Supply Chain and Revenue Management for Online Retailing

by

Yuchen Jiang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2018

Doctoral Committee:

Associate Professor Siqian Shen, Co-Chair
Assistant Professor Cong Shi, Co-Chair
Professor Roman Kapuscinski
Assistant Professor Viswanath Nagarajan

Yuchen Jiang
ycjiang@umich.edu
ORCID iD: 0000-0003-3591-7970

ACKNOWLEDGEMENTS

First of all, I would like to thank my advisors Prof. Cong Shi and Prof. Siqian Shen for their years of guidance and mentorship throughout my Ph.D. studies. Without their support and help, this dissertation would not have been possible. I would also like to thank my committee members Prof. Roman Kapuscinski and Prof. Viswanath Nagarajan, for their helpful discussions and feedback.

Moreover, I would like to express my thanks to the faculty members and fellow Ph.D. students in the Department of Industrial and Operations Engineering. Your lectures, seminars, individual discussions and support help me in various ways throughout my Ph.D. studies. I also owe my thanks to staff members of the department for their assistance and help.

In addition, I appreciate the friendship from Prof. Siqian Shen and Prof. Cong Shi's research group, including former members Dr. Huanan Zhang, Dr. Yan Deng, Dr. Zhihao Chen, Juan Xu and current members Weidong Chen, Hao Yuan and Yiling Zhang. I have learned many research ideas through our discussions. I would also like to thank all my other friends for their advice and support, especially Dr. Yuhui Shi, Yanzhe Lei and Duyi Li.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

1.1 Background

Over the past 20 years, online retailing has been growing at an astonishing rate, with online sales accounting for around one-quarter of the total retail market. According to the U.S. Department of Commerce, total e-commerce sales in 2016 reached \$394.9 billion, a 15.6% increase compared with \$341.7 billion in 2015. E-commerce offers many ways for retailers to reach consumers and conduct business without the need of a physical store. It is easier than ever for businesses to have a digital presence, which offers 24-hour purchasing opportunities all year round with minimum maintenance cost. It also allows online retailers to display their merchandise in any part of the world with little additional expenses. This advantage enables online retailers to expand their market globally and target an extremely focused segment.

With the rapid growth of e-commerce, effective inventory control and revenue management strategies can help many small and large businesses significantly reduce their supply chain costs and increase their total revenue. Comparing with the traditional supply chain management, there are several important features that differentiate the online retail environment.

Quality of service (QoS). In today's customer-driven business environment, it is vital for companies to focus on the QoS. Since the early 2000s, firms started to put tremendous efforts and resources in understanding the customers and the markets. It usually costs five to twenty-five times more to attract a new customer than to retain an old one (cf. Gallo 2014). Customers facing stockouts have been observed abandoning their purchases, switching retailers, substituting similar items and have seldom gone back (see, e.g., Fitzsimons 2000). One of the most common challenges in making supply chain decisions, at its most fundamental, boils down to minimizing the supply chain costs while still delivering great customer service. Providing excellent customer service not only increases sales and profits, but also helps companies stand out in the marketplace.

In make-to-stock inventory management systems where retailers place orders before demands arrive, service-level is widely used, both in theory and in practice, to measure the performance of inventory replenishment policies (cf. Ghiani et al. 2005a). It directly characterizes the QoS, which is closely related to customer satisfaction and firms revenue. By enforcing a service-level requirement, companies are able to improve the QoS by guaranteeing a small stockout rates. In comparison, traditional inventory models often assume linear cost functions to penalize backorders or lost sales, which is primarily for analytical tractability rather than an accurate representation of reality (see, e.g., Bertsimas and Paschalidis 2001). The mechanism of varying unit penalty costs can hardly take effect on the QoS performance of a system. The service-level requirements presented in our models are more applicable in online retail settings.

In addition, after-sales services could also be crucial in delivering great customer services (cf. Cohen et al. 2006). Companies have to handle the return, repair, and disposal of failed components. The returned products, though some parts may be damaged, can be remanufactured and resold. The remanufacturing process includes repair or replacement of worn-out or obsolete components and modules, which has a lower production cost than the manufacturing process. The dual production methods further complicate the inventory control system, but it helps firms retain their old customers by providing great after-sales services.

Seasonal demands. One of the most important features in online retailing is the demand seasonality. For example, the holiday season brings immense opportunities for ecommerce since their demands are often five to ten times higher than usual demands (cf. Blogger 2016). According to Hsu (2017), retailers on Alibaba's platforms had recorded 25.3 billion worth of gross merchandise volume (GMV) on Singles' Day in 2017. Hence, rather than assuming stationary demand in the classical inventory control literature, it is important to consider a practical demand process that is seasonal, forecast-based, and driven by the state of the economy. In the existing literature, several demand models are widely used for different applications, such as Markov-modulated demand process (MMDP) (see, e.g., Sethi and Cheng 1997) and autoregressive demand (see, e.g., Mills 1991). In practice, martingale model of forecast evolution (MMFE for short, see, e.g., Graves et al. 1986, Heath and Jackson 1994a) and advance demand information (ADI) (see, e.g., Gallego and Özer 2001) are often used to forecast the future demand. These demand models will complicate the mathematical formulation and may increase the complexity of finding an optimal solution. Therefore, in order to find good inventory and pricing policies, more efficient algorithms are needed to solve these complex models.

Sales history and customer ratings. With the rapid growth of the Internet, more information is available in the e-commerce marketplace. Online buyers are able to do their own research before purchases. For example, due to the instantaneous pricing information available to online buyers, they are able to compare offerings of sellers worldwide. Moreover, most online retailers (such as Amazon) also provides abundant product information to buyers such as sales rank, customers' rating and customers' review. Some e-retailers also allow customers to sort (or refine) their search results based on these metrics. Consequently, the product information exerts a huge influence on customers' buying choices. For example, Chevalier and Mayzlin (2006) found that online customer ratings significantly impact product sales by analyzing data from online bookstores. Given thousands of different products, customers are usually non-expert and they tend to have the belief that bestsellers with higher star ratings generally have better quality guarantees.

The extra product information can change the traditional pricing strategies. According to Remy et al. (2010), many of today's successful online retailers use adaptive pricing strategy in order to utilize customers' feedbacks. By collecting and analyzing Amazon's prices for bestsellers in the camera and video category every hour for three months, they observed that the price of the same product is changed constantly over time. The changes in prices and sales ranks strongly indicate that online retailers are willing to offer price discounts to attract more price-sensitive customers and improve their sales rank and customers' rating.

1.2 Research Goals

In general, the goal of this research is to study the inventory control and pricing problems arise in today's online retail environment. Specifically, our goals of this research are summarized as follows.

1. We focus on the key features of the online retailing described in Section 1.1 by developing realistic and complex mathematical models. We propose several mathematical formulations to solve inventory control and pricing problems that arise in online retail marketplace. Compared with traditional models, our proposed models are more applicable in reality and can be generalized to other domains.

2. Based on the proposed models, we use mathematical tools to characterize the structure of optimal policies. We provide an analysis of the structure of optimal policies for our proposed models. For some complex models, it turns out that optimal policies follow a simple structure. For instance, we show that an optimal policy for the stochastic inventory system under service-level constraints is a total-base-stock policy. We also show that with sales rank information, the optimal pricing policy follows a cyclic structure. These special structures can be utilized in practice to help retail companies designing heuristic policies to make their inventory decisions and pricing decisions.

3. Due to the complexity of mathematical models, it is usually computationally intractable to find optimal policies when problem size grows larger. Therefore, one of the research goals is to design efficient algorithms to find near-optimal policies. For instance, we develop approximation algorithms for stochastic inventory control models with service-level constraints and non-stationary seasonal demands, which can be applied to provide provably-good inventory control strategies in online retail business. Our proposed algorithms are easy-to-implement and computationally efficient. In addition, we also provide theoretical worst-case analysis to our proposed algorithms, which is often very challenging for multistage stochastic optimization models. The theoretical performance analysis provides insights into the algorithms and motivates an empirical study of their typical performance through extensive computational experiments.

1.3 Contributions

This dissertation studies inventory control and pricing decisions in the context of online retail. It consists of three essays, each analyzing a different problem in the area of supply chain management and revenue management. All three essays contribute to the arising area of today's online retail environment.

In Chapter 2, we study a multi-period production planning problem under joint service-level constraint. The joint service-level constraint is enforced to improve the QoS of inventory management system. Instead of assuming known demand distribution, we incorporate the historical demand data into our model and study a data-driven optimization problem. We also consider a variant in which pricing decisions are required to be made together with the inventory decisions. Via computations of diverse instances, we demonstrate the effectiveness of our approach by analyzing the solution feasibility and objective bounds and conducting sensitivity analysis.

In Chapter 3, we study a dynamic inventory control problem with enforced service-level constraint in each period. In addition to the periodic-review inventory control systems, we also study a variant model, in which retailer has returned products and need to make remanufacturing decisions. Such model is capable of providing after-sales services, which can help retailers to improve their QoS. Moreover, we also consider seasonal demands presented in the system and model the demand with a generic demand structure that generalizes the correlated demand models in existing literature. We formulate both problems using dynamic programs, and propose 2-approximation algorithms for both models. The core concept developed is called the *delayed forced holding and production cost*, which is proven effective in dealing with service-level constrained inventory systems. Our extensive computational experiments show that the proposed algorithms on average perform within 2% error of optimality. The techniques developed in this work can lend themselves to important problems in other domains, such as resource allocation, appointment scheduling, etc. We believe these easy-to-implement and efficient algorithms can be widely applied in industry that typically requires realistic assumptions and large-scale data sets.

In Chapter 4, we study periodic-review dynamic pricing problems under the effect of sales ranking. The demand in each period is a deterministic function of price and sales rank. The goal is to find the optimal price in each period so that retailers can maximize the total profit. We consider both single-product model and multi-product model in which an online retailer manages multiple substitutable products. We show that the cyclic pricing policy is optimal for both models. For the single-product model, we also quantify the length of the optimal pricing cycle. Our results show that the optimal cycle length depends on the impact of the sales rank as well as the expected profit for each price choices. In addition, we conduct numerical studies for both models to show the benefit of cyclic pricing policy.

In general, the outcome of this research will help companies (especially online retailers) manage their inventory and pricing decisions in order to minimize their supply chain costs or maximize their total revenue. The proposed models in this research project can be widely used in practice for different applications and generate effective managerial decisions. In addition, our proposed models and methodologies can be applied to a broader class of traditional or emerging supply chain applications.

CHAPTER II

Production Planning Problem with Joint Service-Level Guarantee

2.1 Abstract

We consider a class of single-stage T-period production planning problems under demand uncertainty. The main feature of this work is to incorporate a joint service-level constraint to restrict the joint probability of having backorders in any period. This is motivated by manufacturing and online retailing applications, in which firms need to decide the production quantities ex ante, and also have stringent service-level agreements. The inflexibility of dynamically altering the pre-determined production schedule may be due to contractual agreement with external suppliers or other economic factors such as enormously large fixed costs and long lead time. We focus on two stochastic variants of this problem, with or without pricing decisions, both subject to a joint service-level guarantee. The demand distribution could be nonstationary and correlated across different periods. Using the sample average approximation (SAA) approach for solving chance-constrained programs, we re-formulate the two variants as mixed-integer linear programs (MILPs). Via computations of diverse instances, we demonstrate the effectiveness of the SAA approach, analyze the solution feasibility and objective bounds, and conduct sensitivity analysis for the two MILPs. The approaches can be generalized to a wide variety of production planning problems, and the resulting MILPs can be efficiently computed by commercial solvers.

2.2 Introductory Remarks

In this chapter, we study a class of production planning problems subject to a joint service-level constraint. The problems fall into the category of single-stage T-period stochastic optimization problems with no recourse decisions (i.e., firms need to plan their production and/or pricing decisions ex ante and cannot change them in subsequent periods). This class of problems is primarily motivated by manufacturing or online retailing applications, in which firms have stringent service-level requirements or need to provide high customer services, but do not have sufficient flexibility of altering their decisions due to the related issues, such as contractual agreement with outside suppliers, enormously large fixed costs and long lead time.

One motivating example of this research comes from online retailing. For example, consider an online apparel store that sells clothes to customers. In apparel business, there are multiple phases (including product development, production, marketing, etc.) before a product can be sold on market. Many apparel stores set their production sites overseas to reduce production cost but increases their production lead-time. Usually, there are months of lead-time for a new collection to be delivered to stores. The retailer has to plan inventory quantities and prices ahead of selling season. Moreover, online digital marketplace offers a more competitive environment compared with traditional offline markets (cf. Brynjolfsson et al. 2011). Customers facing stockouts have been observed abandoning their purchases, switching retailers, substituting similar items and have seldom gone back (see, e.g., Fitzsimons 2000). This motivates us to impose a service-level requirement to restrict the joint stockout probabilities in each period which improves customers' shopping experience and helps firms to maintain excellent reputation. In addition, demands for the online apparel store can be seasonal, non-stationary without known distribution. During holiday season, demand increases dramatically (five to ten times higher). It is critical for retailer to rely on historical demand data and plan production quantities and prices accordingly.

Other than e-commerce applications, this research is also related to other emerging inventory control and pricing problems. For example, CEMEX, a multinational building materials company that often signs contracts with large event organizers. The firm was the key supplier of cement for multiple large events in 2014, including notably a contract with FIFA to supply 28000 tons of cement for the new soccer stadium in Manaus (cf. CEMEX 2015). The firm has to plan production quantities and prices for the committed projects long before starting the project. Moreover, guaranteeing an adequate service-level is absolutely essential for successful completion of these projects. Another motivating example arises in the U.S. automotive industry. The lead time for building a new car model is typically 52 months (cf. Fine et al. 1996) due to lengthy design and testing cycles, and the lead time for manufacturing an existing car model is typically 10 to 18 months (cf. ARI 2015). The fixed cost related to altering a production schedule is also quite high. Due to these reasons, car manufacturers usually make their production plans ahead of the selling season and their planning decisions cannot be easily changed afterwards. Meanwhile, an adequate service-level is important to maintain firms' revenue and goodwill.

To capture the aforementioned applications, we propose two single-stage T-period models (with or without pricing decisions), subject to a joint service-level guarantee. The first model considers the production decisions while the second model concerns both the production and pricing decisions. For the first model, the demands are random, which can be non-stationary and correlated across different periods. Our goal is to minimize the expected total cost, including (linear) production costs, inventory holding costs and backorder penalty costs, subject to a joint service-level constraint over the T-periods to restrict the probability of having unmet demands during the planning horizon. For the second model, we assume a classical additive demand model in which the demand depends linearly on the price plus a random disturbance term (see, e.g., Chen and Simchi-Levi 2012). We consider discrete pricing and continuous pricing options, in which prices are chosen from a given finite set of values or from a bounded price range, respectively. Our goal is to maximize the total expected profit, also subject to a joint service-level constraint.

We remark again that both problems considered in this chapter belong to the category of single-stage stochastic optimization problems with no recourse decisions, which should be distinguished from dynamic inventory control problems considered in the literature (see, e.g., Zipkin 2000). The main feature of our models is to incorporate a joint service-level constraint, which is practically relevant but computationally intractable. Our approach employs the *sample average approximation* (SAA for short) method (see, e.g., Luedtke and Ahmed 2008) to reformulate our chance constrained problems as mixed-integer linear programming (MILP) models and to compute upper and lower bounds of the optimal objective values as well as feasible solutions with certain confidence levels.

2.2.1 Relevant Literature

The traditional study of production planning problem has been focused on deterministic models with known demand. Zangwill (1969) developed a deterministic lot-sizing model that allowed for backlogged demand and proposed a network approach. They further proposed dynamic programming algorithms to compute optimal planning policies based on network formulations. Pochet and Wolsey (1988) studied several strong MILP reformulations of the uncapacitated lot-sizing problem with backlogging. They also described a family of strong valid inequalities that can be effectively used in a cut generation algorithm. Florian et al. (1980) studied capacitated lot-sizing problem and showed that the deterministic problem is NP-hard. Recently, Absi et al. (2011) studied the single item uncapacitated lot-sizing problem with production time windows, lost sales, early productions and backlogs. They presented MILP formulations of these models and developed dynamic programming algorithms to solve them. González-Ramírez et al. (2011) proposed a heuristic algorithm to solve a multi-product, multi-period capacitated lot-sizing problem with pricing, where the deterministic demand was assumed to be linear in price.

In this chapter, we focus on stochastic variants of production planning problems, subject to a joint service-level constraint. There has been limited literature on this topic, among which Bookbinder and Tan (1988) studied a multi-period lot-sizing problem that imposed individual service-level constraint in each period and their demand distributions were known. In contrast, our model considers a joint service-level constraint that poses more computational challenges, and empirical demand samples are given instead of an explicit demand distribution function. We reformulate our problem as an MILP model using the SAA approach, which is based on Monte Carlo simulation of random samples, to approximate the expected value function by the corresponding sample average. Kleywegt et al. (2002) studied the convergence rates, stopping rules and computational complexity of the SAA method. They also presented a numerical example for solving the stochastic knapsack problem using the SAA method. Verweij et al. (2003) formulated stochastic routing problems using the SAA approach. They applied decomposition and branch-and-cut techniques to numerically solve the approximating problems. Pagnoncelli et al. (2009) applied the SAA method to solve two chance constrained problems, namely, linear portfolio selection problem and blending problem with a joint chance constraint. Recently, Mancilla and Storer (2012) formulated a stochastic scheduling problem using the SAA approach and proposed a heuristic method based on Benders decomposition.

Our main methodology builds upon the theory developed in Luedtke and Ahmed (2008). They first proposed to use the SAA approach to find feasible solutions and lower bounds on the optimal objective value of a general chance-constrained program. Keeping the same required risk level, they showed that the corresponding SAA counterpart yields a lower bound of the optimal objective value. To find a feasible solution, they showed that it suffices to solve a sample-based approximation with a smaller risk level. They also mathematically derived the required sample sizes in theory for having a lower bound or a feasible solution with high confidence. Our work contributes to the literature by being first to employ the SAA approach to solve a class of production planning problems subject to a joint service-level constraint, which is typically computationally intensive.

2.2.2 Contributions

The main contributions of this chapter are summarized as follows.

1. From the modeling perspective, we propose two new production planning models (with and without pricing decisions) subject to a joint service-level constraint. In classical stochastic production planning problems, unsatisfied demands are typically penalized by a linear backorder cost only. However, it is usually important for firms to maintain their reliability or credibility by persistently satisfying all the market demands in each period with high probability. Some firms use *α-service-level* (defined as the probability that the demand is fully satisfied) to measure their QoS. This metric is yet neglected in most classical production planning models in the literature. This motivates us to incorporate a joint α-service-level constraint that ensures the market demands being satisfied in each period with a sufficiently high probability, so that the related firms can remain competitive and profitable.

   Moreover, in most stochastic production planning models in the existing literature, the demand distributions in each period are given explicitly, while in real-life applications, it is usually difficult to deduce the true underlying demand distribution. The SAA reformulation can be done without knowing the exact demand distribution; however, a large amount of empirical data (more than 5000) is needed to solve the SAA reformulation under the nominal risk level and such amount of data may not be available in reality. We show that if we tune the risk parameter in the SAA reformulation smaller, we are able to obtain good feasible solutions (within 5% of optimality) using much less empirical data (around 300), which makes the data-collection work less demanding. Also, our proposed models allow for nonstationary and generally correlated demands.

2. From the computational perspective, we employ the SAA method for chance-constrained programming and reformulate the two production planning models as MILPs using finite samples. However, due to the large amount of samples needed in the SAA reformulation, solving the resulting MILPs exactly is computationally intensive. Tuning the risk parameter in the SAA reformulation smaller than the required service-level (i.e., more conservative), we achieve feasible solutions by solving the resulting MILPs via much fewer samples. The feasible solutions provide an upper (lower) bound on the cost-minimization (profit-maximization) problem. On the other hand, we also compute a lower (upper) bound on the cost-minimization (profit-maximization) problem by setting the risk parameter to be at least equal to the service-level and solving multiple SAA counterparts with fewer number of samples.

A comprehensive numerical study has been conducted using three popular demand models with different patterns of demand correlations among periods (i.e., identical and independent demand distributions, Markov modulated demand process (MMDP) and autoregressive demands (AR models)), which are extensively used in theory and practice. For each problem instance with different demand model and different required service-level, both upper and lower bounds are computed and validated. We then compare our bounds with the optimal solutions (for reasonable problem sizes). It can be observed that the more samples we use, the less gap it has from our bounds to the optimal solutions. Also, the actual sample size needed to achieve at a given confidence level for both upper and lower bound solutions is a magnitude smaller than the theoretic bounds proposed in Luedtke and Ahmed (2008). We conduct sensitivity analysis for production planning with pricing and our numerical results show that when the demand is less sensitive to the price, the firm tends to increase the price while keeping the demand at the same level, in order to obtain a better profit.

2.2.3 Structure and General Notation

The remainder of this chapter is organized as follows. In Section 2.3, we introduce the notation and formulate the joint service-level constrained stochastic production planning problem. Section 2.4 formulates the production planning problem with pricing options. The computational results and insights for both problems with/without pricing options are given in Section 2.5. Finally, Section 2.6 concludes the chapter and gives future research directions.

Throughout the chapter, for notational convenience, we use a capital letter and its lowercase form to distinguish between a random variable and its realization. We use $\triangleq$ to indicate "is defined as", and $\mathbb{1}(A)$ is the indicator function taking value 1 if statement A is true and 0 otherwise. For any $x \in \mathbb{R}$, we denote $x^+ = \max\{x, 0\}$. We also use $\lceil x \rceil$ to denote the smallest integer that is no less than x and use $\lfloor x \rfloor$ to denote the largest integer not greater than x.

2.3 Production Planning with Joint Service-Level Constraint

2.3.1 Mathematical Formulation

Consider a finite horizon of T periods. The classic production planning problem decides the production quantities for each period simultaneously at the beginning of the planning horizon (denoted as $q_1, q_2, \ldots, q_T$). During each period t ($t = 1, \ldots, T$), demands are realized and three types of cost are incurred: production cost (with a per-unit cost $c_t > 0$), holding cost for on-hand inventories from period t to $t+1$ (with a per-unit cost $h_t > 0$), and penalty cost for backlogged demand (with a per-unit cost $p_t > 0$). The objective is to minimize the total cost over the T periods. Let $D_1, \ldots, D_T$ denote the random demands over the T periods and they may be independently distributed or correlated. As a general convention, we use a capital letter to denote a random variable and a lower case letter to denote its realization.

Let X_t and B_t be random variables that denote on-hand inventories and backorders at the end of period $t = 1, 2, \ldots, T$, respectively. Clearly, both X_t and B_t must be nonnegative. The initial inventory and backorder levels are denoted by x_0 and b_0, respectively. We formulate the production planning problem under a joint service-level constraint as

$$\begin{aligned}
\text{(PP)} \quad \min \quad & \sum_{t=1}^{T} (c_t q_t + h_t \mathbb{E}[X_t] + p_t \mathbb{E}[B_t]) && (2.1)\\
\text{s.t.} \quad & X_{t-1} + q_t + B_t = D_t + X_t + B_{t-1}, \quad \forall t = 1, \ldots, T, && (2.2)\\
& \mathbb{P}(X_t - B_t \geqslant 0, \forall t = 1, \ldots, T) \geqslant 1 - \theta, && (2.3)\\
& q_t \geqslant 0, \quad \forall t = 1, \ldots, T. && (2.4)
\end{aligned}$$

The objective (2.1) minimizes the total ordering cost, expected inventory cost and expected backlogging cost. In each period t, the incoming items are X_{t-1}, q_t and B_t while the outgoing items are X_t, B_{t-1} and D_t. To balance them, we formulate (2.2) as the *flow-balance constraints*. Constraint (2.3) requires that the probability of satisfying the demands in all T periods is at least $1 - \theta$, which defines the service-level.

2.3.2 Reformulation Using the SAA Approach

Consider N samples of demands over T periods denoted by $d^{(i)} = (d_1^{(i)}, \ldots, d_T^{(i)})$ ($i = 1, 2, \ldots, N$) where each sample i is equally likely to occur with probability $1/N$. The on-hand inventories and backorders vary according to demand samples, denoted by $x^{(i)} = (x_0^{(i)}, \ldots, x_T^{(i)})$ and $b^{(i)} = (b_0^{(i)}, \ldots, b_T^{(i)})$, respectively. The initial on-hand inventory and backorder are pre-determined regardless of the realization of random demands, i.e., $x_0^{(i)} = x_0$ and $b_0^{(i)} = b_0$ for all $i = 1, 2, \ldots, N$. The ordering quantities $q_1, \ldots q_t$ are decided before knowing the demand realizations, and thus do not depend on the specific samples.

In each sample i, the balance constraint (2.2) is presented as

$$x_{t-1}^{(i)} - x_t^{(i)} - b_{t-1}^{(i)} + b_t^{(i)} + q_t = d_t^{(i)}, \quad \forall t = 1, \ldots, T. \tag{2.5}$$

We compute the total expected cost as:

$$\sum_{t=1}^{T} c_t q_t + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(h_t x_t^{(i)} + p_t b_t^{(i)} \right).$$

The joint service-level constraint (2.3) is equivalent to

$$\sum_{i=1}^{N} \mathbb{1} \left\{ x_t^{(i)} \geqslant b_t^{(i)}, \forall t = 1, 2, \ldots, T \right\} \geqslant \lceil (1-\theta) N \rceil. \tag{2.6}$$

Define binary variables $y^{(i)}$ such that if we choose to violate the i-th sample then $y^{(i)} = 1$ and $y^{(i)} = 0$ otherwise. We then replace the joint service-level constraint (2.6) by:

$$\begin{cases} x_t^{(i)} - b_t^{(i)} \geqslant -M_t^{(i)} y^{(i)}, \quad \forall t = 1, \ldots, T, & (2.7) \\ \sum_{i=1}^{N} y^{(i)} \leqslant \lfloor \theta N \rfloor, & (2.8) \\ y \in \{0, 1\}^N. & (2.9) \end{cases}$$

Note that for each period t and sample i, the big-M coefficient $M_t^{(i)} = -x_0 + b_0 + \sum_{s=1}^{t} d_s^{(i)}$ is a valid upper bound for $-x_t^{(i)} + b_t^{(i)}$, because

$$\begin{aligned} b_t^{(i)} - x_t^{(i)} &= (d_t^{(i)} - q_t) + (b_{t-1}^{(i)} - x_{t-1}^{(i)}) \\ &= \sum_{s=1}^{t} (d_s^{(i)} - q_s) + (b_0 - x_0) \\ &\leqslant (b_0 - x_0) + \sum_{s=1}^{t} d_s^{(i)}. \end{aligned}$$

When $y^{(i)} = 0$, the constraint $x_t^{(i)} - b_t^{(i)} \geqslant 0$ is enforced for each $t = 1, 2, \ldots, T$. When $y^{(i)} = 1$, the joint service-level constraint in the i-th sample can be violated and the total number of violated samples is no more than $\lfloor \theta N \rfloor$ ensured by the constraint $\sum_{i=1}^{N} y^{(i)} \leqslant \lfloor \theta N \rfloor$.

