our knowledge. Such a theory would require a detailed scenario describing when we observe the leadtimes, or more generally how we obtain information about them... The optimal policy is probably different in each case, but complicated in everyone."

Most of the existing literature that treats order crossovers or independent leadtimes considers specific order-up to policies, such as base-stock policies with fixed base-stock levels (see for example Kulkarni and Yan (2012), Robinson and Bradley (2008), Bradley and Robinson (2005), Robinson et al. (2001), He et al. (1998) and the references therein; see also Muthuraman et al. (2014) for a review of literature on inventory systems with stochastic leadtimes). As we show, in this paper such policies are sub-optimal and can perform poorly.

In this paper, we consider the specific setting of a continuous review inventory system where demand arises according to a Poisson process. Inventory can be stocked ahead of demand but incurs a holding cost. A replenishment order for one or more units can be placed at any time. The leadtime for each unit ordered is a random variable. The random variables that describe these leadtimes are i.i.d. and exponentially distributed. We allow for a constraint on the total number of units that can be on order at any time, so that there are at most m units ordered at any time but not yet received. However, m can be arbitrarily large. Thus, the setting we consider is one where there is independence among the leadtimes of units ordered at the same time as well as independence among the leadtimes of units that have been ordered at different times. Such a setting arises naturally when the supply process consists of m production facilities (or operators), with each facility capable of producing one unit at a time with exponentially distributed production times (ordering decisions in this case correspond to determining how many of the m facilities should be put into production)[1]. Such a system generalizes the widely studied integrated production-inventory system with a single facility (see for example, Ha (1997)), De Vericourt et al. (2002), Benjaafar and ElHafsi (2006), Benjaafar et al. (2011), and the references therein). Most of the existing literature on integrated production-inventory systems deals with settings where the production process is a single facility. To our knowledge, there are no known results for the optimal control policy for systems with multiple facilities. Therefore, this paper makes a contribution to that literature as well.

We characterize the structure of the optimal control policy and show that it can be specified by a threshold function $r(x)$ where x is the net inventory level. Specifically, we show that it is optimal to order if the number of units on order is less than $r(x)$ and not to order otherwise. More significantly, we show that $r(x)$ is non-increasing in x and that, once $r(x)$ starts to decrease, it continues to do so at a rate that is greater than or equal to one. This implies that the threshold function can be fully described by at most m parameters: a parameter s, corresponding to the largest inventory level for which the number of units on order is at its maximum value m, and at most $m-1$ additional parameters $k_0 = m \geq k_1 \geq \cdots \geq k_{m-1}$, corresponding to the optimal number of units on order at inventory levels $s+i$, for $i = 1, \cdots, m-1$.

[1] It may also arise in settings where each unit corresponds to a batch (e.g., a truckload) that is handled and shipped individually and that experiences independent random delays along the way; see Zipkin (2000) for further discussion.

A base-stock policy with a fixed base-stock level can of course be expressed in terms of a similar threshold function. However, in the case of a base-stock policy, the function $r(x)$ is linear and given by $r(x) = s - x$, where s is the base-stock level. In our case, the function $r(x)$ is non-linear (and in all cases tested observed to be concave). To our knowledge, such a feature has not been documented previously in the literature. This feature appears to be a consequence of the fact that, with i.i.d. leadtimes, the time until the next replenishment arrives is decreasing in the number of units on order (the more units on order, the shorter the time until the next shipment). This is not the case in systems where leadtimes are sequential and independent or in systems where leadtimes are sequential and increasing in the number of units on order, or production-inventory systems with a single facility.

To characterize the structure of the optimal policy, we employ two forms of the optimality equation. We use each to show that the optimal cost function satisfies certain properties, which together imply properties for the optimal policy. We rely on an application of the Banach fixed point theorem to prove some of these properties, which are difficult to prove using standard induction arguments. The application of the Banach fixed point theorem is novel and potentially useful to other optimal control problems.

Uncovering the structure of the optimal policy allows us to develop an efficient algorithm for computing the optimal value of the policy parameters and the corresponding optimal cost. Also, inspired by the structure of the optimal policy, we investigate two plausible simple heuristics, each specified by a single parameter, and examine their performance for a wide range of parameter values. We find that although the two heuristics can be effective for certain ranges of parameters, they can perform poorly when either the ratio of inventory holding to shortage costs is high or demand rate is low. Finally, we extend our analysis to systems with lost sales and systems where order cancellations are possible. In both cases, we characterize the structure of the optimal policy. For systems with order cancellation, we show that the optimal policy reduces to a *bang-bang* policy with $k_0 = m$ and $k_i = 0$ for $i \neq 0$. We show that order cancellation is particularly beneficial when either the ratio of inventory holding to shortage (backorder or lost sales) costs is high or demand rate is low.

The rest of this chapter is organized as follows. In Section 3.2, we describe the problem and formulation. In Section 3.3, we characterize the structure of the optimal policy. In Section 3.4, we describe an efficient algorithm for computing the parameters of the optimal policy. In Section 3.5, we consider heuristic policies and evaluate their performance. In section 3.6, we extend the analysis to related problems.

3.2 Problem Formulation

We consider a single item inventory system where demand arises continuously over time according to a Poisson process. Inventory can be stocked ahead of demand but incurs a holding cost h per unit per unit of time. Demand that cannot be immediately fulfilled from inventory is backordered and incurs a backorder cost b per unit per unit of time. The system is continuously reviewed and a replenishment order for one or more units can be placed at any time. A cost c is incurred when each unit is received. Procurement leadtimes for units ordered are i.i.d. and exponentially-distributed with mean $1/\mu$ with the realization of the leadtime occurring at the time a unit is received. We allow for a constraint on the total number of units that can be on order, so that there are at most m units, at any time, that have been ordered but not yet received (m can be arbitrarily large). The constraint on the number of units on order arises naturally in some settings. As mentioned in Section 3.1, this includes integrated production-inventory systems where m corresponds to the number of available production facilities and the number of units on order corresponds to the number of facilities that are currently producing.

The state of the system at time t can be described by the pair $(X(t), Y(t))$, where $X(t) \in \mathbb{Z}$, the set of integers, and $Y(t) \in \{0, 1, \ldots, m\}$, with $X(t)$ denoting net inventory at time t and $Y(t)$ the number of units on order at time t (in an integrated production-inventory system, this corresponds to the number of facilities currently producing an item). Net inventory can be either positive or negative, with $X(t)^+ = \max\{0, X(t)\}$ corresponding to on-hand inventory and $X(t)^- = \max\{0, -X(t)\}$ corresponding to backorder level. Because both leadtimes and times between consecutive demand orders are exponentially distributed, the system is memoryless and decision epochs can be restricted to only the times when the state changes (state changes are triggered by the delivery of a unit that was on order or the arrival of new demand). The memoryless property allows us to formulate the problem as a Markov Decision Process (MDP) and to restrict our attention to the class of Markovian policies for which actions taken at a particular decision epoch depend only on the current state of the system.

In each state, the system manager must decide whether or not to place additional orders and if so how many. This is equivalent to deciding on the number

k $(0 \leq k \leq m-y)$ of additional orders that must be placed where y is the number of current orders that have been placed but have not been received yet. Let $\mathbb{S} = \mathbb{Z} \times \{0, 1, \ldots, m\}$ denote the state space and let $\mathbf{A}(x, y) = \{0, 1, \ldots, m-y\}$ denote the action set. An action $\mathbf{a}^\pi(x, y)$, in the set $\mathbf{A}(x, y)$, corresponds to the number of units to order. Thus, a policy π specifies, for each state (x, y), an action $\mathbf{a}^\pi(x, y)$ in the set $\mathbf{A}(x, y)$. For instance, action $\mathbf{a}^\pi(x, y) = k$ corresponds to ordering k units in addition to the y units currently on order. In the case of an integrated production-inventory system, this corresponds to deciding on how many additional production facilities to activate, in addition to those that are currently producing.

The expected discounted cost (the sum of inventory holding, backorder, and ordering costs) over an infinite planning horizon, $v^\pi(x, y)$, obtained under a policy π and a starting state (x, y), can be written as:

$$v^\pi(x, y) = E^\pi_{(x,y)} \left\{ \int_0^\infty e^{-\alpha t} \left(hX^+(t) + bX^-(t) \right) dt + \int_0^\infty e^{-\alpha t} c dN(t) \right\}, \qquad (3.1)$$

where $N(t)$ is the cumulative number of units that have been received up to time t and $\alpha > 0$ is the discount rate. Our objective is to choose a policy π^* that minimizes the expected discounted cost.

Let $0 = t_0 \leq t_1 \leq t_2 \leq \ldots$ be the transition times when the system changes from one state to a different state. Then $t_{j+1} - t_j$ is exponentially distributed with rate $\beta_{x,y}(\mathbf{a}(x, y)) = \lambda + \mu(y + \mathbf{a}(x, y))$ if the state of the system at time t_j is (x, y) and action $\mathbf{a}(x, y)$ is selected in this state. The state of the system remains the same between transitions. Therefore, $\{Z_j = (X(t_j), Y(t_j)) : j \geq 0\}$ is a Markov chain, with transition probabilities given by:

$$p_{(x_1,y_1),(x_2,y_2)}(\mathbf{a}_1) = \begin{cases} \dfrac{\lambda}{\beta_{x,y}(\mathbf{a}_1)} & \text{if}(x_2, y_2) = (x_1 - 1, y_1 + \mathbf{a}_1), \\ \dfrac{\mu(y_1 + \mathbf{a}_1)}{\beta_{x,y}(\mathbf{a}_1)} & \text{if}(x_2, y_2) = (x_1 + 1, (y_1 + \mathbf{a}_1 - 1)^+), \\ 0 & \text{otherwise}, \end{cases}$$

where $\mathbf{a}_1 = \mathbf{a}(x_1, y_1)$. This allows us to transform the continuous time decision process into a discrete time one. Let $N(t_j)$ denote the cumulative number of units that have been received by the jth transition. Then, $v^\pi(x, y)$ in (3.1) can be rewritten in the equivalent form:

$$\begin{aligned} v^\pi(x, y) =& E^\pi_{(\mathbf{x},\mathbf{y})} \left[\sum_{i=0}^{\infty} \left(\prod_{k=0}^{i} \frac{\beta_{x_k,y_k}(\mathbf{a}(x_k, y_k))}{\alpha + \beta_{x_k,y_k}(\mathbf{a}(x_k, y_k))} \right) \frac{hx_i^+ + bx_i^-}{\alpha + \beta_{x_k,y_k}(\mathbf{a}(x_k, y_k))} \right. \\ & \left. + c \sum_{i=0}^{\infty} \left(\prod_{k=0}^{i} \frac{\beta_{x_k,y_k}(\mathbf{a}(x_k, y_k))}{\alpha + \beta_{x_k,y_k}(\mathbf{a}(x_k, y_k))} \right) (N(t_i) - N(t_{i-1})) \right], \end{aligned}$$

where $(x_i, y_i) = (x(t_i), y(t_i))$, $\mathbf{x} = (x_1, x_2, \ldots)$ and $\mathbf{y} = (y_1, y_2, \ldots)$. The optimal cost function, v^*, can be shown to satisfy tho following optimality equation:

$$\begin{aligned} v^*(x, y) = \min_{\mathbf{a}(x,y)} & \left\{ \frac{hx^+ + bx^-}{\alpha + \beta_{x,y}(\mathbf{a}(x, y))} + \frac{\mu(y + \mathbf{a}(x, y))}{\alpha + \beta_{x,y}(\mathbf{a}(x, y))} c \right. \\ & \left. + \frac{\beta_{x,y}(\mathbf{a}(x, y))}{\alpha + \beta_{x,y}(\mathbf{a}(x, y))} \sum_{(x',y')} p_{(x,y),(x',y')}(\mathbf{a}(x, y)) v^*(x', y') \right\} \end{aligned}$$

or equivalently

$$v^*(x, y) = \min_{y \le u \le m} \left\{ \frac{g(x) + \lambda v^*(x - 1, u) + u\mu \left(v^*(x + 1, (u - 1)^+) + c \right)}{\alpha + \lambda + u\mu} \right\}, \quad (3.2)$$

where $g(x) = hx^+ + bx^-$. Letting

$$w^*(x, u) = \frac{g(x) + \lambda v^*(x - 1, u) + u\mu \left(v^*(x + 1, (u - 1)^+) + c \right)}{\alpha + \lambda + u\mu},$$

optimality equation (3.2) can be rewritten as

$$v^*(x,y) = \min_{y\le u\le m} \{w^*(x,u)\} = \min_{0\le k\le m-y} \{w^*(x,y+k)\}.$$

The control variable u specifies the total number of units on order at each state transition given the y units already on order; while the control variable k specifies the number of units to be ordered in addition to the y units already on order. Let $u^*(x,y) = \max \operatorname*{argmin}_{y\le u\le m}\{w^*(x,u)\}$ and $k^*(x,y) = u^*(x,y) - y$. Then $u^*(x,y)$ denotes the optimal number of units on order in state (x,y) and $k^*(x,y)$ denotes the optimal number of units to order, in addition to y, in state (x,y).

We will find it useful to work with a *uniformized* version of the problem (see Lippman 1975), in which the transition rate in each state under any action is $\beta = m\mu + \lambda$ so that the transition times $0 = \hat{t}_0 \le \hat{t}_1 \le \hat{t}_2 \le \dots$ are such that the times between transitions $\{\hat{t}_{j+1} - \hat{t}_j : j \ge 0\}$ form a sequence of *i.i.d.* exponentially distributed random variables, each with mean $1/\beta$. This leads to a Markov chain defined by $\{\hat{Z}_j = (X(\hat{t}_j), Y(\hat{t}_j)) : j \ge 0\}$ with transition probabilities given by:

$$\hat{p}_{(x_1,y_1),(x_2,y_2)}(\mathbf{a}_1) = \begin{cases} \dfrac{\lambda}{\beta} & \text{if } (x_2,y_2) = (x_1-1, y_1+\mathbf{a}_1), \\ \dfrac{\mu(y_1+\mathbf{a}_1)}{\beta} & \text{if } (x_2,y_2) = (x_1+1, (y_1+\mathbf{a}_1-1)^+), \\ \dfrac{\beta-\lambda-\mu(y_1+\mathbf{a}_1)}{\beta} & \text{if } (x_2,y_2) = (x_1,y_1), \\ 0 & \text{otherwise.} \end{cases}$$

where $\mathbf{a}_1 = \mathbf{a}(x_1,y_1)$. Let $\hat{N}(t_j)$ denote the cumulative number of units that have been received by the jth transition. Then, $v^\pi(x,y)$ in (3.1) can be rewritten as:

$$v^\pi(x,y) = E^\pi_{(x,y)}\left[\frac{hx^+ + bx^-}{\alpha+\beta}\sum_{i=0}^{\infty}\left(\frac{\beta}{\alpha+\beta}\right)^i + c\sum_{i=1}^{\infty}\left(\frac{\beta}{\alpha+\beta}\right)^i\left(\hat{N}(t_i) - \hat{N}(t_{i-1})\right)\right].$$

Hence, the optimal cost function, v^*, satisfies the optimality equation,

$$v^*(x,y) = \min_{\mathbf{a}(x,y)} \left\{ \frac{hx^+ + bx^-}{\alpha+\beta} + \frac{\mu(y+\mathbf{a}(x,y))}{\alpha+\beta} c \right.$$
$$\left. + \frac{\beta}{\alpha+\beta} \sum_{(x',y')} \hat{p}_{(x,y),(x',y')}(\mathbf{a}(x,y)) v^*(x',y') \right\}.$$

Without loss of generality, we rescale time by letting $\alpha + \beta = 1$. Using the corresponding transition probabilities and time scaling, we can rewrite the optimality equation as follows:

$$v^*(x,y) = \min_{0 \le k \le m-y} \{ g(x) + \lambda v^*(x-1, y+k) + (m-y-k)\mu v^*(x,y)$$
$$+ (y+k)\mu \left(v^*(x+1, (y+k-1)^+) + c \right) \}. \tag{3.3}$$

In Lemma 8 below, we show that another form of the optimality equation is given by a modified version of (3) in which the arguments of v^* on the right hand side of the equality all consist of $y + k$. Such a form will prove to be useful in Section 3 in characterizing the structure of the optimal policy, as it significantly simplifies the analysis. More specifically, we show in Lemma 1 that the optimal cost function satisfies the following optimality equation:

$$v^*(x,y) = \min_{0 \le k \le m-y} \{ g(x) + \lambda v^*(x-1, y+k) + (m-y-k)\mu v^*(x, y+k)$$
$$+ (y+k)\mu \left(v^*(x+1, (y+k-1)^+) + c \right) \},$$

which can also be written as

$$v^*(x,y) = \min_{y \le u \le m} \{ g(x) + \lambda v^*(x-1, u) + (m-u)\mu v^*(x,u)$$
$$+ u\mu \left(v^*(x+1, (u-1)^+) + c \right) \}. \tag{3.4}$$

Lemma 8. *The optimal cost function v^* satisfies optimality equation (3.4).*

Proof. Using optimality equation (3.2), we have

$$\begin{aligned} v^*(x,y) &= \min_{y\le u\le m}\left\{\frac{g(x)+\lambda v^*(x-1,u)+u\mu(v^*(x+1,(u-1)^+)+c)}{\alpha+\lambda+u\mu}\right\} \\ &\le \min_{y+k\le u\le m}\left\{\frac{g(x)+\lambda v^*(x-1,u)+u\mu\left(v^*(x+1,(u-1)^+)+c\right)}{\alpha+\lambda+u\mu}\right\} \\ &= v^*(x,y+k). \end{aligned}$$

Noting that

$$\begin{aligned} u^*(x,y) &= \operatorname*{argmin}_{y\le u\le m}\left\{\frac{g(x)+\lambda v^*(x-1,u)+u\mu\left(v^*(x+1,(u-1)^+)+c\right)}{\alpha+\lambda+u\mu}\right\} \\ &= \operatorname*{argmin}_{u^*(x,y)\le u\le m}\left\{\frac{g(x)+\lambda v^*(x-1,u)+u\mu\left(v^*(x+1,(u-1)^+)+c\right)}{\alpha+\lambda+u\mu}\right\} \end{aligned}$$

leads to $v^*(x,y)=v^*(x,u^*(x,y))=v^*(x,y+k^*(x,y))$. Let $f(x,y,k)=g(x)+\lambda v^*(x-1,y+k)+(y+k)\mu\left(v^*(x+1,(y+k-1)^+)+c\right)+(m-y-k)\mu v^*(x,y)$, and $\hat{f}(x,y,k)=g(x)+\lambda v^*(x-1,y+k)+(y+k)\mu\left(v^*(x+1,(y+k-1)^+)+c\right)+(m-y-k)\mu v^*(x,y+k)$. Then, for any $k\in[0,m-y]$, we have $\hat{f}(x,y,k)\ge f(x,y,k)\ge f(x,y,k^*(x,y))=\hat{f}(x,y,k^*(x,y))$. Consequently, $v^*(x,y)=\min_{0\le k\le m-y}\{f(x,y,k)\}=f(x,y,k^*(x,y))=\hat{f}(x,y,k^*(x,y))=\min_{0\le k\le m-y}\{\hat{f}(x,y,k)\}$, which completes the proof of Lemma 8. □

Define $\hat{w}^*(x,u)=g(x)+\lambda v^*(x-1,u)+u\mu\left(v^*(x+1,(u-1)^+)+c\right)+(m-u)\mu v^*(x,u)$. Then optimality equation (3.4) can be more compactly expressed as

$$v^*(x,y)=\min_{y\le u\le m}\{\hat{w}^*(x,u)\}.$$

Since the optimal cost function, v^*, satisfies both optimality equations (3.2) and (3.4), it follows that

$$u^*(x,y) = \max \operatorname*{argmin}_{y \le u \le m} \{w^*(x,u)\} = \max \operatorname*{argmin}_{y \le u \le m} \{\hat{w}^*(x,u)\}.$$

Finally, for any function v defined on $\mathbb{S}$, we define operators T and $\hat{T}$ as follows:

$$Tv(x,y) = \min_{y \le u \le m} \left\{ \frac{g(x) + \lambda v(x-1,u) + u\mu \left(v(x+1,(u-1)^+) + c \right)}{\alpha + \lambda + u\mu} \right\},$$

and

$$\begin{aligned} \hat{T}(x,y) = \min_{y \le u \le m} & \{ g(x) + \lambda v(x-1,u) + (m-u)\mu v(x,u) \\ & + u\mu \left(v(x+1,(u-1)^+) + c \right) \}. \end{aligned}$$

Then, by virtue of (3.2) and (3.4), we have $v^* = Tv^*$ and $v^* = \hat{T}v^*$.

3.3 The Structure of the Optimal Policy

In this section, we characterize the structure of the optimal policy. To do so, we identify a set of properties specified in Definition 1 below and show that the optimal cost function satisfies these properties. Then, we show that these properties imply specific rules for the optimal action in each state.