Therefore, we approximate a multi-period plan of optimal ordering quantities by using the following MILP model.

$$
\begin{aligned}
\text{(SAA-PP)} \quad \min \quad & \left\{ \sum_{t=1}^{T} c_t q_t + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \left(h_t x_t^{(i)} + p_t b_t^{(i)} \right) \right\} \\
\text{s.t.} \quad & (2.5), (2.7)\text{–}(2.9), \\
& x_0^{(i)} = x_0,\ b_0^{(i)} = b_0, \quad \forall i = 1, \ldots, N, && (2.10) \\
& x_t^{(i)}, b_t^{(i)}, q_t \geqslant 0, \quad \forall t = 1, \ldots, T,\ i = 1, \ldots, N. && (2.11)
\end{aligned}
$$

We present a necessary condition for any optimal production plan in the following proposition.

Proposition 2.1. *For any optimal solutions* $\{q_t^*, x_t^{(i)^*}, b_t^{(i)^*}\}_{i,t}$ *to the (SAA-PP), we have either* $x_t^{(i)^*} = 0$ *or* $b_t^{(i)^*} = 0$ *holds for all* $i = 1, 2, \ldots, N$ *and* $t = 1, 2, \ldots, T$*.*

Proof. We show the results by contradiction. Suppose that there exists an optimal solution with $x_t^{(i)} > 0$ and $b_t^{(i)} > 0$ for some $t \in \{1, 2, \ldots, T\}$ and $i \in \{1, 2, \ldots, N\}$. We can replace $x_t^{(i)}$ and $b_t^{(i)}$ by $\tilde{x}_t^{(i)} = x_t^{(i)} - \min\{x_t^{(i)}, b_t^{(i)}\}$ and $\tilde{b}_t^{(i)} = b_t^{(i)} - \min\{x_t^{(i)}, b_t^{(i)}\}$ while keeping the other decisions the same. Since the inventory level remains unchanged, i.e., $\tilde{x}_t^{(i)} - \tilde{b}_t^{(i)} = x_t^{(i)} - b_t^{(i)}$, all the constraints are satisfied under the new solution. Meanwhile, this new feasible solution decreases the total cost by $\frac{1}{N}(h_t + p_t) \cdot \min(x_t^{(i)}, b_t^{(i)}) > 0$. This contradicts with the fact that the original solution is optimal and thus completes the proof. □

The above proposition is also true when no service-level requirement is present (see, e.g., Zipkin 2000). It asserts that even in the case of having a joint service-level constraint, there is no incentive to have backorders while holding positive inventories. This result is also true when no service-level requirement.

2.4 Model Variant with Pricing Options

2.4.1 Notation and Problem Formulation

We consider pricing decisions in the above production planning problem. In this variant, besides the ordering quantity q_t for each period $t = 1, 2, \ldots, T$, the manager also decides the price r_t for each period $t = 1, 2, \ldots, T$ at the beginning of the whole time horizon. The price r_t set by the manager affects the underlying demand distribution $D(t)$ and thus the realization d_t. The goal is to maximize the total expected profit over the T periods.

We interpret the random demand by a deterministic linear function in price r_t plus a noise term, i.e., $D_t(r_t) = -a_t r_t + \beta_t + \tilde{\epsilon}_t$, where $\tilde{\epsilon}_t$ is a random variable with $\mathbb{E}[\tilde{\epsilon}_t] = 0$ for

$t = 1, \ldots, T$, and both $a_t, \beta_t > 0$. This demand model is well known as the additive demand model in the literature (see, e.g., Mills (1959)). It allows for correlated demands over periods, which indicates that $\{\tilde{\epsilon}_t\}_{t=1}^T$ are not necessarily independent random variables.

We further assume that once the demand is realized in period t, it establishes a contract between the buyer and the retailer with a unit price of r_t. In other words, given that the realized demand $D_t = d_t$, it immediately incurs a revenue of $r_t d_t$, no matter when the demand is satisfied. The price $r = (r_1, r_2, \ldots, r_T)^\mathsf{T}$ is chosen from a given set $P \subseteq \mathbb{R}^T$. We can then formulate the production planning problem with pricing options as

$$
\begin{aligned}
\text{(PO)} \quad \max \quad & \sum_{t=1}^{T} \Big(r_t \mathbb{E}[D_t] - c_t q_t - h_t \mathbb{E}[X_t] - p_t \mathbb{E}[B_t] \Big) \\
\text{s.t.} \quad & \text{(2.2)–(2.4)}, \\
& D_t = -a_t r_t + \beta_t + \tilde{\epsilon}_t, \quad \forall t = 1, \ldots, T, \qquad (2.12) \\
& r \in P.
\end{aligned}
$$

In the formulation above, the constraints (2.2)–(2.4) are carried from the model without pricing options. The constraint (2.12) shows the relationship between price and demand in each period.

2.4.2 Sample Average Approximation Reformulation

We reformulate the (PO) model using N sample data, namely $d^{(i)} = (d_1^{(i)}, d_2^{(i)}, \ldots, d_T^{(i)})^\mathsf{T}$, $i = 1, 2, \ldots, N$. For the i-th sample, we use $\epsilon_t^{(i)}$ to denote the realizations of $\tilde{\epsilon}_t$ and other notations remain the same. Then the total expected profit is computed by

$$
-\sum_{t=1}^{T} c_t q_t + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \left[r_t d_t^{(i)} - h_t x_t^{(i)} - p_t b_t^{(i)} \right].
$$

Hence, using the linear relationship between demand and price, i.e., $d_t^{(i)} = -a_t r_t + \beta_t + \epsilon_t^{(i)}$, we reformulate the (PO) model as:

$$
\begin{aligned}
&\text{(SAA-PO)} \\
&\max \quad \sum_{t=1}^{T} \big(- c_t q_t - a_t r_t^2 + \beta_t r_t \big) + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \big(\epsilon_t^{(i)} r_t - h_t x_t^{(i)} - p_t b_t^{(i)} \big) \\
&\text{s.t.} \quad x_{t-1}^{(i)} - x_t^{(i)} - b_{t-1}^{(i)} + b_t^{(i)} + q_t = -a_t r_t + \beta_t + \epsilon_t^{(i)}, \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
&\qquad x_t^{(i)} - b_t^{(i)} \geqslant -\Big(- x_0 + b_0 + \sum_{s=1}^{t} \big(- a_s r_s + \beta_s + \epsilon_s^{(i)} \big) \Big) y^{(i)}, \\
&\qquad\qquad \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \qquad (2.13) \\
&\qquad (2.8) - (2.11), \\
&\qquad r \in P.
\end{aligned}
$$

Here, the probabilistic constraint (2.3) in the (PO) model is equivalent to constraints (2.13), (2.8) and (2.9), for which we define new binary variables $y^{(i)}$ ($i = 1, \ldots, N$). However, the above model still involve nonlinear terms r_t^2 in the objective function and bilinear terms $r_s y^{(i)}$ in constraint (2.13). We reformulate it as a linear model for two specific price sets: discrete price set and continuous price set, where price decisions p_t, $t = 1, ..., T$, are independently determined for each period. For the former, the price is drawn from a set of finite possible prices, denoted by set $R_t = \{\gamma_1^t, \ldots, \gamma_{m_t}^t\}$. In this case, the price set can be written as $P = \times_{t=1}^{T} R_t$. For the latter, we consider the possible price r_t chosen from a price interval $[L_t, U_t]$ given for each period t and the price set is specified as $P = \times_{t=1}^{T} [L_t, U_t]$. We will give MILP models for each price set in the following subsections. Note that our model is also capable of describing the relationship among the prices in each period. For example, if the prices are not allowed to increase from periods to periods, we can simply add the constraint $r_1 \geqslant r_2 \geqslant \cdots \geqslant r_T$ to our model while the complexity of the resulting model remains the same.

2.4.2.1 Discrete price set

Consider a finite set $R_t = \{\gamma_1^t, \ldots, \gamma_{m_t}^t\}$ from which product price is chosen in each period. Define a binary decision variable u_{jt} to indicate whether using the j-th price option, i.e., γ_j^t ($j = 1, \ldots, m_t$). In each period, $u_{jt} = 1$ if $r_t = \gamma_j^t$ and $u_{jt} = 0$ otherwise. To ensure only one price is used from the set $(\gamma_1^t, \gamma_2^t, \ldots, \gamma_{m_t}^t)$ in each period t, we require $\sum_{j=1}^{m_t} u_{jt} = 1$ for each

t. Then, the quadratic term r_t^2 in objective function can be expressed by a linear term:

$$r_t^2 = \sum_{j=1}^{m_t} (\gamma_j^t)^2 u_{jt}. \tag{2.14}$$

Similarly, for any $t = 1, \ldots, T$ and $i = 1, \ldots, N$, the nonlinear term $r_t y^{(i)}$ in the joint service-level constraint can be rewritten as

$$r_t y^{(i)} = \sum_{j=1}^{m_t} \gamma_j^t u_{jt} y^{(i)}. \tag{2.15}$$

To further linearize the term $u_{jt} y^{(i)}$ in (2.15), we introduce another decision variable $v_{jt}^{(i)}$ to replace $u_{jt} y^{(i)}$ and add the following McCormick inequalities to force $v_{jt}^{(i)} = u_{jt} y^{(i)}$ (see McCormick (1976)):

$$\begin{cases} v_{jt}^{(i)} \leqslant u_{jt} \\ v_{jt}^{(i)} \leqslant y^{(i)} \\ v_{jt}^{(i)} \geqslant u_{jt} + y^{(i)} - 1 \\ v_{jt}^{(i)} \geqslant 0. \end{cases} \tag{2.16}$$

Here both u_{jt} and $y^{(i)}$ are binary variables. If $u_{jt} = 1$ and $y^{(i)} = 1$, then $v_{jt}^{(i)} = 1$; otherwise, the constraints enforce $v_{jt}^{(i)} = 0$. Therefore, for discrete pricing options, using equalities (2.14) and (2.15), we can formulate the following (PP-D) model:

(PP-D)

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T} \Bigg(- c_t q_t - a_t \sum_{j=1}^{m_t} (\gamma_j^t)^2 u_{jt} + \beta_t \sum_{j=1}^{m_t} \gamma_j^t u_{jt} \Bigg) + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \Bigg(\epsilon_t^{(i)} \sum_{j=1}^{m_t} \gamma_j^t u_{jt} - h_t x_t^{(i)} - p_t b_t^{(i)} \Bigg) \\
\text{s.t.} \quad & x_{t-1}^{(i)} - x_t^{(i)} - b_{t-1}^{(i)} + b_t^{(i)} + q_t = -a_t \sum_{j=1}^{m_t} \gamma_j^t u_{jt} + \beta_t + \epsilon_t^{(i)}, \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
& x_t^{(i)} - b_t^{(i)} \geqslant (x_0 - b_0) y^{(i)} + \sum_{s=1}^{t} \Bigg(a_s \sum_{j=1}^{m_s} \gamma_j^s v_{js}^{(i)} - \beta_s y^{(i)} - \epsilon_s^{(i)} y^{(i)} \Bigg), \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
& (2.8)\text{–}(2.11), \\
& (2.16), \quad \forall t = 1, \ldots, T,\ i = 1, \ldots, N,\ j = 1, \ldots, m_t, \\
& \sum_{j=1}^{m_t} u_{jt} = 1, \quad \forall t = 1, \ldots, T, \\
& u_{jt} \in \{0, 1\}, v_{jt}^{(i)} \geqslant 0, \quad \forall t = 1, \ldots, T,\ i = 1, \ldots, N,\ j = 1, \ldots, m_t.
\end{aligned}$$

2.4.2.2 Continuous price set

For continuous pricing options, the price in period t is chosen from the set $P_t = [L_t, R_t]$. We first note that the price r_t must be non-negative, hence $L_t \geqslant 0$. Also, r_t should be bounded above by β_t/a_t, otherwise the expected demand $-a_t r_t + \beta_t < 0$, which is unlikely to happen in reality. Hence, the retailer will never set a price higher than β_t/a_t, and therefore, $U_t \leqslant \beta_t/a_t$.

For the nonlinear term $r_t y^{(i)}$ presented in the joint service-level constraint, we linearize it by introducing a new decision variable $w_t^{(i)}$. The following sets of linear inequalities enforce $w_t^{(i)} = r_t y^{(i)}$ when $y^{(i)}$ is a binary:

$$\begin{cases} w_t^{(i)} \leqslant r_t \\ w_t^{(i)} \leqslant U_t y^{(i)} \\ w_t^{(i)} \geqslant r_t + U_t(y^{(i)} - 1) \\ w_t^{(i)} \geqslant 0. \end{cases} \tag{2.17}$$

Finally we can formulate the following (PP-C) model with quadratic objective for production planning problem with continuous pricing options as follows

$$\begin{aligned}
&\text{(PP-C)} \\
&\max \quad \sum_{t=1}^{T} \big(-c_t q_t - a_t r_t^2 + \beta_t r_t \big) + \frac{1}{N} \sum_{t=1}^{T} \sum_{i=1}^{N} \big(\epsilon_t^{(i)} r_t - h_t x_t^{(i)} - p_t b_t^{(i)} \big) \\
&\text{s.t.} \quad x_{t-1}^{(i)} - x_t^{(i)} - b_{t-1}^{(i)} + b_t^{(i)} + q_t = -a_t r_t + \beta_t + \epsilon_t^{(i)}, \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
&\qquad x_t^{(i)} - b_t^{(i)} \geqslant (x_0 - b_0) y^{(i)} + \sum_{s=1}^{t} \big(a_s w_s^{(i)} - \beta_s y^{(i)} - \epsilon_s^{(i)} y^{(i)} \big), \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
&\qquad \text{(2.8)–(2.11)}, \\
&\qquad \text{(2.17)}, \quad \forall t = 1, \ldots, T,\ i = 1, \ldots, N, \\
&\qquad y^{(i)} \in \{0, 1\}, \quad \forall i = 1, \ldots, N, \\
&\qquad L_t \leqslant r_t \leqslant U_t, \quad \forall t = 1, \ldots, T.
\end{aligned}$$

2.5 Computational Results

2.5.1 Solution Methods

In general, an MILP reformulation of a chance-constrained program is computationally intractable since it usually requires a large number of Monte Carlo samples to attain solution accuracy. Luedtke and Ahmed (2008) suggested using the SAA approach for solving general chance-constrained programs, and derived theocratical sample-size bounds for obtaining solutions that satisfy the chance constraints with certain confidence for specific risk levels. Specifically, consider a generic chance-constrained program:

$$(\mathrm{P}_\theta): \quad z_\theta^* = \min\{f(x) : x \in X_\theta\},$$

where $X_\theta = \Big\{x \in X : \mathbb{P}\left\{G(x,\xi) \geqslant \mathbf{0}\right\} \geqslant 1-\theta\Big\}$. Here $X \subset \mathbb{R}^n$ represents a deterministic feasible region, $f : \mathbb{R}^n \to \mathbb{R}$ represents the objective to be minimized, ξ is a random vector with support $\Xi \subseteq \mathbb{R}^d$, $G : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^m$ is a given constraint mapping and θ is a risk parameter of service-level. We assume that z_θ^* exists and is finite.

The SAA counterpart of the chance-constrained problem (P_θ) with risk parameter α is defined as

$$(\mathrm{P}_\alpha^N): \quad \hat{z}_\alpha^N = \min\{f(x) : x \in X_\alpha\},$$

where $X_\alpha = \Big\{x \in X : \frac{1}{N}\sum_{i=1}^N \mathbb{1}\Big(G(x,\xi^i) \geqslant \mathbf{0}\Big) \geqslant 1-\alpha\Big\}$.

Feasible solutions: To obtain feasible solutions of (P_θ), we can choose a smaller risk parameter $\alpha < \theta$ and solve the SAA counterpart (P_α^N). As shown in Luedtke and Ahmed (2008), if the sample size

$$N \geqslant \frac{1}{2(\theta-\alpha)^2}\log\left(\frac{|X \backslash X_\theta|}{\delta}\right), \tag{2.18}$$

then solving (P_α^N) will yield a feasible solution to (P_θ) with probability at least $1-\delta$. This gives a theoretical sample size to guarantee a feasible solution for P_θ using the SAA approach with a confidence level $1-\delta$.

Lower bounds: To obtain lower bounds on the original optimization problem (P_θ), we set $\alpha = \theta$ and to generate M SAA problems, namely, $(\mathrm{P}_{\theta,i}^N)$ $(i = 1, 2, \ldots, M)$. Then we solve each sample-based problem and obtain a set of optimal objective values, denoted by $\hat{z}_{\theta,i}^N$ $(i = 1, 2, \ldots, M)$. The L-th minimum value among all M optimal objective value is denoted by $\hat{z}_{\theta,[L]}^N$. Then, according to Luedtke and Ahmed (2008), the following result holds:

$$\mathbb{P}(\hat{z}_{\theta,[L]}^N \leqslant z_\theta^*) \geqslant 1 - \sum_{i=0}^{L-1}\binom{M}{i}(1/2)^M \tag{2.19}$$

for large enough N relative to ϵ (e.g., $N\epsilon \geqslant 10$). Hence, we can say that $\hat{z}_{\theta,[L]}^N$ is a lower bound of the objective value with a confidence level $1 - \sum_{i=0}^{L-1}\binom{M}{i}(1/2)^M$.

We test the effectiveness of the SAA approach applied to both models in this chapter with and without the pricing option. We use CPLEX 12.5.1 for solving all MILP models. All the computations are performed on a 3.40GHz Intel(R) Xeon(R) CPU.

2.5.2 Stochastic Production Planning Problem

We present numerical results on the stochastic production planning problem. We numerically solve the appropriate SAA counterpart problems and compute both the upper bounds and the lower bounds for optimal objective values. We also compute the required sample size in practice to show the effectiveness of our approach.

2.5.2.1 Parameter setting

We use randomly generated instances to demosntrate the general features of our models and approaches. We randomly generate instances based on some specific demand distributions in all our test. Companies can cast their own problems with empirical data and specific parameters.

We test both i.i.d. demand and correlated demand. We consider the total number of periods $T = 5$, and assume stationary unit ordering cost, unit holding cost and unit backlogging cost in each period, which are $c = 5$, $h = 1$ and $p = 10$, respectively.

The i.i.d demand in each period follows Poisson distribution with mean value 20. For the correlated demand, we consider both Markov Modulated Demand Process (MMDP) (see, e.g., Chen and Song (2001)) and Autoregressive Model of Order 1 (AR(1)) (see, e.g., Mills (1991)).

The demands generated from MMDP have three states corresponding to the *state of economy*: poor (1), fair (2) and good (3). In each period t, given that the current state is $i_t \in \{1, 2, 3\}$, we test cases where the demand distributions are Poisson with mean value $10i_t$. We also assume that the state of the economy follows a Markov chain with the initial state 1 (i.e., poor) and the transition probability matrix

$$\mathcal{P} = \begin{pmatrix} 0.2 & 0.5 & 0.3 \\ 0.4 & 0.2 & 0.4 \\ 0.1 & 0.6 & 0.3 \end{pmatrix}.$$

For the AR(1) demand case, the demands in period t satisfy $d_t = d_{t-1} + \eta$, where the noise term η is normally distributed with mean 0 and standard deviation 1. We set the initial demand $d_0 = 20$. In all our computations, we test the problems with required service-level $\theta = 0.02$ and 0.05.

2.5.2.2 Feasible solutions

We aim at demonstrating the effectiveness of the SAA approach for finding feasible solutions. To compute for feasible solutions, we set the risk level $\alpha = 0$. This gives us a more conservative SAA counterpart problem and hence, a relatively small sample size N is needed to compute feasible solutions. Also, we numerically compute the solution to the SAA counterpart which uses the required service-level as the risk parameter (i.e., $\alpha = \theta$). We compare the statistics of using these two different methods.

Our numerical test consists of two parts. First, we generate N samples and solve the corresponding SAA instances each time. The above process is repeated $M = 10$ times so that we obtain 10 solutions for the same problem. The second part is to validate if all these 10 solutions are feasible. We conduct a *posteriori* check to compute the risk for each solution: we generate a simulation sample with $N' = 10,000$ scenarios, and check the number of scenarios that are violated under the larger problem for each given solution. The solution risk is then given by

$$\mathcal{R} = \frac{\text{number of violated scenarios}}{N'}.$$

If $\mathcal{R} < \theta$, the service-level requirements are satisfied; otherwise, the solution is not feasible. For solution risk, we report the average (Avg), minimum (Min), maximum (Max), and sample standard deviation (σ) over the solutions given by the 10 SAA problems. We also report the number of feasible solutions as well as the average, minimum, maximum, and sample standard deviation of the cost over these feasible solutions.

Tables 2.1–2.6 summarize our computational results for finding feasible solutions to the stochastic production planning problem. When not applicable, we indicate *** in the corresponding entry of each table. Our observations are summarized as follows:

1. From each table, we observe that as the sample size N grows, the average solution risk decreases and the number of feasible solutions increases. This is because as more samples are used, more constraints are being enforced into the model, which leads to a smaller feasible region. Hence, as the sample size grows, we can obtain more conservative solutions by solving the SAA problems which have lower solution risks and a higher likelihood to be feasible at the nominal risk level θ.

2. We observe that using $\alpha = 0$ requires much less samples to achieve a feasible solution than using $\alpha = \theta$. For example, in Table 2.1, we only need 300 samples to get a feasible solution with confidence level 90% by using $\alpha = 0$. However, solving the SAA reformulation at the nominal risk level $\alpha = 0.02$ requires at least 3000 samples to

guarantee a confidence level of 80%. Note that the problem size grows as the number of samples increases, we conclude that solving the SAA problem by setting risk level $\alpha = 0$ is more efficient than solving the original SAA problem at the nominal risk level in terms of obtaining feasible solutions. We further notice that the required sample sizes in our tests are also smaller than the theoretical bound given in Luedtke and Ahmed (2008). For instance, in Tables 2.3 and 2.4, the required sample size to achieve 90% confidence level is $N = 300$ and $N = 100$, respectively; on the other hand, the theoretical required sample sizes can be calculated by (2.18), which are 2500 and 400, respectively. The smaller sample size not only makes the computation more efficient, but also makes the data-collection work less demanding.

3. In terms of the costs for feasible solutions, we observe that using $\alpha = 0$ yields a higher average cost and a higher variance among all feasible solutions than using $\alpha = \theta$. For example, in Table 2.6 the average cost for feasible solutions is 640.7 using $\alpha = 0$ and $N = 50$ samples, as compared to the average cost of 605.58 by setting $\alpha = \theta$ and $N = 3000$. Hence, although using a smaller risk level $\alpha = 0$ is more efficient to compute a feasible solution under a given confidence level, it might yield a more conservative solution that has a higher cost than solving the SAA problems under the nominal risk level. Therefore, using a smaller risk level $\alpha = 0$ gives a feasible solution and an upper bound on the objective values efficiently.

Table 2.1: Solution results of i.i.d. demand for stochastic production planning problem without pricing for $\theta = 0.02$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	50	0.054	0.003	0.146	0.043	1	782.64	782.64	782.64	***
	100	0.025	0.003	0.054	0.015	4	764.37	715.65	818.83	40.12
	200	0.016	0.004	0.033	0.008	8	752.89	724.38	788.75	22.23
	300	0.011	0.005	0.021	0.005	9	771.58	744.26	808.36	19.02
	400	0.007	0.003	0.012	0.002	10	771.47	742.12	804.69	17.89
0.02	250	0.033	0.014	0.046	0.153	1	725.75	725.75	725.75	***
	500	0.028	0.017	0.044	0.152	3	713.58	709.00	716.47	3.28
	1000	0.025	0.016	0.035	0.132	3	715.34	714.16	716.43	0.93
	2000	0.021	0.017	0.029	0.004	4	712.39	710.30	718.24	3.38
	3000	0.019	0.017	0.024	0.002	8	711.26	707.06	714.28	2.36

Table 2.2: Solution results of i.i.d. demand for stochastic production planning problem without pricing for $\theta = 0.05$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	20	0.113	0.017	0.287	0.076	2	765.78	721.8	809.75	43.97
	60	0.051	0.008	0.084	0.024	4	714.18	691.37	740.97	18.71
	100	0.025	0.003	0.054	0.015	9	743.15	707.62	818.83	35.37
	140	0.021	0.003	0.041	0.011	10	733.92	698.41	790.60	25.23
0.05	500	0.064	0.048	0.080	0.010	1	671.83	671.83	671.83	***
	750	0.055	0.041	0.066	0.008	3	675.08	673.71	677.44	1.67
	1000	0.054	0.047	0.066	0.006	3	675.72	674.03	677.61	1.47
	2000	0.048	0.039	0.055	0.004	8	675.16	671.48	680.13	2.66

Table 2.3: Solution results of MMDP for stochastic production planning problem without pricing decisions for $\theta = 0.02$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	50	0.051	0.008	0.118	0.033	2	990.17	946.06	1034.28	44.11
	100	0.033	0.012	0.065	0.016	2	897.52	895.45	899.58	2.07
	200	0.019	0.003	0.036	0.011	5	919.38	846.43	980.29	45.91
	300	0.012	0.004	0.027	0.006	9	898.96	862.15	940.07	28.33
	400	0.006	0.001	0.012	0.004	10	944.01	896.75	1016.14	40.52
0.05	500	0.031	0.016	0.044	0.008	1	845.08	845.08	845.08	***
	1000	0.025	0.017	0.032	0.005	3	842.58	839.77	847.47	3.48
	2000	0.021	0.016	0.027	0.003	4	843.95	839.18	849.02	3.76
	3000	0.020	0.017	0.024	0.002	3	844.72	842.03	847.13	2.09

Table 2.4: Solution results of MMDP for stochastic production planning problem without pricing decisions for $\theta = 0.05$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	25	0.124	0.039	0.274	0.089	2	912.14	877.28	947.00	34.86
	50	0.064	0.025	0.129	0.034	4	853.81	828.44	885.00	22.60
	75	0.036	0.009	0.090	0.025	8	862.35	806.27	959.87	42.71
	100	0.033	0.012	0.065	0.016	9	857.77	807.34	899.58	35.42
	125	0.025	0.011	0.039	0.009	10	878.36	851.79	922.89	22.17
0.05	500	0.061	0.046	0.075	0.009	1	797.85	797.85	797.85	***
	1000	0.054	0.048	0.065	0.005	2	793.13	792.99	793.26	0.14
	2000	0.052	0.045	0.060	0.004	2	794.11	793.42	794.80	0.69

Table 2.5: Solution results of AR(1) demand for stochastic production planning problem without pricing decisions for $\theta = 0.02$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	25	0.057	0.012	0.094	0.025	1	640.96	640.96	640.96	***
	50	0.041	0.008	0.113	0.033	3	650.48	642.68	664.28	9.79
	75	0.023	0.010	0.048	0.012	6	647.59	638.01	665.27	10.11
	100	0.012	0.001	0.029	0.008	9	659.12	628.93	718.99	29.51
	125	0.010	0.001	0.019	0.006	10	659.67	630.07	720.3	28.09
0.05	500	0.025	0.015	0.042	0.008	2	632.46	630.9	634.02	***
	1000	0.02	0.013	0.026	0.005	5	631.43	626.43	634.18	2.94
	2000	0.018	0.016	0.022	0.002	8	627.75	626.10	633.12	2.17
	3000	0.018	0.016	0.022	0.002	7	627.94	626.14	629.87	1.03
	4000	0.018	0.016	0.02	0.001	10	627.31	625.81	628.49	0.82

Table 2.6: Solution results of AR(1) demand for stochastic production planning problem without pricing decisions for $\theta = 0.05$

		Solution Risk				Feasible Solutions Cost				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	20	0.080	0.017	0.157	0.053	4	641.93	624.05	682.55	23.65
	30	0.044	0.007	0.094	0.031	6	633.26	616.77	676.87	20.89
	40	0.035	0.003	0.091	0.030	7	643.36	614.33	672.73	18.30
	50	0.039	0.002	0.094	0.033	7	640.70	613.06	674.98	22.12
	60	0.020	0.001	0.037	0.012	10	638.14	610.63	685.65	21.62
0.05	500	0.055	0.038	0.091	0.016	5	606.78	604.76	609.82	1.68
	1000	0.049	0.041	0.058	0.006	5	606.33	604.60	608.62	1.56
	2000	0.048	0.042	0.055	0.004	7	605.51	604.57	607.32	0.87
	3000	0.046	0.043	0.051	0.003	8	605.58	604.42	606.82	0.65
	4000	0.045	0.043	0.048	0.002	10	605.33	604.06	606.36	0.66

2.5.2.3 Lower bounds

We study the lower bounds for the stochastic production planning problem by using $\alpha = \theta$. We take $L = 1, \ldots, 4$ to generate lower bounds by using the optimal objective values of the $M = 10$ SAA problems. According to (2.19), the confidence levels of using $L = 1, \ldots, 4$ are 0.999, 0.989, 0.945 and 0.828, respectively. In addition to the lower bounds computed at each confidence level, we also report optimality gaps, defined as the percentage that the lower bound is below the cost of best feasible solution (i.e., the minimum cost among all feasible solutions, given by Tables 2.1–2.6).

Tables 2.7–2.12 report the test results. Combining the test results on the lower bounds and the results of feasible solutions in Section 2.5.2.2, we can obtain the range of the optimal cost. For example, in the i.i.d. demand case with service-level $\theta = 0.02$, solving $M = 10$ SAA instances with sample size $N = 250$ yields a feasible solution of cost 725.75 shown in Table 2.1 while getting a lower bound 675.24 with a confidence level 0.999 shown in Table 2.7. This means, we have at least 99.9% confidence to say that the optimal solution is at most $(725.75 - 675.24)/725.75 \times 100\% \approx 6.96\%$ less costly than the best feasible solution we get. Similarly, we can analyze the problem with other demand cases and different service-level parameters using corresponding tables.

From these results, we observe that as sample size N becomes larger, the lower bound becomes larger and the gap becomes smaller at each confidence level. When the gap reaches zero, we come to a conclusion that the best feasible solution is the optimal solution with the corresponding confidence level. For example, as we notice from Table 2.12, when the sample size $N = 2000$, we have confidence at least 82.8% that the feasible solution of cost 604.57 is optimal; when the sample size raises to $N = 3000$, our confidence increases from 82.8% to 99.9%. Thus, for a certain confidence level, we can make a better estimation of the optimal solution of a SAA problem as we increase the sample size N.

Table 2.7: Lower bounds of i.i.d. demand for stochastic production planning problem without pricing for $\alpha = \theta = 0.02$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
250	675.24	683.10	683.35	684.52	6.96%	5.88%	5.84%	5.68%
500	675.83	691.38	691.94	692.90	4.68%	2.49%	2.41%	2.27%
1000	687.04	695.08	695.46	700.28	3.80%	2.67%	2.62%	1.94%
2000	694.73	698.35	699.20	702.01	2.19%	1.68%	1.56%	1.17%
3000	701.10	702.49	707.06	709.11	0.84%	0.65%	0.00%	–0.29%

Table 2.8: Lower bounds of i.i.d. demand for stochastic production planning problem without pricing for $\alpha = \theta = 0.05$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	655.94	655.97	658.63	660.89	2.36%	2.36%	1.96%	1.63%
750	662.20	663.93	664.02	667.83	1.71%	1.45%	1.44%	0.87%
1000	661.15	662.96	664.37	667.49	1.91%	1.64%	1.43%	0.97%
2000	668.37	669.04	671.48	673.31	0.46%	0.36%	0.00%	−0.27%

Table 2.9: Lower bounds of MMDP for stochastic production planning problem without pricing for $\alpha = \theta = 0.02$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	801.63	803.84	808.33	813.20	5.14%	4.88%	4.35%	3.77%
1000	812.40	813.36	817.94	824.20	3.26%	3.14%	2.60%	1.85%
2000	824.67	830.47	830.98	832.43	1.73%	1.04%	0.98%	0.80%
3000	830.88	835.12	836.84	837.06	1.32%	0.82%	0.62%	0.59%

Table 2.10: Lower bounds of MMDP for stochastic production planning problem without pricing for $\alpha = \theta = 0.05$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	768.09	770.16	770.58	774.45	3.73%	3.47%	3.42%	2.93%
1000	775.65	780.50	780.80	782.01	2.19%	1.57%	1.54%	1.38%
2000	778.77	780.37	784.50	784.76	1.85%	1.64%	1.12%	1.09%

Table 2.11: Lower bounds of AR(1) demand for stochastic production planning problem without pricing for $\alpha = \theta = 0.02$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	612.87	614.25	615.61	617.04	2.86%	2.64%	2.42%	2.20%
1000	620.22	620.25	620.81	621.43	0.99%	0.99%	0.90%	0.80%
2000	622.00	625.37	626.10	626.41	0.66%	0.12%	0.00%	−0.05%
3000	621.54	625.20	625.62	626.14	0.73%	0.15%	0.08%	0.00%
4000	625.81	626.44	626.92	626.95	0.00%	−0.10%	−0.18%	−0.18%

Table 2.12: Lower bounds of AR(1) demand for stochastic production planning problem without pricing for $\alpha = \theta = 0.05$

	LB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	584.34	597.29	600.79	600.91	3.38%	1.23%	0.66%	0.64%
1000	599.22	600.58	601.15	602.62	0.89%	0.66%	0.57%	0.33%
2000	600.56	602.23	603.20	604.57	0.66%	0.39%	0.23%	0.00%
3000	604.42	604.80	605.14	605.28	0.00%	−0.06%	−0.12%	−0.14%
4000	604.06	604.44	604.98	605.14	0.00%	−0.06%	−0.15%	−0.18%

2.5.3 Production Planning with Pricing Options

In this section, we report the computational results of multi-period joint service-level constrained production planning with pricing options. We also conduct sensitivity analysis for this model.

2.5.3.1 Parameter setting

Consider the continuous pricing in the test instances. The setting of cost parameters is the same as those in Section 2.5.2.1. Moreover, we set $a_t = -5$ and $\beta_t = 200$ in the function $d_t(r_t) = a_t r_t + \beta_t + \tilde{\epsilon}_t$ for all $t = 1, \dots, T$. The noise term $\tilde{\epsilon}_t$ follows normal distribution with mean 0 and standard deviation 22 for all $t = 1, \dots, T$. We also assume that the pricing range in each period t is between $W_t^L = 18$ and $W_t^U = 40$. We fix the required service-level $\theta = 0.02$.

2.5.3.2 Feasible solutions

Table 2.13 reports statistics of the solutions of i.i.d. demand for production planning with pricing options. Table 2.14 reports statistics of the solutions of AR(1) demand for production planning with pricing options. The insights of our numerical results are summarized as follows:

1. As the sample size N grows, the average solution risk decreases and the number of feasible solutions increases since more constraints are being enforced into the model, which leads to a smaller feasible region. Hence, as the sample size increases, solving the SAA counterpart under any fixed risk level yields a lower solution risk and a higher likelihood to be feasible at the nominal risk level θ.

2. From Table 2.13 and Table 2.14, we observe that using $\alpha = 0$ requires much fewer

samples to achieve a feasible solution than using $\alpha = \theta$. For example, in Table 2.13, we generate 300 samples to obtain a feasible solution with confidence level 100% by using $\alpha = 0$. However, even using 2500 samples in the SAA reformulation, with risk level $\alpha = 0.02$, can only find a feasible solution at a confidence level of 60%. In our numerical test, solving a SAA reformulation that involves more than 2500 samples is computationally intractable (more than three CPU minutes for each instance). Therefore, an efficient way to compute a feasible solution is to solve the SAA reformulation with a more conservative risk level $\alpha = 0$. The smaller sample size not only makes the computation more efficient, but also makes the data-collection work less demanding.

3. In terms of the profit for feasible solutions, we observe that using $\alpha = 0$ yields a lower average profit and a higher variance among all feasible solutions than using $\alpha = 0.02$. For example, in Table 2.14 the average profit for feasible solution is 6365.02 using $\alpha = 0$ and $N = 200$ samples, as compared to the average profit of 6567.83 by setting $\alpha = 0.02$ and $N = 2000$. Hence, although using a smaller risk level $\alpha = 0$ is more efficient to compute a feasible solution under a given confidence level, it will yield a more conservative solution that has a lower profit than solving the SAA problems under the nominal risk level. Therefore, using a smaller risk level $\alpha = 0$ helps us to find a feasible solution and a lower bound on the total profit efficiently.

Table 2.13: Solution results of i.i.d. demand for production planning with pricing for $\theta = 0.02$

		Solution Risk				Profit for Feasible Solutions				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0.00	50	0.047	0.020	0.148	0.036	1	6652.42	6652.42	6652.42	***
	100	0.024	0.008	0.045	0.010	4	6479.34	6308.48	6581.04	104.64
	200	0.017	0.008	0.027	0.006	8	6534.51	6362.91	6678.65	100.52
	300	0.010	0.005	0.020	0.004	10	6438.01	6261.56	6647.29	119.30
0.02	500	0.026	0.018	0.035	0.005	1	6655.84	6655.84	6655.84	***
	1000	0.024	0.016	0.034	0.004	2	6677.97	6644.61	6711.33	33.36
	1500	0.021	0.014	0.026	0.004	2	6614.26	6612.54	6615.97	1.72
	2000	0.022	0.018	0.027	0.003	3	6645.73	6627.55	6658.49	13.20
	2500	0.021	0.017	0.027	0.003	6	6658.29	6634.35	6685.84	16.75

2.5.3.3 Upper bounds

We check the upper bounds for production planning with pricing options when $\alpha = \theta = 0.02$. The gaps are defined as the percent by which the upper bound is above the best feasible solution (i.e., the maximum profit among all feasible solutions in this case). We

Table 2.14: Solution results of AR(1) demand for production planning with pricing for $\theta = 0.02$

		Solution Risk				Profit for Feasible Solutions				
α	N	Avg	Min	Max	σ	#	Avg	Min	Max	σ
0	50	0.038	0.011	0.090	0.023	2	6425.48	6215.96	6635.01	209.52
	100	0.024	0.006	0.049	0.014	6	6453.05	6242.66	6632.97	154.77
	200	0.018	0.006	0.035	0.009	7	6365.02	6203.31	6576.07	123.50
	300	0.008	0.002	0.017	0.004	10	6261.79	5955.75	6414.22	127.06
0.02	1000	0.024	0.018	0.032	0.003	1	6602.79	6602.79	6602.79	***
	1500	0.022	0.018	0.026	0.002	2	6561.49	6546.36	6576.63	15.14
	2000	0.022	0.018	0.024	0.002	3	6567.83	6536.66	6610.64	31.30
	2500	0.022	0.018	0.028	0.003	4	6566.65	6529.12	6594.29	28.15

use $L = 1, \ldots, 4$ to generate bounds by the optimal objective values of the $M = 10$ SAA problems and the corresponding confidence levels given by (2.19) are 0.999, 0.989, 0.945, 0.828, respectively.

Table 2.15 shows the upper bounds for i.i.d. demand for production planning with pricing options when $\alpha = \theta = 0.02$. Table 2.16 shows the upper bounds for AR(1) demand for production planning with pricing options when $\alpha = \theta = 0.02$.

Combining the test results on the upper bounds and the results of feasible solutions in Section 2.5.3.2, we can obtain the range of the optimal profit. For example, in the i.i.d. demand case, solving $M = 10$ SAA instances with sample size $N = 1000$ yields a best feasible solution with profit 6711.33, as shown in Table 2.13; we also get an upper bound of 6740.62 with confidence 98.9% (shown in Table 2.15). This means, we have at least 98.9% confidence to say that the optimal profit is at most $(6740.62 - 6711.33)/6740.62 \approx 0.44\%$ greater than the best feasible solution 6740.62. We can make a similar analysis with other demand cases using corresponding tables.

From these results, we observe that as sample size N becomes larger, the upper bound becomes smaller and the gap becomes smaller at each confidence level. When the gap reaches zero, we can conclude that the best feasible solution is the optimal solution with the corresponding confidence level. For example, as we observe from Table 2.16, when the sample size $N = 1000$, we have confidence at least 94.5% to say that the optimal profit is at most 0.1% greater than 6609.35; when the sample size increases to $N = 2000$, we have confidence at least 94.5% to say that the feasible solution with profit 6610.64 is optimal. Thus, for a certain confidence level, we can make a better estimation of the optimal solution of a SAA problem as we increase the sample size N.

Table 2.15: Upper bounds of i.i.d. demand for production planning with pricing for $\alpha = \theta = 0.02$

	UB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
500	6741.58	6723.42	6695.46	6692.89	1.29%	1.02%	0.60%	0.56%
1000	6767.01	6740.62	6711.33	6688.30	0.83%	0.44%	0.00%	−0.34%
1500	6702.55	6693.99	6689.16	6668.40	1.31%	1.18%	1.11%	0.79%
2000	6687.78	6687.74	6670.77	6658.49	0.44%	0.44%	0.18%	0.00%
2500	6696.60	6696.56	6686.90	6685.84	0.16%	0.16%	0.02%	0.00%

Table 2.16: Upper bounds of AR(1) demand for production planning with pricing for $\alpha = \theta = 0.02$

	UB with confidence at least				Gap with confidence at least			
N	0.999	0.989	0.945	0.828	0.999	0.989	0.945	0.828
1000	6649.49	6621.39	6609.35	6606.17	0.71%	0.28%	0.10%	0.05%
1500	6659.58	6628.80	6623.04	6582.9	1.26%	0.79%	0.71%	0.10%
2000	6653.04	6636.38	6610.64	6605.67	0.64%	0.39%	0.00%	-0.08%
2500	6629.96	6613.79	6594.29	6593.37	0.54%	0.30%	0.00%	-0.01%

2.5.3.4 Sensitivity analysis

We conduct sensitivity analysis for production planning with pricing options. We focus on parameters a_t and β_t in the function $d_t(r_t) = a_t r_t + \beta_t + \tilde{\epsilon}_t$ for $t = 1, \ldots, T$. To better demonstrate our sensitivity results, we assume that the demand function is time invariant, i.e., $a_t = a$ and $\beta_t = \beta$ for all $t = 1, \ldots, T$. The noise term $\tilde{\epsilon}_t$ follows normal distribution with mean 0 and standard deviation 22 for all $t = 1, \ldots, T$. We use $\theta = 0.02$, $\alpha = 0$ and $N = 300$ since they can yield a feasible solution with high confidence level to the model, as shown in Table 2.13.

Sensitivity analysis on the slope a. We fix $\beta = 200$ and vary a in the set $\{-1, \ldots, -10\}$ in each period. Our test result is shown in Figure 2.1.

As shown in Figure 2.1, the total profit increases as the absolute value of a decreases. Moreover, in each period, the optimal order quantity and the optimal price increase as the absolute value of a decreases.

As the absolute value of a decreases, the demand will be less sensitive to the price set by the retailer; hence, the retailer has the motivation to increase the price and still keep the demand to the same level; as a result, the total profit increases.

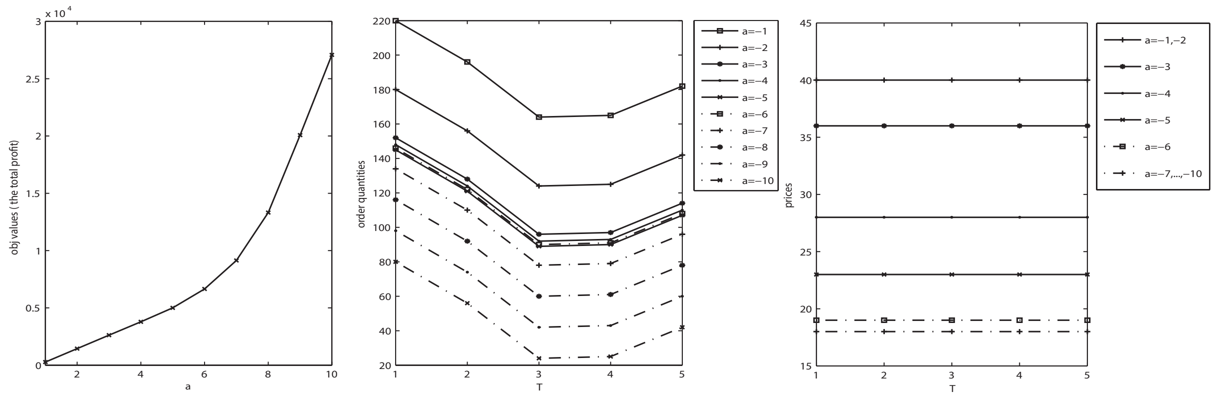

Figure 2.1: Sensitivity analysis on the slope a

Sensitivity analysis on β. We fix $a = -5$ and vary β in the set $\{160, \ldots, 250\}$. Our test result is shown in Figure 2.2.

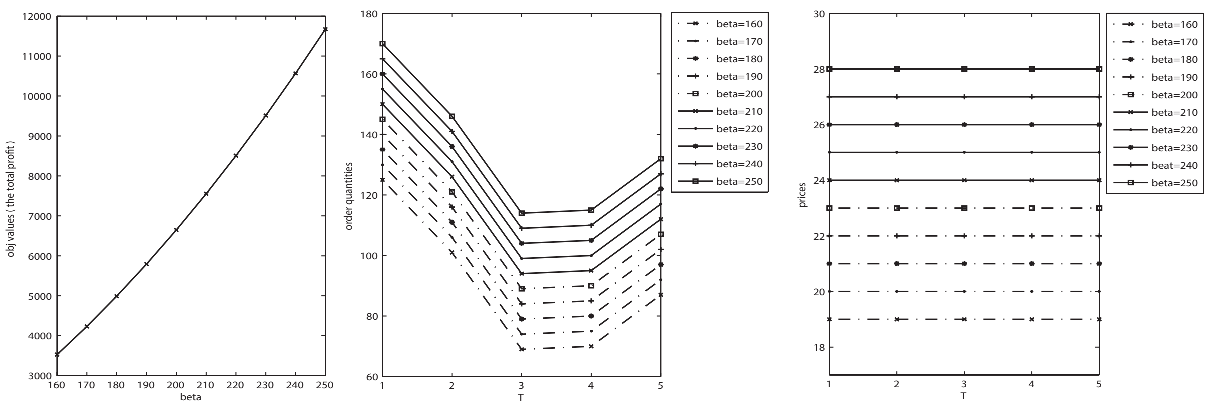

Figure 2.2: Sensitivity analysis on β

As shown in Figure 2.2, the total profit, the optimal order quantity and the optimal price in each period all increase as the value of β increases.

As the value of β increases, the basic demand increases; hence, the retailer has motivation to increase the price and still keep the demand in a higher level. The higher level demand leads to more order quantities. Consequently, the total profit increases.

2.6 Concluding Remarks

In this chapter, we propose two models of production planning problem under a probabilistic service-level guarantee (interpreted as stockout probabilities) over the entire planning horizon. The first model is an inventory management model while the second one also involves pricing decisions. We reformulate these two models as mixed-integer programs based on a finite set of discrete samples of the uncertainty, and solve them by using the SAA method. However, the resulting MILP models are computationally intractable since the SAA method requires a very large sample size; we computationally obtain feasible solutions and lower bounds on these models by adjusting the risk parameter, which gives us an efficient way to bound the optimal cost and the optimal profit. We conduct extensive computational tests under different service-level requirements and demand cases, so as to demonstrate the feasible solutions and lower bounds as well as to suggest reasonable sample sizes in practice.

An interesting direction for future research is to study the dynamic version of our model in which the price and inventory in each period can be changed dynamically at the beginning of each period. Such model solves dynamic pricing problems by using dynamic programming formulation rather than static MILP formulation. This is mainly due to the reason that MILP formulations for a dynamic problem require to grow a scenario tree to represent decisions based on the system states, which could be intractable to compute. (see, e.g., Huang and Ahmed 2005, 2008).

CHAPTER III

Stochastic Inventory Control Under Service-Level Constraints

3.1 Abstract

Motivated by the importance of service quality in today's customer business environment, we consider two periodic-review stochastic inventory models with *probabilistic service-level guarantees* for restricting stockout probabilities: (i) the classical inventory control model with backlogging and (ii) the remanufacturing inventory control model with random product returns. We assume that demands are stochastic nonstationary, evolving and correlated over time. We establish the optimality of generalized base-stock policies, and propose new 2-approximation algorithms for both models. The core concept developed in this chapter is called *the delayed forced holding and production cost*, which is proven effective in dealing with service-level constrained inventory systems. Our extensive computational experiments show that the proposed algorithms on average perform within 2% error of optimality.

3.2 Introductory Remarks

In today's customer-driven business environment, it is vital for companies to focus on the QoS. Since the early 2000s, firms started to put tremendous effort and resource into understanding the customers and the markets. Those who could consistently provide a superior service to their customers would remain an excellent reputation and keep most of its loyal buyers. Customers facing stockouts have been observed abandoning their purchases, switching retailers, substituting similar items and have seldom gone back (see, e.g., Fitzsimons (2000)). One of the most common challenges in making supply chain decisions, at its most fundamental, boils down to minimizing the supply chain costs while still delivering great customer service.

On one hand, the notion of service-level requirement has been widely used both in theory and in practice to measure the performance of inventory replenishment policies (cf. Ghiani et al. 2005b). It is typically defined as a probabilistic constraint so that the demand is satisfied with a high probability. By enforcing a service-level requirement, companies are able to improve the QoS by guaranteeing a small stockout rates. There are several empirical studies of the sensitivity of inventory service-level on demand in business-to-consumer settings (cf. Fitzsimons 2000, Anderson et al. 2006, Jing and Lewis 2011). In particular, according to Jing and Lewis (2011), stockout rates have a dramatic impact on the firm's profitability and the firm can achieve many of the benefits through small decreases in stockout rates.

There are abundant examples in practice where the service-level plays an important role in firms' supply chain management. For example, online grocery store (such as AmazonFresh) generally has a very high service-level expectation, especially for its dairy product section. When a customer wants to buy 2% reduced-fat milk, it must be available with a very high probability. If not, the store runs the risk of losing the sale as well as the customer. Clearly, customers are more willing to buy from those grocery stores who always have enough stock. This directly explains why many grocery stores generally enforce a high service-level (from 85% to 98%) on dairy products. Their optimal inventory replenishment policy must meet the service-level requirement while minimizing the total inventory cost over the planning horizon. Likewise, many other industries, such as food and fashion, also set a high service-level to maintain a high in-stock inventory, which helps satisfy customers' demands and avoid stockouts. In addition, service-level agreements (SLA) are usually enforced in some industries such as the semi-conductor industry to guarantee the delivery of manufactured products. As pointed out by Katok et al. (2008), SLAs are used to improve supply chain coordination and there are contractual financial penalties and rewards associated with failing or achieving a target service-level. In general, having a service-level requirement helps firms maintain their reputation and increase their revenue in the long run (see Chen and Krass (2001) for more examples and discussions).

On the other hand, after-sales services could also be crucial in delivering great customer services (cf. Cohen et al. 2006). Companies have to handle the return, repair, and disposal of failed components. The returned products, though some parts may be damaged, can be remanufactured and resold. The remanufacturing process includes repair or replacement of worn-out or obsolete components and modules, which has a lower production cost than the manufacturing process. Examples of remanufacturing occur in many industries, such as personal computers, cell phones, automotive parts, etc. For example, Apple Inc., one of the world largest technology companies, produces several electronic products including iPhone, iPad and Mac. The firm provides excellent after-sales customer service by allowing customers return their products within two weeks of purchase without paying re-stocking fee. Each year, Apple Inc. receives millions of product returns and they remanufactured the returned products by replacing defective parts with new parts. In addition, Apple also imposes a high service-level requirement on their long-term supply to retain loyal customers and providing products to them immediately. The company is much more willing to make extra productions and hold them as inventory than having a lost-sales which may potentially jeopardize their reputation. Both the service-level requirement and its remanufacturing process help Apple build up an excellent customer service and earn them a great reputation among other online competitors, which benefits the company in the long run (cf. Gallo 2012).

To address all the aforementioned issues in inventory management, we study periodic-review *service-level constrained* stochastic inventory systems where the stockout probability is lower bounded by a threshold value in each period. This type of service-level constraint is commonly known as the α-service-level in the literature (see, e.g., Simchi-Levi et al. (2014), Snyder and Shen (2011), Chen and Krass (2001)). We consider two fundamental stochastic inventory models with α-service-level constraints: the *multi-period backlogging model* and the *multi-period backlogging model with remanufacturing*, with a general stochastic demand process (i.e., correlated, nonstationary and evolving demand). In the service-level constrained backlogging model, the firm makes a production decision in each period to minimize the total expected production, holding and backlogging costs over a finite planning horizon, subject to a given service-level requirement. In the counterpart model with remanufacturing, in addition to the regular production, there are some products being returned at the beginning of each period (commonly referred to as cores, see, e.g., Tao and Zhou (2014)), which can be remanufactured into regular products at a lower cost. The objective is to decide the manufacturing and remanufacturing quantities in each period so as to minimize the total expected costs, subject to a given service-level requirement.

As seen from our literature review below, there has also been growing research on both the theoretical and computational aspects of service-level constrained inventory systems. There are mainly two sources of motivation. First, traditional inventory models usually assume linear cost functions to penalize inventory, backorders, or lost sales. However, the assumption of linear backlogging or lost-sale penalty is primary for analytical tractability rather than an accurate representation of reality (see Bertsimas and Paschalidis (2001) for a detailed discussion). The mechanism of varying unit penalty costs can hardly take effect on the QoS performance of a system, mainly due to the difficulty of quantifying customer satisfaction. In this regard, imposing a target service-level is a much more direct way to quantify and improve the QoS performance of an inventory system. Second, as extensively discussed in Chen and Krass (2001), the backlogging cost is often very difficult to quantify in practice. Hence, a target service-level constraint is thus considered as an effective (if not more so) alternative performance measure.

In this chapter, we consider a generalized model that incorporates both the service-level constraints and the penalty cost for each unsatisfied demand, which has several advantages. First, when the firm does not have a good estimate of the backlogging cost, the firm can simply set the per-unit penalty cost to be zero in our model, which then reduces to the conventional model with service-level constraints only. Our algorithms, analysis, and results hold for this special case. Second, when the firm does have a good estimate of the backlogging cost (e.g., from historical data), considering backlogging cost together can significantly alleviate the problem of suffering from severe backlogs in the worst-case scenarios (since the service-level constraints only guarantee that the demand will be met with a certain positive probability in each period).

3.2.1 Main Results and Contributions

We consider a *service-level constrained* backlogging model and a remanufacturing model with random product returns. The demand and return processes can be non-stationary and generally correlated. We summarize the key results as follows.

(a) We establish the optimality of a base-stock policy for the backlogging model (Proposition 3.1). We also establish the optimality of a total base-stock policy for the remanufacturing model (Proposition 3.6); more specifically, in each period, we bring the total inventory position (after production) to an optimal threshold level by remanufacturing as many returned products as possible.