In order to simplify the notation, we introduce the difference operators

$$\Delta_x v(x,y) = v(x+1,y) - v(x,y),$$

$$\Delta_y v(x,y) = v(x,y+1) - v(x,y),$$

for real valued functions v defined on $\mathbb{S}$ and combinations of such operators, including

$$\Delta_{x,y} v(x,y) = \Delta_x \Delta_y v(x,y) = \Delta_x v(x,y+1) - \Delta_x v(x,y),$$

$$\Delta_{y,y}v(x,y) = \Delta_y \Delta_y v(x,y) = \Delta_y v(x,y+1) - \Delta_y v(x,y).$$

Definition 1. *Let $\mathcal{V}$ be the set of real valued functions defined on $\mathbb{S}$, such that if $v \in \mathcal{V}$, we have:*

P1: $\Delta_{y,y}v(x,y) \geq 0$ *for* $y < m-1$,

P2: $\Delta_{x,y}v(x,y) \geq 0$ *for* $y < m$, *and*

P3: $\Delta_y v(x+1,y) - \Delta_y v(x,y+1) \geq 0$ *for* $y < m-1$.

Property P1 implies that $v(x,y)$ is convex in y. Property P2 implies that $v(x,y)$ is supermodular in (x,y). Property P3 implies that the cost difference $v(x+1,y) - v(x,y+1)$ is non-decreasing in y.

In what follows, we show that the optimal cost function v^* is in the set $\mathcal{V}$, which in turn implies a specific structure for the optimal policy, as described in Theorem 6. Unfortunately, the standard approach of showing, via induction, that a single optimality operator preserves these properties is difficult to use here (as neither T nor $\hat{T}$ can be shown to preserve the properties in Definition 1). Therefore, we resort to a different approach that employs both operators T and $\hat{T}$, showing that each preserves certain properties. We use these results to construct an operator that we show to be a contraction mapping with zero as a fixed point. This then allows us to apply the Banach fixed point theorem to show that v^* is in the set $\mathcal{V}$.

In preparation for Lemma 9, we introduce the following definitions.

Definition 2. *Let $\mathcal{V}_1$ be the set of real valued functions defined on $\mathbb{S}$, such that if $v \in \mathcal{V}_1$, then v satisfies property P2 and property P4 defined below:*

P4: $\Delta_x g(x) + \lambda \Delta_x v(x-1,y_1) - (\alpha+\lambda)\Delta_x v(x+1,y_2) \leq 0,.$ *for* $y_2 \leq y_1 \leq m$.

Definition 3. *Let $\mathcal{V}_2$ be the set of real valued functions defined on $\mathbb{S}$, such that if $v \in \mathcal{V}_2$, then v satisfies:*

P5: $0 \leq \Delta_y v(x,y) \leq \frac{(h+\alpha c)\mu}{\alpha^2}$ *for* $y < m$, *and*

P6: $v(x+1,y) - v(x,y+1) \leq \frac{h}{\alpha}$ *for* $y < m$.

Lemma 9. *The optimal cost function* v^* *satisfies properties P1-P3. That is,* $v^* \in \mathcal{V}$.

Proof. We divide the proof into four parts.

(i) We show that if $v \in V_1$, then $Tv \in \mathcal{V}_1$. This implies that $v^* \in \mathcal{V}_1$.

(ii) We show that if $v \in \mathcal{V}_2$, then $\hat{T}v \in \mathcal{V}_2$. This implies that $v^* \in \mathcal{V}_2$.

(iii) Using the results of parts (i) and (ii), we construct an operator $\mathcal{O}$ that is a contraction mapping with a fixed point at zero, such that $\Phi(x,y,k) \geq \mathcal{O}\Phi(x,y,k)$, where $\Phi(x,y,k) = \Delta_y v^*(x,y+k) - \Delta_y v^*(x,y)$. Then, we apply the Banach fixed point theorem to show that $\Phi(x,y,k) \geq 0$ for all x, y and k. This implies that $\Delta_{y,y} v^*(x,y) = \Phi\left(x,y,1\right) \geq 0$.

(iv) Using a similar approach to the one used in part (iii), we show that $\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1) \geq 0$.

We begin with part (i). For any $v \in \mathcal{V}_1$, let

$$w(x,y) = \frac{g(x) + \lambda v(x-1,y) + y\mu\left(v(x+1,(y-1)^+) + c\right)}{\alpha + \lambda + y\mu},$$

and $u(x,y) = \max \operatorname*{argmin}_{y \leq u' \leq m} w(x,u')$. It is not difficult to verify that

$$\begin{aligned}\Delta_{x,y} w(x,y) = &- \mu \frac{\Delta_x g(x) + \lambda \Delta_x v(x-1,y) - (\alpha+\lambda)\Delta_x v(x+1,(y-1)^+)}{(\alpha+\lambda+y\mu)(\alpha+\lambda+(y+1)\mu)} \\ &+ \frac{(y+1)\mu\Delta_{x,y} v(x+1,(y-1)^+)}{\alpha+\lambda+(y+1)\mu} + \frac{\lambda\Delta_{x,y} v(x-1,y)}{\alpha+\lambda+(y+1)\mu} \geq 0.\end{aligned}$$

The first term is nonnegative due to the fact that v satisfies property P4. The last two terms are nonnegative due to the fact that v satisfies property P2.

Noting that $Tv(x,y) = \min_{y \le u' \le m} w(x,u') = \min\{w(x,y), \min_{y+1 \le u' \le m} w(x,u')\} = \min\{w(x,y), Tv(x,y+1)\}$, we have

$$\begin{aligned}
\Delta_y Tv(x+1,y) &= Tv(x+1,y+1) - \min\left\{w(x,y), Tv(x,y+1)\right\} \\
&= \max\left\{Tv(x+1,y+1) - w(x+1,y), 0\right\} \\
&= \max\left\{\min_{y+1 \le u' \le m}\left\{w(x+1,u')\right\} - w(x+1,y), 0\right\} \\
&= \max\left\{\min_{y+1 \le u' \le m}\left\{\sum_{j=y}^{u'-1} \Delta_y w(x+1,j)\right\}, 0\right\} \\
&\ge \max\left\{\min_{y+1 \le u' \le m}\left\{\sum_{j=y}^{u'-1} \Delta_y w(x,j)\right\}, 0\right\} \\
&= \Delta_y Tv(x,y).
\end{aligned}$$

Therefore, Tv satisfies property P2 for $v \in \mathcal{V}_1$.

Next, we show that Tv satisfies property P4. First, we show that $u(x,y) \ge u(x+1,y)$. To this end, note that $w(x,y) - w(x,u(x,y)) > 0$ for all $y > u(x,y)$. Since $\Delta_{x,y} w(x,y) \ge 0$, we have $\Delta_x w(x,y) - \Delta_x w(x,u(x,y)) \ge 0$ for all $y > u(x,y)$. Equivalently,

$$w(x+1,y) - w(x+1,u(x,y)) \ge w(x,y) - w(x,u(x,y)) > 0,$$

for all $y > u(x,y)$, which implies $u(x,y) \ge u(x+1,y)$. Next, since

$$\begin{aligned}
\Delta_x Tv(x,y) &= \min_{y \le u' \le m} w(x+1,u') - \min_{y \le u' \le m} w(x,u') \\
&\ge w(x+1,u(x+1,y)) - w(x,u(x+1,y)) \\
&= \Delta_x w(x,u(x+1,y)),
\end{aligned}$$

and $\Delta_x Tv(x,y) \le w(x+1,u(x,y)) - w(x,u(x,y)) = \Delta_x w(x,u(x,y))$, it follows that

$$
\begin{aligned}
&\Delta_x g(x) + \lambda\Delta_x Tv(x-1, y_1) - (\alpha+\lambda)\Delta_x Tv(x+1, y_2) \\
&\qquad \le \Delta_x g(x) + \lambda\Delta_x w(x-1, u(x-1, y_1)) - (\alpha+\lambda)\Delta_x w(x+1, u(x+2, y_2)).
\end{aligned}
$$

Let $\hat{y}_1 = u(x-1, y_1)$ and $\hat{y}_2 = u(x+2, y_2)$. For $y_1 \ge y_2$, we have

$$\hat{y}_1 = u(x-1, y_1) \ge u(x+2, y_1) \ge u(x+2, y_2) = \hat{y}_2,$$

where the second inequality is due to the fact that

$$u(x, y+1) = \max \operatorname*{argmin}_{y+1 \le u' \le m} w(x, u') \ge \max \operatorname*{argmin}_{y \le u' \le m} w(x, u') = u(x, y).$$

Therefore, for $y_1 \ge y_2$, we have

$$
\begin{aligned}
&\Delta_x g(x) + \lambda\Delta_x Tv(x-1, y_1) - (\alpha+\lambda)\Delta_x Tv(x+1, y_2) \\
&\qquad \le \Delta_x g(x) + \lambda\Delta_x w(x-1, \hat{y}_1) - (\alpha+\lambda)\Delta_x w(x+1, \hat{y}_2) \\
&\qquad = \Delta_x g(x) + \lambda\frac{\Delta_x g(x-1) + \lambda\Delta_x v(x-2, \hat{y}_1) + \hat{y}_1\mu\Delta_x v(x, (\hat{y}_1-1)^+)}{\alpha+\lambda+\hat{y}_1\mu} \\
&\qquad\quad - (\alpha+\lambda)\frac{\Delta_x g(x+1) + \lambda\Delta_x v(x, \hat{y}_2) + \hat{y}_2\mu\Delta_x v(x+2, (\hat{y}_2-1)^+)}{\alpha+\lambda+\hat{y}_2\mu} \\
&\qquad = -\Delta_{x,x} g(x) + \lambda\frac{\Delta_x g(x-1) + \lambda\Delta_x v(x-2, \hat{y}_1) - (\alpha+\lambda)\Delta_x v(x, \hat{y}_2)}{\alpha+\lambda+\hat{y}_1\mu} \\
&\qquad\quad + \hat{y}_2\mu\frac{\Delta_x g(x+1) + \lambda\Delta_x v(x, (\hat{y}_1-1)^+) - (\alpha+\lambda)\Delta_x v(x+2, (\hat{y}_2-1)^+)}{\alpha+\lambda+\hat{y}_2\mu} \\
&\qquad \le 0.
\end{aligned}
$$

As a result, Tv satisfies property P4 for $v \in \mathcal{V}_1$. Therefore, $Tv \in \mathcal{V}_1$. To show that $v^* \in \mathcal{V}_1$, we use the fact that (1) $v^* = \lim_{n\to\infty} T^n v$ for any $v \in \mathcal{V}_1$ (see Proposition 3.1.5 and 3.1.6, Bertsekas (2007)), and (2) $T^n v \in \mathcal{V}_1$ since $Tv \in \mathcal{V}_1$.

Next, we consider paet (ii). For any $v \in \mathcal{V}_2$, let

$$\hat{w}(x,u) = g(x) + \lambda v(x-1,u) + u\mu\left(v(x+1,(u-1)^+) + c\right) + (m-u)\mu v(x,u),$$

and $\hat{u}(x,y) = \max \operatorname*{argmin}_{y \le u \le m} \hat{w}(x,u)$. Note that

$$\begin{aligned}
\Delta_y \hat{w}(x,u) &= \lambda \Delta_y v(x-1,u) + u\mu \Delta_y v(x+1,(u-1)^+) \\
&\quad + (m-u)\mu \Delta_y v(x,u) + \mu\left(v(x+1,u) - v(x,u+1) + c\right) \\
&\le \lambda \frac{(h+\alpha c)\mu}{\alpha^2} + u\mu \frac{(h+\alpha c)\mu}{\alpha^2} + (m-u)\mu \frac{(h+\alpha c)\mu}{\alpha^2} + \left(\frac{h}{\alpha} + c\right)\mu \\
&= \frac{(h+\alpha c)\mu}{\alpha^2}.
\end{aligned}$$

Hence, if $\hat{u}(x,y) > y$, we have $\Delta_y \hat{T}v(x,y) = \hat{w}(x,\hat{u}(x,y)) - \hat{w}(x,\hat{u}(x,y)) = 0$, and if $\hat{u}(x,y) = y$, we have

$$\Delta_y \hat{T}v(x,y) = \min_{y+1 \le u \le m} \hat{w}(x,u) - \hat{w}(x,y) \le \Delta_y \hat{w}(x,y) \le \frac{(h+\alpha c)\mu}{\alpha^2}.$$

On the other hand, we have

$$\Delta_y \hat{T}v(x,y) = \min_{y+1 \le u \le m} \hat{w}(x,u) - \min_{y \le u \le m} \hat{w}(x,u) \ge 0.$$

Therefore, $\hat{T}v$ satisfies property P5 for $v \in \mathcal{V}_2$.

Noting that

$$\begin{aligned}
\hat{w}(x+1,u) - \hat{w}(x,u+1) &= g(x+1) - g(x) + \lambda(v(x,u) - v(x-1,u+1)) \\
&\quad + (m-u-1)\mu(v(x+1,u) - v(x,u+1)) \\
&\quad + u\mu(v(x+2,(u-1)^+) - v(x+1,u)) \\
&\le h + \lambda\frac{h}{\alpha} + u\mu\frac{h}{\alpha} + (m-u-1)\mu\frac{h}{\alpha} \\
&\le h + \frac{(1-\alpha)}{(\lambda+m\mu)}\frac{h}{\alpha} = \frac{h}{\alpha}
\end{aligned}$$

leads to

$$\begin{aligned}
\hat{T}v(x+1,y) - \hat{T}v(x,y+1) &= \hat{w}(x+1,\hat{u}(x+1,y)) - \hat{w}(x,\hat{u}(x,y+1)) \\
&\le \hat{w}(x+1,\hat{u}(x,y+1)-1) - \hat{w}(x,\hat{u}(x,y+1)) \\
&\le \frac{h}{\alpha}.
\end{aligned}$$

That is, $\hat{T}v$ satisfies property P6 for $v \in \mathcal{V}_2$. Therefore, $\hat{T}v \in \mathcal{V}_2$. To show that $v^* \in \mathcal{V}_2$, we use the fact that (1) $v^* = \lim_{n\to\infty}\hat{T}^n v$ for any $v \in \mathcal{V}_2$ (see Proposition 3.1.5 and 3.1.6, Bertsekas (2007)), and (2) $\hat{T}^n v \in \mathcal{V}_2$ since $\hat{T}v \in \mathcal{V}_2$.

We now consider part (iii). Note teat if $u^*(x,y) > y$, then $\Delta_y v^*(x,y) = 0$ and if $u^*(x,y) = y$, then $\Delta_y v^*(x,y) \le \Delta_y \hat{w}^*(x,y)$. Therefore, if $u^*(x,y) > y$, we have

$$\Delta_y v^*(x,y+j) - \Delta_y v^*(x,y) = \Delta_y v^*(x,y+j) \ge 0.$$

If $u^*(x,y) = y$, then

$$\begin{aligned}
\Delta_y v^*(x,y+j) - \Delta_y v^*(x,y) &= v^*(x,y+j+1) - v^*(x,y+j) - \Delta_y v^*(x,y) \\
&\ge \hat{w}^*(x,u^*(x,y+j+1)) - \Delta_y \hat{w}^*(x,y) \\
&\quad - \hat{w}^*(x,u^*(x,y+j+1)-1) \\
&= \Delta_y \hat{w}^*(x,\hat{u}(x,y+j+1)-1) - \Delta_y \hat{w}^*(x,y).
\end{aligned}$$

Now let

$$\Phi(x,y,j) = \begin{cases} \Delta_y v^*(x,y+j) - \Delta_y v^*(x,y) & \text{if } y+j < m\ , \\ 0 & \text{otherwise.} \end{cases}$$

Then, it can be verified that

$$\Delta_y \hat{w}^*(x,y+j) - \Delta_y \hat{w}^*(x,y)$$

$$
\begin{aligned}
&= \lambda\left(\Delta_y v^*(x-1,y+j) - \Delta_y v^*(x-1,y)\right)\\
&\quad + (m-y-j-1)\mu\left(\Delta_y v^*(x,y+j) - \Delta_y v^*(x,y)\right)\\
&\quad + y\mu\left(\Delta_y v^*(x+1,y+j-1) - \Delta_y v^*(x+1,(y-1)^+)\right)\\
&\quad + j\mu\left(\Delta_y v^*(x+1,y+j-1) - \Delta_y v^*(x+1,y)\right)\\
&\quad + \mu\left(\Delta_x v^*(x,y+j) - \Delta_x v^*(x,y)\right) + j\mu\Delta_{x,y} v^*(x,y)\\
&\geq \lambda\Phi(x-1,y,j) + (m-y-j-1)\mu\Phi(x,y,j)\\
&\quad + y\mu\Phi(x+1,(y-1)^+,j) + j\mu\Phi(x+1,y,(j-1)^+),
\end{aligned}
$$

where the inequality is due to the fact that $\Delta_{x,y} v^*(x,y) \geq 0$.

Let $L(x,y,j) = \hat{u}(x,y+j+1) - y - 1 \geq 0$. Then, if $\hat{u}(x,y) = y$, we have

$$
\begin{aligned}
\Phi(x,y,j) \geq &\lambda\Phi(x-1,y,L(x,y,j)) + (m-y-L(x,y,j)-1)\mu\Phi(x,y,L(x,y,j))\\
&+ y\mu\Phi(x+1,(y-1)^+,L(x,y,j))\\
&+ L(x,y,j)\mu\Phi(x+1,y,(L(x,y,j)-1)^+).
\end{aligned}
$$

Let $\hat{\mathbb{S}} = \{v : \mathbb{Z} \times \{0,1,\ldots,m\} \times \{0,1,\ldots,m\} \to \mathbb{R}, v \text{ is bounded}\}$. Then, $\Phi \in \hat{S}$, since

$$
|\Phi(x,y,j)| = |\Delta_y v^*(x,y+j) - \Delta_y v^*(x,y)| \leq \frac{(h+\alpha c)\mu}{\alpha^2}.
$$

For any function $v \in \hat{\mathbb{S}}$, define operator $\mathcal{O}$ as follows,

$$
\mathcal{O}v(x,y,j) = \begin{cases} \hat{\mathcal{O}}v(x,y,j) & \text{if } \hat{u}(x,y) = y \text{ and } y+j < m,\\ 0 & \text{otherwise.} \end{cases}
$$

where $\hat{\mathcal{O}}v(x,y,j) = \lambda v(x-1,y,L(x,y,j)) + (m-y-L(x,y,j)-1)\mu v(x,y,L(x,y,j)) + y\mu v(x+1,(y-1)^+,L(x,y,j)) + L(x,y,j)\mu v(x+1,y,(L(x,y,j)-1)^+)$. Hence, based on the above analysis, we have

$\Phi \geq \mathcal{O}\Phi$. Also, for any $v_1, v_2 \in \hat{\mathbb{S}}$, it is not difficult to verify that $\|\mathcal{O}v_1 - \mathcal{O}v_2\|_1 \leq (\lambda + (m-1)\mu)\|v_1 - v_2\|_1$. Therefore, $\mathcal{O}$ is a contraction mapping on a metric space $\left(\hat{\mathbb{S}}, \|\cdot\|_1\right)$. Clearly, $\left(\hat{\mathbb{S}}, \|\cdot\|_1\right)$ is a complete metric space and $\mathcal{O}\mathbf{0} = \mathbf{0}$. By the Banach fixed point theorem, $\mathcal{O}$ has a unique fixed point $\mathbf{0}$, and for any function $v \in \hat{\mathbb{S}}$, $\lim_{n\to\infty}\mathcal{O}^n v = \mathbf{0}$. Also, since $\Phi \geq \mathcal{O}\Phi \geq \mathcal{O}^2\Phi \geq \cdots \geq \mathcal{O}^n\Phi$, we have $\Phi \geq \lim_{n\to\infty}\mathcal{O}^n\Phi = \mathbf{0}$. This implies that

$$\Delta_{y,y}v^*(x,y) = \Delta_y v^*(x,y+1) - \Delta_y v^*(x,y) = \Phi(x,y,1) \geq 0.$$

Lastly, we consider part (iv). Since $v^* \in \mathcal{V}_1$ and $\Delta_{y,y}v^*(x,y) \geq 0$, we have

$$\begin{aligned}\Delta_{y,y}\hat{w}^*(x,y) &= \lambda\Delta_{y,y}v^*(x-1,y) + (m-y-2)\Delta_{y,y}v^*(x,y) \\ &\quad + y\mu\Delta_{y,y}v^*(x+1,(y-1)^+) + 2\mu\Delta_{x,y}v^*(x,y) \geq 0.\end{aligned}$$

Therefore, $u^*(x,y) = \max\{u^*(x,0), y\}$, and

$$\Delta_y v^*(x,y) = \begin{cases} 0 & \text{if } y < u^*(x,0), \\ \Delta_y\hat{w}^*(x,y) & \text{otherwise.} \end{cases}$$

Noting that $\Delta_y\hat{w}(x,y) < 0$, if $y < u^*(x,0)$ we have $\Delta_y v^*(x,y) \geq \Delta_y\hat{w}^*(x,y)$. Also, since $v^*(x,y)$ satisfies $\Delta_x g(x) + \lambda\Delta_x v^*(x-1,y) - (\alpha+\lambda)\Delta_x v^*(x+1,(y-1)^+) \leq 0$, then, if $y \geq u^*(x,0)$, we have

$$\begin{aligned}\Delta_x v^*(x,y) &= \frac{\Delta_x g(x) + \lambda\Delta_x v^*(x-1,y) + y\mu\Delta_x v^*(x+1,(y-1)^+)}{\alpha+\lambda+y\mu} \\ &\leq \Delta_x v^*(x+1,(y-1)^+).\end{aligned}$$

As a consequence, if $y+1 < u^*(x,0)$, we have

$$\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1) = \Delta_y v^*(x+1,y) \geq 0,$$

and, if $y+1 \geq u^*(x,0)$, we have

$$\begin{aligned}
&\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1) \\
&\qquad \geq \Delta_y \hat{w}^*(x+1,y) - \Delta_y \hat{w}^*(x,y+1) \\
&\qquad = \lambda \left(\Delta_y v^*(x,y) - \Delta_y v^*(x-1,y+1)\right) + \mu \left(\Delta_x v^*(x+1,y) - \Delta_x v^*(x,y+1)\right) \\
&\qquad\quad + (m-y-2)\mu \left(\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1)\right) \\
&\qquad\quad + y\mu \left(\Delta_y v^*(x+2,(y-1)^+) - \Delta_y v^*(x+1,y)\right) \\
&\qquad \geq \lambda \left(\Delta_y v^*(x,y) - \Delta_y v^*(x-1,y+1)\right) \\
&\qquad\quad + (m-y-2)\mu \left(\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1)\right) \\
&\qquad\quad + y\mu \left(\Delta_y v^*(x+2,(y-1)^+) - \Delta_y v^*(x+1,y)\right). \qquad (3.5)
\end{aligned}$$

Using a similar approach to the proof of part (iii) and invoking the Banach fixed point theorem, it is not difficult to verify that $\Delta_y v^*(x+1,y) - \Delta_y v^*(x,y+1) \geq 0$. This completes the proof of Lemma 9. □

We are now ready to state the main result of the paper.