(b) Finding the exact optimal policy via a brute-force dynamic programming (DP) approach is computationally intractable, despite its simple form. We propose two new approximation algorithms (named *Split-Merge-Balance (SMB) algorithms*) for both backlogging and remanufacturing models to efficiently compute near-optimal solutions. Our analysis shows that both algorithms admit a worst-case performance of two, i.e., the expected cost of our algorithm is at most two times the expected cost of an optimal solution (Theorems 3.5 and 3.10). Through testing a large set of demand instances, we demonstrate via extensive simulation that the SMB policies perform near-optimal (within 2% error of optimality) and also yield significant reduction of solution time.

The core new concept developed in this chapter is the notion of *delayed forced costs* in designing provably-good policies for service-level constrained inventory systems. The major difficulty of designing approximation algorithms for service-level constrained models is the impossibility of balancing the expected marginal holding costs against the expected marginal backlogging costs, which is the dominated technique (i.e., *cost-balancing technique*) in some related literature (see, e.g., Levi et al. 2007, 2008a,b). Our algorithms first split the marginal costs into two parts (i.e., *forced costs* and *future costs*) and shift all the forced costs to one period later (called *delayed forced costs*); after regrouping the future costs and delayed costs according to their monotonicity, our algorithms balance the expected overage cost against the expected underage cost associated with each period. The introduction of the delayed forced costs makes the worst-case analysis invariably harder – one needs to group consecutive intervals together to amortize the sum of future cost and delayed forced cost against the optimal policy (Lemma 3.4). This is in sharp contrast to the aforementioned studies where period-by-period amortization is sufficient for the classical backlogging model.

For the remanufacturing model, the amortization of production costs becomes non-trivial and our worst-case analysis builds upon on the elegant partitioning technique introduced in Tao and Zhou (2014). There is a challenge we need to overcome: due to the service-level constraints in our model, we split each holding cost and production cost into two parts and use the delayed holding cost and delayed production cost when designing a modified SMB algorithm. Consequently, our worst-case analysis needs to bound both parts of production costs in different sets of periods, which cannot be readily adapted from their results (see the detailed discussions before Lemma 3.9).

In general, we believe that the concept of delayed forced costs could be widely applied in designing algorithms for any general service-level constrained inventory systems.

3.2.2 Literature Review

This chapter is closely related to the following research domains and the related literature.

Stochastic inventory system with service-level constraints. Despite the huge body of literature on stochastic inventory theory (cf. Zipkin 2000), surprisingly, only a few papers studied stochastic inventory systems with probabilistic constraints that incorporate service-level guarantees. Among them, Bookbinder and Tan (1988) studied a probabilistic lot-sizing problem using a "static-dynamic uncertainty" strategy. In their two-stage model, a retailer first decided a schedule (or replenishment periods) to place orders. Then, she made adjustments to the planned orders when demand was released. Chen and Krass (2001) showed that the (s, S) policy is optimal under independently and identically distributed (i.i.d.) demands for an infinite-horizon stationary setting. Boyaci and Gallego (2001) proposed effective heuristic procedures to serial inventory systems with service-level constraints on stockout probabilities. Shang and Song (2006) also studied a serial base-stock inventory model under simple Poisson demands and the same type of service-level constraints. They developed a closed-form approximation for the optimal base-stock levels. Bertsimas and Paschalidis (2001) considered a multiclass make-to-stock manufacturing system with probabilistic service-level guarantees, and devised a production policy that minimized inventory costs under a stockout probability guarantee using queueing methods. Goh and Sim (2011) carried out a computational study (using a software called ROME) of a distributionally robust periodic-review inventory problem with fill rate constraints. More recently, Wei et al. (2017) studied a periodic-review service-level constrained inventory system with lost-sales and lead times. They proposed a simple heuristic by solving a linear programming (LP) problem derived from a deterministic inventory model with backlogging, and showed that the proposed heuristic is asymptotically optimal as the service-level grows to 100%, and derived a simple and explicit bound on the optimality gap. The probabilistic constraints that impose service-level guarantees in each period (i.e., α-service-levels) are used in the majority of the literature, which is also the primary focus of this chapter.

Stochastic inventory system with remanufacturing. Zhou et al. (2011) studied the structure of optimal policies for the remanufacturing inventory system with multiple types of returned products. They showed that the optimal policy is a modified base-stock policy, which can be completely characterized by a sequence of control parameters. Tao and Zhou (2014) proposed an approximation algorithm for the stochastic inventory system with remanufacturing. They also proved that the cost of their proposed algorithm is at most twice of the optimal cost. Gong and Chao (2013) focused on the capacitated inventory systems with remanufacturing. Building upon the preservation result by Chen et al. (2013), they showed that the optimal remanufacturing policy is a modified remanufacture-down-to policy and the optimal manufacturing policy is a modified total-up-to policy. Our remanufacturing model differs from all of the aforementioned models by incorporating service-level constraints in each period.

Approximation algorithms on inventory systems. The DP approach is effective in characterizing the structural properties of optimal policies. However, the computational complexity is very sensitive to the dimension of the state space. In fact, it has been shown in Halman et al. (2009) that the stochastic lot-sizing model (without service-level constraints) is NP-hard. Our work is closely related to recent studies of approximation algorithms for stochastic periodic-review inventory systems pioneered by Levi et al. (2007). They first introduced the marginal cost accounting scheme, which associated a cost with each decision made by a particular policy. They proposed a dual-balancing policy which admitted a worst-case performance guarantee of 2 for the backlogging model. Subsequently, Levi et al. (2008a,b) and Levi and Shi (2013) proposed approximation algorithms for the lost-sales, capacitated, and lot-sizing models, respectively. More recently, Truong (2014) re-derived the 2-approximation ratio for the backlogging model via a look-ahead optimization approach. Tao and Zhou (2014) proposed an approximation algorithm for a remanufacturing system with a worst-case performance guarantee of two. Cheung et al. (2016), Nagarajan and Shi (2016) proposed approximation algorithms for submodular joint replenishment problems. There has also been a series of studies on perishable/fresh inventory systems (see, e.g., Chao et al. (2015, 2017), Zhang et al. (2016)). However, none of these papers imposed service-level constraints in their models while our work focuses on designing approximation algorithms for inventory models with service-level guarantees.

3.2.3 Organization and General Notation

We organize the remainder of the chapter as follows. In Section 3.3, we formulate the service-level constrained backlogging model as a dynamic program and present the structural properties of an optimal policy. In Section 3.4, we derive a 2-approximation algorithm for the classical backlogging model. Section 3.5 extends our structural result and approximation algorithm to the remanufacturing system. Section 3.6, we carry out numerical experiments and demonstrate the effectiveness of our proposed policy. Finally, Section 3.7 concludes the chapter and presents future research avenues.

Throughout the chapter, we use increasing and decreasing in a non-strict sense. For notational convenience, we use a capital letter and its lower-case form to distinguish between a random variable and its realization. We use $\triangleq$ to mean "is defined as", and $\mathbb{1}(A)$ is the indicator function taking value 1 if statement "A" is true and 0 otherwise. For any $x \in \mathbb{R}$, we denote $x^+ = \max\{x, 0\}$. For any sequence $x_i,\ i = 1, 2, \ldots$, we let $x_{[i,j]} = \sum_{k=i}^{j} x_k$ and $x_{[i,j)} = \sum_{k=i}^{j-1} x_k$, where the summation over an empty set is defined as 0. For any $a,\ b \in \mathbb{R}$, we denote $a \wedge b = \min\{a, b\}$.

3.3 Periodic-Review Inventory Systems with Service-Level Constraints

We formally describe a periodic-review service-level constrained inventory system. Consider a finite planning horizon of T periods indexed by $t = 1, \ldots, T$. The production lead time is assumed to be $L \geqslant 0$.

Demand structure. We adopt the same demand structure as in Levi et al. (2007) and Tao and Zhou (2014). An information set f_t is observed at the beginning of each period t. It contains all the available information that can be used to predict future demands, such as the realized demands $(d_1, \ldots, d_{t-1})$ and possibly some other exogenous information (denoted by ρ_t at period t). For example, when the state of economy is observed at each period, ρ_t corresponds to the state of economy at period t. The conditional joint distribution of the future demands $(D_t, \ldots, D_T)$ is determined by $f_t = (d_1, \ldots, d_{t-1}, \rho_1, \ldots, \rho_t)$. We denote $\mathcal{F}_t$ as the set of all the possible realizations of the information set f_t. Our demand model generalizes the existing correlated demand models in the literature, such as Markov-modulated demand process (MMDP) (Sethi and Cheng 1997), autoregressive demand (Mills 1991), which will be described in Section 3.6. Our demand model is also useful in practice, in which martingale model of forecast evolution (MMFE for short, see, e.g., Graves et al. 1986, Heath and Jackson 1994b) and advance demand information (ADI) (see, e.g., Gallego and Özer 2001) are used to forecast the future demand.

Service-level requirements. Our model incorporates the service-level requirement, which is defined as a probabilistic constraint for each period t. Following Chen and Krass (2001), the constraint enforces that the demand in each period $t + L$ is satisfied by a certain probability θ_t. (Note that θ_t is in fact associated with period $t + L$.) Mathematically, it is given by

$$\mathbb{P}(NI_{t+L} \geqslant 0 \mid f_t) \geqslant \theta_t, \quad \forall t = 1, \cdots, T, \tag{3.1}$$

where NI_t denotes the net inventory at the end of period t, which can be either positive (in the presence of on-hand inventory) or negative (in the presence of backorders). This type of service-level constraint is commonly referred to as the α-service-level in the literature (see §2.3.1. for a detailed discussion in Chen and Krass (2001)). This type of α-service-level has also been discussed in Bookbinder and Tan (1988), Nahmias (1993) and Sethi and Cheng

(1997). Note that our production decisions will not affect net inventories for the first L periods, and hence we enforce the service-level requirements from period $L+1$ to $L+T$.

System dynamics. In each period t, events occur in the following sequence: First, the manufacturer determines the production quantity (denoted by q_t) in period t. The planned production quantity should satisfy the service-level constraint (3.1). Then d_t units of demands are received. As a consequence, unused products are stored as inventory, which incurs the holding cost; on the other hand, unsatisfied demands incur the backlogging cost and are carried to the next period. The production, holding and backlogging cost functions are assumed to be linear, with per-unit costs c_t, h_t and b_t, respectively. The goal is to decide production quantities that achieve the required service-levels in each period and to minimize total expected cost at the same time.

A dynamic programming formulation. We formulate the problem using dynamic programming (DP) approach. Since no products will arrive in the first L periods, it suffices to consider the total cost from period $L+1$ to period $L+T$. We first calculate the immediate cost associated with decisions in period t. Define x_t as the inventory position at the beginning of period t, which equals to the current on-hand inventory plus the pipeline inventory minus the backorders, i.e., $x_t = NI_{t-1} + q_{[t-L,t)}$. Then the inventory position in the next period equals to the current inventory position plus production quantity minus the demand in the current period, i.e., $x_{t+1} = x_t + q_t - d_t$. Let $y_t = x_t + q_t \geqslant x_t$ be the inventory position in period t after production. Then the net inventory at the end of period $t+L$ can be written as $NI_{t+L} = y_t - D_{[t,t+L]}$. Therefore, the total holding and backlogging cost is given by

$$G_t(y_t, f_t) \triangleq h_{t+L}\mathbb{E}[(y_t - D_{[t,t+L]})^+ \mid f_t] + b_{t+L}\mathbb{E}[(D_{[t,t+L]} - y_t)^+ \mid f_t].$$

It is clear that G_t is continuous and convex in its first component. Adding the production cost, the total cost in period t is therefore $G_t(y_t, f_t) + c_t q_t$. We can also rewrite the service-level constraint defined in (3.1) as

$$\mathbb{P}(y_t - D_{[t,t+L]} \geqslant 0 \mid f_t) \geqslant \theta_t, \quad \forall t = 1, \cdots, T. \tag{3.2}$$

To simplify the above constraint, we define a threshold value

$$r_t = r_t(f_t) \triangleq \inf\{r \in \mathbb{R}^+ : \mathbb{P}(D_{[t,t+L]} \leqslant r \mid f_t) \geqslant \theta_t\},$$

i.e., $r(f_t)$ is the θ_t-quantile of the random variable $D_{[t,t+L]}$ given information f_t. Then the service-level constraint (3.2) is equivalent to $y_t \geqslant r_t$ for all $t = 1, \ldots, T$. In some special cases, r_t is very easy to compute. For example, when demands are independent Poisson or Normal random variables, the aggregated demand $D_{[t,t+L]}$ follows the same distribution type. For the demands following general distributions, we can use the Monte Carlo sampling method Glasserman (2004) to empirically obtain r_t.

Let $v_t(x_t, f_t)$ be the minimal expected cost from period $t + L$ to period $T + L$ given the inventory position x_t and the information set $f_t \in \mathcal{F}_t$ at the beginning of period t. The Bellman's equations are

$$\begin{aligned} &v_{T+1}(x_{T+1}, f_{T+1}) = 0, \quad \forall x_{T+1} \in \mathbb{R}, f_{T+1} \in \mathcal{F}_{T+1}, \\ &v_t(x_t, f_t) = \min_{y_t \geqslant \max\{r_t, x_t\}} \Big\{ G_t(y_t, f_t) + c_t q_t + \mathbb{E}[v_{t+1}(y_t - D_t, \mathcal{F}_{t+1}) \mid f_t] \Big\}, \quad t = 1, \ldots, T. \end{aligned} \tag{3.3}$$

Structure of optimal policies. Using (3.3), the structure of optimal policies is characterized in Proposition 3.1.

Proposition 3.1. *For the inventory control problem defined in* (3.3)*, an optimal policy is a modified base stock policy. More specifically, there exists* $\{s(f_t)\}_{t=1}^T$ *such that*

$$y_t^*(x_t) = \begin{cases} \max\{r_t, s(f_t)\}, & \text{if } x_t < s(f_t); \\ \max\{r_t, x_t\}, & \text{if } x_t \geqslant s(f_t). \end{cases}$$

Proof. This is a special case of Proposition 3.6. □

Proposition 3.1 asserts that any optimal policy has the following structure: if the inventory position in period t is no less than the threshold $s(f_t)$, an optimal policy produces up to the service-level r_t; otherwise, it brings the inventory position to $\max\{r_t, s(f_t)\}$. Therefore, the higher the service-level, the more orders are placed by the optimal policy. In particular, when there is no service-level requirement presented in the model (i.e., $r_t = 0$), the structure of optimal policy reduces to the well-known *base stock policy* (see, e.g., Zipkin (2000)).

3.4 Provably-Good Policies for Service-Level Constrained Inventory Systems

Computing an exact optimal policy through a brute-force DP model is generally intractable under correlated demand structures, despite the simple structure of optimal policies. To this end, we provide an approximation algorithm, called *Split-Merge-Balance policy* (denoted by the SMB policy), for practically solving the service-level constrained inventory problem. We show that the SMB policy admits a worst-case performance guarantee of 2, i.e., the expected cost of the policy is at most twice the expected cost of an optimal policy, regardless of any distributions of the random demands and choices of the cost parameters. Moreover, the SMB policy performs empirically near-optimal, demonstrated by extensive numerical tests in §3.6.

The traditional inventory cost accounting scheme (in dynamic programming) decomposes the total costs by periods. In the following, we present a new marginal cost accounting scheme for our model under service-level constraints: it decomposes the total cost in terms of the marginal costs of individual decisions and these marginal costs may include costs in both the current and subsequent periods. This extends and generalizes the marginal cost accounting discussed by Levi et al. (2007).

3.4.1 Review of the Dual-Balancing Policy

The underlying idea of the SMB policy is based on the *dual-balancing policy* proposed by Levi et al. (2007). The traditional inventory cost accounting scheme (in dynamic programming) decomposes the total costs by periods; Levi et al. (2007) proposes a marginal cost accounting scheme and computes the marginal holding cost by

$$H_t(q_t) = \sum_{j=t+L}^{T+L} h_j\big((X_t + q_t - D_{[t,j]})^+ - (X_t - D_{[t,j]})^+\big), \tag{3.4}$$

where X_t denotes a random inventory position which realizes at the beginning of period t.

The marginal backlogging cost is the same as the classical per-period backlogging cost, i.e.,

$$\Pi_t(q_t) = b_{t+L}(D_{[t,t+L]} - (q_t + X_t))^+. \tag{3.5}$$

The dual balancing policy balances the marginal holding cost in (3.4) against the marginal backlogging cost (3.5) and it admits a worst-case performance guarantee of two.

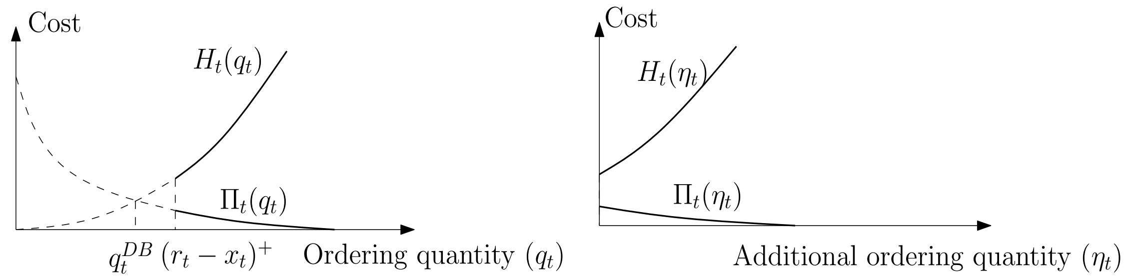

Figure 3.1: Dual-balancing policy: no balancing point

However, the dual balancing policy cannot be directly applied to our model because the balancing quantity for the marginal holding cost and the marginal backlogging cost may not exist due to the service-level constraints. In periods where the constrained service-level is much higher than the current inventory position, the expected marginal holding cost is always greater than the expected marginal backlogging cost. In such a case, one cannot find a feasible production quantity which perfectly balances the expected marginal holding cost against the expected marginal backlogging cost. (see Figure 3.1).

3.4.2 Split-Merge-Balance Policy

Without loss of generality, we assume that the unit production cost in each period is zero following a standard cost transformation in the literature (see, e.g., Zipkin (2000)). That is, for any system with positive unit production cost c_t, there is an equivalent system with revised costs $c'_t = 0$, $h'_{t+L} = h_{t+L} + c_t - c_{t+1}$ and $b'_{t+L} = b_{t+L} - c_t + c_{t+1}$. This allows us to only consider holding costs and backlogging costs.

Marginal cost accounting scheme (Split). We first present a new marginal cost accounting scheme for our inventory model under service-level requirements, which generalizes the marginal cost accounting scheme discussed by Levi et al. (2007). In the presence of service-level constraints, we split the marginal holding cost into two parts. The first part is called *forced holding cost* (denoted by $\tilde{H}_t$), which accounts for the holding cost from producing up to the service-level $\bar{X}_t = \max\{X_t, r_t\}$ in period t. The forced holding cost is unavoidable and it is independent of the current decision. The second part of the marginal holding cost is an *additional future holding cost* (denoted by $\hat{H}_t$) incurred by producing additional (controllable) $\eta_t = X_t + q_t - \bar{X}_t$. The reason behind this split is that the forced marginal holding cost $\tilde{H}_t$ is fixed given production decisions in previous periods, and hence only the additional marginal holding cost $\hat{H}_t$ is affected by the current decision η_t.

Suppose that X_t is the inventory position at the beginning of period t. We compute the forced holding cost $\tilde{H}_t$ by

$$\tilde{H}_t \quad = \quad \sum_{j=t+L}^{T+L} h_j\big((\bar{X}_t - D_{[t,j]})^+ - (X_t - D_{[t,j]})^+\big), \tag{3.6}$$

where $h_j(\bar{X}_t - D_{[t,j]})^+$ is the marginal holding cost in period j for producing up to $\bar{X}_t$ in period t and $h_j(X_t - D_{[t,j]})^+$ is the marginal holding cost in period j for producing nothing in period t. Similarly, the additional future holding cost $\hat{H}_t$ is computed as

$$\hat{H}_t(\eta_t) = \sum_{j=t+L}^{T+L} h_j\big((\bar{X}_t + \eta_t - D_{[t,j]})^+ - (\bar{X}_t - D_{[t,j]})^+\big), \tag{3.7}$$

where $h_j(\bar{X}_t + \eta_t - D_{[t,j]})^+$ stands for the marginal holding cost in period j for producing an additional η_t in period t and $h_j(\bar{X}_t - D_{[t,j]})^+$ is the marginal holding cost in period j for producing nothing additionally in period t. The backlogging cost in period t is the same as the classical per-period backlogging cost, i.e.,

$$B_t(\eta_t) = b_{t+L}(D_{[t,t+L]} - (\eta_t + \bar{X}_t))^+. \tag{3.8}$$

The left graph in Figure 3.2 shows the split marginal costs.

Regrouping the marginal costs (Merge). After splitting the marginal costs, we next regroup them. The marginal costs fall into two categories. One is called *overage cost* and it includes the marginal costs incurred due to production. Specifically, the additional holding cost $\hat{H}_t$ is overage cost since it will increase when an additional production is made. We name the other category *underage cost*, which includes the marginal costs incurred due to lack of productions, such as the backlogging cost $B_t(\eta_t)$. For the forced holding cost $\tilde{H}_t$, however, it is not affected by the decision in the same period because it is pre-determined by the production made in the last period. For this reason, we compute the *delayed holding cost* in the subsequent period (i.e., $\tilde{H}_{t+1}$) as soon as the production is made in period t. Specifically, once the additional production quantity η_t is determined, we compute the *delayed holding cost* by

$$\tilde{H}_{t+1}(\eta_t) \quad = \quad \sum_{j=t+L+1}^{T+L} h_j\big((\bar{X}_{t+1} - D_{[t,j]})^+ - (X_{t+1} - D_{[t,j]})^+\big). \tag{3.9}$$

Note that the *delayed holding cost* requires to compute r_{t+1} based on f_t rather than f_{t+1}. Hence, the term $\bar{X}_{t+1}$ in (3.9) should be treated as a random variable depending on the realization of D_t.

As we can see from (3.9), the more we produce in period t, the more inventory position we have at the beginning of period $t+1$ and the less delayed holding cost will be incurred. Thus, by shifting all the marginal forced holding costs to one period later, we conclude that $\tilde{H}_{t+1}$ is decreasing in η_t and hence, it belongs to underage cost. The right graph in Figure 3.2 illustrates the shifted cost.

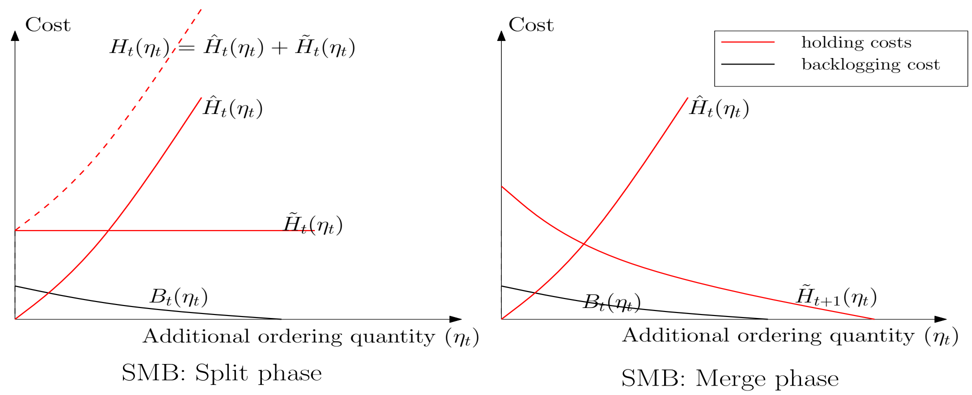

Figure 3.2: Marginal costs under the SMB policy

SMB policy (Balance). We describe the SMB policy as follows: At the beginning of each period t, we first calculate the *balancing quantity* η_t^{SMB}, which balances the conditional expected overage cost against the conditional expected underage cost. In other words, η_t^{SMB} solves

$$\mathbb{E}[\hat{H}_t(\eta_t) \mid f_t] = \mathbb{E}[\tilde{H}_{t+1}(\eta_t) + B_t(\eta_t) \mid f_t]. \tag{3.10}$$

Then the SMB policy produces $q_t^{SMB} = \eta_t^{SMB} + \bar{x}_t - x_t$ in period t when $X_t = x_t$ is realized. Here in the SMB policy, the balancing quantity must exist due to the following facts: (i) $\hat{H}_t(\eta_t)$, $\tilde{H}_{t+1}(\eta_t)$ and $B_t(\eta_t)$ are continuous; (ii) $\hat{H}_t(\eta_t)$ is non-decreasing in η_t while $\tilde{H}_{t+1}(\eta_t) + B_t(\eta_t)$ are non-increasing in η_t; (iii) As η_t increases from 0 to $+\infty$, the left hand side of (3.10) also increases from 0 to $+\infty$ while the right hand side of (3.10) decreases from a positive number to 0. Moreover, the balancing quantity η_t^{SMB} can be computed efficiently using a bisection search method.

To evaluate the total cost of a given policy P in a convenient way, we define the required service-level in period $T+L+1$ to be zero, i.e., $r_{T+1} = 0$. Under this convention, the forced costs must vanish in period $t+1$, i.e., $\tilde{H}_{T+1}^P = 0$. Then for a given policy P, the total cost

$\mathcal{C}(P)$ is given by

$$\mathcal{C}(P) \;=\; \sum_{t=1}^{T}((\tilde{H}_t^P + \hat{H}_t^P) + B_t^P) = \sum_{t=1}^{T}(\hat{H}_t^P + (\tilde{H}_{t+1}^P + B_t^P)) + \tilde{H}_1^P \qquad (3.11)$$

We note that the delayed holding cost $\tilde{H}_{t+1}^P$ for any policy P is computed using the demand information f_t we obtained in the previous period. Also, note that the forced cost $\tilde{H}_1^P$ can be calculated without knowing specific stochastic demand information and the policy P we use. Hence, they are realized at the beginning of the planning horizon and are fixed in any policy P we refer to.

3.4.3 Worst-Case Analysis of the SMB Policy

Now we establish the worst-case guarantee of 2 for the proposed SMB policy, which is the key result of this chapter. Due to service-level constraints, the forced holding cost and additional holding cost components need to be considered separately. To this end, we use an algebraic method to prove our desired results, which departs from the unit-matching techniques used in Levi et al. (2007).

To begin with, we define the following random sets of periods:

- $\mathcal{T}_H = \{t \mid Y_t^{SMB} < Y_t^{OPT}\}$ denotes the set of periods in which the optimal policy yields more ending inventory in period $t + L$ than the SMB policy;
- $\mathcal{T}_B = \{t \mid Y_t^{SMB} \geqslant Y_t^{OPT}\}$ denotes the set of periods in which the optimal policy yields less or equal ending inventory in period $t+L$ compared to the SMB policy; it is evident that $\mathcal{T}_H$ and $\mathcal{T}_B$ are disjoint sets and $\mathcal{T}_H \cup \mathcal{T}_B = \{1, 2, \ldots, T\}$.

Our main results are based on the following lemmas. The key idea is to calculate the total cost of the SMB policy using periods in sets $\mathcal{T}_H$ and $\mathcal{T}_B$. Then in each period, we bound either the overage cost or the underage cost by the corresponding cost for the optimal policy, according to which set the current period belongs to.

Lemma 3.2. $\mathbb{E}[\mathcal{C}(SMB)] = 2\mathbb{E}\Bigg[\sum_{t\in\mathcal{T}_H} \hat{H}_t^{SMB} + \sum_{t\in\mathcal{T}_B}(\tilde{H}_{t+1}^{SMB} + B_t^{SMB})\Bigg] + \tilde{H}_1^{SMB}.$

Proof. Let ζ_t be the random balanced cost by the SMB policy in period t, i.e., ζ_t =

$\mathbb{E}[\hat{H}_t^{SMB}|F_t] = \mathbb{E}[\tilde{H}_{t+1}^{SMB} + B_t^{SMB}|F_t]$. According to (3.11), we have

$$\begin{aligned}
\mathbb{E}[\mathcal{C}(SMB)] &= \sum_{t=1}^{T} \mathbb{E}\big[\mathbb{E}[\hat{H}_t^{SMB} + (\tilde{H}_{t+1}^{SMB} + B_t^{SMB}) \mid F_t]\big] + \tilde{H}_1^{SMB} \\
&= 2\sum_{t=1}^{T} \mathbb{E}[\zeta_t] + \tilde{H}_1^{SMB} \\
&= 2\sum_{t} \mathbb{E}\big[\big(\mathbb{1}(t \in \mathcal{T}_H) + \mathbb{1}(t \in \mathcal{T}_B)\big) \cdot \zeta_t\big] + \tilde{H}_1^{SMB} \\
&= 2\sum_{t} \mathbb{E}\big[E[\hat{H}_t^{SMB} \cdot \mathbb{1}(t \in \mathcal{T}_H) \mid F_t]\big] \\
&\quad +2\sum_{t} \mathbb{E}\big[\mathbb{E}[(\tilde{H}_{t+1}^{SMB} + B_t^{SMB}) \cdot \mathbb{1}(t \in \mathcal{T}_B) \mid F_t]\big] + \tilde{H}_1^{SMB} \\
&= 2\mathbb{E}\Bigg[\sum_{t \in \mathcal{T}_H} \hat{H}_t^{SMB} + \sum_{t \in \mathcal{T}_B} (\tilde{H}_{t+1}^{SMB} + B_t^{SMB})\Bigg] + \tilde{H}_1^{SMB}.
\end{aligned}$$

□

First, we consider the case when $t \in \mathcal{T}_B$. In this case, the ending inventory position in period t for the optimal policy is lower than that of the SMB policy, so it yields more backlogging cost in the current period. Moreover, given the relatively lower ending inventory position for the optimal policy, the inventory position at the beginning of the next period for the optimal policy must also be lower. Thus, the optimal policy must yield a larger forced holding cost in period $t+1$. We summarize these observations in the following lemma.

Lemma 3.3. *For any $t \in \mathcal{T}_B$, we have:*

1. $B_t^{SMB} \leqslant B_t^{OPT}$;
2. $\tilde{H}_{t+1}^{SMB} \leqslant \tilde{H}_{t+1}^{OPT}$.

Proof. Suppose $t \in \mathcal{T}_B$, then $Y_t^{SMB} \geqslant Y_t^{OPT}$. Therefore,

$$B_t^{SMB} = b_{t+L}(D_{[t,t+L]} - Y_t^{SMB})^+ \leqslant b_{t+L}(D_{[t,t+L]} - Y_t^{OPT})^+ = B_t^{OPT}.$$

Moreover, since $X_{t+1}^{SMB} = Y_t^{SMB} - D_t \geqslant Y_t^{OPT} - D_t = X_{t+1}^{OPT}$, we have

$$\bar{X}_{t+1}^{SMB} - X_{t+1}^{SMB} = (r_{t+1} - X_{t+1}^{SMB})^+ \leqslant (r_{t+1} - X_{t+1}^{OPT})^+ = \bar{X}_{t+1}^{OPT} - X_{t+1}^{OPT}.$$

Since Equation (3.6) is equivalent to

$$\tilde{H}_t = \sum_{j=t+L}^{T+L} h_j\big(\bar{X}_t - X_t - (D_{[t,j]} - X_t)^+\big)^+,$$

we conclude that

$$\begin{aligned}\tilde{H}_{t+1}^{SMB} &= \sum_{j=t+1+L}^{T+L} h_j((\bar{X}_{t+1}^{SMB} - X_{t+1}^{SMB} - (X_{t+1}^{SMB} - D_{[t+1,j]})^+))^+ \\ &\leqslant \sum_{j=t+1+L}^{T+L} h_j((\bar{X}_{t+1}^{OPT} - X_{t+1}^{OPT} - (X_{t+1}^{OPT} - D_{[t+1,j]})^+))^+ = \tilde{H}_{t+1}^{OPT}.\end{aligned}$$

□

Second, for any period $t \in \mathcal{T}_H$, the ending inventory of the SMB policy is lower than that of the optimal policy. Consider consecutive periods $[t^1, t^2] \subseteq \mathcal{T}_H$. At the beginning of period t^1, the inventory position of the SMB policy is higher while at the end of period t^2 the inventory of the SMB policy is lower. This implies that the SMB policy must make fewer additional productions than the optimal policy. As a result, the additional holding cost of the SMB policy is dominated by the additional holding cost of the optimal policy. We summarize this result in the following lemma:

Lemma 3.4. *For $t \in \mathcal{T}_H$, $\sum_{t\in\mathcal{T}_H} \hat{H}_t^{SMB} \leqslant \sum_{t\in\mathcal{T}_H} \hat{H}_t^{OPT}$.*

Proof. Proof of Lemma 3.4. To show this, we prove the following inequality:

$$\sum_{t\in\mathcal{T}_H} (\hat{H}_t^{SMB} + \tilde{H}_{t+1}^{SMB}) \leqslant \sum_{t\in\mathcal{T}_H} (\hat{H}_t^{OPT} + \tilde{H}_{t+1}^{OPT}). \tag{3.12}$$

Since $\tilde{H}_{t+1}^{SMB} \geqslant \tilde{H}_{t+1}^{OPT}$ for any $t \in \mathcal{T}_H$ (following a similar argument in Lemma 3.3), we conclude that (3.12) implies our desired inequality.

Notice that for any policy P, we have

$$\begin{aligned}\hat{H}_t^P + \tilde{H}_{t+1}^P &= \sum_{j=t+L}^{T+L} h_j((\bar{X}_t^P + \eta_t - D_{[t,j]})^+ - (\bar{X}_t^P - D_{[t,j]})^+)\Big\} \\ &\quad + \sum_{j=t+L}^{T+L} h_j\Big\{((\bar{X}_{t+1}^P - D_{[t+1,j]})^+ - (X_{t+1}^P - D_{[t+1,j]})^+)\Big\} \\ &= \sum_{j=t+L}^{T+L} h_j\Big\{(\bar{X}_{t+1}^P - D_{[t+1,j]})^+ - (\bar{X}_t^P - D_{[t,j]})^+\Big\},\end{aligned}$$

where the first equality applies the definition of $\hat{H}_t^P$ and $\tilde{H}_{t+1}^P$ (see (3.6) and (3.7)) and the second one cancels the first and the last terms using the system dynamic $X_{t+1}^P = \bar{X}_t^P + \eta_t^P - D_t$.

Suppose that $\mathcal{T}_H$ has n intervals, i.e., $\mathcal{T}_H = I_1 \cup I_2 \cup \cdots \cup I_n$, where $I_s = [\xi_s^1, \xi_s^2]$. Then it suffices to show the desired inequality on each interval, i.e.,

$$\sum_{t=\xi_s^1}^{\xi_s^2} (\hat{H}_t^{SMB} + \tilde{H}_{t+1}^{SMB}) \leqslant \sum_{t=\xi_s^1}^{\xi_s^2} (\hat{H}_t^{OPT} + \tilde{H}_{t+1}^{OPT}).$$

Let $\gamma_s(j) = \min\{j - L, \xi_s^2\}$. By interchanging the order of summation, we conclude that for each interval I_s and any policy P,

$$\begin{aligned}
\sum_{t=\xi_s^1}^{\xi_s^2} (\hat{H}_t^P + \tilde{H}_{t+1}^P) &= \sum_{t=\xi_s^1}^{\xi_s^2} \sum_{j=t+L}^{T+L} h_j \Big\{ (\bar{X}_{t+1}^P - D_{[t+1,j]})^+ - (\bar{X}_t^P - D_{[t,j]})^+ \Big\} \\
&= \sum_{j=\xi_s^1+L}^{T+L} h_j \sum_{t=\xi_s^1}^{\gamma_s(j)} \Big\{ (\bar{X}_{t+1}^P - D_{[t+1,j]})^+ - (\bar{X}_t^P - D_{[t,j]})^+ \Big\} \\
&= \sum_{j=\xi_s^1+L}^{T+L} h_j \Big\{ (\bar{X}_{\gamma_s(j)+1}^P - D_{[\gamma_s(j)+1,j]})^+ - (\bar{X}_{\xi_s^1}^P - D_{[\xi_s^1,j]})^+ \Big\}. \qquad (3.13)
\end{aligned}$$

For all $j = \xi_s^1 + L, \ldots, T + L$, $\gamma_s(j) \in I_s \subseteq \mathcal{T}_H$, thus,

$$X_{\gamma_s(j)+1}^{SMB} = Y_{\gamma_s(j)}^{SMB} - D_{\gamma_s(j)} < Y_{\gamma_s(j)}^{OPT} - D_{\gamma_s(j)} = X_{\gamma_s(j)+1}^{OPT},$$

and consequently,

$$(\bar{X}_{\gamma_s(j)+1}^{SMB} - D_{[\gamma_s(j)+1,j]})^+ \leqslant (\bar{X}_{\gamma_s(j)+1}^{OPT} - D_{[\gamma_s(j)+1,j]})^+. \qquad (3.14)$$

Also, since $\xi_s^1 - 1 \notin \mathcal{T}_H$, we obtain

$$X_{\xi_s^1}^{SMB} = Y_{\xi_s^1-1}^{SMB} - D_{\xi_s^1-1} \geqslant Y_{\xi_s^1-1}^{OPT} - D_{\xi_s^1-1} = X_{\xi_s^1}^{OPT},$$

(this inequality also holds when $\xi_s^1 = 1$ since $X_1^{SMB} = X_1^{OPT}$) and hence

$$(\bar{X}_{\xi_s^1}^{SMB} - D_{[\xi_s^1,j]})^+ \geqslant (\bar{X}_{\xi_s^1}^{OPT} - D_{[\xi_s^1,j]})^+. \qquad (3.15)$$

Therefore,

$$\begin{aligned}
\sum_{t=\xi_s^1}^{\xi_s^2}(\hat{H}_t^{SMB}+\tilde{H}_{t+1}^{SMB}) &= \sum_{j=\xi_s^1+L}^{T+L} h_j\Big\{(\bar{X}^{SMB}_{\gamma_s(j)+1}-D_{[\gamma_s(j)+1,j]})^+ - (\bar{X}^{SMB}_{\xi_s^1}-D_{[\xi_s^1,j]})^+\Big\}\\
&\leqslant \sum_{j=\xi_s^1+L}^{T+L} h_j\Big\{(\bar{X}^{OPT}_{\gamma_s(j)+1}-D_{[\gamma_s(j)+1,j]})^+ - (\bar{X}^{OPT}_{\xi_s^1}-D_{[\xi_s^1,j]})^+\Big\}\\
&= \sum_{t=\xi_s^1}^{\xi_s^2}(\hat{H}_t^{OPT}+\tilde{H}_{t+1}^{OPT}),
\end{aligned}$$

where the first and the third equalities follow from (3.13), and the second inequality follows from (3.14) and (3.15). □

Combining Lemma 3.2 to Lemma 3.4 together, we have

$$\begin{aligned}
\mathbb{E}[\mathcal{C}(SMB)] &= 2\mathbb{E}\Big[\sum_{t\in\mathcal{T}_H}\hat{H}_t^{SMB}+\sum_{t\in\mathcal{T}_B}(\tilde{H}_{t+1}^{SMB}+B_t^{SMB})\Big]+\tilde{H}_1^{SMB}\\
&\leqslant 2\mathbb{E}\Big[\sum_{t\in\mathcal{T}_H}\hat{H}_t^{OPT}+\sum_{t\in\mathcal{T}_B}(\tilde{H}_{t+1}^{OPT}+B_t^{OPT})\Big]+\tilde{H}_1^{OPT}\\
&\leqslant 2\mathbb{E}\Big[\sum_{t=1}^{T}(\hat{H}_t^{OPT}+\tilde{H}_{t+1}^{OPT}+B_t^{OPT})+\tilde{H}_1^{OPT}\Big]=2\mathbb{E}[\mathcal{C}(OPT)].
\end{aligned}$$

Hence, we have proved the following theorem, which provides a worst-case performance guarantee on the result of the SMB policy.

Theorem 3.5. *The SMB policy has a worst-case performance guarantee of two, i.e., for each instance of the backlogging model under service-level constraints, the expected cost of the SMB policy is at most two times the expected cost of an optimal solution, i.e.,* $\mathbb{E}[\mathcal{C}(SMB)] \leqslant 2\mathbb{E}[\mathcal{C}(OPT)]$.

3.5 Remanufacturing System with Service-Level Requirements

We consider a remanufacturing system with general random demands and product returns. In the remanufacturing system, the manufacturer receives a random number of returned products at the beginning of each period. The returned products received during each period can be remanufactured to a new product at a lower cost by replacing some components. The major difference between the remanufacturing model and our previous basic model is the dual modes of production, i.e., the manufacturer can either produce by remanufacturing a returned product or by regular way using raw materials. This leads to the nonlinear production cost, which makes the model even harder.

In the following, we will first formulate our model using dynamic programming and show the structure of optimal policies. Then we will generalize the proposed SMB policy to solve the remanufacturing model, which also guarantee us a worst-case performance of two. Our technique is based on Tao and Zhou (2014), in which they proposed a two approximation algorithm for the remanufacturing system without service-level requirement. However, our results are different in the following ways. First, due to the service-level requirements presented in our model, our algorithm departs from the one proposed in Tao and Zhou (2014) (details are provided in §3.5.2). Second, in worst case analysis, amortizing the production costs of the modified SMB policy is different since we need to handle two parts of split production cost, i.e., the forced production cost and the additional production cost. We will provide more details in §3.5.3.

3.5.1 Model and DP Formulation

We adopt most of the notations described in §3.3. At each period t, the manufacturer first receives a random number of returned products (denoted by U_t). The manufacturer then decides the remanufacturing quantity q_t^1 and the manufacturing quantity q_t^2. We assume that both production methods have the same lead time, denoted by L (see, e.g., Zhou et al. (2011)). The total number of productions at period t are computed by $q_t = q_t^1 + q_t^2$. We use c_t^1 and c_t^2 to denote the remanufacturing cost and the manufacturing cost with $c_t^1 < c_t^2$. We also assume that $c_t^2 - c_t^1$ is non-increasing (see, e.g., Zhou et al. (2011)). This assumption holds in practice where manufacturing costs can be reduced significantly over time while remanufacturing costs are lower and hard to be reduced. Finally, demand realizes and corresponding cost occurs.

In the remanufacturing model, the information set f_t is realized at the beginning of period t, which consists of the realized demands $(d_1, \ldots, d_{t-1})$, the realized returns $(u_1, \ldots, u_t)$ and some exogenous information $(\rho_1, \ldots, \rho_t)$ such as the state of economy. The conditional joint distribution of future demand and returns $(D_t, \ldots, D_T, U_t, \ldots, U_T)$ is determined by the information set f_t.

To derive a dynamic programming formulation, we first describe the state vector as follows. The state consists of a time period t, inventory position x_t at the beginning of period t, total number of returned products w_t at the beginning of period t, and information

$f_t \in \mathcal{F}_t$. The system dynamics are

$$x_{t+1} = x_t + q_t - d_t, \quad \forall t = 1, \cdots, T-1, \tag{3.16}$$

$$w_{t+1} = u_{t+1} + w_t - q_t^1, \quad \forall t = 1, \cdots, T-1, \tag{3.17}$$

with $x_1 = 0$ and $w_1 = u_1$. Let the value function $v_t(x_t, w_t, f_t)$ be the minimal expected cost from period $t + L$ to period $T + L$. In each period t, given the state vector (x_t, w_t, f_t), we need to decide the remanufacturing quantity q_t^1 and the manufacturing quantity q_t^2. The remanufacturing quantity q_t^1 is bounded above by w_t, while the service-level constraint enforces $y_t = x_t + q_t^1 + q_t^2$ bounded below by r_t. Hence, feasible choices of the two types of quantities are in the set

$$\mathcal{Q}(x_t, w_t, r_t) = \{(q_t^1, q_t^2) \mid 0 \leqslant q_t^1 \leqslant w_t,\ q_t^2 \geqslant 0,\ q_t^1 + q_t^2 \geqslant r_t - x_t\}.$$

The Bellman's equations are

$$\begin{aligned}
&v_{T+1}(x_{T+1}, w_{T+1}, f_{T+1}) = 0, \quad \forall x_{T+1} \in \mathbb{R}, w_{T+1} \in \mathbb{R}^+ \cup \{0\}, f_{T+1} \in \mathcal{F}_{T+1}, \\
&v_t(x_t, w_t, f_t) = \min_{q_t^1, q_t^2 \in \mathcal{Q}(x_t, w_t, r_t)} \Big\{ G_t(x_t + q_t^1 + q_t^2, f_t) + c_t^1 q_t^1 + c_t^2 q_t^2 \\
&\qquad\qquad + \mathbb{E}[v_{t+1}(y_t - D_t, U_{t+1} + w_t - q_t^1, \mathcal{F}_{t+1}) \mid f_t] \Big\}, t = 1, \ldots, T.
\end{aligned} \tag{3.18}$$

Using the above DP formulation (3.18), the structure of optimal policies is characterized in the following proposition.

Proposition 3.6. *For the inventory control problem defined in* (3.18)*, an optimal policy is a total base stock (including both manufacturing and remanufacturing) policy. More specifically, there exists* $\{s(w_t, f_t)\}_{t=1}^T$ *such that*

$$y_t^*(x_t, w_t) = \begin{cases} \max\{r_t, s(w_t, f_t)\}, & \text{if } x_t < s(w_t, f_t); \\ \max\{r_t, x_t\}, & \text{if } x_t \geqslant s(w_t, f_t). \end{cases}$$

and

$$q_t^{1*}(x_t, w_t) = \min\{w_t, y_t^* - x_t\},\ q_t^{2*}(x_t, w_t) = y_t^* - x_t - q_t^{1*}(x_t).$$

Proof. For simplicity, we will omit the information f_t in the proof. We change the decision variable from (q_t^1, q_t^2) to (q_t^1, y_t) where $y_t = x_t + q_t^1 + q_t^2$ and define

$$\begin{aligned} J_t(x_t, w_t, y_t, q_t^1) &= G_t^h(y_t) + c_t^1 q_t^1 + c_t^2 q_t^2 + \mathbb{E}[v_{t+1}(y_t - D_t, U_{t+1} + w_t - q_t^1)] \\ &= G_t^h(y_t) + c_t^2 y_t - (c_t^2 - c_t^1) q_t^1 - c_t^2 x_t \\ &\quad + \mathbb{E}[v_{t+1}(y_t - D_t, U_{t+1} + w_t - q_t^1)] \end{aligned} \tag{3.19}$$

for $t = 1, 2, \ldots, T$. Then the value function can be computed by

$$v_t(x_t, w_t) = \min_{q_t^1, y_t \in \mathcal{Y}(x_t, w_t, r_t)\}} \{J_t(x_t, w_t, y_t, q_t^1)\},$$

where $\mathcal{Y}(x_t, w_t, r_t) = \{(q_t^1, y_t) \mid 0 \leqslant q_t^1 \leqslant \min\{w_t, y_t - x_t\}, y_t \geqslant \max\{x_t, r_t\}\}$.

We first show that for all $f_t \in \mathcal{F}_t$, $v_t(x_t, w_t)$ is separable, convex in (x_t, w_t) and linear and non-increasing in w_t with rate less than $c_t^2 - c_t^1$, i.e., there exists a convex function $z_t(\cdot)$ and a coefficient $\kappa_t \in [0, c_t^2 - c_t^1]$ such that $v_t(x_t, w_t) = z_t(x_t) - \kappa_t w_t$.

Clearly, when $t = T + 1$, $v_{T+1}(x_{T+1}, w_{T+1}) = 0$ satisfies these conditions.

Suppose the statement is true for $v_{t+1}(x_{t+1}, w_{t+1})$, then we can write

$$\begin{aligned} \mathbb{E}[v_{t+1}(y_t - D_t, U_{t+1} + w_t - q_t^1)] &= \mathbb{E}[z_{t+1}(y_t - D_t) - \kappa_{t+1}(U_{t+1} + w_t - q_t^1)] \\ &= \tilde{z}_{t+1}(y_t) - \kappa_{t+1}(w_t - q_t^1) - \kappa_{t+1}\mathbb{E}[U_{t+1}], \end{aligned}$$

where $\tilde{z}_{t+1}(\cdot)$ is still a convex function. Hence, applying (3.19), we conclude that

$$\begin{aligned} J_t(x_t, w_t, y_t, q_t^1) &= G_t^h(y_t) + c_t^2 y_t - (c_t^2 - c_t^1) q_t^1 - c_t^2 x_t + \tilde{z}_{t+1}(y_t) - \kappa_{t+1}(w_t - q_t^1) \\ &\quad - \kappa_{t+1}\mathbb{E}[U_{t+1}] \\ &= (\kappa_{t+1} - (c_t^2 - c_t^1)) q_t^1 + (G_t^h(y_t) + \tilde{z}_{t+1}(y_t) + c_t^2 y_t) - c_t^2 x_t \\ &\quad - \kappa_{t+1}(w_t + \mathbb{E}[U_{t+1}]) \end{aligned}$$

Now we can compute $v_t(x_t, w_t)$ by first optimizing with respect to q_t^1. Since $\kappa_{t+1} \leqslant c_{t+1}^2 - c_{t+1}^1 \leqslant c_t^2 - c_t^1$, the coefficient of q_t^1 is non-positive and hence $J_t(x_t, w_t, y_t, q_t^1)$ is non-increasing in q_t^1. Observing the feasible set $\mathcal{Y}(x_t, w_t, r_t) = \{(q_t^1, y_t) \mid 0 \leqslant q_t^1 \leqslant w_t, y_t \geqslant x_t + q_t^1, y_t \geqslant r_t\}$, we conclude that ${q_t^1}^* = \min\{w_t, y_t - x_t\}$.

- If $w_t \geqslant y_t - x_t$, ${q_t^1}^* = w_t$. Then

$$J_t(x_t, w_t, y_t, {q_t^1}^*) = (G_t^h(y_t) + \tilde{z}_{t+1}(y_t) + c_t^2 y_t) - (c_t^2 - c_t^1) w_t - c_t^2 x_t - \kappa_{t+1}\mathbb{E}[U_{t+1}].$$

Since both $G_t^h(y_t)$ and $\tilde{z}_{t+1}(y_t)$ are convex in y_t we conclude that $J_t(x_t, w_t, y_t, q_t^{1*})$ is a convex function in y_t. Thus, if we define

$$s_t = s(w_t, f_t) = \arg\min_{y_t}\{J_t(x_t, w_t, y_t, q_t^{1*})\} = \arg\min_{y}\{G_t^h(y) + \tilde{z}_{t+1}(y) + c_t^2 y\},$$

we conclude that $y_t^* = \max\{s_t, x_t, r_t\}$ minimizes $J_t(x_t, w_t, y_t, q_t^{1*})$ with respect to the feasible set $\mathcal{Y}(x_t, w_t, r_t)$. Moreover, $z_t(x_t) = (G_t^h(y_t^*) + \tilde{z}_{t+1}(y_t^*) + c_t^2 y_t^*) - c_t^2 x_t - \kappa_{t+1}\mathbb{E}[U_{t+1}]$ must be a convex function. Hence, $v_t(x_t, w_t) = J_t(x_t, w_t, y_t^*, q_t^{1*}) = z_t(x_t) - (c_t^2 - c_t^1)w_t$, where $z_t(\cdot)$ is a convex function.

- If $w_t < y_t - x_t$, $q_t^{1*} = y_t - x_t$. In this case,

  $$J_t(x_t, w_t, y_t, q_t^{1*}) = (G_t^h(y_t) + \tilde{z}_{t+1}(y_t) + (\kappa_{t+1} + c_t^1)y_t) - \kappa_{t+1}w_t - (\kappa_{t+1} + c_t^1)x_t - \kappa_{t+1}\mathbb{E}[U_{t+1}],$$

  which is also convex in y_t. Hence, by defining

  $$s_t = s(w_t, f_t) = \arg\min_{y_t}\{J_t(x_t, w_t, y_t, q_t^{1*})\} = \arg\min_{y}\{G_t^h(y) + \tilde{z}_{t+1}(y) + (\kappa_{t+1} + c_t^1)y\},$$

  we have $y_t^* = \max\{s_t, x_t, r_t\}$ minimizes $J_t(x_t, w_t, y_t, q_t^{1*})$. Moreover, $z_t(x_t) = (G_t^h(y_t^*) + \tilde{z}_{t+1}(y_t^*) + (\kappa_{t+1} + c_t^1)y_t^*) - (\kappa_{t+1} + c_t^1)x_t - \kappa_{t+1}\mathbb{E}[U_{t+1}]$ is convex in x_t and $v_t(x_t, w_t) = J_t(x_t, w_t, y_t^*, q_t^{1*}) = z_t(x_t) - \kappa_{t+1}w_t$. Note that $\kappa_{t+1} \leqslant c_{t+1}^2 - c_{t+1}^1 \leqslant c_t^2 - c_t^1$.

Combining the two cases discussed above, we conclude that $v_t(x_t, w_t) = z_t(x_t) - \kappa_t w_t$ where $z_t(\cdot)$ is a convex function and $\kappa_t \in [0, c_t^2 - c_t^1]$. Moreover, the minimizer (y_t^*, q_t^{1*}) must satisfy

$$y_t^* = \max\{x_t, s_t, r_t\}; \quad q_t^{1*} = \min\{w_t, y_t^* - x_t\},$$

where $s_t = s(w_t, f_t) = \arg\min_{y_t}\{J_t(x_t, w_t, y_t, q_t^{1*})\}$. □

Proposition 3.6 asserts that any optimal policy has the following structure: if the inventory position in period t is no less than the threshold $s(w_t, f_t)$, an optimal policy produces up to the required service-level r_t; otherwise, it brings the total inventory position (after production) to $\max\{r_t, s(w_t, f_t)\}$. Therefore, the higher the service-level, the more orders are placed by the optimal policy. Moreover, optimal policy will remanufacture returned products as much as possible before manufacturing new products.

3.5.2 Modified Split-Merge-Balance Policy

Without loss of generality, we assume $c_t^1 = 0$ following a standard cost transformation in the literature (see, e.g., Zipkin (2000)). That is, for any system with positive unit remanufacturing cost c_t^1, there is an equivalent system with revised unit remanufacturing cost $c_t^{1\prime} = 0$ and revised unit manufacturing cost $c_t^{2\prime} = c_t^2 - c_t^1$. Note that c_t^2 cannot be normalized to zero and we use $\bar{c}_t = c_t^2 - c_t^1$ to denote the unit production cost in the transformed system with $c_t^1 = 0$ in the following discussion.