Theorem 6. *The optimal control policy is specified by an inventory level-dependent threshold $r^*(x)$ such that, when the system is in state (x,y), it is optimal to order $r^*(x) - y$ if $y < r^*(x)$ and not to order otherwise. Furthermore, the threshold $r^*(x)$ has the following properties:*

(1) $r^*(x)$ *is non-increasing in* x,

(2) $0 \leq r^*(x) \leq m$.

(3) $r^*(x+1) \leq \max\{0, r^*(x)-1\}$, *for* $x \geq s^*$, *where* $s^* = \max\{x | r^*(x) = m\}$.

Proof. To show that the optimal policy is specified by a threshold $r^*(x)$, it is sufficient to show $\hat{w}^*(x,u)$ is convex in u. To see that this is the case, note that

$$\begin{aligned}\Delta_{y,y}\hat{w}^*(x,u) &= \lambda\Delta_{y,y}v^*(x-1,u) + (m-u)\mu\Delta_{y,y}v^*(x,u)) \\ &\quad + u\mu\Delta_{y,y}v^*(x+1,(u-1)^+ + 2\mu(\Delta_y v^*(x+1,u) - \Delta_y v^*(x,u+1)) \\ &\geq 0,\end{aligned}$$

where the inequality is due to Properties P1 and P3. Let $r^*(x) = \max \operatorname*{argmin}_{0\leq u\leq m}\{\hat{w}^*(x,u)\} = u^*(x,0)$. Then, $u^*(x,y) = \max\{u^*(x,0),y\} = \max\{r^*(x),y\}$. Therefore, the optimal policy is to order $r^*(x)-y$ if $y < r^*(x)$ and not to order otherwise.

Next, we prove that $r^*(x)$ satisfies properties 1-3. First, note that

$$\begin{aligned}\Delta_{x,y}w^*(x,y) = -\mu&\frac{\Delta_x g(x) + \lambda\Delta_x v^*(x-1,y) - (\alpha+\lambda)\Delta_x v^*(x+1,(y-1)^+)}{(\alpha+\lambda+y\mu)(\alpha+\lambda+(y+1)\mu)} \\ &+ \frac{(y+1)\mu\Delta_{x,y}v^*(x+1,(y-1)^+)}{\alpha+\lambda+(y+1)\mu} + \frac{\lambda\Delta_{x,y}v^*(x-1,y)}{\alpha+\lambda+(y+1)\mu} \geq 0,\end{aligned}$$

where the first term on the right-hand side of the equality is nonnegative because v^* satisfies property P4. The last two terms are nonnegative due to the fact that v^* satisfies property P2. This implies that

$$r^*(x+1) = u^*(x+1,0) \leq u^*(x,0) = r^*(x).$$

Hence, $r^*(x)$ is non-increasing in inventory level x. As such, if $r^*(x) = 0$, then $r^*(x+1) = 0$, and if $r^*(x+1) = m$, then $r^*(x) = m$.

The fact that $0 \leq r^*(x) \leq m$ follows immediately from the definition of $r^*(x)$. To show that the last property holds, note that by virtue of Property P3 and inequality (3.5), we have $\Delta_y\hat{w}^*(x+1,y) - \Delta_y\hat{w}^*(x,y+1) \geq 0$, if $y+1 \geq r^*(x)$.

Note also that $r^*(x)$ can be written as $r^*(x) = \min\{y, 0 \le y \le m : \Delta_y \hat{w}^*(x, y) > 0\}$. Therefore, $\Delta_y \hat{w}^*(x+1, r^*(x)-1) \ge \Delta_y \hat{w}^*(x, r^*(x)) > 0$, which implies that $r^*(x) - 1 \ge r^*(x+1)$. This completes the proof of Theorem 6. □

Theorem 6 reveals an important feature of the optimal policy. The optimal policy is not in general a base-stock policy (see for example Figure 3.1). Although the optical policy can be expressed in terms of a threshold function, it is different from a policy that follows a fixed base-stock level. Under a base-stock policy with a fixed base-stock level$s$, the threshold function is linear and given by $r(x) = s - x$ (the presence of a capacity constraint modifies this slightly so that $r(x) = \min\{m, s - x\}$). In our case, the threshold function can be non-linear (see Figure 3.1). Moreover, the inventory position, in contrast to the one under a fixed base stock level, is not constant and in fact is path-dependent and can be non-monotonic as a function of x. so our knowledge, an optimal policy with such a structure has not been documented previously in the literature.

The fact that $r^*(x+1) \le \max\{0, r^*(x) - 1\}$, for $x \ge s^*$, is another important feature. It implies that, once $r^*(x)$ starts to strictly decrease, it will continue to do so with each unit increase in x. This means that if $r^*(x)$ starts decreasing at $x = s^*$, it would reach $r^*(x) = 0$ for $x \le s^* + m$. In other words, the decrease in $r^*(x)$ takes place over at most m steps (see Figure 3.1 for an illustration). As a result, the optimal policy can be fully specified by the vector $(s^*, \mathbf{k}^*)$ where s^* is as defined in Theorem 6 and $\mathbf{k}^* = (k_0^*, k_1^*, \ldots, k_{m-1}^*)$. Here, s^* represents the largest inventory level for which the maximum m orders have beer placed and k_j^* represents the optimal number of units to order at inventory level $s^* + j$, for $j = 0, \ldots, m-1$. Note that since there is no limit on backlogs there always exists

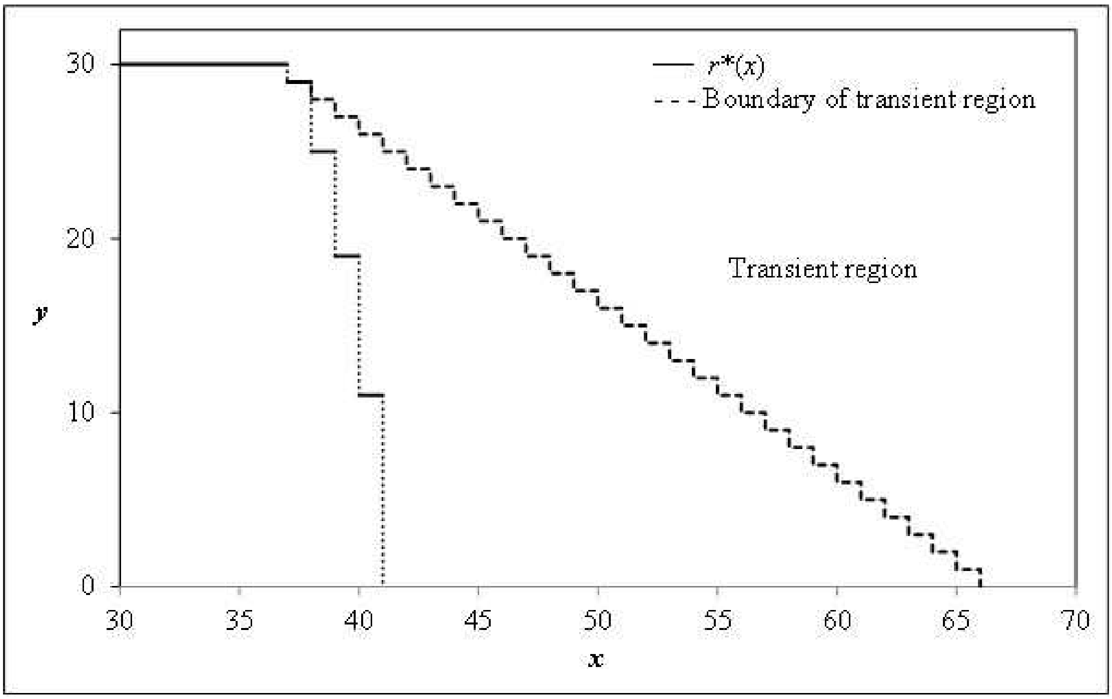

Figure 3.1: Illustration of the optimal policy ($m = 30$, $\mu = 1$, $\lambda = 28.5$, $c = 0$, $h = 2$, and $b = 15$)

an inventory level at which it is optimal to have m units on order ($k_0^* = m$).

Furthermore, if we let $M(x) = m - (x - s^*)$ for $x \in [s^*, s^* + m]$, we note that all states (x, y) such that $x > s^*$ and $y > M(x)$ are transient. This can be verified by noting that once the system reaches inventory level s^*, and since cancellations are not allowed, orders are delivered one at a time, along $M(x)$, until eventually all orders are received at which time the system enters state $(s^* + m, 0)$. At this point, no orders can be placed until inventory is depletes back to level $s^* + l$ for $l = \text{argmin}\{k_l | k_l > 0\}$. Hence, y never exceeds $\min\{m, M(x)\}$. In section 3.4, we show how we can exploit all the above features of the optimal policy to construct a Markov chain model that allows us to efficiently compute the parameters of the optimal policy.

We conclude this section by noting that the above results extend to two important cases: (1) the case where decisions are made based on the average cost criterion and (2) the case where demand that cannot be fulfilled immediately from on-hand inventory is lost and not backlogged (the case of lost sales).

For systems where the objective is to minimize the long run average cost, we can show that, given a police π, the average cost rate is given by:

$$J^{\pi}(x,y) = \lim_{T\to\infty} \sup \frac{1}{T} E^{\pi}_{(x,y)} \left\{ \int_0^T (hX(t) + bY(t))\, dt + \int_0^T c dN(t) \right\}. \tag{3.6}$$

A policy π^* that yields $J^*(x,y) = \inf_{\pi} J^{\pi}(x,y)$ for all states (x,y) is said to be optimal for the average cost criterion. In the following theorem, we show that the optimal policy retains all of the properties observed in Theorem 6 under the expected discounted cost criterion.

Theorem 7. *The optimal policy under the average cost criterion retains all the properties of the optimal policy under the discounted cost criterion, namely that there exists an inventory level-dependent threshold $r^*(x)$ such that it is optimal to order $r^*(x) - y$ if $y < r^*(x)$ when the system is in state (x,y) and not to order otherwise. Furthermore, $r^*(x)$ satisfies properties 1-3 in Theorem 6.*

Proof. The existence of an optimal policy for the average cost, and for this average cost to be finite and independent of the starting state, can be proven via an argument involving taking the limit as $\alpha \to 0$ in the discounted cost problem. However, in order to apply this argument, we must show that the following two conditions hold (see Cavazos-Cadena and Sennott (1992) and

Weber and Stidham Jr (1987)): (1) there exists a stationary policy π that induces an irreducible positive recurrent Markov chain with finite average cost J^π, and (2) the number of states for which one-stage cost $hx^+ + bx^- \leq J^\pi$ is finite.

In order to prove condition 1, we can choose a special case of the optimal policy with parameters s^* and $k^* = (m, 0, \cdots, 0)$ resulting in a maximum reachable inventory level of $s^* + m$ (provided we start the system at an inventory level $x \leq s^* + m$). In this case, it is straightforward to show that this policy yields an irreducible positive recurrent Markov Chain with a finite average cost (provided $\lambda/m\mu < 1$ as shown in Section 3.5 below). It is easy to verify that condition 2 holds since $hx^+ + bx^-$ is convex in x and goes to infinity when $x \to \pm\infty$. Hence, for any positive value γ, the number of states for which $hx^+ + bx^- \leq \gamma$ is always finite. Under these conditions, Weber and Stidham Jr (1987) showed that there exists a constant J^* and a function $f(x, y)$ such that

$$\begin{aligned} f(x,y) + J^* \geq hx^+ + bx^- + \min_{0\leq k\leq m-y} \{&\lambda f(x-1, y+k) + (m-y-k)\mu f(x,y) \\ &+(y+k)\mu \left[f(x+1, [y+k-1]^+) + c\right]\}. \end{aligned} \tag{3.7}$$

Furthermore, the stationary policy that minimizes the fight hand side of the above equation for each state (x, y) it an optimal policy for the average cost criterion and yields a constant average cost J^*. Hence, properties of the average cost optimal policy are the same as and determined through function $f(x, y)$ in much the same way as were properties of the discounted cost optimal policy determined by $v^*(x, y)$. $\square$

Next, we consider the case of systems with lost sales. For systems with lost sales, demand that cannot be immediately fulfilled is considered lost and incurs a lost sale cost L per unit of unfulfilled demand. In this case, the optimal cost function v^* can be shown to satisfy the following optimality equations:

$$v^*(x,y) = \begin{cases} \min\limits_{y \le u \le m} \dfrac{hx + \lambda v^*(x-1,u) + u\mu\left(v^*(x+1,(u-1)^+) + c\right)}{\alpha + \lambda + u\mu}, & \text{if } x > 0 \ , \\ \min\limits_{y \le u \le m} \dfrac{\lambda(v^*(x,u) + L) + u\mu\left(v^*(x+1,(u-1)^+) + c\right)}{\alpha + \lambda + u\mu}, & \text{if } x = 0 \ , \end{cases}$$

and

$$v^*(x,y) = \begin{cases} \min\limits_{0 \le k \le m-y} \{hx + \lambda v^*(x-1,y+k) + (m-y-k)\mu v^*(x,y) & \\ \qquad +(y+k)\mu\left[v^*(x+1,[y+k-1]^+) + c\right]\} & \text{if } x > 0 \ , \\ \min\limits_{0 \le k \le m-y} \{\lambda\left(v^*(x,y+k) + L\right) + (m-y-k)\mu v^*(x,y) & \\ \qquad +(y+k)\mu\left[v^*(x+1,[y+k-1]^+) + c\right])\} & \text{if } x = 0 \ . \end{cases}$$

Similar to the backlog case, it is not difficult no show that v^* satisfies properties P1-P3 of definition 1. Hence, Theorems 6 and 7 apply to the lost sales case as well. Here, we point out though, that depending on the parameters of the system, an inventory level for which the on order level is equal to m may not be optimal. This is, in contrast with the backlog case where there exists an inventery/backlog level for which the maximum number of orders, m, is used. In other words, for the lost sale case, it may be optimal to have $r^*(x) = \bar{m} < m$ and $r^*(x+\bar{m}) = 0$.

In the next section, we use the structural properties of the optimal policy to devise an algorithm which allows us to obtain the parameters of the optimal policy efficiently (the dynamic programming approach suffers from the curse of dimensionality when the state space is large, which would be the case when the demand rate is high or the holding cost rate is low).

3.4 Computing the Parameters of the Optimal Policy

In this section, we show how we can use the properties of the optimal policy to model the dynamics of the system (under the optimal policy) as a Markov Chain. In turn, this allows us to derive expressions for the probability distribution of system states that we use to devise an algorithm to compute the parameters of the optimal policy. We focus on the average cost criterion because it is independent of the initial state and because it is easier to compute once the system state probabilities have been determined.

In section 3.3, we showed that the optimal policy belongs to the class of policies that are fully characterized by a pair $(s, \mathbf{k})$ where s represents the largest inventory level for which we set the number of units on order to its maximum feasible value and $\mathbf{k} = (k_0, k_1, \ldots, k_{m-1})$ is a vector with elements k_i such that k_i specifies the optimal number of units on order when $x = s + i$. The objective of this section is to determine the values of s and $\mathbf{k}$, which we denote by s^* and $\mathbf{k}^*$, that minimize the average cost.

Focusing on the recurrent region of the state space under the optimal policy, we restrict ourselves to states (x, y) that fall into one of the following sub-regions: (1) $x \leq s$ and $y = m$, (2) $s < x \leq s + m - 1$ and $k_{x-s} \leq y \leq m - (x - s)$, and (3) $x \geq s + m$ and $y = 0$. State transitions occur with rates λ and $\eta(x, y)$, where

$$\eta(x,y) = \begin{cases} m\mu & \text{if } x \in (-\infty, s] \text{ and } y = m, \\ y\mu & \text{if } x = s+i,\ k_i \leq y \leq m-i,\ i = 1, \ldots, m-1, \\ 0 & \text{otherwise.} \end{cases}$$

Figure 3.2 depicts the state transition diagram for the special case $\mathbf{k} = (m, \ldots, m-n, k_{n+1}, \ldots, 2, 0, 0)$.

Obtaining closed form expressions for the probability $\pi(x, y)$ of being in state (x, y) is difficult in general. What we propose instead is an algorithm that uses expressions obtained recursively to efficiently compute the probabilities $\pi(x, y)$. First, note that the Markov chain in Figure 3.2 follows a simple birth-death process for spates (x, m), where $x \in (-\infty, s]$ and for states (x, y) where $x \in (s, s+n]$ and $y = k_{x-s}$, with $n = \max\{i \in [1, m-1] | k_{i-1} - k_i = 1\}$. Note also that $\pi(x, y) = 0$ for (1) $x \in (-\infty, s]$ and $y < m$ and (2) $x \in (s+1, s+m)$ and $y < k_{x-s}$.

Define $\pi'(x, y) = \pi(x, y)/\pi(s+n, k_n)$. Then $\pi(x, y) = \pi'(x, y) \times \pi(s+n, k_n)$. Upon normalization (i.e., using the fact that the sum of the all probabilities is equal to 1), we obtain

$$\pi(x, y) = \frac{\pi'(x, y)}{\sum_{i=-\infty}^{s+m} \sum_{j=0}^{m} \pi'(i, j)}. \tag{3.8}$$

Noting that $\pi'(s+n, k_n) = 1$, we have

$$\pi'(s+n-i, k_{n-i}) = \left(\prod_{l=1}^{i} \frac{\lambda}{k_{n-l}\mu} \right) \pi'(s+n, k_n), \tag{3.9}$$

and

$$\pi'(s-i, m) = \rho_m^i \left(\prod_{l=1}^{n} \frac{\lambda}{\mu k_{n-l}} \right) \pi'(s+n, k_n), \tag{3.10}$$

where $\rho_m = \lambda/m\mu$ (for the case of backlogs, we assume $\rho_m < 1$ to ensure system stability). In the case of an integrated production-inventory system, ρ_m corresponds to the utilization of the facilities.

For states (x, y) such that $s+n < x \leq s+m$ and $k_{x-s} \leq y \leq k_n + n - x + s$ (where we define $k_m = 0$) the remaining $\pi'(x, y)$'s can be calculated in sequence

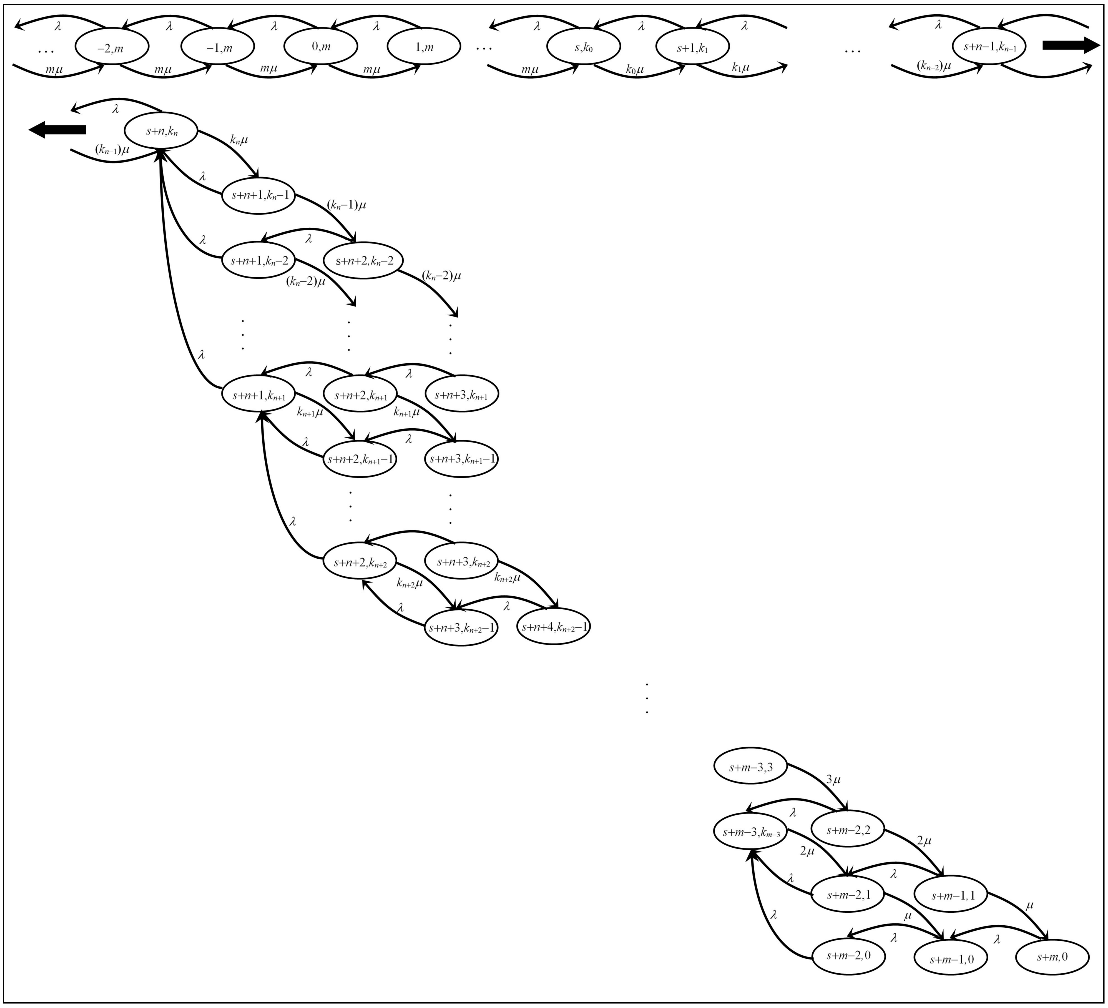

Figure 3.2: The state transition diagram under a policy specified by parameters s and $\mathbf{k}$

as follows:

$$\pi'(s+n+i, k_l-j) = A(k_l-j)\pi'(s+n+i-1, k_l-j+1)+B(k_l-j)\pi'(s+n+i+1, k_l-j), \tag{3.11}$$

for $l = n, \dots, m-1$, $j = 1, \dots, k_l - k_{l+1}$ and $i = j, j-1, \dots, 1$, and

$$\pi'(s+n+a, k_{n+a}) = \begin{cases} \dfrac{k_n\mu}{\lambda}\pi'(s+n, k_n) - \sum\limits_{j=k_{n+1}+1}^{k_n-1} \pi'(s+n+1, j) & \text{if } a = 1, k_{n+a} \neq 0, \\ \dfrac{\pi'(s+n+a-1, k_{n+a-1})}{B(k_{n+a-1})} - \sum\limits_{j=k_{n+a}+1}^{k_{n+a-1}} \pi'(s+n+a, j) & \text{if } a > 1, k_{n+a} \neq 0. \end{cases} \tag{3.12}$$

where $A(t) = (t+1)\mu/(\lambda + t\mu)$ and $B(t) = \lambda/(\lambda + t\mu)$.

Note that for states (x, y) such that $s + n < x \leq s + m$ and $k_{x-s} \leq y \leq m - (x - s)$, the computations of the $\pi'(x, y)$'s are carried out in a specific order. For each row, where a row corresponds to a value of y (see Figure 3.2), the $\pi'(x, y)$'s are computed in decreasing values of x. Once all the values of a row have been computed, computations for row $y-1$ begin and so on until all $\pi'(x, y)$'s have been computed. It is worth mentioning that $\pi'(s+n+a, k_{n+a})$, for $a = 1, \dots, m-n-2$, is computed from the balance equation of state $(s + n + a - 1, k_{n+a-1})$ (given by (3.12)) since at this stage all $\pi'(x, y)$'s involved in the balance equation of state $(s + n + a - 1, k_{n+a-1})$ have been computed except for state $(s + n + a, k_{n+a})$.

Given the steady state probabilities, $\pi(x, y)$, we obtain the marginal distribution of the inventory level as follows

$$p_{m-i} = \begin{cases} \sum\limits_{j=k_i}^{m-i} \pi(s+i, j) & \text{if } i = 0, \dots, m, \\ \pi(s+i, m) & \text{if } i = -1, -2, \dots, -\infty, \end{cases} \tag{3.13}$$

where p_i is the probability that the net inventory level is $s + m - i$. The average total cost $J(s,k)$ can then be written as

$$J(s,\mathbf{k}) = h\sum_{i=0}^{s+m}(s+m-i)p_i + b\sum_{i=s+m+1}^{\infty}(i-s-m)p_i + \lambda c, \tag{3.14}$$

where the three terms represent the expected inventory holding cost, the expected backorder cost, and the expected production cost, respectively. It is not difficult to show that $J(s,\mathbf{k})$ is convex in s. Therefore, the optimal stock level s^* is given by the smallest integer s for which

$$J(s,\mathbf{k}) - J(s+1,\mathbf{k}) \leq 0. \tag{3.15}$$

Here, it is important to note that the distribution specified by $\pi(x,y)$ is independent of the choice of the value s since none of the expressions (3.9)-(3.12) involves s in the computation of the value of the probabilities. This is valuable because it allows for the algorithm to be run only once using a large enough value of s (any $s \geq -m$ will do). The parameter s^* can then be obtained as follows:

$$s^* = \max\left\{ s \geq -m | \sum_{i=0}^{s+m-1} p_i \leq \frac{b}{h+b} \right\} - m.$$

To determine the optimal vector $\mathbf{k}$, we carry out an exhaustive search of all feasible vectors $\mathbf{k}$. This search is significantly expedited by taking advantage of Property 3 of Theorem 6 and the fact that the decrease of $r^*(x)$ to zero takes place over no more than m steps. Table 3.1 shows the number of feasible $\mathbf{k}$ vectors, using full enumeration (a total of m^{m-1} possible $\mathbf{k}$ vectors) and using Property 3 of Theorem 6. As we can see, using Property 3 dramatically reduces the number of feasible $\mathbf{k}$ vectors. In Table 3.2, we compare the computational performance of our search algorithm to the performance of a standard value iteration algorithm (see for example Puterman (2014, Sec. 8.5.1)) for solving the dynamic program in

(3.3). We do so for an illustrative range of values of m and $\lambda/m\mu$. As we can see, our algorithm can significantly outperform the standard value iteration algorithm. This is particularly the case when the recurring region of the state space is large. This is true when $\lambda/m\mu$ is large, leading to high backorder levels (a high $\lambda/m\mu$ corresponds to either high demand or long lead time). In fact, the computational effort for the value iteration algorithms grows exponentially with $\lambda/m\mu$ while it grows more modestly for our algorithm.

Table 3.1: Number of feasible **k** vectors

m	Full enumeration	Enumeration using Property 3 of Theorem 6
5	625	16
10	109	512
15	$2.9193 \times 10^{1}6$	16384
20	$5.2429 \times 10^{2}4$	524288

The above approach can also be adapted to the lost sales case (details are omitted for brevity). Note that in the case of lost sales, the optimal policy requires specifying an additional parameter $\bar{m} \leq m$ which corresponds to the optimal maximum number of units on order. In other words, the optimal policy may never need to place an order of size m. The average total cost, given parameters $\bar{m}, s$, and $\mathbf{k}$, can then be expressed as follows

$$J^L(\bar{m}, s, k) = \lambda L p_0 + h \sum_{i=0}^{s+\bar{m}} i p_i + \lambda c(1 - p_0), \tag{3.16}$$

where the three terms in the above expression correspond, respectively, to the expected lost sales cost, the expected inventory holding cost and the expected production cost.

We conclude this section by noting that we observed numerically, and in all the cases tested, that $r^*(x)$ is concave in x (this is the case in both the backorder and

Table 3.2: Computational performance comparisons

m	λ/μ	CPU time (seconds)	
		Value iteration algorithm	Proposed algorithm
5	0.6	0.415	0.002
	0.7	0.71	0.002
	0.75	0.994	0.002
	0.8	1.487	0.002
	0.85	7.29	0.002
	0.9	42.678	0.003
	0.95	633.844	0.003
10	0.6	0.819	0.055
	0.7	1.406	0.065
	0.75	1.96	0.074
	0.8	2.928	0.086
	0.85	14.384	0.088
	0.9	79.017	0.094
	0.95	1194.145	0.111
15	0.6	1.257	0.208
	0.7	2.148	0.245
	0.75	2.993	0.275
	0.8	4.461	0.312
	0.85	21.114	0.321
	0.9	123.098	0.338
	0.95	1953.74	0.389
20	0.6	1.888	7.61
	0.7	2.983	8.791
	0.75	4.16	9.776
	0.8	6.198	11.043
	0.85	30.511	11.332
	0.9	180.525	11.85
	0.95	2886.58	13.477

lost sale cases). That is, $r^*(x+2) - r^*(x+1) \leq r^*(x+1) - r^*(x)$. In terms of the k_l parameters, the concavity od $r^*(x)$ translates into to the following constraints:

$$0 \leq k_l - k_{l+1} \leq k_{l+1} - k_{l+2}, \text{ for } l = 0, \ldots, m-3,\ k_l \leq m, \text{ and } k_0 = m. \quad (3.17)$$

These constraints, if they were to be included in our search algorithm, would further reduce the number of feasible **k** vectors (see Table 3.3). This would also lead to further improvements in the computational performance of our algorithm (see Table 3.4).

Table 3.3: Number of feasible **k** vectors with concave $r^*(x)$

m	Number of feasible **k** vectors
5	12
10	97
15	508
20	2087

3.5 Heuristics

In this section, we evaluate the performance of two heuristic policies that are simpler to implement and communicate than the optimal policy. Both policies belong to the same class of policies as the optimal one. Namely, both can be specified in terms of a threshold s on the inventory level and a vector of thresholds $\mathbf{k} = (k_0, \ldots, k_{m-1})$ on the inventory on order.

Heuristic H1: Under this heuristic, we set $k_0 = m$ and $k_j = 0$ for $j \neq 0$. In other words, if $x < s$, we order the maximum number of units to bring the number of units on order to its maximum value m; otherwise, we do not order. Hence,

Table 3.4: Computational performance of the proposed algorithm with concave $r^*(x)$

m	λ/μ	CPU time (seconds)
5	0.6	0.001
	0.7	0.001
	0.75	0.001
	0.8	0.002
	0.85	0.002
	0.9	0.002
	0.95	0.002
10	0.6	0.01
	0.7	0.012
	0.75	0.014
	0.8	0.016
	0.85	0.017
	0.9	0.018
	0.95	0.021
15	0.6	0.065
	0.7	0.076
	0.75	0.085
	0.8	0.097
	0.85	0.1
	0.9	0.105
	0.95	0.121
20	0.6	0.303
	0.7	0.35
	0.75	0.389
	0.8	0.44
	0.85	0.451
	0.9	0.472
	0.95	0.536

the policy is specified in terms on the single parameter s. For a given s, the expected average cost can be obtained using the approach described in Section 3.4 as illustrated in Figure 3.3. This leads to an average cost given by:

$$J^{\mathrm{H1}}(s) = h \sum_{i=0}^{s+m} (s+m-i) p^{\mathrm{H1}}_{s+m-i} + b \sum_{i=s+m+1}^{\infty} (i-s-m) p^{\mathrm{H1}}_{s+m-i} + \lambda c. \quad (3.18)$$

Noting that the average cost is convex in s, the value of s that minimizes this cost, s^{H1}, can be obtained as described in Section 3.4 using an equivalent expression to (15).

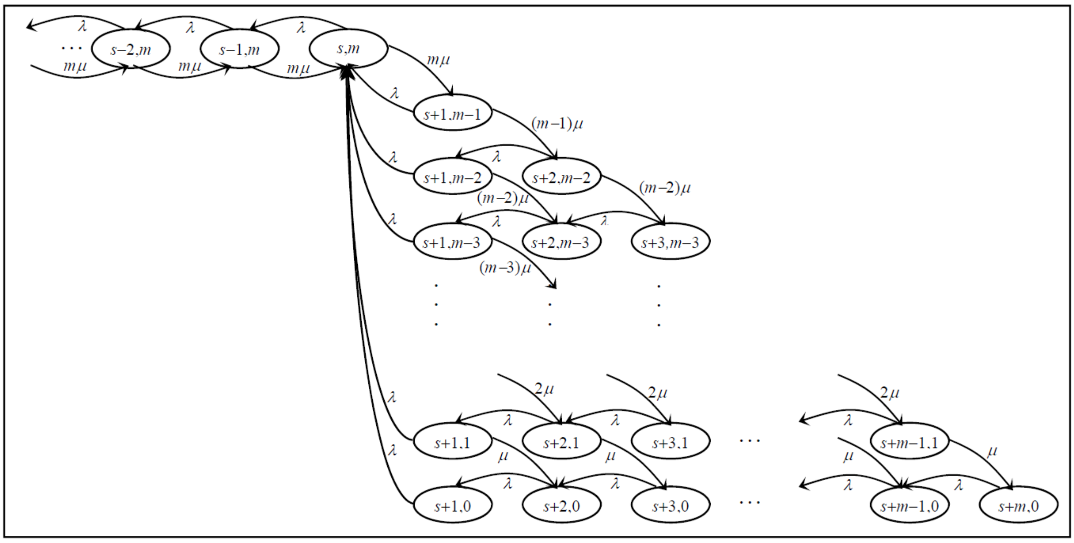

Figure 3.3: State transition diagram under Heuristic H1

Heuristic H2: Under this heuristic, wi set $k_0 = m$ and $k_j = m - j$ for $j \neq 0$. In other words, if $x < s$, we set the number of units on order to its maximum value m as in heuristic H1; if $x = s + j$, we bring the number of unite on order to $m - j$; otherwise, we do not order. This also means that once the threshold function starts decreasing from its maximum value m, it continues to decrease by one unit for each unit increase in inventory. Hence, under this heuristic, the inventory position stays constant and equals to $s+m$ if $s \leq x \leq s+m$. This implies that the policy is a modified base-stock policy with base-stock level $s+m$ (ordering in the way such that the inventory position is as close to $s+m$ as possible). As in heuristic H1, the policy here is specified by the single parameter s.

For a given s, we can again follow the approach described in Section 3.4 to obtain the average cost. However, in this case, the analysis simplifies. In particular, the state of the system can be described by the net inventors level. This allows us to characterize, in closed from, the probabilities p_i, where p_i is the probability that the net inventory level is $s+m-i$:

$$p_i = \begin{cases} (\lambda/\mu)^i \dfrac{p_0}{i!} & \text{for } i = 1, \ldots, m, \\ (\lambda/\mu)^m \dfrac{\rho_m^{i-m}}{m!} p_0 & \text{for } i = m+1, \ldots, \end{cases} \tag{3.19}$$

and

$$p_0 = \left(1 + \sum_{i=1}^{m} \frac{(\lambda/\mu)^i}{i!} + \frac{(\lambda/\mu)^m}{m!} \frac{\rho_m}{(1-\rho_m)} \right)^{-1}. \tag{3.20}$$

where $\rho_m = \lambda/m\mu$. The expected total cost of the system under this policy is given as follows

$$J^{\mathrm{H2}}(s) = h \sum_{i=0}^{s+m} (s+m-i)p_i + b \sum_{i=s+m+1}^{\infty} (i-s-m)p_i + \lambda c. \tag{3.21}$$

Noting again that the average cost function is convex in s, the optimal value of the threshold s^{H2} can be determined as follows. Let

$$s^+ = \left\lceil \log \left(\left(\frac{1-\rho_m}{p_0 (\lambda/\mu)^m / m!} \right) \left(\frac{h}{h+b} \right) \right) / \log(\rho_m) \right\rceil,$$

where the notation $\lceil w \rceil$ indicates the smallest integer that is greater than or equal to w. If $s^+ \geq 0$ then $s^{\mathrm{H2}} = s^+$. Otherwise, $s^{\mathrm{H2}} = \max\{s \geq$

$-m|\sum_{i=0}^{s+m-1} p_i \leq b/(h+b)\} - m$. Note that when $m = 1$, $p_0 = 1 - \lambda/\mu$, the optimal base-stock level $s^{\text{H2}} + m$ reduces to $\lceil \log(h/(h+b))/\log(\lambda/\mu) \rceil$. For an integrated production-inventory system, this corresponds to the optimal base-stock level for a system with a single facility (see Buzacott and Shanthikumar for a similar result).

To test the performance of heuristics H1 and H2, we choose a base system with parameters $m = 20$, $\mu = 1.0$, $\lambda = 18$, $h = 2$, $b = 15$, and $c = 0$. We vary parameter values one at a time and obtain the percentage difference between the average cost of the heuristic and that of the optimal policy:

$$\textit{Percentage diff.} = \frac{\textit{Heuristic policy average cost} - \textit{Optimal policy average cost}}{\textit{Optimal policy average cost}} \times 100.$$

Representative results are shown in Table 3.5. Note that we set the procurement cost c to zero since it is always incurred in the case of backorders and can be incorporated into the lost sales cost in the case of lost sales. From Table 3.5, we first note that $s^{\text{H1}} \geq s^* \geq s^{\text{H2}}$. This means that Heuristic H1 is associated with the highest maximum attainable inventory level $s^{\text{H1}} + m$, and Heuristic H2 with the lowest level $s^{\text{H2}} + m$. This is because heuristic H1 lacks the ability to adjust the number of units on order in the way heuristic H2 and the optimal policy do. Heuristic H2 must adjust its threshold for the number of units on order one unit at a time and lacks the flexibility of multiple unit increase or decrease that the optimal policy has. These results can also be explained by the shape of the threshold function $r^*(x)$ (illustrated by the vector $\mathbf{k}^*$ as shown in Table 3.5. Note that when the demand rate is low, the optimal vector $\mathbf{k}^*$ tends to have more nonzero values and approaches the corresponding vector for Heuristic H2. On the other hand, for systems with high utilization, $\mathbf{k}^*$ tends to have fewer

Table 3.5: Performance of Heuristics H1 and H2 in the case of backorders

		$s^*, \mathbf{k}^*$	s^{H1}	s^{H2}	Percentage difference H1	Percentage difference H2
h	2	16, (20, 17, 12, 5, 0, . . . , 0)	17	14	0.045	0.991
	5	9, (20, 16, 11, 4, 0, . . . , 0)	10	7	0.071	1.512
	7	7, (20, 15, 9, 1, 0, . . . , 0)	8	5	0.083	1.837
	10	4, (20,19,14,8,0, . . . , 0)	6	3	0.105	2.368
	12	3, (20, 19, 14, 8, 0, . . . , 0)	5	2	0.118	2.669
	13	3, (20, 17, 12, 5, 0, . . . , 0)	4	1	0.122	2.715
	15	2, (20, 18, 13, 7, 0, . . . , 0)	4	0	0.127	3.088
	17	2, (20, 16, 10, 2, 0, . . . , 0)	3	0	0.153	3.267
	20	1, (20, 17, 12, 5, 0, . . . , 0)	2	-1	0.165	3.697
	50	-2, (20, 18, 13, 6, 0, . . . , 0)	-1	-4	0.327	6.873
b	2	2, (20, 18, 13, 7, 0, . . . , 0)	4	0	0.127	3.088
	5	8, (20, 15, 9, 1, 0, . . . , 0)	9	6	0.077	1.68
	7	10, (20, 17, 12, 5, 0,..., 0)	11	8	0.063	1.4
	10	13,(20, 16, 10, 2, 0, . . . , 0)	14	11	0.056	1.175
	15	16, (20, 17, 12, 5, 0,. . . , 0)	17	14	0.045	0.991
	20	18, (20, 19, 14, 8, 0, . . . , 0)	20	17	0.04	0.902
	25	20, (20, 19, 14, 8, 0, . . . , 0)	22	19	0.038	0.844
	30	22, (20, 17, 12, 5, 0, . . . , 0)	23	20	0.035	0.766
	35	23, (20, 19, 14, 8, 0, . . . , 0)	25	22	0.034	0.755
	40	25, (20, 15, 9, 1, 0, . . . , 0)	26	23	0.032	0.698
	50	27, (20, 15, 9, 1, 0, . . . , 0)	28	25	0.03	0.652
	75	30,(20, 19, 14, 8, 0, . . . , 0)	32	28	0.027	0.609
	100	33, (20, 17, 12, 5, 0, . . . , 0)	34	31	0.025	0.542
λ	4	-7, (20, 19, 17, 15, 13, 11, 9, 6, 3, 0, . . . , 0)	-2	-14	96.896	29.265
	6	-5, (20, 19, 17, 15, 12, 9, 6, 2, 0, . . . , 0)	-1	-11	43.244	32.535
	8	-3, (20, 18, 15, 12, 8, 4, 0, . . . , 0)	-1	-9	17.134	36.441
	10	-2, (20, 18, 15, 11, 7, 2, 0, . . . , 0)	0	-6	7.426	33.331
	12	-1, (20, 19, 16, 12, 7, 1, 0, . . . , 0)	1	-4	2.844	24.242
	14	1, (20, 19, 15, 11, 5, 0, . . . , 0)	3	-1	1.058	12.716
	16	5, (20, 18, 13, 8, 0, . . . , 0)	7	3	0.285	4.807
	18	16, (20, 17, 12, 5, 0, . . . , 0)	17	14	0.045	0.991
	19	37, (20, 19, 14, 8, 0, . . . , 0)	39	36	0.011	0.223

nonzero values and approaches the corresponding vector for Heuristic H1.