The key idea of modified SMB policy (MSMB for short) is similar to the SMB policy proposed in §3.4.2, which has three phases, namely, split, merge and balance. However, in the remanufacturing system, we have to consider production cost in addition to the holding and backlogging costs. Recall that $\bar{X}_t = \max\{X_t, r_t\}$ is the required service-level, $\eta_t = y_t - \bar{X}_t$ is the controllable producing quantity in period t, and the marginal production cost in period t is given by

$$E_t(\eta_t) = \bar{c}_t(\eta_t + \bar{X}_t - X_t - W_t)^+.$$

Similar as splitting the holding cost in §3.4.2, we also split the marginal production cost into two parts: the *forced production cost* (denoted by $\tilde{E}_t$) which accounts for the cost of producing up to the required service-level $\bar{X}_t = \max\{X_t, r_t\}$ in period t and the *additional production cost* (denoted by $\hat{E}_t$) determined by the amount of additional (controllable) produces η_t. Specifically, if the number of returned products in period t is denoted by W_t, the forced production cost $\tilde{E}_t$ is computed by

$$\tilde{E}_t \;\; = \;\; \bar{c}_t(\bar{X}_t - X_t - W_t)^+, \tag{3.20}$$

and the additional production cost is

$$\hat{E}_t(\eta_t) = \bar{c}_t\big((\eta_t + \bar{X}_t - X_t - W_t)^+ - (\bar{X}_t - X_t - W_t)^+\big). \tag{3.21}$$

Next, we regroup the costs based on whether it belongs to overage cost or underage cost. It is evident that the additional production cost $\hat{E}_t(\eta_t)$ is overage cost since it increases when η_t increases (i.e., more productions are made). For the forced production cost $\tilde{E}_t$, although it does not depend on the decision η_t in the current period, it occurs due to the lack of productions in the previous period. Hence, by shifting the cost to one period later, we conclude that

$$\tilde{E}_{t+1}(\eta_t) = \bar{c}_{t+1}(\bar{X}_{t+1} - X_{t+1} - W_{t+1})^+ \tag{3.22}$$

is non-increasing in η_t and it belongs to underage cost. Similar as the delayed holding cost, the *delayed production cost* requires to compute r_{t+1} based on f_t rather than f_{t+1}. Hence, the term $\bar{X}_{t+1}$ in (3.22) should be treated as a random variable depending on the realization of D_t.

To summarize, if we use Φ_t and Ψ_t to denote the total overage cost and the total underage cost in period t, we have $\Phi_t(\eta_t) = \hat{H}_t(\eta_t) + \hat{E}_t(\eta_t)$ and $\Psi_t(\eta_t) = \tilde{H}_{t+1}(\eta_t) + \tilde{E}_{t+1}(\eta_t) + B_t(\eta_t)$. Figure 3.3 illustrates the split phase and the merge phase of the MSMB policy.

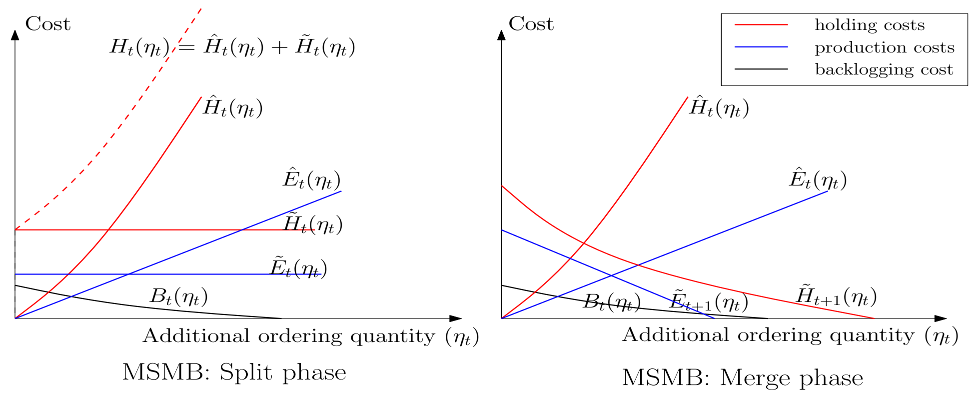

Figure 3.3: Marginal costs under the SMB policy

Finally, we balance the overage cost against the underage cost, i.e., η_t^{MSMB} solves

$$\mathbb{E}[\Phi_t(\eta_t) \mid f_t] = \mathbb{E}[\Psi_t(\eta_t) \mid f_t]. \tag{3.23}$$

Thus, the MSMB policy produces $q_t^{MSMB} = \eta_t^{MSMB} + \bar{x}_t - x_t$ in period t. Moreover, it fully utilizes the returned products to remanufacture, i.e., $q_t^{1,MSMB} = \min\{w_t, q_t^{MSMB}\}$ and $q_t^{2,MSMB} = (q_t^{MSMB} - w_t)^+$.

3.5.3 Worst-Case Analysis of the MSMB Policy

Now we establish the worst-case guarantee of two for the proposed MSMB policy. The main difficulty in our analysis is to amortize the production costs of the MSMB policy against that of the optimal policy. Our proof techniques are based on Tao and Zhou (2014), in which the authors constructed a set of periods such that the total production costs of the balancing policy are dominated by that of the optimal policy. They further showed the same inequality holds for the set of periods in which the optimal policy yields less or equal inventory compared to the balancing policy (see Lemma 4 of Tao and Zhou (2014)). However, our proof is different since we need to bound both the forced production cost and the additional production cost at the same time. As a result, the same inequality does not hold any more; instead, we delay the forced production cost to one period later and bound the total additional production costs in periods $\mathcal{T}_\Phi$ plus the total delayed production costs in periods $\mathcal{T}_\Psi$, which is crucial in our analysis (see Lemma 3.9).

In the following discussion, we will only focus on a particular type of policy, namely the *rational policies*. These policies will not manufacture at a higher per-unit cost unless there are no returned products to remanufacture. As we have already shown in Proposition 3.6, any optimal policy (indicated by OPT) is a rational policy and the MSMB policy described above is also rational. We will use superscripts to specify which policy we refer to.

We define new variables and introduce new notation. For any given policy P, let W_t^P be the total number of returned products in period t, and S_t^P be the remaining number of returned products after producing up to the service-level $\bar{X}_t^P$ in period t. We also split the additional production quantity η_t into two parts, denoted by $\eta_t^{1,P}$ and $\eta_t^{2,P}$, representing the additional remanufacturing quantity and additional manufaturing quantity of a given policy P. Because we only consider rational policies, we must have $\eta_t^{1,P} \leqslant S_t^P$ and $\eta_t^{2,P} = (\eta_t^P - S_t^P)^+$. Moreover, the system dynamics follow

$$\begin{cases} W_t^P = U_t + S_{t-1}^P - \eta_{t-1}^{1,P}; \\ S_t^P = (W_t^P - (\bar{X}_t^P - X_t^P))^+; \\ X_t^P = \bar{X}_{t-1}^P + \eta_{t-1}^{1,P} + \eta_{t-1}^{2,P} - d_{t-1}. \end{cases} \tag{3.24}$$

We rewrite the additional production cost as $\hat{E}_t^P = \bar{c}_t \eta_t^{2,P}$, and define three random sets of periods as follows.

- $\mathcal{T}_\Phi = \{t \mid Y_t^{MSMB} < Y_t^{OPT}\}$ denotes the set of periods t in which the optimal policy yields more ending inventory in periods $t+L$ than the MSMB policy;
- $\mathcal{T}_\Psi = \{t \mid Y_t^{MSMB} \geqslant Y_t^{OPT}\}$ denotes the set of periods t in which the optimal policy yields less or equal ending inventory in period $t+L$ compared to the MSMB policy; it is evident that $\mathcal{T}_\Phi$ and $\mathcal{T}_\Psi$ are disjoint sets and $\mathcal{T}_\Phi \cup \mathcal{T}_\Psi = \{1, 2, \ldots, T\}$.
- $\mathcal{T}_c = \{t \mid \bar{X}_t^{MSMB} + S_t^{MSMB} + \eta_t^{2,MSMB} \geqslant \bar{X}_t^{OPT} + S_t^{OPT} + \eta_t^{2,OPT}\}$. Following the system dynamics in (3.24), we can equivalently present the set as $\mathcal{T}_c = \{t \mid W_{t+1}^{MSMB} + X_{t+1}^{MSMB} \geqslant W_{t+1}^{OPT} + X_{t+1}^{OPT}\}$. The quantity $W_t^P + X_t^P$ stands for the maximum producing-up-to level without having any production cost in period t (which we refer to as the *free-production level*). The set $\mathcal{T}_c$ can be interpreted as the periods in which the free-production level in the next period for the MSMB policy is lower than that of the optimal policy.

Our main results are based on the following lemmas.

Lemma 3.7. $\mathbb{E}[\mathcal{C}(MSMB)] = 2\mathbb{E}\Big[\sum_{t\in\mathcal{T}_\Phi}\Phi_t^{MSMB} + \sum_{t\in\mathcal{T}_\Psi}\Psi_t^{MSMB}\Big] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB}$.

Proof. Proof of Lemma 3.7. Let ξ_t be the random balanced cost by the MSMB policy in period t, i.e., $\xi_t = \mathbb{E}[\Phi_t^{MSMB}|F_t] = \mathbb{E}[\Psi_t^{MSMB}|F_t]$. The total cost is computed by

$$\begin{aligned}
\mathbb{E}[\mathcal{C}(MSMB)] &= \sum_{t=1}^{T} E[\Phi_t^{MSMB} + \Psi_t^{MSMB}] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB} \\
&= \sum_{t=1}^{T} \mathbb{E}\big[\mathbb{E}[\Phi_t^{MSMB} + \Psi_t^{MSMB} \mid F_t]\big] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB} \\
&= 2\sum_{t=1}^{T} \mathbb{E}[\xi_t] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB} \\
&= 2\sum_{t} \mathbb{E}\big[\big(\mathbb{1}(t\in\mathcal{T}_\Phi) + \mathbb{1}(t\in\mathcal{T}_\Psi)\big)\cdot\xi_t\big] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB} \\
&= 2\Bigg(\sum_{t} \mathbb{E}\big[E[\Phi_t^{MSMB}\cdot\mathbb{1}(t\in\mathcal{T}_\Phi) \mid F_t]\big] + \sum_{t} \mathbb{E}\big[\mathbb{E}[\Psi_t^{MSMB}\cdot\mathbb{1}(t\in\mathcal{T}_\Psi) \mid F_t]\big]\Bigg) \\
&\quad + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB} \\
&= 2\mathbb{E}\Bigg[\sum_{t\in\mathcal{T}_\Phi}\Phi_t^{MSMB} + \sum_{t\in\mathcal{T}_\Psi}\Psi_t^{MSMB}\Bigg] + \tilde{H}_1^{MSMB} + \tilde{E}_1^{MSMB}.
\end{aligned}$$

□

We note that the both forced holding cost $\tilde{H}_1^{MSMB}$ and forced production cost $\tilde{E}_1^{MSMB}$ can be computed without knowing specific stochastic demand information and the policy we use. Hence, they are realized at the beginning of the planning horizon and are fixed in any policy P we refer to.

The following lemma restates the results proved in Lemma 3.3 and Lemma 3.4, under the modified SMB policy. The proof is identical to the proofs of Lemma 3.3 and Lemma 3.4, and thus omitted here.

Lemma 3.8. *1. For any $t\in\mathcal{T}_\Psi$, we have $B_t^{MSMB} \leqslant B_t^{OPT}$ and $\tilde{H}_{t+1}^{MSMB} \leqslant \tilde{H}_{t+1}^{OPT}$.*

2. For $t\in\mathcal{T}_\Phi$, we have $\sum_{t\in\mathcal{T}_\Phi}\hat{H}_t^{MSMB} \leqslant \sum_{t\in\mathcal{T}_\Phi}\hat{H}_t^{OPT}$.

Proof. The proof is identical to that in Lemmas 3.3 and 3.4. □

The next lemma is crucial in our analysis, which deals with production costs. The difficulty of the analysis lies in the fact that the production cost does not only depend on the ending inventory level but also depends on the number of returned products W_t.

For this reason, we first compare the production cost of the MSMB policy and that of the optimal policy for sets $\mathcal{T}_c$ and $\mathcal{T}_c^c$. When $t \in \mathcal{T}_c$, the free-production level for the MSMB policy is higher than that of the optimal policy in the next period $t+1$. Therefore, the forced production cost $\tilde{E}_{t+1}$ for the MSMB policy must be lower. For set $\mathcal{T}_c^c$, consider any consecutive time interval $[t_1, t_2]$ that belongs to $\mathcal{T}_c^c$. Compared with the optimal policy, the free-production level for the MSMB policy is higher at the beginning of period t_1, while it becomes lower at the end of period t_2. This can happen only when the MSMB policy uses more free productions. As a result, the total production cost for the MSMB policy during this time interval must be less than that of the optimal policy. Finally, we extend the results to sets $\mathcal{T}_\Phi$ and $\mathcal{T}_\Psi$ using the fact that $\eta_t^{2,MSMB} = 0$ for all $t \in \mathcal{T}_\Phi \cap \mathcal{T}_c$ and $\eta_t^{2,OPT} = 0$ for all $t \in \mathcal{T}_\Psi \cap \mathcal{T}_c^c$.

Connection and comparison with Tao and Zhou (2014). Our construction of the set $\mathcal{T}_c$ is based on the technique used in Tao and Zhou (2014), but the analysis is different in the following aspects. First, they showed that the total production costs of the balancing policy in periods $\mathcal{T}_c$ are no more than that of the optimal policy, i.e.,

$$\sum_{t \in \mathcal{T}_c^c} (\hat{E}_t^{MSMB} + \tilde{E}_t^{MSMB}) \leqslant \sum_{t \in \mathcal{T}_c^c} (\hat{E}_t^{OPT} + \tilde{E}_t^{OPT}).$$

However, the above inequality does not hold in our model since the forced production cost $\tilde{E}_t^{MSMB}$ is pre-determined by the previous decisions. This motivates us to consider a delayed production cost which shifts the forced production cost to one period later. The reason behind this is that the delayed production cost is determined as soon as the productions are made in the current period and it can be treated as a penalty for not producing enough in the current period. We show that the total additional production costs plus the total delayed production costs of our MSMB policy in periods $\mathcal{T}_c$ are no more than that of the optimal policy. (see the second inequality in Lemma 3.9).

Secondly, after comparing the total production costs in periods $\mathcal{T}_c$, Tao and Zhou (2014) proved the same inequality holds in periods $\mathcal{T}_\Phi$ (see Lemma 4 in Tao and Zhou (2014)), i.e.,

$$\sum_{t \in \mathcal{T}_\Phi} (\hat{E}_t^{MSMB} + \tilde{E}_{t+1}^{MSMB}) \leqslant \sum_{t \in \mathcal{T}_\Phi} (\hat{E}_t^{OPT} + \tilde{E}_{t+1}^{OPT}).$$

Again, this inequality does not hold in our case; instead, we show that the total additional costs in periods $\mathcal{T}_\Phi$ plus the total delayed production costs in periods $\mathcal{T}_\Psi$ are dominated by that of the optimal policy (see the third inequality in Lemma 3.9). The idea is to bound the overage cost in periods $\mathcal{T}_\Phi$ and the underage cost in periods $\mathcal{T}_\Psi$. We summarize our results in the following lemma.

Lemma 3.9. *For the production costs, we have*

1. *For* $t \in \mathcal{T}_c$, $\tilde{E}_{t+1}^{MSMB} \leqslant \tilde{E}_{t+1}^{OPT}$*; For* $t \in \mathcal{T}_c^c$, $\tilde{E}_{t+1}^{MSMB} \geqslant \tilde{E}_{t+1}^{OPT}$*;*
2. $\sum_{t\in\mathcal{T}_c^c}(\hat{E}_t^{MSMB} + \tilde{E}_{t+1}^{MSMB}) \leqslant \sum_{t\in\mathcal{T}_c^c}(\hat{E}_t^{OPT} + \tilde{E}_{t+1}^{OPT})$*;*
3. $\sum_{t\in\mathcal{T}_\Phi} \hat{E}_t^{MSMB} + \sum_{t\in\mathcal{T}_\Psi} \tilde{E}_{t+1}^{MSMB} \leqslant \sum_{t\in\mathcal{T}_\Phi} \hat{E}_t^{OPT} + \sum_{t\in\mathcal{T}_\Psi} \tilde{E}_{t+1}^{OPT}$.

Proof. Proof of Lemma 3.9. For the first part, we have

$$\begin{aligned}
X_{t+1}^{MSMB} + W_{t+1}^{MSMB} &= u_{t+1} + \bar{X}_t^{MSMB} + S_t^{MSMB} + \eta_t^{2,MSMB} - D_t \\
&\geqslant u_{t+1} + \bar{X}_t^{OPT} + S_t^{OPT} + \eta_t^{2,OPT} - D_t \\
&= X_{t+1}^{OPT} + W_{t+1}^{OPT}
\end{aligned}$$

whenever $t \in \mathcal{T}_c$. Note that (3.20) is equivalent to $\tilde{E}_t = \bar{c}_t(r_t - X_t - W_t)^+$. Therefore,

$$\begin{aligned}
\tilde{E}_{t+1}^{MSMB} &= c_{t+1}(r_{t+1} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+ \\
&\leqslant c_{t+1}(r_{t+1} - X_{t+1}^{OPT} - W_{t+1}^{OPT})^+ \\
&= \tilde{E}_{t+1}^{OPT},
\end{aligned}$$

for all $t \in \mathcal{T}_c$. Similarly, the inequality reverses when $t \in \mathcal{T}_c^c$.

Our second inequality is different from the one proved in Tao and Zhou (2014), in which they showed that the total production cost in periods $\mathcal{T}_c^c$ is less than that of the optimal policy,

For the second inequality, suppose that $\mathcal{T}_c^c$ has n intervals, namely, $\mathcal{T}_c^c = I_1 \cup I_2 \cup \cdots \cup I_n$ where $I_a = [\xi_a^1, \xi_a^2]$. Then, we only need to show the inequality holds for each interval I_a. To see this, we compare the free-production level $\bar{X}_t + S_t$ between the MSMB and OPT policies. Given a policy P, following the system dynamics in (3.24), we have

$$\begin{aligned}
\bar{X}_t^P + S_t^P &= \bar{X}_t^P + ((u_t + S_{t-1}^P - \eta_{t-1}^{1,P}) - (\bar{X}_t^P - X_t^P))^+ \\
&= \max\{\bar{X}_t^P, u_t + S_{t-1}^P - \eta_{t-1}^{1,P} + X_t^P\} \\
&= \max\{\bar{X}_t^P, u_t + S_{t-1}^P + \bar{X}_{t-1}^P + \eta_{t-1}^{2,P} - D_{t-1}\} \\
&= \max\{r_t, u_t + S_{t-1}^P + \bar{X}_{t-1}^P + \eta_{t-1}^{2,P} - D_{t-1}\},
\end{aligned}$$

where the last equality follows from $S_{t-1} + \bar{X}_{t-1}^P + \eta_{t-1}^{2,P} - D_{t-1} \geqslant X_t^P$ and $\bar{X}_t^P = \max\{X_t^P, r_t\}$.

Therefore, for $t \in \mathcal{T}_e^c$,

$$\begin{aligned}
\bar{X}_{t+1}^{MSMB} + S_{t+1}^{MSMB} &= \max\{r_{t+1}, u_{t+1} + S_t^{MSMB} + \bar{X}_t^{MSMB} + \eta_t^{2,MSMB} - D_t\} \\
&\leqslant \max\{r_{t+1}, u_{t+1} + S_t^{OPT} + \bar{X}_t^{OPT} + \eta_t^{2,OPT} - D_t\} \\
&= \bar{X}_{t+1}^{OPT} + S_{t+1}^{OPT}
\end{aligned}$$

and similarly, $\bar{X}_{t+1}^{MSMB} + S_{t+1}^{MSMB} \geqslant \bar{X}_{t+1}^{OPT} + S_{t+1}^{OPT}$ for all $t \in \mathcal{T}_e$. Using the system dynamics, for any given policy P, the additional manufacturing quantity is given by

$$\eta_t^{2,P} = (X_{t+1}^P + W_{t+1}^P - \bar{X}_t^P - S_t^P) + (D_t - u_{t+1}).$$

Thus for any $\beta \in I_a$,

$$\begin{aligned}
\sum_{t=\xi_a^1}^{\beta} \eta_t^{2,P} &= \sum_{t=\xi_a^1}^{\beta} (X_{t+1}^P + W_{t+1}^P - \bar{X}_t^P - S_t^P) + \sum_{t=\xi_a^1}^{\beta} (D_t - u_{t+1}) \qquad (3.25) \\
&= \sum_{t=\xi_a^1}^{\beta} (X_{t+1}^P + W_{t+1}^P - \bar{X}_{t+1}^P - S_{t+1}^P) + (\bar{X}_{\beta+1}^P + S_{\beta+1}^P) - (\bar{X}_{\xi_a^1}^P + S_{\xi_a^1}^P) \\
&\quad + \sum_{t=\xi_a^1}^{\beta} (D_t - u_{t+1}) \\
&= -\sum_{t=\xi_a^1}^{\beta} (\bar{X}_{t+1}^P - X_{t+1}^P - W_{t+1}^P)^+ + (\bar{X}_{\beta+1}^P + S_{\beta+1}^P) - (\bar{X}_{\xi_a^1}^P + S_{\xi_a^1}^P) \\
&\quad + \sum_{t=\xi_a^1}^{\beta} (D_t - u_{t+1}).
\end{aligned}$$

Since $\beta \in \mathcal{T}_c^c$ and $\xi_a^1 - 1 \in \mathcal{T}_c$, we have $\bar{X}_{\beta+1}^{MSMB} + S_{\beta+1}^{MSMB} \leqslant \bar{X}_{\beta+1}^{OPT} + S_{\beta+1}^{OPT}$ and $\bar{X}_{\xi_a^1}^{MSMB} +$

$S_{\xi_a^1}^{MSMB} \geqslant \bar{X}_{\xi_a^1}^{OPT} + S_{\xi_a^1}^{OPT}$. Thus, from equality (3.25), we obtain

$$\begin{aligned}
& \sum_{t=\xi_a^1}^{\beta} \left[\eta_t^{2,MSMB} + (\bar{X}_{t+1}^{MSMB} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+ \right] \\
= & \ (\bar{X}_{\beta+1}^{MSMB} + S_{\beta+1}^{MSMB}) - (\bar{X}_{\xi_a^1}^{MSMB} + S_{\xi_a^1}^{MSMB}) + \sum_{t=\xi_a^1}^{\beta} (D_t - u_{t+1}) \\
\leqslant & \ (\bar{X}_{\beta+1}^{OPT} + S_{\beta+1}^{OPT}) - (\bar{X}_{\xi_a^1}^{OPT} + S_{\xi_a^1}^{OPT}) + \sum_{t=\xi_a^1}^{\beta} (D_t - u_{t+1}) \\
= & \ \sum_{t=\xi_a^1}^{\beta} \left[\eta_t^{2,OPT} + (\bar{X}_{t+1}^{OPT} - X_{t+1}^{OPT} - W_{t+1}^{OPT})^+ \right].
\end{aligned}$$

Note that $\eta_t^{2,MSMB}$ is the additional manufacturing quantity and $(\bar{X}_{t+1}^{MSMB} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+$ is the forced manufacturing quantity, the above inequality allows us to compare the cumulative manufacturing quantity.

Following the assumption that the unit production cost $\bar{c}_t$ is non-increasing in t, we define $\Delta_t = \bar{c}_t - c_{t+1}$ for all $t \in [\xi_a^1, \xi_a^2)$ and $\Delta_{\xi_a^2} = c_{\xi_a^2}$. Then since $\Delta_t \geqslant 0$ for all $t \in [\xi_a^1, \xi_a^2]$, by interchanging the order of summation, we conclude that

$$\begin{aligned}
\sum_{t \in I_a} (\hat{E}_t^{MSMB} + \tilde{E}_{t+1}^{MSMB}) &= \sum_{t=\xi_a^1}^{\xi_a^2} \bar{c}_t \left[\eta_t^{2,MSMB} + (\bar{X}_{t+1}^{MSMB} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+ \right] \\
&= \sum_{t=\xi_a^1}^{\xi_a^2} \left(\sum_{\beta=t}^{\xi_a^2} \Delta_\beta \right) \left[\eta_t^{2,MSMB} + (\bar{X}_{t+1}^{MSMB} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+ \right] \\
&= \sum_{\beta=\xi_a^1}^{\xi_a^2} \Delta_\beta \sum_{t=\xi_a^1}^{\beta} \left[\eta_t^{2,MSMB} + (\bar{X}_{t+1}^{MSMB} - X_{t+1}^{MSMB} - W_{t+1}^{MSMB})^+ \right] \\
&\leqslant \sum_{\beta=\xi_a^1}^{\xi_a^2} \Delta_\beta \sum_{t=\xi_a^1}^{\beta} \left[\eta_t^{2,OPT} + (\bar{X}_{t+1}^{OPT} - X_{t+1}^{OPT} - W_{t+1}^{OPT})^+ \right] \\
&= \sum_{t \in I_a} (\hat{E}_t^{OPT} + \tilde{E}_{t+1}^{OPT}),
\end{aligned}$$

which proves the second equality.