As shown in Table 3.5 for the case of backorders, the heuristics perform well except when the holding cost is high, the backorder cost is low, or the demand rate is high. The heuristics lack the flexibility of the optimal policy to adjust the ordering threshold levels. This can lead to higher inventory levels when the heuristics are used, with the associated costs increasing with higher holding costs, lower backorder costs, or lower demand rates. The effect of the holding and backorder costs is more pronounced for heuristic H2 because the heuristic cannot adjust down the ordering thresholds sufficiently quickly. The effect of the low demand rate is more pronounced for heuristic H1 because, under H1, the order up to level cannot be smaller than m. Once delivered, ordered units that are not used to fulfill demand immediately tend to be held in inventory longer leading to higher holding costs.

Heuristics H1 and H2 can easily be adapted to the case of lost sales using the results of Section 3.3. Numerical results comparing the performance of the heuristics to that of the optimal policy are shown in Table 3.6 (in this case, the base system has parameters $m = 20$, $\mu = 1.0$, $\lambda = 19$, $h = 5$, $L = 150$, and $c = 0$). The results can be explained similarly to the case with backorders. Some of the differences in the relative performance of the two heuristics appear to be due the fact that parameters s^{H1} and s^{H2} can no longer be negative as in the backorder case.

Table 3.6: Performance of Heuristics H1 and H2 in the case of lost sales

		$s^*, \mathbf{k}^*$	s^{H1}	s^{H2}	Percentage difference H1	H2
h	5	17, (20, 15, 10, 2, 0, . . . , 0)	17	15	0.0487	1.062
	10	12, (20, 18, 13, 6, 0, . . . , 0)	13	10	0.061	1.49
	15	10, (20, 16, 11, 3, 0, . . . , 0)	10	8	0.082	1.79
	20	8, (20, 19, 14, 8, 0, . . . , 0)	9	7	0.087	2.194
	25	7, (20, 18, 13, 7, 0, . . . , 0)	8	5	0.102	2.5
	30	6, (20, 19, 14, 8, 0, . . . , 0)	7	5	0.121	2.728
	35	6, (20, 16, 10, 3, 0, . . . , 0)	6	4	0.136	2.927
	40	5, (20, 18, 13, 6, 0, . . . , 0)	6	3	0.137	3.335
	45	5, (20, 18, 13, 6, 0, . . . , 0)	5	3	0.159	3.435
	50	4, (20, 19, 14, 7, 0, . . . , 0)	5	3	0.153	4.851
	70	3, (20, 18, 13, 6, 0, . . . , 0)	4	1	0.196	4.851
	80	3, (20, 16, 10, 2, 0, . . . , 0)	3	1	0.234	5.067
	90	2, (20, 19, 14, 8, 0, . . . , 0)	3	1	0.243	5.614
	100	2, (20, 17, 12, 5, 0, . . . , 0)	3	0	0.26	6.175
L	25	2, (20, 15, 9, 1, 0, . . . , 0)	2	0	0.296	6.68
	50	4, (20, 15, 10, 2, 0, . . . , 0)	4	2	0.191	4.149
	75	5, (20, 18, 13, 6, 0, . . . , 0)	6	3	0.137	3.335
	100	6, (20, 19, 14, 8, 0, . . . , 0)	7	5	0.121	2.728
	150	8, (20, 19, 14, 7, 0, . . . , 0)	9	7	0.087	2.194
	175	9, (20, 18, 12, 6, 0, . . . , 0)	10	7	0.081	1.966
	200	10, (20, 16, 11, 3, 0, . . . , 0)	10	8	0.082	1.79
	250	11, (20, 17, 12, 6, 0, . . . , 0)	12	9	0.067	1.618
	300	12, (20, 18, 13, 6, 0, . . . , 0)	13	10	0.061	1.49
	350	14, (20, 18, 12, 6, 0, . . . , 0)	14	11	0.056	1.317
	400	14, (20, 17, 11, 4, 0, . . . , 0)	14	12	0.056	1.274
	500	15, (20, 19, 14, 8, 0, . . . , 0)	16	14	0.052	1.181
l	4	0, (11, 8, 5, 2, 0, . . . , 0)	1	0	5.538	61.392
	6	0, (15, 12, 9, 5, 1, 0, . . . , 0)	1	0	6.223	74.245
	8	0, (18, 15, 12, 9, 5, 0, . . . , 0)	1	0	7.909	88.66
	10	1, (20, 16, 12, 8, 4, 0, . . . , 0)	2	0	4.358	56.123
	12	2, (20, 17, 13, 8, 3, 0, . . . , 0)	3	0	1.95	22.314
	14	3, (20, 19, 15, 10, 5, 0, . . . , 0)	4	1	0.858	10.331
	16	5, (20, 19, 15, 10, 0, . . . , 0)	6	3	0.314	5.256
	18	8, (20, 19, 14, 7, 0, . . . , 0)	9	7	0.087	2.194
	19	10, (20, 18, 13, 6, 0, . . . , 0)	11	9	0.043	1.205
	21	16, (20, 16, 10, 0, . . . , 0)	16	15	0.008	0.233
	23	25, (20, 18, 11, 1, 0, . . . , 0)	25	24	0	0.01
	25	38, (20, 15, 7, 0, . . . , 0)	38	37	0.022	0.798
	30	64, (20, 0, . . . , 0)	64	71	0	0

3.6 Systems with Order Cancellation

In this section, we consider a system similar in all aspects to the original model described in Sections 3.2 and 3.3, except that now we allow orders to be cancelled at no cost after they have been placed. This means that, at each decision epoch, we can choose how many units to have on order by arbitrarily increasing or decreasing the number of such orders. In other words, at each decision epoch, we decide on an order quantity k, such that $0 \leq k \leq m$, where k is unconstrained by the number of previously placed orders. Because leadtimes are memoryless, it is not necessary to keep track of the number of units on order. The state of the system, at time t, can be fully described by only the inventory level $X(t)$. The assumption of order cancellation is a common assumption in much of the literature on integrated-production inventory systems (see for example Ha (1997), De Vericourt et al. (2002), and Benjaafar and ElHafsi (2006)).

Using the same methodology as in the case of no order cancellation, we can show that the optimal cost function v^* satisfies, upon uniformization, the following optimality equation (the details are omitted for brevity):

$$v^*(x) = g(x) + \lambda v^*(x-1) + \min_{0 \leq k \leq m} \{k\mu\,(v^*(x+1)+c) + (m-k)\mu v^*(x)\}\,, \quad (3.22)$$

which we can rewrite as

$$v^*(x) = g(x) + \lambda v^*(x-1) + m\mu v^*(x) + \min_{0 \leq k \leq m} \{k\mu\,(v^*(x+1) - v^*(x) + c)\}\,. \quad (3.23)$$

Theorem 8. *The optimal control policy is a base-stock policy with base-stock level s^* such that if $x < s^*$, it is optimal to bring the number of units on order to m, otherwise it is optimal to bring this number to zero.*

Proof. It is easy to show that the optimal cost function is convex in x. That is, the difference $\Delta v^*(x) = v^*(x+1) - v^*(x)$ is non-decreasing in x. In turn, this implies that when $\Delta v^*(x) + c \geq 0$, it is optimal to bring the number of units on order to zero (by cancelling pending orders if necessary). On the other hand, when $\Delta v^*(x) + c < 0$, it is optimal to bring the number of units on order up to m. Thus, in each decision epoch, the optimal policy is a so-called bang-bang policy, where the optimal number of orders to place is either 0 or m. The convexity of $v^*(x)$ also implies that the optimal policy is a base-stock policy with base-stock level s^*, where $s^* = \min\{x|\Delta v^*(x) + c \geq 0\}$, such that it is optimal to bring the number of units on order to m if $x < s^*$ and to bring it to 0 otherwise. $\square$

Noting that once one unit is produced, it is possible to cancel the production of all remaining ones. It is not difficult to see that the dynamics of the system are the same as those where leadtime is exponentially distributed with rate $m\mu$ and only one unit can be on order at any time. In the case of an integrated production-inventory system, this means that the system is equivalent to one with a single facility with a production rate $m\mu$. The dynamics of such a system can be described by a simple Markov chain and various performance measures can be obtained in this case in closed form. In particular, given a base-stock level s, the average cost is given by (we again omit the details for the sake of brevity):

$$J(s) = \lambda c + h\left(s - \rho_m \frac{(1-\rho_m^s)}{(1-\rho_m)}\right) + b\frac{\rho_m^{s+1}}{(1-\rho_m)}s, \tag{3.24}$$

where $\rho_m = \lambda/m\mu$. Noting that the average cost is convex in s, the optimal base-stock level is given by the smallest integer for which $J(s+1) - J(s) \geq 0$. It is not difficult to show that the base-stock level, s^*, is given by:

$$s^* = \left\lceil \frac{\log\left(\frac{h}{h+b}\right)}{\log\left(\lambda/m\mu\right)} \right\rceil. \tag{3.25}$$

Similarly, for systems with lost sales, we can show that the optimal cost function v^* satisfies the following optimality equation:

$$v^*(x) = \begin{cases} g(x) + \lambda v^*(x-1) + m\mu v^*(x) & \\ \quad + \min\limits_{0 \le k \le m} \left\{k\mu\left(v^*(x+1) - v^*(x) + c\right)\right\} & \text{if } x > 0, \\ g(x) + \lambda\left(v^*(x) + L\right) + m\mu v^*(x) & \\ \quad + \min\limits_{0 \le k \le m} \left\{k\mu\left(v^*(x+1) - v^*(x) + c\right)\right\} & \text{otherwise.} \end{cases} \tag{3.26}$$

As in the corresponding backlog case, the optimal policy is bang-bang. It is optimal to bring the number of units on order to m if $x < s^*$ and to zero otherwise. Here too, the dynamics of the system can be modeled using a Markov Chain. Various performance measures can be obtained in closed form. In particular, for a given base-stock level s, the average cost is given by

$$J(s) = \lambda \frac{(1-\rho_m)\rho_m^s}{(1-\rho_m^{s+1})} L + \frac{(1-\rho_m)(s+\rho_m^{s+1}) - \rho_m(1-\rho_m^{s+1})}{(1-\rho_m)(1-\rho_m^{s+1})} h + \lambda c \frac{1-\rho_m^s}{(1-\rho_m^{s+1})}, \tag{3.27}$$

where $\rho_m = \lambda/m\mu$. Noting again that the average cost is convex in s, the optimal base-stock level can be easily computed.

We conclude this section by providing numerical results that examine the benefit from order cancellation. To do so, we compare the performance of the original model (a system with backlog and no order cancellation) and its lost sale counterpart to a system that allows for order cancellation. Using a base system with parameters $m = 20$, $\mu = 1.0$, $\lambda = 18$, $h = 2$, $b = 15$, and $c = 0$, Tables 3.7 and 3.8 show the percentage difference in average cost between systems without

Table 3.7: Percentage difference in average cost between systems without and with order cancellation (the backorder case)

h	2	0.881
	5	1.328
	7	1.629
	10	2.093
	12	2.365
	13	2.45
	15	2.783
	17	2.816
	20	3.375
	50	7.346
b	2	2.783
	5	1.478
	7	1.248
	10	0.994
	15	0.881
	20	0.795
	25	0.736
	30	0.68
	35	0.657
	40	0.607
	50	0.562
	75	0.527
	100	0.48
λ	4	141.049
	6	98.15
	8	65.426
	10	37.318
	12	21.346
	14	9.901
	16	4.095
	18	0.881
	19	0.21

Table 3.8: Percentage difference in average cost between systems without and with order cancellation (the lost sale case)

h	5	0.914
	10	1.329
	15	1.552
	20	1.955
	25	2.22
	30	2.498
	35	2.577
	40	2.98
	45	3.078
	50	3.51
	70	4.373
	80	4.705
	90	5.4
	100	5.6
L	25	6.74
	50	3.85
	75	2.98
	100	0.025
	150	0.02
	175	0.018
	200	0.016
	250	0.0144
	300	0.0133
	350	0.0122
	400	0.011
	500	0.01
λ	4	59.662
	6	49.969
	8	35.418
	10	23.105
	12	14.43
	14	8.067
	16	4.356
	18	1.955
	19	1.136
	21	0.221
	23	0.012
	25	0
	30	0

and with order cancellation for the backlog and lost sales cases respectively, where

$$Percentage\ diff. = \frac{Average\ cost\ without\ cancellation - Average\ cost\ with\ cancellation}{Average\ cost\ with\ cancellation} \times 100.$$

Results from a more extensive set of experiments reveal similar observations.

As expected, a system in which order cancellations are possible results in lower costs since it has the ability to quickly adjust the number of orders (or, in the case of a production-inventory system, the production capacity). As shown in Table 3.7, for the case of backlogs, the benefit of order cancellation increases as the holding to the backorder cost ratio increases. Without order cancellation, all placed orders eventually show up in inventory. The cost implication of the resulting inventory is higher with higher inventory holding cost or with lower backorder cost. The benefit of order cancellation increases with decreases in the demand rate. This is because, when the demand rate is low, any inventory held tends to be held for longer periods of time. Systems with order cancellation can mitigate the need for holding inventory by placing multiple orders when demand arises, thereby expediting deliveries, but then cancelling pending orders once demand is satisfied. This ability to expedite deliveries without repercussion on inventory holding cost is not available to the system without order cancellations. Such systems end up carrying more inventory on average than systems without order cancellation. Table 3.8 tells a similar story for systems with lost sales.

Chapter 4

Optimal Policies for Inventory Systems with Concave Ordering Costs

4.1 Introduction

Most of the literature on inventory systems usually assumes a linear ordering cost or a linear ordering cost with a setup cost. As Scarf (1963) argues, "This type of cost functions has appeared in inventory theory not necessarily because of its realism, but because it provides one of the few examples of cost functions with a decreasing average cost for which the analysis of inventory policies is relatively easy." In this paper, we consider inventory systems with general concave ordering cost functions, where the type of ordering costs described in Scarf (1959) is a special case under our setting. The class of concave ordering cost functions is a special type of a decreasing average cost and there are many examples of concave ordering costs in practice. Consider the following examples.

Quantity Discounts: Quantity discounts provide a practical foundation for coordinating inventory decisions in supply chains. Sellers usually employ quantity discount schemes or contracts to give buyers the incentive to buy more. That is, the larger the order is, the lower the marginal price will be. Ordering costs with quantity discounts can usually be expressed by piecewise linear concave functions: first, there is a setup cost; then the first few items have the same per-unit cost; the next few items have a lower per-unit cost, and so on.

The Effect of Economies of Scale: In economics, one of the common assumptions on production functions is that they have the economies of scale feature. Basically, the more a firm produces the same item, the more efficient the production technology will be. The transportation costs in supply chains also exhibit economies of scale: the more volume of goods to be shipped, the cheaper the marginal cost will be. Concave functions are an important class of functions exhibiting the feature of economies of scale.

Procurement with Multiple Suppliers: In many cases in practice, there are multiple suppliers available for a buyer. Usually local suppliers offer relatively lower setup costs but with higher per-unit costs and overseas or distant suppliers offer higher setup costs but with lower per-unit costs. Hence, if the buyer chooses the suppliers optimally, the resulting ordering cost is a piecewise linear concave function. A related case of a buyer that purchases from both long-term suppliers and spot markets is treated in Yi and Scheller-Wolf (2003) and the references therein. There are many other examples of concave costs due to the availability of multiple choices of labor and production, see Fox et al. (2006) for examples and references therein.

Scarf (1959) proves that the (s, S) policy is optimal for an inventory system with a fixed ordering cost and a unit ordering cost and does it by introducing the notion of K-convexity. This type of ordering cost is a special case of concave ordering costs. Karlin (1958) analyzes the optimal ordering policy for a one-period inventory problem with concave ordering costs. Scarf (1963) points out that it is difficult to generalize the result to the dynamic multiperiod setting. There has been only limited research on stochastic inventory systems with concave ordering costs. Porteus (1971) analyzes inventory systems with piecewise linear concave ordering costs. He shows that a generalized (s, S) policy is optimal for a multi-period periodic review inventory system under some mild assumption on cost functions and that demand has a one-sided Polya density. He does it by introducing a generalized notion of K-convexity called quasi-K-convexity. However, the class of one-sided Polya densities does not include many densities encountered in practice, for example, the normal distribution, beta distribution and most gamma distributions, although it does include the exponential distribution and all its finite convolutions. Porteus (1972) also shows that the generalized (s, S) policy is optimal for uniform demand distributions.

Fox et al. (2006) consider the optimal policy for an inventory system with two suppliers: the buyer incurs a high variable cost but negligible fixed cost for the first supplier (HVC) and a lower variable cost but a substantial fixed cost for the second supplier (LVC). The resulting ordering cost is a two-piece linear concave function.

They show that the optimal policy is a (s, S_{HVC}, S_{LVC}) policy, which is a special case of the generalized (s, S) policy, under the condition that the demand density is log-concave. Their proof relies on K-convexity and quasi-convex properties since they consider a two-piece linear concave function. Although the class of log-concave densities is less restrictive than the class of one-sided Polya densities, it still only covers a limited range of distributions. Furthermore, their results do not cover general piecewise linear concave ordering costs. Hence, whether or not the generalized (s, S) policy is optimal for general demand distributions remained an open question.

Recently, Chen et al. (2010) consider joint pricing and inventory control for inventory systems with concave ordering costs. They utilize quasi-K-convexity to show that the optimal policy is a generalized (s, S, p) policy when demand distributions are Polya or uniform. Another related paper is Yi and Scheller-Wolf (2003), where they also consider a two-supplier inventory problem: the buyer has a long-term contract from a regular supplier with a minimum and maximum purchasing quantity, and the buyer can also purchase from a spot market that has no quantity limitation but with a fixed entry fee. They partially characterize the structure of the optimal policy and their proof relies on a closure property of K-convexity. Note that the ordering cost in their case is no longer concave since they assume a limited capacity for the regular supplier and the corresponding optimal policy is not a generalized (s, S) policy.

Chen and Simchi-Levi (2004) consider the joint inventory-pricing control problem with fixed ordering costs. They introduce the concept of sym-K-convexity, which is a generalization of K-convexity, and show that the optimal policy can be fully characterized except for a bounded interval for the multiplicative demand model. Chao and Zipkin (2008) study a model with a fixed cost function that is neither convex nor concave: the fixed cost is incurred only if the order quantity exceeds a threshold, and hence the cost function can be written as $c(x) = K\delta(x - C)$ for some constant C. They apply the property of K-convexity and partially characterize the optimal policy with three critical points which divide the state space into five regions.

In contrast to the existing literature, we characterize the structure of optimal policies for inventory systems with concave ordering costs with general demand distributions. In order to analyze the structure of the optimal policy, we first introduce a monotone condition that ensures the optimality of a generalized (s, S) policy. We then introduce the concept of c-convexity, a generalization of K-convexity, and use it to show that the value function for this problem is c-convex with respect to a modified ordering cost function. Based on the c-convexity of the value function, we show that, except for a bounded region of the state space, the generalized (s, S) policy is optimal. We also provide conditions under which the generalized (s, S) policy is optimal for all regions of the state space. Our results can be readily extended to systems with time-varying cost parameters, systems with fixed leadtimes and to systems with lost sales. The notion of c-convexity we introduce in this paper may also have usefulness to other inventory control problems.

4.2 Inventory Systems with Concave Ordering Costs

We consider a single product single stage inventory problem with multiple periods, stochastic demands, and zero leadtime. The assumption on zero leadtime is not critical and is made for ease of exposition (see Section 4.6 for extensions). Demand ξ_t in each period t is a continuous random variable with $E[\xi_t] < \infty$ and distribution function $F_t(x), x \geq 0$, where $t = 1, \cdots, T$ and T corresponds to the length of the planning horizon. Demands in different periods are independent but not necessarily identically distributed (i.e., demand can be time-varying). Inventory is replenished from an outside supplier immediately (i.e., with zero leadtime) with ample stock. Demand is satisfied from on-hand inventory, if any is available; otherwise it is backordered. In each period, the inventory manager must decide on the quantity to order to minimize the expected discounted cost over the entire planning horizon. There are three types of costs in each period t: (1) an ordering cost $c(z)$ if the order quantity is $z, z \geq 0$, (2) a holding cost $h_t(x^+)$ and (3) a backordering cost $b_t(x^-)$ given the inventory level x in period t, where $x^+ = \max\{0, x\}$ and $x^- = \max\{0, -x\}$. Finally, we allow a discount factor $\alpha \in (0, 1]$.

In order to simplify our presentation, we first consider a piecewise linear concave ordering cost $c(\cdot)$ with n linear pieces. Specifically, we can express

$$c(x) = \min_{i=1,\cdots,n} \{K_i \delta(x) + c_i x\},$$

with $0 \leq K_1 < K_2 < \cdots < K_n$ and $c_1 > c_2 > \cdots > c_n \geq 0$, where δ is defined as follows

$$\delta(x) = \begin{cases} 1 & \text{if } x > 0, \\ 0 & \text{if } x \leq 0. \end{cases} \tag{4.1}$$

We will discuss how we can deal with general concave ordering costs at the end of Section 4.3 and time-varying costs in the section on extensions (see Section 4.6).

Let

$$g_t(x - \xi_t) = \begin{cases} h_t(x - \xi_t) & \text{if } x \geq \xi_t, \\ b_t(\xi_t - x) & \text{otherwise.} \end{cases}$$

Let x_t be the starting inventory level and y_t be the post-ordering inventory level for period t, with $x_{t+1} = y_t - \xi_t$. Given $x_1, \cdots, x_T, y_1, \cdots, y_T$, i.e., the ordering quantities being $q_t = y_t - x_t, t = 1, \cdots, T$, the expected discounted total cost is given by

$$E\left\{\sum_{t=1}^{T} \alpha^t [c(y_t - x_t) + g_t(y_t - \xi_t)]\right\}. \tag{4.2}$$

Let $v_t^*(x)$ be the value function (the optimal expected discounted cost) in period t when the inventory level in period t is x. Then the corresponding dynamic programming formulation is given by

$$v_t^*(x) = \min_{y \geq x} \{c(y - x) + Eg_t(y - \xi_t) + \alpha E v_{t+1}^*(y - \xi_t)\}. \tag{4.3}$$

Finally we let $v_{T+1}^*(x) = 0$ for all x.

Assumption 1. *We assume that $L_t(y) = Eg_t(y - \xi_t)$ is convex in y and finite for any y.*

For example, this assumption is satisfied if h_t and b_t are linear and $E[\xi_t] < \infty$. The finiteness of L_t ensures that L_t is continuous on $(-\infty, \infty)$ (by the dominated convergence theorem).

Let

$$H_t(y) = Eg_t(y - \xi_t) + \alpha E v^*_{t+1}(y - \xi_t).$$

Then the optimality equation is given by

$$v^*_t(x) = \min_{y \geq x} [c(y - x) + H_t(y)]. \tag{4.4}$$

Given x, let $y_t(x)$ be the smallest minimizer of $c(y - x) + H_t(y)$, i.e.,

$$y_t(x) = \min \arg \min_{y \geq x} \{c(y - x) + H_t(y)\}. \tag{4.5}$$

Hence, given the current inventory level is x, it is optimal to order $y_t(x) - x$ quantity in period t.

4.3 The Structure of the Optimal Policy

In this section, we show that, except for a bounded region, the optimal policy can be described by a generalized (s, S) policy.

First, we show a conditional monotone property for $y_t(x)$ for any concave function c.

Theorem 9. *Suppose that $y_t(x) > x$, then $y_t(z) \leq y_t(x)$ for $z \in (x, y_t(x))$.*

Proof. Suppose that for some x, we have $y_t(x) - x > 0$. Let $y_t(x) > z > x$. Since we know that $c(x)$ is concave in x, for $\omega > 0$ we have

$$c(y_t(x) + \omega - z) - c(y_t(x) - z) \geq c(y_t(x) + \omega - x) - c(y_t(x) - x),$$

which implies that

$$\begin{aligned}
&c(y_t(x)+\omega-z)+H_t(y_t(x)+\omega)-[c(y_t(x)-z)+H_t(y_t(x))] \\
\geq & c(y_t(x)+\omega-x)+H_t(y_t(x)+\omega)-[c(y_t(x)-x)+H_t(y_t(x))] \geq 0.
\end{aligned} \tag{4.6}$$

Since inequality (4.6) is true for all $\omega > 0$ and c is continuous in $(0, \infty)$, it follows that $y_t(z) \leq y_t(x)$. □

As far as we know, this is a new result in the literature. The interesting aspect of this result is that the conditional monotone property in period t holds for general concave ordering costs. However we do not know what will happen for $z \notin (x, y_t(x))$. Next, we describe a monotone condition that is the key to characterizing the structure of optimal policies.

Condition 1. *$y_t(x_2) > x_2$ implies that $y_t(x_1) > x_1$ for any $x_1 < x_2$. In words, if it is optimal to order a positive amount when the starting inventory level is x_2, then it must be optimal to order a positive amount when the starting inventory level is less than x_2 in period t.*

It turns out that if we know that $y_t(x_2) > x_2$ implies that $y_t(x_1) > x_1$ for any $x_1 < x_2$, then coupled with Theorem 9, we can show that $y_t(x_1) \geq y_t(x_2)$ for all $x_1 < x_2$ such that $y_t(x_2) > x_2$.

Lemma 10. *Under Condition 1, we have*

(1) $y_t(x_1) \geq y_t(x_2)$ for all $y_t(x_2) > x_2$ and $x_1 < x_2$.

(2) There exists some x_0 such that $y_t(x) = x$ for all $x \geq x_0$ and $y_t(x)$ is non-increasing in x for $x \in (-\infty, x_0]$.

Proof. We prove the first part by contradiction. Suppose we have $y_t(x_1) < y_t(x_2)$ given that $y_t(x_2) > x_2$, $y_t(x_1) > x_1$ and $x_1 < x_2$. We differentiate two cases. (1) $y_t(x_1) \leq x_2$. This case is impossible, since under Condition 1, we must have $y_t(y_t(x_1)) > y_t(x_1)$, i.e., $y_t(x_1)$ is not an optimal order-up-to level for x_1, which violates the optimality of $y_t(\cdot)$. (2) $y_t(x_1) \in (x_2, y_t(x_2))$. This case is impossible since it violates Theorem 9. We know that by Theorem 9, we must have $y_t(x_2) < y_t(x_1)$ since $x_2 \in (x_1, y_t(x_1))$.

Since we have $\lim_{x\to\infty} H_t(x) = \infty$, it follows that for sufficiently large x, we must have $y_t(x) = x$. Let x_0 the smallest value such that $y_t(x) = x$. Then $y_t(x) > x$ for all $x < x_0$ by Condition 1. It follows that $y_t(x)$ is non-increasing in x in that domain by part (1) of this lemma. It can also be shown that $y_t(x) = x$ for all $x > x_0$, otherwise if $y_t(x) > x > x_0$ then we must have $y_t(x_0) > x_0$. This contradicts the definition of x_0. □

Theorem 10. *If Condition 1 is satisfied, then the optimal inventory policy in period t is a generalized (s, S) policy, i.e., there exists $(s_{m,t}, \cdots, s_{1,t}, S_{1,t}, \cdots, S_{m,t})$ with $s_{m,t} < s_{m-1,t} < \cdots < s_{1,t} \leq S_{1,t} < S_{2,t} < \cdots < S_{m,t}$ for some $m \leq n$ such that if $x < s_{m,t}$ then we order up to $S_{m,t}$ and if $x \in [s_{i,t}, s_{i-1,t})$ then we order up to $S_{i-1,t}$ for $i = 2, \ldots, m$, and finally we order nothing for $x \geq s_{1,t}$. Hence, we have at most n distinctive such order-up-to levels $S_{i,t}$.*

Proof. Let $s_{1,t} = \min\{x : H_t(x) \leq c(y-x)+H_t(y), y > x\}$, i.e., $s_{1,t}$ is the minimum starting inventory level such that it is optimal to order nothing (the existence of $s_{1,t}$ is due to $\lim_{x\to\infty} H_t(x) = \infty$ and H_t is continuous). It follows that if $x > s_{1,t}$, then it is also optimal to order nothing, since otherwise it would violate Condition 1. Also if $x < s_{1,t}$, then it must be optimal to order a positive quantity and the post-ordering inventory level must be greater than or equal to $s_{1,t}$, since otherwise it would violate the definition of $s_{1,t}$.

Let

$$\hat{S}_{i,t} = \min \arg \min_{y \geq s_{1,t}} \{H_t(y) + c_i y\},$$

i.e., $\hat{S}_{i,t}$ is the minimum of $H_t(y)+c_i y$ on $[s_{1,t}, \infty)$ (the existence of $\hat{S}_{i,t}$ is due to the continuity of H_t). Since $c_1 > c_2 > \cdots > c_n$, it follows that $\hat{S}_{1,t} \leq \hat{S}_{2,t} \leq \cdots \leq \hat{S}_{n,t}$. For $x \in (-\infty, s_{1,t})$, let

$$\begin{aligned} v_{i,t}(x) &= \min_{y_i \geq x} [c_i y_i + K_i \delta(y_i - x) + H_t(y_i)] - c_i x \\ &= \min\{\min_{y_i > x} [c_i y_i + K_i + H_t(y_i)], c_i x + H_t(x)\} - c_i x. \end{aligned}$$

We have

$$\min_{y \geq x} \{c(y - x) + H_t(y)\} = \min_{i=1,\cdots,n} \{v_{i,t}(x)\}.$$

It follows that for any starting inventory level $x \in (-\infty, s_{1,t})$ (note that it is optimal to order a positive quantity for such starting inventory level x), it must be optimal to order to one of the levels in $\{\hat{S}_{1,t}, \hat{S}_{2,t}, \cdots \hat{S}_{n-1,t}\}$. Ties can be broken by choosing the smallest solution.

Suppose that for small enough δ it is optimal to order up to $\hat{S}_{i_1,t}$ for some $i_1 \in \{1, \cdots, n\}$ for starting inventory level $x \in [s_{1,t} - \delta, s_{1,t})$. If $i_1 = n$, then we are done. Otherwise, let $s_{2,t}$ be the smallest value such that it is optimal to order up to $\hat{S}_{i_1,t}$, i.e., $y_t(x) = \hat{S}_{i_1,t}$ for $x \in [s_{2,t}, s_{1,t})$. We define $S_{1,t} \equiv \hat{S}_{i_1,t}$. Since it is also optimal to order a positive quantity for $x < s_{2,t}$, by Lemma 10, we have $y_t(x) > \hat{S}_{i_1,t}$ for $x < s_{2,t}$. Again, suppose for some small enough δ it is optimal to order-up-to $\hat{S}_{i_2,t}$ for some $i_2 \in \{1, \cdots, n\}$ for $x \in [s_{2,t} - \delta, s_{1,t})$. Obviously, we have

$i_2 > i_1$ by the conditional monotone property. We define $S_{2,t} \equiv \hat{S}_{i_2,t}$. If $S_{2,t} = \hat{S}_{n,t}$, then we are done. Otherwise by a similar argument, we can iteratively define $s_{i,t}$ and $S_{i,t}$ $(i > 3)$ such that it is optimal to order-up-to $S_{i,t}$ for $x \in [s_{i+1,t}, s_i)$ until we have some $s_{m,t}$ and $S_{m,t} = \hat{S}_{n,t}$. Then it follows that if $x < s_{m,t}$, it is optimal to order up to $S_{m,t} = \hat{S}_{n,t}$. It is also clear that $m \leq n$. □

To analyze the structure of the optimal policy, we can rewrite the ordering cost as follows: $c(x) = K_i + c_i x$ for $x \in [z_{i-1}, z_i], i = 2, \cdots, n-1$, $c(x) = K_1\delta(x) + c_1 x$ for $x \in [0, z_1]$ and $c(x) = K_n + c_n x$ for $x \geq z_{n-1}$, where $0 < z_1 < \cdots < z_{n-1}$. Note that $c(x) - c_n x \geq 0$ for all $x \geq 0$. We first use the following transformation. We define $\bar{c}(x) \equiv c(x) - c_n x$. It is clear that $\bar{c}(x) = K_n$ for $x \geq z_{n-1}$. Then we can reformulate the dynamic recursion in (4.3) as follows.

$$\bar{v}_t^*(x) = \min_{y \geq x} E\{\bar{c}(y-x) + (1-\alpha)c_n[y-\xi_t] + c_n\xi_t + g_t(y-\xi_t) + \alpha E\bar{v}_{t+1}^*(y-\xi_t)\}, \quad (4.7)$$

with $\bar{v}_{T+1}(x) = c_n x$. Let

$$\bar{G}_t(y) = (1-\alpha)c_n[y - E\xi_t] + c_n E\xi_t + Eg_t(y-\xi_t) + \alpha E\bar{v}_{t+1}^*(y-\xi_t),$$

and

$$\bar{L}_t(y) = (1-\alpha)c_n[y - E(\xi_t)] + c_n E\xi_t + Eg_t(y-\xi_t).$$

Then we have

$$\bar{v}_t^*(x) = \min_{y \geq x}\{\bar{c}(y-x) + \bar{G}_t(y)\}. \quad (4.8)$$

We can show that $v_t^*(x) = \bar{v}_t^*(x) - c_n x$.

Next, we introduce a new generalized convexity notion to which we refer as c-convexity.

Definition 4. *A function f is said to be c-convex for any nonnegative nondecreasing concave function c if for any $x_1 < x_2$ and $\theta \in (0,1)$ the following inequality holds*

$$\theta f(x_1) + (1-\theta)[f(x_2) + c(x_2 - \theta x_1 - (1-\theta)x_2)] \geq f(\theta x_1 + (1-\theta)x_2).$$

One can view c-convexity as a generalization of K-convexity.

Based on the definition of c-convexity, we can show that the following lemma holds.

Lemma 11. *c-convex functions have the following properties. Assuming $c^i, i = 1,2$ are nonnegative nondecreasing concave functions.*

1. *Convexity is equivalent to 0-convexity, where 0 denotes that $c(x) \equiv 0$ for all $x \geq 0$.*

2. *If g is c-convex, then $g(x+a)$ is also c-convex for any a.*

3. *If g is c^1-convex, then it is also c^2-convex if $c^2(x) \geq c^1(x)$ for all $x \geq 0$.*

4. *If g_i is c^i-convex for $i = 1,2$, then $a_1 g_1 + a_2 g_2$ is $a_1 c^1 + a_2 c^2$-convex for all non-negative a_i.*

5. *If g is c-convex and $f(x) = E[g(x-\xi)] < \infty$, where ξ is a random variable, then $f(x)$ is also c-convex.*

The proof of the lemma is straightforward and hence omitted.

Next, we show that the value function is indeed $\bar{c}$-convex where $\bar{c}(x) = c(x) - c_n x$.

Lemma 12. *$\bar{v}_t^*(x)$ is $\bar{c}$-convex for all t.*

Proof. We use induction to show that $\bar{v}_t^*(x)$ is $\bar{c}$-convex for all t. It is clear that $\bar{v}_T^*$ is $\bar{c}$-convex since it is a linear function. Suppose that $\bar{v}_{t+1}$ is $\bar{c}$-convex. Since $\bar{L}_t$ is convex, then $\bar{G}_t(x) = \bar{L}_t(x) + \alpha E\bar{v}_{t+1}^*(x - \xi_t)$ must be $\bar{c}$-convex according to Lemma 11. For any $x_1 < x_2$ and $\theta \in (0,1)$, we differentiate two possible cases: (1) $y_t(x_1) \geq \theta x_1 + (1-\theta)x_2$ and (2) $y_t(x_1) < \theta x_1 + (1-\theta)x_2$ (recall that $y_t(x)$ is the optimal order-up-to level for x).

For case (1), we have

$$\begin{aligned}
&\theta\bar{v}_t^*(x_1) + (1-\theta)[\bar{v}_t^*(x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
=&\theta[\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1)] + (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) \\
&+ \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
\geq&\theta[\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1)] + (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - \theta x_1 - (1-\theta)x_2)] \\
\geq&\theta[\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - \theta x_1 - (1-\theta)x_2)] + (1-\theta)[\bar{G}_t(y_t(x_2)) \\
&+ \bar{c}(y_t(x_2) - \theta x_1 - (1-\theta)x_2)] \\
\geq&\bar{v}_t^*(\theta x_1 + (1-\theta)x_2).
\end{aligned}$$

The first inequality is due to the subadditivity of $\bar{c}$ since $\bar{c}$ is concave. The second inequality is due to the fact that $\bar{c}$ is nondecreasing and $y_t(x_1) \geq \theta x_1 + (1-\theta)x_2$.

For case (2), we differentiate two subcases: (2a) $\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1) > \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)$ and (2b) $\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1) \leq \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)$.

First, we consider subcase (2a). In this subcase, we have

$$\begin{aligned}
&\theta\bar{v}_t^*(x_1) + (1-\theta)[\bar{v}_t^*(x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
=&\theta[\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1)] + (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) \\
&+ \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)]
\end{aligned}$$

$$
\begin{aligned}
>&\theta[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
&+ (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - \theta x_1 - (1-\theta)x_2)] \\
\geq& \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - \theta x_1 - (1-\theta)x_2) \\
\geq& \bar{v}_t^*(\theta x_1 + (1-\theta)x_2).
\end{aligned}
$$

The first inequality is due to the subadditivity of $\bar{c}$ and the assumption of subcase (2a) and the second inequality is due to the subadditivity of $\bar{c}$.

Next, we consider subcase (2b). Since $y_t(x_1) < \theta x_1 + (1-\theta)x_2$ and $\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1) \leq \bar{G}_t(x_1)$, it follows that there exists some $\hat{x}_1$ such that $x_1 \leq \hat{x}_1 \leq y_t(x_1)$ and $\bar{G}_t(\hat{x}_1) = \bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1)$ since $\bar{G}_t$ is continuous. Then there exists $1 > \rho \geq \theta$ such that $\rho\hat{x}_1 + (1-\rho)y_t(x_2) = \theta x_1 + (1-\theta)x_2$. In this subcase we have

$$
\begin{aligned}
&\theta\bar{v}_t^*(x_1) + (1-\theta)[\bar{v}_t^*(x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
=&\theta[\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1)] + (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) \\
&+ \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
=&\theta[\bar{G}_t(\hat{x}_1)] + (1-\theta)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
\geq&\rho[\bar{G}_t(\hat{x}_1)] + (1-\rho)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)] \\
\geq&\rho[\bar{G}_t(\hat{x}_1)] + (1-\rho)[\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - \rho\hat{x}_1 - (1-\rho)y_t(x_2))] \\
\geq&\bar{G}_t(\rho\hat{x}_1 + (1-\rho)y_t(x_2)) \\
=&\bar{G}_t(\theta x_1 + (1-\theta)x_2) \\
\geq&\bar{v}_t^*(\theta x_1 + (1-\theta)x_2),
\end{aligned}
$$

where the first inequality is due to $\rho \geq \theta$ and $\bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1) \leq \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - \theta x_1 - (1-\theta)x_2)$. The second inequality is due to the subadditivity of $\bar{c}$. The third inequality is due to $\bar{G}_t$ being $\bar{c}$-convex. This completes the inductive proof. $\square$

Remark 1. *In contrast to the proof of K-convexity in Scarf (1959), our proof of the $\bar{c}$-convexity of the value function does not rely on any structural properties of the optimal policy.*

Based on the $\bar{c}$-convexity of the value function, we can characterize the structure of the optimal policy as follows. Define

$$\hat{S}_{i,t} = \min \arg \min_{y} [\bar{G}_t(y) + (c_i - c_n)y].$$

Let

$$\hat{s}_{n,t} = \max\{x | \bar{G}_t(x) > \bar{G}_t(\hat{S}_{n,t}) + K_n, \hat{S}_{n,t} - x \geq z_{n-1}\},$$

i.e., $\hat{s}_{n,t}$ is the largest value such that ordering up to $\hat{S}_{n,t}$ is preferable to not ordering (there always exists such $\hat{s}_{n,t}$ since $\bar{G}_t$ is $\bar{c}$-convex and $L_t(x) \to \infty$ as $x \to -\infty$). Also let $\underline{s}_{0,t}$ be the maximum z such that it is optimal to order a positive quantity for all $x \in [\hat{s}_{n,t}, z]$ in period t.

Theorem 11. *The optimal policy has the following properties.*

(1) The generalized (s, S) policy is optimal for $x < \underline{s}_{0,t}$.

(2) It is optimal not to order for $x \geq \hat{S}_{n,t}$.

(3) If it is optimal to order for $x \in (\underline{s}_{0,t}, \hat{S}_{n,t})$, then its optimal order-up-to level is less than $\hat{S}_{n,t}$.

Thus, except for the interval $(\underline{s}_{0,t}, \hat{S}_{n,t})$, the generalized (s, S) policy is optimal.