To show the last inequality, we first claim that for $\eta_t^{2,MSMB} = 0$ for all $t \in \mathcal{T}_\Phi \cap \mathcal{T}_c$. Otherwise, suppose $\eta_t^{2,MSMB} > 0$, then $\eta_t^{r,MSMB} = S_t^{MSMB}$. Thus,

$$Y_t^{MSMB} = \bar{X}_t^{MSMB} + S_t^{MSMB} + \eta_t^{2,MSMB} \geqslant \bar{X}_t^{OPT} + S_t^{OPT} + \eta_t^{2,OPT} \geqslant Y_t^{OPT},$$

which contradicts with $t \in \mathcal{T}_\Phi$. Similarly, we can show that $\eta_t^{2,OPT} = 0$ holds for all $t \in \mathcal{T}_\Psi \cap \mathcal{T}_c^c$. Hence, using the second inequality we have proved above, we have

$$\sum_{t \in \mathcal{T}_\Phi} \bar{c}_t \eta_t^{2,MSMB} = \sum_{t \in \mathcal{T}_\Phi \cap \mathcal{T}_c^c} \bar{c}_t \eta_t^{2,MSMB} \leqslant \sum_{t \in \mathcal{T}_c^c} \bar{c}_t \eta_t^{2,MSMB} \tag{3.26}$$

and

$$\sum_{t \in \mathcal{T}_\Phi} \bar{c}_t \eta_t^{2,OPT} = \sum_{t \in \mathcal{T}_\Phi \cup \mathcal{T}_c^c} \bar{c}_t \eta_t^{2,OPT} \geqslant \sum_{t \in \mathcal{T}_c^c} \bar{c}_t \eta_t^{2,OPT}. \tag{3.27}$$

Therefore, using the above inequalities together with the two inequalities we have already showed in the lemma, we conclude that

$$\begin{aligned}
\sum_{t \in \mathcal{T}_\Phi} \hat{E}_t^{MSMB} + \sum_{t \in \mathcal{T}_\Psi} \tilde{E}_{t+1}^{MSMB} &\leqslant \sum_{t \in \mathcal{T}_c^c} \hat{E}_t^{MSMB} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c^c} \tilde{E}_{t+1}^{MSMB} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c} \tilde{E}_{t+1}^{MSMB} \\
&= \sum_{t \in \mathcal{T}_c^c} (\hat{E}_t^{MSMB} + \tilde{E}_{t+1}^{MSMB}) - \sum_{t \in \mathcal{T}_\Phi \cap \mathcal{T}_c^c} \tilde{E}_{t+1}^{MSMB} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c} \tilde{E}_{t+1}^{MSMB} \\
&\leqslant \sum_{t \in \mathcal{T}_c^c} (\hat{E}_t^{OPT} + \tilde{E}_{t+1}^{OPT}) - \sum_{t \in \mathcal{T}_\Phi \cap \mathcal{T}_c^c} \tilde{E}_{t+1}^{OPT} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c} \tilde{E}_{t+1}^{OPT} \\
&= \sum_{t \in \mathcal{T}_c^c} \hat{E}_t^{OPT} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c^c} \tilde{E}_{t+1}^{OPT} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c} \tilde{E}_{t+1}^{OPT} \\
&\leqslant \sum_{t \in \mathcal{T}_\Phi} \hat{E}_t^{OPT} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c^c} \tilde{E}_{t+1}^{OPT} + \sum_{t \in \mathcal{T}_\Psi \cap \mathcal{T}_c} \tilde{E}_{t+1}^{OPT} \\
&= \sum_{t \in \mathcal{T}_\Phi} \hat{E}_t^{OPT} + \sum_{t \in \mathcal{T}_\Psi} \tilde{E}_{t+1}^{OPT},
\end{aligned}$$

where the first inequality follows from (3.26), the third inequality follows from the results proved in part 1 and part 2, and the fifth inequality follows from (3.27). □

Combining Lemma 3.7 to Lemma 3.9 together, we have

$$\begin{aligned}
\mathbb{E}[\mathcal{C}(MSMB)] &= 2\mathbb{E}\Bigg[\sum_{t\in\mathcal{T}_\Phi}\Phi_t^{MSMB}+\sum_{t\in\mathcal{T}_\Psi}\Psi_t^{MSMB}\Bigg]+\tilde{H}_1^{MSMB}+\tilde{E}_1^{MSMB}\\
&= 2\mathbb{E}\Bigg[\sum_{t\in\mathcal{T}_\Phi}\hat{H}_t^{MSMB}+\sum_{t\in\mathcal{T}_\Psi}(\tilde{H}_{t+1}^{MSMB}+B_t^{MSMB})\\
&\quad +(\sum_{t\in\mathcal{T}_\Phi}\hat{E}_t^{MSMB}+\sum_{t\in\mathcal{T}_\Psi}\tilde{E}_{t+1}^{MSMB})\Bigg]+\tilde{H}_1^{MSMB}+\tilde{E}_1^{MSMB}\\
&\leqslant 2\mathbb{E}\Bigg[\sum_{t\in\mathcal{T}_\Phi}\hat{H}_t^{OPT}+\sum_{t\in\mathcal{T}_\Psi}(\tilde{H}_{t+1}^{OPT}+B_t^{OPT})+(\sum_{t\in\mathcal{T}_\Phi}\hat{E}_t^{OPT}+\sum_{t\in\mathcal{T}_\Psi}\tilde{E}_{t+1}^{OPT})\Bigg]\\
&\quad +\tilde{H}_1^{OPT}+\tilde{E}_1^{OPT}\\
&\leqslant 2\mathbb{E}\Bigg[\sum_{t=1}^{T}(\Phi_t^{OPT}+\Psi_t^{OPT})+\tilde{H}_1^{OPT}+\tilde{E}_1^{OPT}\Bigg]\\
&= 2\mathbb{E}[\mathcal{C}(OPT)].
\end{aligned}$$

Hence, we have proved the following theorem, which provides a worst-case performance guarantee on the result of the MSMB policy.

Theorem 3.10. *The MSMB policy has a worst-case performance guarantee of two, i.e., for each instance of the backlogging model under service-level constraints, the expected cost of the MSMB policy is at most two times the expected cost of an optimal solution, i.e.,* $\mathbb{E}[\mathcal{C}(MSMB)] \leqslant 2\mathbb{E}[\mathcal{C}(OPT)]$.

3.6 Numerical Experiments

Since the remanufacturing model generalizes the classical backlogging model, we only focus on testing the MSMB policy and compare with the optimal policies derived through DP (for small problem sizes). Our numerical results show that the MSMB policy performs near-optimal for a set of instances with diverse demand and parameter settings. Moreover, the performance of the MSMB policy improves as we increase the levels of the QoS guarantee.

Demand process. We consider the following three demand settings:

1. **Independent and identically distributed (i.i.d.) demands**. We test three specific demand distributions, namely, Exponential, Erlang-2, and Poisson all with mean values equal to 10.

2. **Markov-modulated demand process (MMDP).** MMDP considers an underlying Markov Chain and assumes that the demand distribution depends on the state of the Markov Chain. The state at period t, denoted by $s_t \in \{1, 2, 3\}$ and is interpreted as the state of the economy (poor, fair or good). Given state s_t at period t, the demand is a random variable with cumulative distribution function $F_t(\cdot)$ and mean value μ_t. The better the state of economy, the larger the mean of the demand, i.e., $\mu_1 < \mu_2 < \mu_3$. The transition probability matrix is defined by $\mathbf{P} = (p_{ij})_{3\times 3}$, where p_{ij} denotes the transition probability from state s_i to state s_j. In our test data, the state of the economy follows a Markov chain with initial probabilities $p_1 = p_2 = p_3 = 1/3$ and transition probabilities

$$\mathbf{P} = \begin{bmatrix} 0.6 & 0.3 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.1 & 0.3 & 0.6 \end{bmatrix}.$$

For each state $s_t \in \{1, 2, 3\}$, we also set the demand mean value as $5s_t$ in period t.

We test three specific demand distributions, namely, Poisson, Uniform and Normal. The parameter of the Poisson distribution is solely governed by the mean value (set as $5, 10$, and 15). For the Uniformly distributed demand, we consider intervals $[0, 10]$, $[5, 15]$ and $[10, 20]$ for the three states, respectively. For the Normal distribution, we set the standard deviation as 2 for all three states.

3. **Autoregressive demands.** For the autoregressive demand, we assume that there is a priori estimation μ_t of the demand at the period t. Besides, the realized demand also depends on the actual sales (or actual deviations from the priori) of previous seasons. Hence, the generic autoregressive demand model with parameter γ has the following form:

$$D_t - \mu_t = \sum_{s=t-\gamma}^{t-1} \psi_{t-s}(D_s - \mu_s) + \epsilon,$$

where ψ_t stands for the extent of correlation for the demand deviations and ϵ is the noise term which is assumed to be a Gaussian white noise (i.e., standard normal distribution). The coefficients $\{\psi_t\}_{t=1}^{\gamma}$ are usually determined by the auto-covariances following the *Yule-Walker Equations* (cf. Hamilton 1994). Our numerical tests cover the following three cases:

- $\gamma = 1$, with coefficients $\psi_1 = 1$;
- $\gamma = 2$, with coefficients $(\psi_1, \psi_2) = (3/4, 1/4)$;

- $\gamma = 3$, with coefficients $(\psi_1, \psi_2, \psi_3) = (1/2, 1/3, 1/6)$;

Return process. We consider two types of return process (a) independent product return with U_t following a Normal distribution with mean $\mu = 20$ and standard deviation $\sigma = 5$; (b) dependent product return on the previous demand with $U_t = 0.3D_{t-1} + \nu_t$, where $\{\nu_t\}$ are i.i.d. Poisson random variables with rate 1.

Paramater settings. We consider a planning horizon $T = 20$ periods and production lead time $L = 2$. We assume that the cost parameters are stationary with a discounted factor $\alpha = 0.99$ and a unit holding cost being normalized to 1. We test different combinations of cost parameters under the three types of demand patterns. Specifically, we assume the unit remanufacturing cost $c_1 = 30$, the unit manufacturing cost $c_2 = 30, 40, 50$, and the unit backlogging cost $p = 50, 70, 90$.

Performance measure. To evaluate the performance of a policy P, we compare it with the results of the optimal policy. We use $\mathcal{C}(P)$ and $\mathcal{C}(OPT)$ to denote the costs given by the two policies, respectively. We define the *performance error* of the policy P as the percentage of increase in the total cost of this policy compared to the optimal cost over the planning horizon, i.e.,

$$\mathcal{E} = \frac{\mathcal{C}(P) - \mathcal{C}(OPT)}{\mathcal{C}(OPT)} \times 100\%.$$

Clearly, the performance error $\mathcal{E}$ is always a positive number. Moreover, a smaller performance error means a better approximation algorithm. We report the values of $\mathcal{E}$ for every testing combination to empirically show that the proposed SMB policy provides close-to-optimal solutions in much more competitive CPU time. All of the numerical experiments are conducted on an Intel(R) Xeon(R) 2.93 GHz PC and we use Matlab R2013a as the solver.

3.6.1 Numerical Results

Tables 3.1–3.6 present all the numerical results: Tables 3.1–3.3 cover the independent product return case and Tables 3.4–3.6 cover the dependent return case. For each instance, we test performance errors of the SMB policy for the i.i.d. demand, MMDP demand and autoregressive demand. Note that the average empirical performance error of the SMB policy is less than 2% in all instances, demonstrating the efficacy of the proposed approximation algorithm. Moreover, comparing the results of i.i.d., MMDP, and autoregressive demands, our algorithm performs consistently better in instances when demands are correlated. This indicates that the SMB policy takes the advantage of given demand correlation information. On the other hand, the average CPU time of the SMB policy is around 1.16 seconds. In contrast, the DP algorithm for finding optimal solutions takes much longer time (173.9 seconds on average) per test instance.

$(c_1,\, c_2,\, p)$	Exponential				Erlang-2				Poisson			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	1.17%	0.95%	1.40%	1.07%	0.52%	0.24%	0.49%	0.81%	0.08%	0.21 %	0.36%	0.57%
$(30, 30, 70)$	0.97%	0.81%	1.11%	1.09%	1.13%	0.48%	0.42%	0.94%	0.04%	0.09%	0.24%	0.49%
$(30, 30, 90)$	1.39%	0.65%	0.77%	1.43%	1.44%	0.69%	0.21%	0.56%	0.11%	0.07%	0.25%	0.43%
$(30, 40, 50)$	0.64%	1.11%	0.75%	0.65%	0.80%	0.46%	0.57%	0.82%	0.06%	0.15%	0.40%	0.58%
$(30, 40, 70)$	0.91%	0.70%	1.24%	1.61%	0.97%	0.56%	0.31%	0.53%	0.08%	0.11%	0.33%	0.51%
$(30, 40, 90)$	0.76%	0.84%	1.35%	1.40%	1.41%	0.72%	0.38%	0.41%	0.10%	0.07 %	0.24%	0.41%
$(30, 50, 50)$	1.28%	0.99%	0.94%	1.01%	0.73%	0.19%	0.46%	0.95%	0.04%	0.11 %	0.37%	0.55%
$(30, 50, 70)$	1.25%	1.00%	1.17%	1.13%	1.29%	0.48%	0.61%	0.63%	0.08%	0.09 %	0.34%	0.52%
$(30, 50, 90)$	1.43%	1.23%	0.92%	1.14%	1.67%	0.53%	0.50%	0.65%	0.18%	0.12%	0.19%	0.40%
max	1.43%	1.23%	1.4%	1.61%	1.67%	0.72%	0.61%	0.95%	0.18%	0.21 %	0.40%	0.58%
mean	1.09%	0.92%	1.07%	1.17%	1.11%	0.48%	0.44%	0.70%	0.09%	0.11%	0.30%	0.50%

Table 3.1: Error $\mathcal{E}$ for i.i.d. demands with different parameters (independent return)

$(c_1,\, c_2,\, p)$	Poisson				Uniform				Normal			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	0.42%	0.16%	0.66%	0.85%	0.39%	0.35%	0.71%	0.90%	0.20%	0.37%	0.65%	0.74%
$(30, 30, 70)$	0.77%	0.29%	0.37%	0.74%	0.55%	0.34%	0.61%	0.83%	0.27%	0.25%	0.55%	0.76%
$(30, 30, 90)$	0.95%	0.39%	0.49%	0.55%	0.72%	0.37%	0.54%	0.82%	0.40%	0.28%	0.49%	0.76%
$(30, 40, 50)$	0.39%	0.30%	0.64%	0.79%	0.40%	0.42%	0.72%	0.88%	0.21%	0.29%	0.61%	0.79%
$(30, 40, 70)$	0.65%	0.31%	0.51%	0.61%	0.56%	0.41%	0.67%	0.87%	0.28%	0.26%	0.56%	0.73%
$(30, 40, 90)$	0.79%	0.30%	0.40%	0.58%	0.65%	0.38%	0.54%	0.84%	0.41%	0.24%	0.43%	0.73%
$(30, 50, 50)$	0.46%	0.25%	0.50%	0.95%	0.40%	0.42%	0.75%	0.87%	0.21%	0.30%	0.62%	0.82%
$(30, 50, 70)$	0.63%	0.27%	0.43%	0.65%	0.56%	0.35%	0.62%	0.84%	0.28%	0.27%	0.61%	0.75%
$(30, 50, 90)$	0.98%	0.39%	0.43%	0.59%	0.64%	0.34%	0.55%	0.77%	0.36%	0.27%	0.49%	0.68%
max	0.98%	0.39%	0.66%	0.95%	0.72%	0.42%	0.75%	0.90%	0.41%	0.37%	0.65%	0.82%
mean	0.67%	0.30%	0.49%	0.70%	0.54%	0.38%	0.63%	0.85%	0.29%	0.28%	0.55%	0.75%

Table 3.2: Error $\mathcal{E}$ for MMDP demands with different parameters (independent return)

$(c_1,\, c_2,\, p)$	AR(1)				AR(2)				AR(3)			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	0.26%	0.30%	0.52%	0.71%	0.22%	0.35%	0.54%	0.71%	0.20%	0.36%	0.58%	0.67%
$(30, 30, 70)$	0.39%	0.33%	0.44%	0.62%	0.33%	0.31%	0.49%	0.64%	0.24%	0.31%	0.50%	0.64%
$(30, 30, 90)$	0.59%	0.38%	0.46%	0.59%	0.51%	0.33%	0.54%	0.62%	0.31%	0.34%	0.43%	0.63%
$(30, 40, 50)$	0.27%	0.34%	0.55%	0.73%	0.24%	0.31%	0.57%	0.68%	0.18%	0.30%	0.57%	0.68%
$(30, 40, 70)$	0.46%	0.30%	0.42%	0.63%	0.30%	0.33%	0.52%	0.66%	0.26%	0.32%	0.46%	0.66%
$(30, 40, 90)$	0.64%	0.38%	0.43%	0.67%	0.41%	0.33%	0.47%	0.58%	0.34%	0.31%	0.44%	0.66%
$(30, 50, 50)$	0.29%	0.30%	0.58%	0.73%	0.24%	0.32%	0.49%	0.68%	0.20%	0.32%	0.60%	0.70%
$(30, 50, 70)$	0.45%	0.32%	0.49%	0.71%	0.31%	0.33%	0.50%	0.68%	0.24%	0.32%	0.50%	0.63%
$(30, 50, 90)$	0.63%	0.35%	0.43%	0.60%	0.42%	0.33%	0.46%	0.63%	0.30%	0.33%	0.46%	0.63%
max	0.64%	0.38%	0.58%	0.73%	0.51%	0.35%	0.57%	0.71%	0.34%	0.36%	0.60%	0.70%
mean	0.44%	0.33%	0.48%	0.66%	0.33%	0.33%	0.51%	0.65%	0.25%	0.32%	0.50%	0.66%

Table 3.3: Error $\mathcal{E}$ for Autoregressive demands with different parameters (independent return)

(c_1, c_2, p)	Exponential				Erlang-2				Poisson			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	0.48%	0.99%	1.54%	1.75%	0.51%	0.97%	1.57%	1.74%	0.41%	0.94%	1.42%	1.59%
$(30, 30, 70)$	0.46%	0.69%	1.33%	1.66%	0.50%	0.68%	1.37%	1.64%	0.31%	0.64%	1.21%	1.53%
$(30, 30, 90)$	0.58%	0.54%	1.14%	1.57%	0.61%	0.55%	1.11%	1.55%	0.34%	0.50%	1.05%	1.47%
$(30, 40, 50)$	0.47%	0.77%	1.29%	1.47%	0.47%	0.98%	1.57%	1.72%	0.35%	0.72%	1.15%	1.34%
$(30, 40, 70)$	0.46%	0.61%	1.14%	1.43%	0.49%	0.68%	1.36%	1.64%	0.32%	0.56%	1.01%	1.32%
$(30, 40, 90)$	0.57%	0.51%	1.00%	1.32%	0.59%	0.53%	1.19%	1.58%	0.34%	0.45%	0.89%	1.23%
$(30, 50, 50)$	0.45%	0.67%	1.10%	1.31%	0.49%	0.97%	1.55%	1.78%	0.32%	0.63%	1.02%	1.17%
$(30, 50, 70)$	0.49%	0.56%	1.02%	1.26%	0.48%	0.70%	1.36%	1.64%	0.32%	0.52%	0.91%	1.14%
$(30, 50, 90)$	0.58%	0.48%	0.88%	1.21%	0.64%	0.55%	1.14%	1.58%	0.34%	0.43%	0.80%	1.07%
max	0.58%	0.99%	1.54%	1.75%	0.64%	0.98%	1.57%	1.78%	0.41%	0.94%	1.42%	1.59%
mean	0.50%	0.65%	1.16%	1.44%	0.53%	0.73%	1.36%	1.65%	0.34%	0.60%	1.05%	1.32%

Table 3.4: Error $\mathcal{E}$ for i.i.d. demands with different parameters (dependent return)

(c_1, c_2, p)	Poisson				Uniform				Normal			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	0.55%	0.60%	0.90%	1.06%	0.50%	0.98%	1.54%	1.73%	0.35%	0.70%	1.32%	1.49%
$(30, 30, 70)$	0.99%	0.51%	0.83%	1.06%	0.50%	0.69%	1.37%	1.66%	0.22%	0.48%	1.14%	1.41%
$(30, 30, 90)$	0.58%	0.53%	0.55%	0.90%	0.58%	0.54%	1.12%	1.55%	0.18%	0.33%	1.01%	1.32%
$(30, 40, 50)$	0.64%	0.56%	0.83%	1.04%	0.44%	0.78%	1.27%	1.48%	0.27%	0.57%	1.10%	1.23%
$(30, 40, 70)$	0.99%	0.52%	0.64%	0.92%	0.48%	0.58%	1.13%	1.41%	0.19%	0.37%	0.94%	1.18%
$(30, 40, 90)$	0.41%	0.64%	0.56%	0.92%	0.61%	0.49%	0.98%	1.34%	0.19%	0.29%	0.81%	1.08%
$(30, 50, 50)$	0.73%	0.53%	0.69%	0.90%	0.45%	0.66%	1.10%	1.31%	0.24%	0.44%	0.93%	1.09%
$(30, 50, 70)$	0.57%	0.51%	0.59%	0.87%	0.50%	0.53%	1.01%	1.27%	0.19%	0.33%	0.82%	1.01%
$(30, 50, 90)$	0.41%	0.63%	0.55%	0.81%	0.55%	0.46%	0.89%	1.22%	0.19%	0.27%	0.71%	0.95%
max	0.99%	0.64%	0.90%	1.06%	0.61%	0.98%	1.54%	1.73%	0.35%	0.70%	1.32%	1.49%
mean	0.65%	0.56%	0.68%	0.94%	0.51%	0.64%	1.16%	1.44%	0.23%	0.42%	0.97%	1.20%

Table 3.5: Error $\mathcal{E}$ for MMDP demands with different parameters (dependent return)

(c_1, c_2, p)	AR(1)				AR(2)				AR(3)			
	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$	$\theta = 0.9$	$\theta = 0.95$	$\theta = 0.98$	$\theta = 0.99$
$(30, 30, 50)$	0.52%	0.89%	1.42%	1.59%	0.51%	0.98%	1.43%	1.59%	0.62%	1.04%	1.46%	1.59%
$(30, 30, 70)$	0.51%	0.69%	1.26%	1.49%	0.48%	0.75%	1.28%	1.48%	0.50%	0.85%	1.33%	1.52%
$(30, 30, 90)$	0.69%	0.59%	1.10%	1.42%	0.49%	0.66%	1.12%	1.42%	0.46%	0.71%	1.19%	1.46%
$(30, 40, 50)$	0.44%	0.70%	1.14%	1.32%	0.42%	0.74%	1.16%	1.32%	0.46%	0.79%	1.19%	1.31%
$(30, 40, 70)$	0.46%	0.56%	1.02%	1.25%	0.39%	0.61%	1.05%	1.23%	0.42%	0.67%	1.10%	1.26%
$(30, 40, 90)$	0.55%	0.51%	0.86%	1.18%	0.44%	0.55%	0.93%	1.19%	0.41%	0.59%	0.96%	1.22%
$(30, 50, 50)$	0.39%	0.57%	0.99%	1.17%	0.38%	0.62%	1.00%	1.14%	0.42%	0.65%	1.01%	1.13%
$(30, 50, 70)$	0.43%	0.50%	0.88%	1.12%	0.37%	0.52%	0.91%	1.10%	0.38%	0.59%	0.94%	1.10%
$(30, 50, 90)$	0.51%	0.46%	0.78%	1.02%	0.43%	0.51%	0.81%	1.06%	0.39%	0.53%	0.85%	1.04%
max	0.69%	0.89%	1.42%	1.59%	0.51%	0.98%	1.43%	1.59%	0.62%	1.04%	1.46%	1.59%
mean	0.50%	0.61%	1.05%	1.29%	0.44%	0.66%	1.08%	1.28%	0.45%	0.71%	1.11%	1.29%

Table 3.6: Error $\mathcal{E}$ for Autoregressive demands with different parameters (dependent return)

3.7 Concluding Remarks

In this chapter, we studied two stochastic inventory systems with probabilistic guarantees of service-levels (interpreted as stockout probabilities) in each period of a planning horizon. In particular, we derived structural properties of optimal policies for both backlogging and remanufacturing models. The chapter also proposed several efficient and easily implementable approximation algorithms for computing near-optimal solutions, of which the efficacy is demonstrated through numerical experiments with diverse demand settings.

One interesting and important future research avenue is to consider a joint service-level constraint for restricting the stockout probability in any period over a finite time horizon. To the best of our knowledge, only Zhang et al. (2014) has considered a related dynamic lot-sizing problem with a joint chance constraint on stockout probability. The authors formulated as a multi-stage stochastic integer programming model solved by cutting-plane algorithms. There have not been existing papers have characterized optimal or near-optimal policies for stochastic inventory models with joint service-level constraints, which we believe is an important future research topic.

CHAPTER IV

Optimal Dynamic Pricing with Sales Rank Information

4.1 Abstract

With increased popularity of online retail, we see exploding use of sales rank information. Such information is critical for customers as it allows them to choose which product to buy. This implies high importance of sales rank information also for retailers. If choosing lower prices, the sales increase and the rank of the product rises. Per its nature, ranking is a latent variable (current rank reflects the previous sales) and, therefore, it allows retailers to strategically influence ranking of the product.

In this chapter, we study periodic-review dynamic pricing problems in presence of sales rank information. The demand in each period is a function of both price and sales rank. The retailer's goal is to find the optimal pricing policy that takes sales rank into account to maximize the total revenue over a finite time horizon T.

To abstract the critical drivers of the retailer's policy, we first consider a deterministic model and then extend it to stochastic settings. With deterministic demand and a single product, we show that the optimal pricing policy is cyclical – the retailer alternates between high and low prices. We show how the the length of the optimal cycle depends on the expected revenue and customers' sensitivity to sales rank. We evaluate the benefit of strategic cycling versus myopic policy of the retailer. For stochastic demand case, we derive the structure of the optimal pricing policy, which generalizes the results in the deterministic case. We show that it is upper bounded by single-period myopic optimal price and is monotonically decreasing in the sales rank. Our numerical experiments illustrate the potential of revenue increases when a strategic-cycling pricing policy is used. We consider the demand and rank sensitivities, as well as different demand patterns. We also observe that penetration policy

is used in optimal policy for products with lower sales rank.

4.2 Introductory Remarks

Many forms of dynamic pricing strategies such as price markdowns, promotions, coupons have been used in practice for many years (cf. Elmaghraby and Keskinocak 2003). In the last decade, the benefits of dynamic pricing strategies have been not only well-acknowledged but also increasingly studied and refined (cf. Talluri and Van Ryzin 2006). It is widely expected that, with increased availability of demand data as well as the flexibility of changing prices, the volume of applications of dynamic pricing will further increase.

The information technology plays an important role for both retailers and consumers. The rapid growth of Internet and e-commerce makes it possible for consumers to search across many online stores, using price-comparison engines available at desktops, mobile sites, and apps. The instantaneous price information increases price competition, dwindles product differentiation, and may decrease brand loyalty (see, e.g., Robert 1998). With consumers having access to real-time information, real-time dynamic pricing becomes necessary for online retailers in order to react in real-time to competitors' price changes.

In addition to providing the instantaneous price information, online channels also make available a wealth of other product sales information such as sales rank, customers' reviews, etc., and it is well documented that customers are aware of that information and take it into account when deciding whether and what product to purchase (cf. Kannan 2017). According to Nielsen (2010), 40% of online shoppers indicate that they would not even buy electronics without consulting online reviews first.

Unfortunately, the sales rank information may be both inaccurate as well as overwhelming. Given thousands of different brands and models, customers are usually not experts in differentiating the quality among similar products. As widely observed in practice (e.g., through use of websites) and in formal studies, customers pay close attention to ranking of products and tend to believe that bestsellers generally have good quality (see, e.g., Chen et al. 2011). As a result, customers often rely on sales history and may be more likely to purchase popular products simply because these products show on the very first page of their search. Not surprisingly, the impact of sales rank on the customers' demand is significant, which could play a crucial role in online retailer revenue management strategies (cf. Chen et al. 2011). A natural question to ask is whether retailer should offer a consistent (stationary) pricing policy, or whether other pricing strategies may be more beneficial to maximize profit.

In practice, prices and sales rank change significantly over time and the changes are related to each other. Many of today's successful companies use ranking information in their dynamic pricing strategies. For example, Amazon prices and sales rank of the Dell Laptop for the past six months can be found online.[1] As shown in Figure 4.1, the price varies and the highest price exceeds the lowest price by 40%. Similarly, see Figure 4.2, the sales rank changes continuously over time. Comparing these two figures, we see that the sales rank increases whenever there is a price drop-down. The similar results are observed by Remy et al. (2010), by collecting and analyzing Amazon's prices for bestsellers in the camera and video categories. The changes in price and sales rank point to the possibility that online retailers may have an interest in offering price discounts to attract more price-sensitive customers and to improve the sales rank of their products.