Proof. (1) First, we show that if $\bar{G}_t(x) > \bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x)$ and $\hat{S}_{n,t} - x \geq z_{n-1}$, then we must have $\bar{G}_t(x - \delta) > \bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x + \delta)$ for all $\delta > 0$. We show this by contradiction. Suppose that $\bar{G}_t(x_0) > \bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0)$ with $\hat{S}_{n,t} - x_0 \in (z_{n-1}, \infty)$ but $\bar{G}_t(x_0 - \delta) \leq \bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0 + \delta)$ for some $\delta > 0$. There exists some $\rho \in (0, 1)$ such that $\rho(x_0 - \delta) + (1 - \rho)\hat{S}_{n,t} = x_0$. Note that

$$\bar{c}(\hat{S}_{n,t} - x_0) = \rho\bar{c}(\hat{S}_{n,t} - x_0 + \delta) + (1 - \rho)\bar{c}(\hat{S}_{n,t} - x_0),$$

since $\bar{c}(\hat{S}_{n,t} - x_0 + \delta) = \bar{c}(\hat{S}_{n,t} - x_0) = K_n$ for $\hat{S}_{n,t} - x_0 \geq z_{n-1}$. Thus, we have

$$\begin{aligned}
&\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0) \\
=&\rho[\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0 + \delta)] + (1 - \rho)[\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0)] \\
\geq&\rho[\bar{G}_t(x_0 - \delta)] + (1 - \rho)[\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0)] \\
\geq&\bar{G}_t(x_0).
\end{aligned}$$

This implies that $\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0) \geq \bar{G}_t(x_0)$, which contradicts the fact that $\bar{G}_t(\hat{S}_{n,t}) + \bar{c}(\hat{S}_{n,t} - x_0) < \bar{G}_t(x_0)$.

Since

$$\hat{s}_{n,t} = \max\{x|\bar{G}_t(x) > \bar{G}_t(S_{n,t}) + K_n, \hat{S}_{n,t} - x \geq z_{n-1}\},$$

it follows that it is optimal to order a positive quantity for $x < \hat{s}_{n,t}$ based on the above result. By the definition of $\underline{s}_{0,t}$, it is optimal to order a positive quantity for all $x \in [\hat{s}_{n,t}, \underline{s}_{0,t})$ in period t. It follows that the generalized (s, S) policy is optimal on $(-\infty, \underline{s}_{0,t})$ based on Theorem 10.

(2) We show this result by contradiction. Suppose that $y_t(x_0) > x_0$ for some $x_0 > \hat{S}_{n,t}$. Then we must have

$$\bar{G}_t(x_0) > \bar{G}_t(y_t(x_0)) + \bar{c}(y_t(x_0) - x_0).$$

But we know that for any $x < x_0$ and $\theta \in (0,1)$ such that $\theta x + (1-\theta)y_t(x_0) = x_0$, we have

$$\theta \bar{G}_t(x) + (1-\theta)[\bar{G}_t(y_t(x_0)) + \bar{c}(y_t(x_0) - x_0)] \geq \bar{G}_t(x_0),$$

since $\bar{G}_t$ is $\bar{c}$-convex. It follows that $\bar{G}_t(x) > \bar{G}_t(x_0)$ and hence $\bar{G}_t(\hat{S}_{n,t}) > \bar{G}_t(x_0)$, which contradicts the fact that $\hat{S}_{n,t}$ minimizes $\bar{G}_t(x)$.

(3) is due to Theorem 9. $\square$

Remark 2. *In general, the generalized (s, S) policy may not be optimal over the interval $(\underline{s}_{0,t}, \hat{S}_{n,t})$. In the Appendix, we provide a counter example.*

We conclude this section by noting that the results regarding the structure of the optimal policy extend to the case where the ordering cost $c(x)$ in each period is a general increasing concave function. This follows from the fact that we can approximate, with arbitrary accuracy, an increasing concave function by a piecewise linear concave function.

4.4 Further Characterization of the Optimal Policy

In this section, we further characterize the optimal policy by showing that (1) the region over which the generalized (s, S) policy may not be optimal can be further reduced, (2) this region is increasing in $c_1 - c_n$, and (3) providing bounds on the optimal order-up-to levels $\hat{S}_{i,t}$.

First, let

$$\eta_{i,t} = \min \arg\min_y [\bar{L}_t(y) + (c_i - c_n)y],$$

i.e., $\eta_{i,t}$ is the global minimum of $\bar{L}_t(y) + (c_i - c_n)y$ for $i = 1, \cdots, n$ assuming it exists. It is clear that $\eta_{1,t} \leq \eta_{2,t} \leq \cdots \leq \eta_{n,t}$ since $c_1 > c_2 > \cdots > c_n$.

Assumption 2. $\eta_{n,1} \leq \eta_{n,2} \leq \cdots \leq \eta_{n,T}$.

Assumption 2 is satisfied if demands and costs are stationary.

First, we state a lemma which is useful in the further characterization of the optimal policy.

Lemma 13. *Under Assumption 2,*

(1) $\bar{v}_t^*(x)$ *is nonincreasing in* x *for any* $x \leq \eta_{n,t}$*, and*

(2) $\bar{G}_t(x_2) - \bar{G}_t(x_1) + (c_i - c_n)(x_2 - x_1) \leq \bar{L}_t(x_2) - \bar{L}_t(x_1) + (c_i - c_n)(x_2 - x_1) \leq 0$ *for* $x_1 \leq x_2 \leq \eta_{i,t}$.

Proof. (1) We show this result by induction. Observe that it is true for period T. Assume that it is true for period $t+1$. It is clear that $\beta E[\bar{v}_{t+1}^*(y - \xi_t)]$ is nonincreasing in y for any $y \leq \eta_{n,t+1}$ by the inductive assumption since $\xi_t \geq 0$. Since $\eta_{n,t} \leq \eta_{n,t+1}$, it follows that $\bar{G}_t(y)$ is also nonincreasing in y for $y \leq \eta_{n,t}$. Let $x_1 < x_2 \leq \eta_{n,t}$. Define $a \vee b = \max\{a, b\}$ for some real numbers a, b. Let $y_t(x)$ be the optimal order-up-to level under state x. We have

$$
\begin{aligned}
\bar{v}_t^*(x_1) =& \bar{G}_t(y_t(x_1)) + \bar{c}(y_t(x_1) - x_1) \\
\geq& \bar{G}_t(y_t(x_1) \vee x_2) + \bar{c}(y_t(x_1) \vee x_2 - x_2) \\
\geq& \min_{y \geq x_2} [\bar{G}_t(y) + \bar{c}(y - x_2)] \\
=& \bar{v}_t^*(x_2).
\end{aligned}
$$

The first inequality is due to the fact that $\bar{G}_t(y)$ is nonincreasing in y for $y \le \eta_{n,t}$ and the fact that $\bar{c}(y_t(x_1) - x_1) \ge \bar{c}(y_t(x_1) \vee x_2 - x_2)$ for any $x_2 \ge x_1$. The second inequality is due to the fact that $y_t(x_1) \vee x_2 \ge x_2$.

(2) For $x_1 \le x_2 \le \eta_{i,t}$, we have

$$\begin{aligned}
&\bar{G}_t(x_2) - \bar{G}_t(x_1) + (c_i - c_n)(x_2 - x_1) \\
=&\bar{L}_t(x_2) - \bar{L}_t(x_1) + (c_i - c_n)(x_2 - x_1) + \beta E[\bar{v}^*_{t+1}(x_2 - \xi_t) - \bar{v}^*_{t+1}(x_1 - \xi_t)] \\
\le&\bar{L}_t(x_2) - \bar{L}_t(x_1) + (c_i - c_n)(x_2 - x_1) \le 0.
\end{aligned}$$

The first inequality is due to the fact that $\bar{v}^*_{t+1}(x)$ is nonincreasing in x for any $x \le \eta_{n,t+1}$. The second inequality is due to the fact that $\bar{L}_t(x) + (c_i - c_n)$ is nonincreasing in x for any $x \le \eta_{i,t}$ and the fact that $\eta_{i,t} \le \eta_{n,t}$. □

Proposition 5. *Under Assumption 2, the optimal policy can be further characterized as follows.*

(1) The generalized (s, S) policy is optimal for $x < \eta_{1,t}$.

(2) It is optimal not to order for $x > \eta_{n,t}$.

(3) $\eta_{1,t} - \eta_{n,t}$ is increasing in $c_1 - c_n$.

(4) $\hat{S}_{i,t} \ge \eta_{i,t}$ for $i = 1, \cdots, n$.

Proof. (1) We show that if it is optimal to order at some $x_2 < \eta_{1,t}$ it must be optimal to order at any x_1 such that $x_1 < x_2$. Then according to Theorem 10, the generalized (s, S) policy is optimal for $x < \eta_{1,t}$. Suppose that $c(y_t(x_2) - x_2) = K_i + c_i(y_t(x_2) - x_2)$. Then we have

$$\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_1) \le \bar{G}_t(y_t(x_2)) + K_i + c_i(y_t(x_2) - x_1) - c_n(y_t(x_2) - x_1)$$

$$
\begin{aligned}
&= \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + (c_i - c_n)(x_2 - x_1) \\
&< \bar{G}(x_2) + (c_i - c_n)(x_2 - x_1) \\
&\leq \bar{G}(x_1).
\end{aligned}
$$

The first inequality is due to the fact that $\bar{c}(x) = \min_i\{K_i + c_i x\} - c_n x$. The second inequality is due to the fact that $\bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) < \bar{G}(x_2)$, i.e., it is optimal to order at x_2. The last inequality is due to $\bar{G}_t(x) + (c_i - c_n)x$ being nonincreasing for $x < \eta_{1,t}$. Thus, it is optimal to order at x_1.

(2) First, we show $\bar{v}_t^*(x_2) + \bar{c}(x_2 - x_1) \geq \bar{v}_t^*(x_1)$ for $x_2 > x_1$. Note that

$$
\begin{aligned}
\bar{v}_t^*(x_2) + \bar{c}(x_2 - x_1) =& \bar{G}_t^*(y_t(x_2)) + \bar{c}(y_t(x_2) - x_2) + \bar{c}(x_2 - x_1) \\
\geq& \bar{G}_t(y_t(x_2)) + \bar{c}(y_t(x_2) - x_1) \\
\geq& \bar{v}_t^*(x_1).
\end{aligned}
$$

The first inequality is due to the subadditivity of $\bar{c}$. Since $\bar{v}_t^*(x_2) + \bar{c}(x_2 - x_1) \geq \bar{v}_t^*(x_1)$ and $\bar{L}_t(x_2) \geq \bar{L}_t(x_1)$ for $x_2 > x_1 \geq \eta_{n,t}$, as a result we must have $\bar{G}_t(x_2) + \bar{c}(x_2 - x_1) \geq \bar{G}_t(x_1)$ for $x_2 > x_1 \geq \eta_{n,t}$. Hence, it is optimal to order nothing for $x \geq \eta_{n,t}$.

(3) This result follows directly from the definition of $\eta_{i,t}$.

(4) This result is due to property (2) of Lemma 13. □

Note that results (2) and (3) of the above proposition hold even without Assumption 2.

As we can see, under Assumption 2, the generalized (s, S) policy is optimal except for the interval $(\eta_{1,t}, \eta_{n,t})$ and the width of this interval is increasing in $c_1 - c_n$, implying that when $c_1 - c_n$ is small, the width of $(\eta_{1,t}, \eta_{n,t})$ is also small.

For $K_1 = 0$, we can further reduce the interval over which the generalized (s, S) policy may not be optimal. Let $\bar{c}_1(x) \equiv c(x) - c_1 x$,

$$\bar{L}_{1,t}(y) = (1-\alpha)c_1[y - E(\xi_t)] + c_1 E\xi_t + Eg_t(y - \xi_t),$$

and

$$\bar{G}_{1,t}(y) = \bar{L}_{1,t}(y) + \alpha E\bar{v}^*_{1,t+1}(y - \xi_t).$$

Then we can reformulate the dynamic recursion as follows.

$$\begin{aligned}\bar{v}^*_{1,t}(x) &= \min_{y \geq x} E\{\bar{c}_1(y-x) + (1-\alpha)c_1[y - \xi_t] + c_1\xi_t + g_t(y - \xi_t) \\ &\qquad + \alpha\bar{v}^*_{1,t+1}(y - \xi_t)\} \\ &= \min_{y \geq x}\{\bar{c}_1(y-x) + \bar{G}_{1,t}(y)\},\end{aligned}$$

with $\bar{v}_{1,T+1}(x) = c_1 x$. Note that $\bar{c}_1(x) \leq 0$ for all $x \geq 0$ and

$$\hat{S}_{1,t} = \min \arg\min_y [\bar{G}_t(y) + (c_1 - c_n)y] = \min \arg\min_y \bar{G}_{1,t}(y).$$

This leads to the following proposition.

Proposition 6. *If $K_1 = 0$, then the region over which we cannot fully characterize the structure of the optimal policy can be reduced to $(\hat{S}_{1,t}, \eta_{n,t})$.*

Proof. We first show that the generalized (s, S) policy is optimal for $x < \hat{S}_{1,t}$. Note that the optimal ordering decision is given by

$$\min_{y \geq x}[\bar{G}_{1,t}(y) + \bar{c}_1(y - x)].$$

Since $\bar{c}_1(y-x) \leq 0$ and $\bar{G}_{1,t}(x) > \bar{G}_{1,t}(\hat{S}_{1,t})$ for $x < \hat{S}_{1,t}$, we have

$$\bar{G}_{1,t}(\hat{S}_{1,t}) + \bar{c}_1(\hat{S}_{1,t} - x) < \bar{G}_{1,t}(x)$$

for $x < \hat{S}_{1,t}$. Hence it is optimal to order a positive quantity under x for all $x < \hat{S}_{1,t}$. In turn, this implies, based on Theorem 10, that the generalized (s, S) policy is optimal for $x < \hat{S}_{1,t}$. By a similar argument as in Proposition 5, we can show that it is optimal not to order for $x > \eta_{n,t}$. □

Note that since $\hat{S}_{1,t} \geq \eta_{1,t}$, the region in which the general (s, S) policy may not be optimal is reduced from $(\eta_{1,t}, \eta_{n,t})$ to $(\hat{S}_{1,t}, \eta_{n,t})$. Also,note that if $\hat{S}_{1,t} \geq \eta_{n,t}$, then a generalized (s, S) policy is optimal over the entire state space. Finally, note that we do not need Assumption 2 for Proposition 6.

4.5 The Optimality of the Generalized (s, S) Policy

In this section, we show that a generalized (s, S) policy is optimal for all regions of the state space if the single period inventory cost satisfies the following assumption.

Assumption 3. *The following inequality holds for all x_1, x_2 such that $|x_1 - x_2| \geq 1$ and $\theta \in (0, 1)$*

$$\theta \bar{L}_t(x_1) + (1-\theta)\bar{L}_t(x_2) \geq \bar{L}_t(\theta x_1 + (1-\theta)x_2) + \theta(1-\theta)(c_1 - c_n)|x_2 - x_1|.$$

Remark 3. *This assumption is related to the concept of strong convexity. A function f is called strongly convex with parameter $m > 0$ if*

$$\theta f(x_1) + (1-\theta)f(x_2) \geq f(x_1 + (1-\theta)x_2) + \frac{1}{2}m|x_1 - x_2|^2,$$

which is equivalent to $f(x) - \frac{1}{2}mx^2$ *being convex (See Rockafellar (2015)). It is clear that Assumption 3 is stronger than regular convexity but weaker than strong convexity with parameter* $2(c_1 - c_n)$ *since* $|x_2 - x_1|^2 \geq |x_2 - x_1|$ *for* $|x_2 - x_1| \geq 1$.

One example that satisfies Assumption 3 is the case in which $g_t(x) = h_t \cdot (x^+)^2 + b_t \cdot (x^-)^2$, where $h_t \geq c_1 - c_n$ and $b_t \geq c_1 - c_n$.

Theorem 12. *If Assumption 3 is true and the minimum ordering quantity is always larger than or equal to 1, then the generalized* (s, S) *policy is optimal.*

Proof. We show this result by contradiction. Suppose that it is optimal to order under state x_0 and it is not optimal to order under state $x_0 - \delta$ for some $\delta > 0$. Let y be the optimal order-up-to level under state x_0. Suppose that $y - x_0 \in (z_{i-1}, z_i]$ for some i. Then we have $\bar{G}_t(y) + \bar{c}(y - x_0) < \bar{G}_t(x_0)$ and $\bar{G}_t(y) + \bar{c}(y - x_0 + \delta) \geq \bar{G}_t(x_0 - \delta)$. There exists some $\rho_1 \in (0, 1)$ such that $\rho_1(x_0 - \delta) + (1 - \rho_1)y = x_0$. Note that for this ρ_1, we have

$$
\begin{aligned}
&\bar{c}(y - x_0 + \delta) - (1 - \rho_1)(c_i - c_n)(y - x_0 + \delta) \\
\leq &K_i + (c_i - c_n)(y - x_0 + \delta) - (1 - \rho_1)(c_i - c_n)(y - x_0 + \delta) \\
= &K_i + (c_i - c_n)(y - x_0) \\
= &\bar{c}(y - x_0),
\end{aligned}
$$

which implies that

$$
\bar{c}(y - x_0) \geq \rho_1[\bar{c}(y - x_0 + \delta) - (1 - \rho_1)(c_i - c_n)(y - x_0 + \delta)] + (1 - \rho_1)[\bar{c}(y - x_0)].
$$

As a result, we have

$$
\bar{G}_t(y) + \bar{c}(y - x_0)
$$

$$
\begin{aligned}
&\geq \rho_1[\bar{G}_t(y) + \bar{c}(y - x_0 + \delta) - (1-\rho_1)(c_i - c_n)(y - x_0 + \delta)] \\
&\quad + (1-\rho_1)[\bar{G}_t(y) + \bar{c}(y - x_0)] \\
&\geq \rho_1[\bar{G}_t(x_0 - \delta) - (1-\rho_1)(c_i - c_n)(y - x_0 + \delta)] + (1-\rho_1)[\bar{G}_t(y) + \bar{c}(y - x_0)] \\
&\geq \rho_1[\bar{L}_t(x_0 - \delta) - (1-\rho_1)(c_1 - c_n)(y - x_0 + \delta)] + (1-\rho_1)\bar{L}_t(y) \\
&\quad + \alpha E[\rho_1 \bar{v}^*_{t+1}(x_0 - \delta - \xi_t) + (1-\rho_1)(\bar{v}^*_{t+1}(y - \xi_t) + \bar{c}(y - x_0))] \\
&\geq \bar{L}_t(x_0) + \alpha E \bar{v}^*_{t+1}(x_0 - \xi_t) \\
&= \bar{G}_t(x_0).
\end{aligned}
$$

The third inequality is due to the fact that $c_1 \geq c_i$ for all i. The last inequality is due to Assumption 3 and the $\bar{c}$-convexity of $\bar{v}^*_{t+1}$. This contradicts the fact that $\bar{G}_t(y) + \bar{c}(y - x_0) < \bar{G}_t(x_0)$.

As a result, if it is optimal to order under x_2, then it must be optimal to order under x_1 for any $x_1 < x_2$. Hence, from Theorem 10, the generalized (s, S) policy is optimal. □

The assumption that the minimum ordering quantity is larger than or equal to 1 applies to the discrete demand case (all related proofs can be modified to accommodate the discrete demand case) and to the case in which K_1 is sufficiently large.

Note that under Assumption 2, Assumption 3 can be relaxed as follows: for all $x_1, x_2 \in (\eta_{1,t}, \eta_{n,t})$ such that $|x_1 - x_2| \geq 1$ and $\theta \in (0, 1)$, the following inequality holds

$$\theta \bar{L}_t(x_1) + (1-\theta)\bar{L}_t(x_2) \geq \bar{L}_t(\theta x_1 + (1-\theta)x_2) + \theta(1-\theta)(c_1 - c_n)|x_2 - x_1|.$$

This is due to the fact that under Assumption 2, we only need to check whether Condition 1 holds for the interval $(\eta_{1,t}, \eta_{n,t})$.

4.6 Extensions to Other Settings

In this section, we briefly explain how our approach can be extended to time-varying ordering costs case, the lost sales case, and the non-zero leadtime case with backordering.

First, our result can be extended to the following time varying piecewise concave ordering costs: $c_t(x) = \min_i\{K_{i,t}\delta(x) + c_{i,t}x\}$ with $\alpha\bar{c}_{t+1}(x) \leq \bar{c}_t(x)$ for all $x \geq 0$ and all t, where subscript t denotes the dependency of the cost parameters on period t. By similar arguments as in the stationary ordering cost case, we can characterize the structure of the optimal policy and show that, except for a bounded region, it is a generalized (s, S) policy.

Next, we consider the case where unfulfilled demand is lost instead of backordered. Let p_t be the unit lost sales cost and h_t be the unit holding cost in period t. Let $v_t^*(x)$ be the value function in period t with starting inventory level x. Then, the function G_t can be modified as follows

$$G_t(y) = E(h_t[y-\xi]^+ + p_t[\xi - y]^+) + \alpha E v_{t+1}^*([y-\xi]^+); \tag{4.9}$$

and the optimality equation rewritten as

$$v_t^*(x) = \min_{y \geq x}[c(y-x) + G_t(y)]. \tag{4.10}$$

Using similar analysis to the one for the backordering case, we can show here too that the structure of the optimal policy, except for a bounded region, is also a generalized (s, S) policy.

Finally, consider the case with fixed leadtime l and backordering, where an order placed in period t is delivered in period $t + l$. Let x be the current inventory in stock, and x_i be the amount of inventory delivered i periods later, where $i = 1, \cdots, l-1$. By a standard transformation as in Zipkin (2000), we can show that the value function in each period only depends on the starting aggregate inventory level $x + x_1 + \cdots + x_{l-1}$. Then, we can carry out similar analysis to the one for zero leadtime case and show that the optimal policy is again a generalized (s, S) except for a bounded region.

4.7 Appendix: A Counter Example

Here we provide a counterexample illustrating that the generalized (s, S) policy may not be optimal over the entire state space. In particular, we show that in this example there exist x_1 and x_2 such that $x_1 < x_2$ and it is optimal to order at x_2 but not optimal to order at x_1.

We consider a 2-period problem with the ordering cost $c(x) = \min\{2x, x+192\}$, i.e., $c_1 = 2$, $c_2 = 1$, $K_1 = 0$ and $K_2 = 192$. Let

$$g_t(x - \xi_t) = \begin{cases} (x - \xi_t), & \text{if } x \geq \xi_t, \\ 3(\xi_t - x), & \text{otherwise,} \end{cases}$$

i.e., the unit holding cost is 1 and the unit backorder cost is 3 in every period. Let $\alpha = 1$. Demands in different periods are *i.i.d.* with density function ϕ given as follows:

$$\phi(x) = \begin{cases} \frac{1}{1280}(x-270), & \text{if } 270 < x \le 280, \\ \frac{1}{1280}[-\frac{8}{5}(x-280)+10], & \text{if } 280 < x \le 285, \\ \frac{1}{640}, & \text{if } 285 < x \le 405, \\ \frac{1}{1280}[2(x-405)+2], & \text{if } 405 < x \le 410, \\ \frac{1}{1280}[-2(x-410)+12], & \text{if } 410 < x \le 415, \\ \frac{1}{640}, & \text{if } 415 < x \le 540, \\ \frac{1}{1280}[-\frac{3}{20}(x-540)+2], & \text{if } 540 < x \le 550, \\ \frac{1}{1280}[\frac{2}{5}(x-550)+\frac{1}{2}], & \text{if } 550 < x \le 555, \\ \frac{5}{2560}, & \text{if } 555 < x \le 803, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

It is clear that ϕ is not a Polya density function. Denote the distribution function of demand by Φ. Then $\Phi^{-1}(0.25) = 405$, $\Phi^{-1}(0.5) = 540$ and $\Phi^{-1}(0.75) = 675$. It is easy to check that in the second period, the optimal policy is a generalized (s, S) policy with $s_1 = S_1 = 405$, $s_2 = 270$ and $S_2 = 540$. Then we have

$$v_2^*(x) = \begin{cases} 192 + (540-x) + E(540-\xi_t), & \text{if } x \le 270, \\ 2(405-x) + E(450-\xi_t), & \text{if } 270 < x \le 450, \\ E(x-\xi_t) & \text{otherwise.} \end{cases}$$

In the first period, we can check that for any $y \ge 540$, $H_1(540) \le c(y-540)+H_1(y)$ and $H(550) > c(555-550)+H_1(555)$. Thus, it is not optimal to order at $x_1 = 540$

and it is optimal to order at $x_2 = 550$. Therefore, the optimal policy in the first period is not a generalized (s, S) policy.

In this example, we can show that the optimal policy in the first period is

$$y_1^*(x) = \begin{cases} 717, & \text{if } x \leq 443, \\ 540, & \text{if } 443 < x \leq 540, \\ x, & \text{if } 540 < x \leq 545, \\ 555, & \text{if } 545 < x \leq 555, \\ x, & \text{otherwise.} \end{cases}$$

Thus, the generalized (s, S) policy is optimal for $x < 540$ and it is optimal to order nothing for $x > 555$. This is consistent with our statement about the optimal policy in Theorem 11. Finally, we note that a counter example for $K_1 > 0$ can also be found with the same demand density function.

Chapter 5

Conclusions and Other Research Projects

In this chapter, we provide conclusions and future research directions on the work presented in Chapters 2 and 3. We also briefly discuss other research projects.

5.1 On the Impact of Input Price Variability and Correlation in Stochastic Inventory Systems

In Section 2, we examined the impact of input price variability on expected cost in inventory systems with stochastic demand and stochastic input prices. For a general class of such systems, we showed that higher input price variability leads to lower expected cost. We showed that this is true for a wide range of assumptions regarding price evolution, including i.i.d. prices and prices that evolve according to a Markovian martingale. We also showed that this is true for systems with both single and multiple periods. We described how the impact of price variability on expected cost can be traced to the concavity of the cost function in input price, which is itself a consequence of the flexibility in adjusting the order quantity as prices vary. In addition, we examined the impact of price correlation over time and across inputs. We found that expected cost is increasing in price correlation over time and decreasing in price correlation across components. Numerical results suggest that higher correlation of prices over time diminishes the benefit derived from price variability while higher correlation of prices across components enhances it.

There are several avenues for future research. It would be useful to extend the analysis to broader classes of systems, including systems with multiple production stages where different components may be needed at different stages. In particular, it would be of interest to investigate how the position of a component in the production process affects the benefit derived from the variability in its input price (e.g., is price variability more beneficial for components that are upstream in the production process or is it more so for components that are downstream?). It would also be useful to consider settings in which there is variability in both the input purchase price and the output selling price. For example, a firm may purchase input from one spot market and sell output to another, with the firm observing both input and output prices at the beginning of each period and then deciding on how much input to buy and how much output to produce and sell. Lastly, it would be valuable to extend our analysis to settings where the firm may not be risk neutral and to account for its attitude toward risk by studying a decision criterion other than expected value.

5.2 Optimal Control of an Inventory System with Stochastic and Independent Leadtimes

In Section 3, we studied an inventory system with stochastic and independent leadtimes. For the case of exponentially distributed leadtimes, we resolved the open question regarding the structure of the optimal policy. In particular, we showed that the optimal policy is specified by a threshold function that is non-increasing in the inventory level. We showed that once the threshold function starts to decrease it continues to do so at a rate that is greater than or equal to one. This implies that the threshold function can be fully described by at most m parameters. Taking advantage of this structure, we provided an efficient algorithm for computing these parameters and the corresponding optimal cost. Also, inspired by the structure of the optimal policy, we investigated two plausible heuristics, as alternatives to the optimal, and examined their performance for a wide range of parameter values. We showed that the heuristics can perform poorly for certain parameter values. Finally, we extended our analysis to systems with lost sales and to systems where order cancellations are possible.

There are several possible avenues for future research. It would be of interest to extend the results to more general settings with respect to the distribution of demand and leadtime. It would also be of interest to extend the results to settings where leadtimes are not identical, as in systems with heterogeneous production facilities or delivery modes. We expect the analysis to be much more difficult in those cases, but there may be special cases for which at least partial characterization of the optimal policy is possible.

5.3 Other Research Projects

5.3.1 Managing Stochastic Inventory Systems with Scarce Resources

We consider a production-inventory system where the input material is scarce and its consumption is subject to a limit over a specified *compliance* period. Examples of such settings are many and include those where limits are imposed on the harvesting of forest products, the hunting and fishing of wild life, and the mining of rare minerals and metals. They also include settings where limits are imposed on the consumption of water or the emission of harmful pollution. In such cases, the amount that can be produced over the compliance period, which may consist of multiple production periods, cannot exceed the specified limit. Imposing such a limit introduces capacity dependencies across production periods, absent from traditional models where capacity constraints are imposed on individual periods. In particular, capacity in each period depends on the production decisions in previous periods and affects the capacity available in future periods. The objective of the system manager, in the face of stochastic demand, is to minimize the sum of inventory holding and shortage costs over a planning horizon consisting of one or more compliance review periods.

We formulated the problem as a stochastic dynamic program with a two-dimensional state space: on-hand inventory level and remaining capacity. We considered an extended state-space version of the problem and showed that this modified version of the problem reduces to a one-dimensional problem. We described various properties of the optimal policy for the modified version of the problem and then showed that these properties also hold for the original problem. We then used these properties to characterize the structure of the optimal policy for the original problem. In particular, we showed that the optimal ordering policy is specified by dynamic thresholds that depend on both the on-hand inventory level and the remaining capacity but only via the sum of these two quantities. In addition, we characterized the impact of the capacity constraint and showed that the expected optimal cost is convex with respect to the remaining capacity, implying that there is diminishing value to capacity. We provided numerical results that examine the tradeoff between the expected optimal cost and the expected cumulative amount ordered, and discussed how both are affected by problem parameters. We evaluated the performance of three plausible heuristics that are simpler to compute and implement. We showed that, although the heuristics can be quite effective under some settings, they can also perform poorly under others. We then considered the problem of jointly optimizing for capacity and inventory control and showed that the associated total cost is convex in capacity and, therefore, the optimal capacity can be computed easily. We also showed that the optimal capacity can be quite sensitive to the price of capacity initially, with even modest prices leading to a significant reduction in the capacity purchased. Finally, we considered various extensions to the original model and show that the optimal policy of the extended models has similar structure.

There are several possible avenues for future research. It would be useful to generalize the results to a broader class of systems, including multi-stage systems where each stage may have its own cumulative ordering/production capacity constraint. It would also be useful to study systems with both cumulative and period capacity constraints, where the period constraint may be due to production capacity limits while the cumulative constraint due to limits on input material availability or negative environmental externalities. Moreover, it would be interesting to compare systems where the cumulative amount ordered over the planning horizon is limited via an explicit constraint (as considered in this paper) to systems where this is achieved via imposing a penalty (or a tax) on ordering, or to systems where there are both a reward and penalty with production depending on whether the cumulative quantity falls below or over a specified threshold. For more details, please refer to Benjaafar et al. (2015b).

5.3.2 Stochastic Inventory Systems with Discount-driven Backorders

Stockouts are quite common for consumer products due to the variability in demand. Most of the inventory literature assumes either backorders or lost sales when stockouts occur. Existing literature that considers both backorders and lost sales assumes that when the on-hand inventory is not available to fulfill current demand, the inventory manager could decide whether to backlog or to reject some or all the demand. However, in practice, when stockouts occur, it is the customer herself who decides whether to wait for the product or walk away. The seller can offer a price discount to incentivize customers to wait for the product in the case of a stockout in order to mitigate lost sales.

In this study, we consider a multi-period stochastic inventory systems with both backorders and lost sales. In the case of a stockout, customers may choose to either wait for the product which corresponds to backorders, or walk away which corresponds to lost sales. We assume that a fraction of the unfulfilled customers are willing to wait and this fraction depends on a discount the seller offers. The higher the discount is, the higher this fraction is, i.e., the more customers are willing to wait. We show that for a given discount, the optimal policy is a base stock policy. The optimal cost is convex in the discount and therefore the optimal discount can be computed easily. We also consider a continuous version of this problem. In this continuous review model, the probability that a customer would wait for the product is increasing in the discount. Again, we characterize the structure of the optimal policy and provide some managerial insights. These results are similar to those in the periodic review model.

We are currently extending the periodic review model by assuming that the backordering process is probabilistic, i.e., there is a range of possible outcomes, including with some positive probability that no customer would be willing to be backordered. In other words, given a discount, there is a distribution for the number of customers who are willing to be backordered. To do so, we introduce a notion of customer valuation of waiting (backordering), which is a random variable. If this value is less than the discount, the customer waits; otherwise the customer does not wait. This allows us to endogenize the probability of backordering and to study how the distribution of valuations affects the optimal discount and the corresponding base stock level. We are also extending the analysis to settings where lost sales in one period affect demand in future periods and those where a customer's probability of waiting is affected by current backorder levels. For more details, please refer to Chen et al. (2015b).

References

R. W. Anderson and J.-P. Danthine. Cross hedging. *The Journal of Political Economy*, 89(6):1182–1196, 1981.

R. W. Anderson and J.-P. Danthine. The time pattern of hedging and the volatility of futures prices. *The Review of Economic Studies*, 50(2):249–266, 1983.

R. N. Batra and A. Ullah. Competitive firm and the theory of input demand under price uncertainty. *The Journal of Political Economy*, 82(3):537–548, 1974.

S. Benjaafar and M. ElHafsi. Production and inventory control of a single product assemble-to-order system with multiple customer classes. *Management Science*, 52(12):1896–1912, 2006.

S. Benjaafar, M. Elhafsi, W. Zhou, and C. Lee. Optimal production and inventory control of multi-stage assembly systems with multiple demand classes. *Operations Research*, 59(2):522–529, 2011.

S. Benjaafar, D. Chen, and M. Elhafsi. Optimal control of an inventory system with stochastic and independent leadtimes. *Working paper*, 2015a.

S. Benjaafar, D. Chen, and R. Wang. Managing stochastic inventory systems with scarce resources. *Working paper, University of Minnesota*, 2015b.

P. Berling and V. Martínez-de-Albéniz. Optimal inventory policies when purchase price and demand are stochastic. *Operations Research*, 59(1):109–124, 2011.

P. Berling and K. Rosling. The effects of financial risks on inventory policy. *Management Science*, 51(12):1804–1815, 2005.

D. P. Bertsekas. Dynamic programming and optimal control, vol. ii. 2007.

O. Boyabatlı, P. R. Kleindorfer, and S. R. Koontz. Integrating long-term and short-term contracting in beef supply chains. *Management Science*, 57(10): 1771–1787, 2011.

J. R. Bradley and L. W. Robinson. Improved base-stock approximations for independent stochastic lead times with order crossover. *Manufacturing & Service Operations Management*, 7(4):319–329, 2005.

J. A. Buzacott and J. G. Shanthikumar. *Stochastic models of manufacturing systems*, volume 4. Prentice Hall Englewood Cliffs, NJ.

J. Casassus and P. Collin-Dufresne. Stochastic convenience yield implied from commodity futures and interest rates. *The Journal of Finance*, 60(5):2283–2331, 2005.

R. Cavazos-Cadena and L. I. Sennott. Comparing recent assumptions for the existence of average optimal stationary policies. *Operations research letters*, 11 (1):33–37, 1992.

X. Chao and P. H. Zipkin. Optimal policy for a periodic-review inventory system under a supply capacity contract. *Operations Research*, 56(1):59–68, 2008.

D. Chen, S. Benjaafar, and W. Cooper. On the impact of input price variability and correlation in stochastic inventory systems. *Working paper, University of Minnesota*, 2015a.

D. Chen, R. Wang, and S. Benjaafar. Stochastic inventory systems with discount-driven backorders. *Working paper, University of Minnesota*, 2015b.

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

X. Chen, Y. Zhang, and S. X. Zhou. Preservation of quasi-k-concavity and its applications. *Operations Research*, 58(4-Part-1):1012–1016, 2010.

Y. Chen, W. Xue, and J. Yang. Technical note-optimal inventory policy in the presence of a long-term supplier and a spot market. *Operations Research*, 61 (1):88–97, 2013.

F. De Vericourt, F. Karaesmen, and Y. Dallery. Optimal stock allocation for a capacitated supply system. *Management Science*, 48(11):1486–1501, 2002.

S. K. Devalkar, R. Anupindi, and A. Sinha. Integrated optimization of procurement, processing, and trade of commodities. *Operations Research*, 59 (6):1369–1381, 2011.

E. J. Fox, R. Metters, and J. Semple. Optimal inventory policy with two suppliers. *Operations Research*, 54(2):389–393, 2006.

S. Gavirneni. Periodic review inventory control with fluctuating purchasing costs. *Operations Research Letters*, 32(4):374–379, 2004.

R. Gibson and E. S. Schwartz. Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance*, 45(3):959–976, 1990.

A. Goel and G. J. Gutierrez. Integrating commodity markets in the optimal procurement policies of a stochastic inventory system. *Available at SSRN 930486*, 2006.

A. Goel and G. J. Gutierrez. Multiechelon procurement and distribution policies for traded commodities. *Management Science*, 57(12):2228–2244, 2011.

K. Golabi. Optimal inventory policies when ordering prices are random. *Operations Research*, 33(3):575–588, 1985.

X. Guo, P. Kaminsky, P. Tomecek, and M. Yuen. Optimal spot market inventory strategies in the presence of cost and price risk. *Mathematical Methods of Operations Research*, 73(1):109–137, 2011.

A. Y. Ha. Optimal dynamic scheduling policy for a make-to-stock production system. *Operations Research*, 45(1):42–53, 1997.

Ç. Haksöz and S. Seshadri. Supply chain operations in the presence of a spot market: a review with discussion. *Journal of the Operational Research Society*, 58(11):1412–1429, 2007.

X. X. He, S. H. Xu, J. K. Ord, and J. C. Hayya. An inventory model with order crossover. *Operations research*, 46(3-supplement-3):S112–S119, 1998.

T.-H. Ho, C. S. Tang, and D. R. Bell. Rational shopping behavior and the option value of variable pricing. *Management Science*, 44(12-Part-2):S145–S160, 1998.

K. Inderfurth and P. Kelle. Capacity reservation under spot market price uncertainty. *International Journal of Production Economics*, 133(1):272–279, 2011.

G. Janakiraman and S. Seshadri. Parametric concavity in stochastic dynamic programs. *Computers & Industrial Engineering*, 61(1):98–102, 2011.

B. A. Kalymon. Stochastic prices in a single-item inventory purchasing model. *Operations Research*, 19(6):1434–1458, 1971.

A. Kamara. Production flexibility, stochastic separation, hedging, and futures prices. *Review of Financial Studies*, 6(4):935–957, 1993.

S. Karlin. One stage inventory models with uncertainty. *Studies in the mathematical theory of inventory and production*, pages 109–134, 1958.

V. Kulkarni and K. Yan. Production-inventory systems in stochastic environment and stochastic lead times. *Queueing Systems*, 70(3):207–231, 2012.

V. Martínez-de-Albéniz and J. V. Simón. A capacitated commodity trading model with market power. *Working Paper, IESE Business School, University of Navarra, Barcelona, Spain*, 2009.

J. Meyer. Two-moment decision models and expected utility maximization. *The American Economic Review*, 77(3):421–430, 1987.

A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks.* Wiley, 2002.

K. Muthuraman, S. Seshadri, and Q. Wu. Inventory management with stochastic lead times. *Mathematics of Operations Research*, 40(2):302–327, 2014.

X. Nie, T. Boyacı, M. Gümüs, S. Ray, and D. Zhang. Joint procurement and demand-side bidding strategies under price volatility. *Working Paper*, 2014.

S. J. Park, G. Lai, and S. Seshadri. The impacts of inventory sharing on commodity procurement. *Available at SSRN 1998914*, 2012.

E. L. Plambeck and T. A. Taylor. On the value of input efficiency, capacity efficiency, and the flexibility to rebalance them. *Manufacturing & Service Operations Management*, 15(4):630–639, 2013.

E. L. Porteus. On the optimality of generalized (s, S) policies. *Management Science*, 17(7):411–426, 1971.

E. L. Porteus. The optimality of generalized (s, S) policies under uniform demand densities. *Management Science*, 18(11):644–646, 1972.

M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming.* John Wiley & Sons, 2014.

L. W. Robinson and J. R. Bradley. Note-further improvements on base-stock approximations for independent stochastic lead times with order crossover. *Manufacturing & Service Operations Management*, 10(2):325–327, 2008.

L. W. Robinson, J. R. Bradley, and L. J. Thomas. Consequences of order crossover under order-up-to inventory policies. *Manufacturing & Service Operations Management*, 3(3):175–188, 2001.

R. T. Rockafellar. *Convex analysis*. Princeton university press, 2015.

A. Sandmo. On the theory of the competitive firm under price uncertainty. *The American Economic Review*, 61(1):65–73, 1971.

H. Scarf. The optimality of (s, S) policies in the dynamic inventory problem. *Mathematical Methods in the Social Sciences*, 1959.

H. E. Scarf. A survey of analytic techniques in inventory theory. *Multistage inventory models and techniques*, 7:185–225, 1963.

E. Schwartz and J. E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.

N. Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467, 2010.

N. Secomandi and S. Kekre. Optimal energy procurement in spot and forward markets. *Manufacturing & Service Operations Management*, 2014.

M. Shaked and J. G. Shanthikumar. *Stochastic Orders*. Springer, 2007.

J.-S. Song and P. Zipkin. Inventory control in a fluctuating demand environment. *Operations Research*, 41(2):351–370, 1993.

R. R. Weber and S. Stidham Jr. Optimal control of service rates in networks of queues. *Advances in applied probability*, pages 202–218, 1987.

J. Yang and Y. Xia. Acquisition management under fluctuating raw material prices. *Production and Operations Management*, 18(2):212–225, 2009.

J. Yi and A. Scheller-Wolf. Dual sourcing from a regular supplier and a spot market. *Working Paper, GSIA-Carnegie Mellon University*, 2003.

Y. Yu, S. Benjaafar, and D. Chen. Optimal policies for inventory systems with concave ordering costs. *Working paper*, 2015.

R. Zhang. Pricing and inventory management under fluctuating costs. *Technical Report, Olin Business School, Washington University in St. Louis*, 2012.

P. H. Zipkin. *Foundations of inventory management*, volume 2. McGraw-Hill New York, 2000.