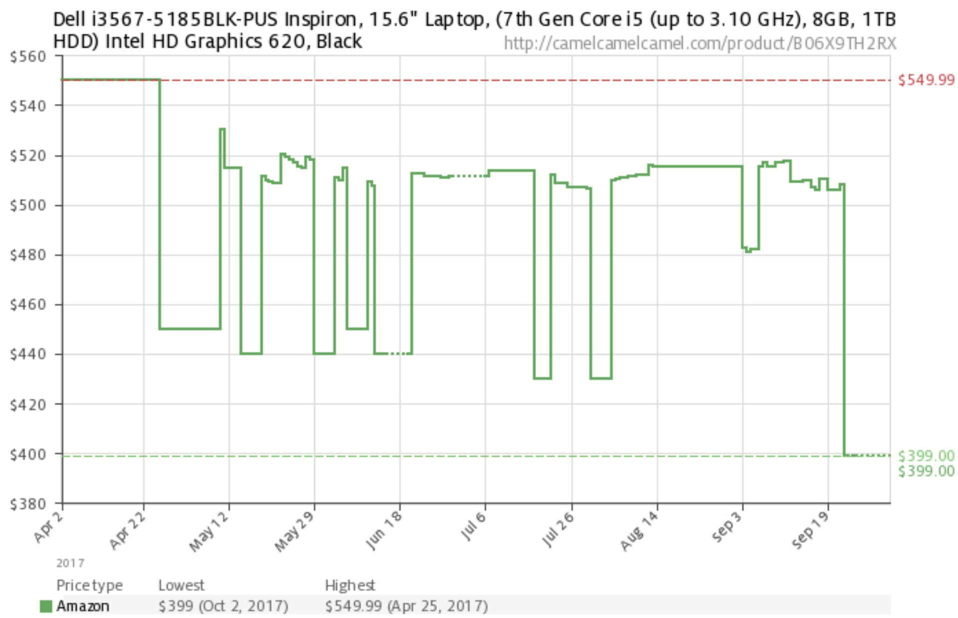

Figure 4.1: Price changes of Dell Laptop from Apr. 2, 2017 to Oct. 2, 2017

In this chapter, we study periodic-review dynamic pricing problems in presence of sales rank information. We assume that the demand in each period depends on current sales rank and current price. Given that product rank is a latent (delayed) variable, it is interesting how pricing policy that is based on sales rank could maximize retailers' total revenue. Our analysis is based on two scenarios where retailer sells either a single product or multiple products with similar characteristics. In the single-product pricing model, the sales rank for the product is either a deterministic or stochastic function of the demand in the previous period. In the multi-product model, we consider the deterministic case with several similar

[1] www.camelcamelcamel.com

Figure 4.2: Sales rank changes of Dell Laptop from Apr. 2, 2017 to Oct. 2, 2017

and substitutable products.

4.2.1 Literature Review

Because our work focuses on the effect of sales rank on product pricing, we review both of these streams of literature.

Sales rank. The literature on sales rank is fairly small, but it belongs to broader literature that studies the effect of historical sales data on current sales. A very significant portion of that literature focuses on the sales volume itself, notably the papers that analyze the diffusion effects (starting with seminal paper by Bass (1969)). Recently, a growing sub-literature focuses on the effect customer reviews (e.g., Chevalier and Mayzlin (2006), Mudambi and Schuff (2010) and Ho-Dac et al. (2013)). Our focus, however, is on a specific substream that analyzes sales *r*ank effects.

The relevance of sales rank is established in several papers. Goolsbee and Chevalier (2002) empirically study price sensitivity of online consumers in presence of sales ranks for two leading online booksellers, Amazon and Barnes and Nobel. They translate the observed sales ranks of each book into sales quantity by assuming the sales follow a Pareto distribution. Using publicly available data on the sales ranks for about 20,000 books, the authors show that there is a significant price sensitivity for online customers, both to a site's own price as well as to some rivals' price. They also show that prices are much more variable online than in retail stores.

There are empirical research papers that focus on distribution of products in online versus traditional retailers, as well as papers that focus on sequencing of the products when presented to customers and on design of appropriate ranking mechanisms.

Specifically, the long tail phenomenon of e-commerce has recently been studied and papers in that subarea find that online retailers sell more products that are less popular than do traditional retailers. For example, Brynjolfsson et al. (2003) investigate how customers benefits from the increased product variety at online booksellers. Brynjolfsson et al. (2011) also investigate the long tail phenomenon of the Internet channel by analyzing data from a multichannel retailer. They conclude that the Internet's long tail not only due to the product variety but also partly driven by lower search costs of the online channel.

It seems that ranking information increases popularity of both the niche products and of best sellers. Tucker and Zhang (2011) find that vendors of niche products benefit from being listed on websites that make popularity information highly important. Contrasting effect (steep tail or superstar) was documented in Rosen (1981) and Noe and Parker (2005).

Recently, Fleder and Hosanagar (2009) examine the rich-get-richer effect for popular products under the impact of recommender systems. Brynjolfsson et al. (2010) argue that both phenomenon (the long tail and the superstar) should be analyzed in an integrated way and identify four major areas of inquiry for future research, which includes its impact on pricing and other marketing strategies. There are attempts to design ranking mechanisms. E.g., Yoo and Kim (2012) study how ranking policy should be set to maximize the value of its online music ranking service. They design a ranking mechanism in which sellers can design the slot size to influence the popularity of music items while consumers will gain indirect benefits through segmented ranking slots and reduced search costs.

We are unaware of any papers that would analyze the effect of taking rank information into pricing policy, which is the focus of this chapter.

Dynamic pricing in revenue management. The study of dynamic pricing in revenue management is pioneered by Thomas (1970). Later the seminal paper Gallego and Van Ryzin (1994) consider a continuous-time formulation with limited inventories over finite-time horizon, where the demand intensity is a function of price. They show that the optimal price is decreasing in the stock level and is increasing the length of the remaining time horizon. The focus is how to use price to encourage customers to buy, given wide or scarce availability of the remaining inventory. Gallego and Van Ryzin (1997) consider a multi-product multi-resource dynamic pricing model and provide two heuristics based on deterministic counterparts. Many extensions of their models are reviewed in Talluri and Van Ryzin (2006) and Özer and Phillips (2012). Most of them assume the demand rate depends only on the current-period price. Recently, this assumption is relaxed by considering customer behaviors.

One major stream of literature focuses on pricing models with customers are forward-looking or strategic. That is, customers strategically timing their purchases, based on factors such as expectations of the future prices, the evolution of valuations, or availability of the product. This body of literature is vast, and we selectively discuss a few that are more closely related to our model (we refer interested readers to Shen and Su (2007) and Aviv et al. (2009) for an extensive overview). Conlisk et al. (1984) were the first to consider the problem of how sellers should price their products in settings where consumers with two possible valuations (high or low) arrive over time. They show that such a seller should use a cyclic policy, pricing high most of the time and dropping the price once in a while to sell to the accumulated mass of low valuation consumers. Besbes and Lobel (2015) then extend the above model to a more general version where customers are characterized by a triplet combination of their arrival time, valuation, and a willingness-to-wait, and show that a cyclic pricing policy is still optimal but often takes the form of a nested policy. More recently, Chen and Farias (2018) and Chen et al. (2018) show the efficacy of static fixed pricing policy with worst-case performance guarantees when customers are strategic.

There is also another growing stream of literature focusing on pricing models with customers having finite patience levels (see, e.g., Ahn et al. (2007), Liu and Cooper (2015), Lobel (2016)). The latter two papers show that optimal policies are periodic and providing a bound on the cycle length, and show that optimal policies are cyclic decreasing in the presence of heterogeneous patience levels. They construct a dynamic programming based algorithms for computing optimal dynamic pricing policies. More recently, Zhang and Jasin (2018) extend their model to a learning setting with the joint distribution of customers valuation and patience level is unknown *a priori*.

Different than the above streams of literature, we focus on role of product rank and its impact on the optimal pricing policies. Our customers are *not* really strategic, and we assume that the demand is affected by product rank (or past sales) in some structured way. In a broad sense, one can view product rank as a active covariate or side-information (see, e.g., recent works by Cohen et al. (2016), Qiang and Bayati (2016)).

Empirical literature studying different angles of dynamic pricing is also growing. Brynjolfsson and Smith (2000) empirically analyze the pricing behavior for the online channel and compare it with the conventional offline retail outlets. Their results indicate that Internet retailers have a higher incentive to make small price changes than conventional retailers. For more literature focus on empirical studies of online price dispersion, one may refer to Pan et al. (2004).

4.2.2 Main Results and Contributions

To study the structure of the optimal pricing policies, we use rank-based multi-period pricing models. Our main contributions are as follows:

- We propose and analyze single-product and multi-product pricing models for e-commerce retailers that incorporate the sales rank information. Sales rank is used by customers to decide whether to purchase the product as well as by retailers who can dynamically adjust their pricing decisions based on the current observed sales rank, to maximize their revenue. To the best of our knowledge, this is the first model that incorporates sales rank into revenue management problems.

- In the single-product model with deterministic customer arrivals, we characterize when cyclic pricing policy is optimal. Moreover, we show how the optimal cycle length depends on the customers' sensitivity to the sales rank and on the expected (retailer's) revenue generated by different prices.

- In a generalized single-product model with stochastic customer arrivals and continuous price set, we characterize the structure of the optimal policy. Under mild concavity assumptions on the revenue function and the sales rank function, we prove that the optimal price in each period is upper bounded by the single-period myopic optimal price. We also show that the optimal price increases as the product has a better sales rank.

- For multi-product rank-based pricing models, we also find that, cyclic policy is optimal. Interestingly, when customers decisions are based on product rankings and prices, retailers can manipulate the demand by increasing the price of the high-rank product and decreasing the price of the low-rank product at the same time, which benefits them due to the boosted demand for the high-rank product.

The rest of the chapter is organized as follows. In Section 4.3, we discuss a single-product model with deterministic customer arrivals. In Section 4.4, we generalize the single-product model by considering continuous price set and stochastic customer arrivals.

4.3 Single-Product Pricing with Deterministic Arrivals

4.3.1 Model Description

In this section, we consider a multi-period rank-based pricing model for a single product with two potential prices managed by an online retailer. The number of periods is T and each period is indexed by $t = 1, 2, \ldots, T$. In period t, n_t customers arrive. Then, the retailer chooses a price p_t from a discrete price set $\mathcal{P} = \{p^1, p^2\}$, with $p^1 < p^2$.

When the valuation of a customer is higher than the price of the product (i.e., $v > p_t$), the customer will buy the product. Each customer's valuation is non-negative and drawn from distribution $\mathcal{F}(\cdot)$. Let $F(\cdot)$ be the cumulative distribution function of the valuation and $\bar{F}(x) = 1 - F(x)$ denotes the proportion of customers who have a valuation at least x. Therefore, if the price in period t is p_t and the number of arrivals is n_t, the total demand in period t is $d_t = n_t\bar{F}(p_t)$.

The rank of the product in the following period is updated based on the quantity of the product sold in the current period. To reflect the reality of rank-based methods, rank with index 1 is the highest ranked product, and larger the index, the lower is the rank of the product. Following this convention, the rank of the product is a non-increasing function of last-period demand $r_{t+1} = h(d_t)$, where $h(\cdot)$ is a non-increasing function. This new rank in period $t + 1$ will affect the number of arriving customers in that period, i.e., $n_{t+1} = g(r_{t+1})$ where $g(\cdot)$ is a non-increasing function with upper bound $C = g(1)$.

The retailer's goal is to choose the prices $p_t \in \mathcal{P}$ such that the total revenue over T periods is maximized. We formulate the corresponding problems as follows:

$$
\begin{aligned}
z_T(r) = \max_{p_1,\ldots,p_T \in \mathcal{P}} \quad & \sum_{t=1}^{T} p_t d_t \\
\text{s.t.} \quad & n_t = \min\{C, g(r_t)\} \quad t = 1, 2, \ldots, T \\
& d_t = n_t \bar{F}(p_t) \quad t = 1, 2, \ldots, T \\
& r_{t+1} = h(d_t) \quad t = 1, 2, \ldots, T \\
& r_0 = r,
\end{aligned}
$$

where r denotes the initial rank of the product. The infinite-horizon objective is then given by

$$
z(r) = \liminf_{T\to\infty} \frac{1}{T} z_T(r) \tag{4.1}
$$

for each initial rank r.

4.3.2 Assumptions and Preliminary Results

Note that in our settings $n_{t+1} = g(h(d_t))$. In order to explicitly compute the optimal pricing policies, we will use fairly standard assumptions used in the literature that implies a specific form of this relationship, $n_{t+1} = \min\{C, \alpha d_t\}$.

Specifically, a standard distributional assumption for this type of rank data is a Pareto distribution (see, e.g., Goolsbee and Chevalier 2002), in which the probability that the sales D exceeds the current level is $\mathbb{P}(D \geqslant d) = (k/d)^{\theta}$. Also the product rank r is defined as satisfying: $\mathbb{P}(D \geqslant d) = r/R_{max}$. The above two equations indicate a log-linear relationship between sales and sales rank, i.e., $\ln d = \tau_0 + \tau_1 \ln r$ where $\tau_1 < 0$ represents the sensitivity of the sales rank to the demand. Therefore, given the rank r_t in period t, we have $\ln d_t = \tau_0 + \tau_1 \ln r_t$ and $\ln n_{t+1} = C + \tau_1' \ln r_t$ (since n_{t+1} is upper bounded by C). Thus, we have $n_{t+1} = \min\{C, \alpha d_t^{\tau_1'/\tau_1}\}$, where $\alpha = e^{C - \tau_0 \tau_1'/\tau_1}$. We assume in the remainder of this section that the sensitivity coefficient τ_1 remains the same for both demand-to-rank function $h(\cdot)$ and rank-to-arrival function $g(\cdot)$. This means $\tau_1 = \tau_1'$ and $n_{t+1} = \min\{C, \alpha d_t\}$. We note here that the parameter α can be interpreted as the customer's sensitivity to the sales rank – larger α implies that customers are more sensitive to the sales rank and past sales demand is more important in determining the current demand. In Section 4.4 for stochastic model, we will consider a generalized transition function.

Using the linear relationship between demand and number of customer arrivals, we formulate the dynamic program as follows:

$$v_{t,T}(n_t) = \max_{p_t \in \mathcal{P}} \left\{ n_t \bar{F}(p_t) p_t + v_{t+1,T}(\min\{C, \alpha n_t \bar{F}(p_t)\}) \right\}, \tag{4.2}$$

where n_t is the number of arrivals in period t and $v_{t,T}(n_t)$ denotes the total revenue from period t to period T.

To provide some preliminary results for the model, we define $\beta_i = \alpha \bar{F}(p^i)$ and $\mu_i = p^i \bar{F}(p^i)$ $(i = 1, 2)$. Since $p^2 > p^1$, we must have $\beta_1 > \beta_2$. As higher sales in period t result in higher arrivals in the following period, we label β_i as the discount factor for the number of arrivals in the next period (even though it is not necessarily smaller than 1). Parameter μ_i denotes the expected revenue for each customer arrival in the current period. Recall that the price set is $\mathcal{P} = \{p^1, p^2\}$. Thus, the dynamic programming formulation can be written as

$$v_{t,T}(n_t) = \max \left\{ n_t \mu_1 + v_{t+1,T}\big(\min\{C, n_t \beta_1\}\big), n_t \mu_2 + v_{t+1,T}\big(\min\{C, n_t \beta_2\}\big) \right\}, \tag{4.3}$$

where $v_{T+1,T}(\cdot) = 0$ and $v_{1,T}(x)$ is the maximum revenue given that the initial number of arrivals is $n_1 = x$. We drop the subscript T and use simplified notation $v_t(n_t)$ if the total number of periods is implied by the context.

We use $p_t^*(n_t)$ to denote the optimal pricing policy in period t, given that the arrivals in period t are $n_t \in (0, C]$. We also use a vector $\mathbf{p}^*(n) = (p_1^*, \ldots, p_T^*)$ to denote the optimal pricing policy with the initial arrivals in period 1, $n \in (0, C]$. Based on the above dynamic programming formulation, we have following results.

Lemma 4.1. *$v_{t,T}(x)$ is a non-decreasing function in $x \in (0, C]$.*

Proof. We show this by induction. First, $v_{T+1}(x)$ is trivially non-decreasing in x. Assuming that $v_{t+1}(x)$ is non-decreasing in x, we are going to show the same for $v_t(x)$. Let $0 \leqslant x < y \leqslant C$, then

$$\begin{aligned} v_t(x) &= \max\Big\{x\mu_1 + v_{t+1}\big(\min\{C, x\beta_1\}\big), x\mu_2 + v_{t+1}\big(\min\{C, x\beta_2\}\big)\Big\} \\ &< \max\Big\{y\mu_1 + v_{t+1}\big(\min\{C, y\beta_1\}\big), y\mu_2 + v_{t+1}\big(\min\{C, y\beta_2\}\big)\Big\} \\ &= v_t(y). \end{aligned}$$

Hence, $v_t(x)$ is a non-decreasing function in $x \in (0, C]$. □

Proposition 4.2. *If $\mu_1 \geqslant \mu_2$, then $p_t^*(n_t) = p^1$ for all $n_t \in (0, C]$.*

Proof. According to Lemma 4.1, we have $v_{t+1}(\min\{C, n_t\beta_1\}) \geqslant v_{t+1}(\min\{C, n_t\beta_2\})$ since $\beta_1 > \beta_2$. Moreover, since $n_t\mu_1 \geqslant n_t\mu_2$, we must have $v_t(n_t) = n_t\mu_1 + v_{t+1}\big(\min\{C, n_t\beta_1\}\big)$, which implies that $p_t^*(n_t) = p^1$ for all $n_t \in (0, C]$. □

The underlying intuition for Proposition 4.2 is straightforward. First, choosing the lower price p^1 will be beneficial in the future periods because there are more customer arrivals due to the higher rank. If the lower price also yields a higher revenue in the current period (i.e., $\mu_1 \geqslant \mu_2$), then there is no incentive to choose a higher price. Hence, imposing the following assumption on the expected revenue allows us to avoid trivial pricing decisions.

Assumption 4.3. *$\mu_2 > \mu_1$.*

Next, we consider the range of parameters β_1 and β_2.

Proposition 4.4. *(a) When $\beta_1 > \beta_2 \geqslant 1$, there exists $t_0 > 0$ such that the optimal pricing policy $p_t^* = p^1$ for all $t < t_0$ and $p_t^* = p^2$ for all $t \geqslant t_0$. Consequently, the optimal long-run average revenue $z(n) = \liminf_{T\to\infty} \frac{1}{T} v_{t,T}(n) = C\mu_2$ for all $n \in (0, C]$.*

(b) *When $\beta_2 < \beta_1 < 1$, the long-run average revenue*

$$z(n_1) = \liminf_{T\to\infty} \frac{1}{T} v_{1,T}(n_1) = 0$$

for all $n_1 \in (0, C]$.

(c) *When $\beta_1 = 1 > \beta_2$, the optimal pricing policy is to always charge low price, i.e., $p_t^* = p^1$ for all $t > 0$. The long-run average revenue*

$$z(n_1) = \liminf_{T\to\infty} \frac{1}{T} v_{t,T}(n_1) = n_1\mu_1$$

for all $n_1 \in (0, C]$.

Proof. (a) Since $\beta_1 > \beta_2 \geqslant 1$, we can always increase the number of arrivals up to C by charging lower price p^1 for a finite number of periods. Suppose $n_{t_0} = C$ for some $t_0 \geqslant 1$, then since $\mu_2 > \mu_1$ and $\beta_2 \geqslant 1$, it is optimal to charge the higher price p^2 in period t_0 as $C\mu_2 + v_{t_0+1}(C) > C\mu_1 + v_{t_0+1}(C)$. Thus, $p_{t_0}^* = p^2$ and $n_{t_0+1} = C$. By induction, we conclude that $p_t^* = p^2$ for all $t \geqslant t_0$. Let R denote the total revenue for the initial t_0 periods. Then, the long-run average revenue can be computed as $z(n) = \liminf_{T\to\infty} \frac{1}{T}(R + (T - t_0)C\mu_2) = C\mu_2$.

(b) Since $\beta_1 < 1$, the number of customer arrivals n_t must decrease in each period. Moreover, for any $\epsilon > 0$, there exists $t_0 > 0$ such that $n_{t_0} \leqslant n\beta_1^{t_0-1} < \epsilon$. Let R denote the total revenue for the initial t_0 periods. Then, the long-run average revenue is upper bounded by $z(n) \leqslant \liminf_{T\to\infty} \frac{1}{T}(R + (T - t_0)\epsilon\mu_2) = \epsilon\mu_2$. Letting $\epsilon \to 0$, we have that $z(n) = 0$ for all $n \in (0, C]$.

(c) First, the high price p^2 may only be optimal for a finite number of periods. Otherwise, following the same argument as in Proposition 4.4, the long-run average revenue would be 0. Now suppose t_0 is the last period p^2 is charged. Then, since $\beta = 1$, the number of arrivals in period $t \geqslant t_0$ must be $n\beta_2^m$, where m is the number of periods using p^2. Denoting the total revenue for the first t_0 periods as R, the long-run average revenue must be $z(n) = \liminf_{T\to\infty} \frac{1}{T}\{R + (T - t_0)n\beta_2^m\mu_2\} = n\beta_2^m\mu_2$. Since $\beta_2 < 1$, $z(n) \leqslant n\mu_2$ and the equality holds when $m = 0$, i.e., no high price is charged in the optimal policy. □

Proposition 4.4 eliminates three extreme cases for parameters β_1 and β_2.

First, when $\beta_1 > \beta_2 \geqslant 1$, the number of arrivals increases until it reaches capacity C. From then on, charging a higher price will not only result in a higher immediate revenue

$C\mu_2$, but also maintains the number of arrivals for the next period. In this case, the long-run average revenue reaches its maximum value.

When $\beta_2 < \beta_1 < 1$, the number of arrivals decreases exponentially over time. Hence, as the time horizon grows, the revenue obtained in each period eventually decays to 0 regardless of the pricing policy. In this case, the long-run average revenue attains its minimum value.

For $\beta_1 = 1$, charging a higher price decreases the number of arrivals which cannot be recovered in the future by charging a lower price, permanently reducing the revenue in each of the following periods. Therefore, in order to maximize the long-run average revenue, it is optimal to choose lower price p^1 in all periods.

Notice that the above proposition rules out the extreme cases, in which the long-run optimal policy is relatively easy to derive. Moreover, in practical applications, the ranking of a product usually increases as a lower price is charged, which results in a higher volume of customers in the next period. Similarly, a higher price decreases the sales rank of the product and yields a lower volume of customers in the subsequent period. Therefore, it is realistic (and also technically-appealing) to impose the Assumption 4.5 in our single-product model.

Assumption 4.5. $\beta_1 > 1 > \beta_2$.

For ease of analysis, we assume that the initial number of customer arrivals is C. This assumption will not change the structure of the optimal policy nor the optimal long-run revenue.[2]

Assumption 4.6. $n_1 = C$

Using Assumptions 4.3–4.6, we can further simplify the dynamic programming formulation (4.3) as

$$v_{t,T}(n_t) = \max\left\{ n_t\mu_1 + v_{t+1,T}\big(\min\{C, n_t\beta_1\}\big), n_t\mu_2 + v_{t+1,T}(n_t\beta_2) \right\}, \qquad (4.4)$$

since we always have $n_t\beta_2 < n_t \leqslant C$.

4.3.3 Structural Results for Two Special Cases

In this section, we consider two special cases that allow us to observe the dynamics of cyclic policy in its easiest (clearest) form. The assumptions on β_1 and β_2 allow us to significantly

[2] This is because the retailer can always charge lower price for a finite number of initial periods to increase the number of customers to C. The revenue for the initial periods will not affect the average long-run revenue.

reduce the number of states in the dynamic programming and also to precisely describe the policy, including features like the cycle length. The same policy will hold in general setting, as we will see in te following section.

- In the first case, we assume that there exists an integer $L \in \mathbb{N}^+$ such that $\beta_2 = \beta_1^{-L}$. This means the negative impact (on the ranking) of charging a higher price overwhelms the positive impact of charging a lower price.

- In the second case, we assume that there exists an integer $L \in \mathbb{N}^+$ such that $\beta_1 = \beta_2^{-L}$. In this case, charging a lower price has higher impact on ranking than charging a higher price.

In both cases, we characterize the structure of the optimal pricing policy.

Case 1: $\beta_2 = \beta_1^{-L}$

As the number of initial arrivals is $n = C$, the assumption on β_1 and β_2 implies that the states can be represented as $\{C/\beta_1^m : m = 0, 1, \dots\}$. Thus, we can simplify our dynamic program (4.4) as follows

$$v_t(m_t) = \max\left\{v_{t+1}((m_t-1)^+), \frac{\xi C}{\beta_1^{m_t}} + v_{t+1}(m_t+L)\right\} + \frac{\mu_1 C}{\beta_1^{m_t}}, \tag{4.5}$$

where $\xi = \mu_2 - \mu_1 > 0$. This allows us to characterize the optimal pricing policy.

Lemma 4.7. *Suppose $\beta_2 = \beta_1^{-L}$ for some $L \in \mathbb{N}^+$. Let*

$$v_{T+1}(m) = \begin{cases} -\mu_1 C(m - \sum_{i=1}^m \beta_1^{-i}) & \textit{if } m = 0, 1, \dots, L \\ v_{T+1}(m-L-1) - \xi C & \textit{if } m > L \end{cases}$$

be the salvage value at the end of the selling horizon.
If $\frac{\mu_2}{\mu_1} \geqslant L + 1 - \sum_{i=1}^L \beta_1^{-i}$, then the optimal policy $\{p_t : t = 1, 2, \dots T\}$ is

$$p_t(m) = \begin{cases} p^2 & \textit{if } m = 0 \\ p^1 & \textit{if } m = 1, 2, \dots \end{cases}$$

otherwise, the optimal policy is $p_t(m) = p^1$ for all $m = 0, 1, \dots$

Proof. Suppose $\frac{\mu_2}{\mu_1} \geqslant L + 1 - \sum_{i=1}^L \beta_1^{-i}$. Then, using formulation (4.5), $p_t(m) = p^2$ for $m = 0$ is equivalent to

$$v_{t+1}(0) \leqslant \xi C + v_{t+1}(L), \tag{4.6}$$

and $p_t(m) = p^1$ for $m \geqslant 1$ is equivalent to

$$v_{t+1}(m-1) > \frac{\xi C}{\beta_1^m} + v_{t+1}(m+L), \quad \forall m = 1, 2, \ldots \tag{4.7}$$

We use induction to show inequalities (4.6) and (4.7) hold for each period $t = T, T-1, \ldots, 1$.

- Base case. When $t = T$, by definition, $v_{T+1}(0) - v_{T+1}(L) = \mu_1 C(L - \sum_{i=1}^{L} \beta_1^{-i})$. Since $\frac{\mu_2}{\mu_1} \geqslant L + 1 - \sum_{i=1}^{L} \beta_1^{-i}$, $\mu_1 C(L - \sum_{i=1}^{L} \beta_1^{-i}) \leqslant \xi C$. Moreover, by the definition of $v_{T+1}(\cdot)$, for all $m \geqslant 1$, we have $v_{T+1}(m-1) - v_{T+1}(m+L) = \xi C > \frac{\xi C}{\beta_1^m}$. Hence, both inequalities (4.6) and (4.7) hold for $t = T$.

- Induction step. Suppose the inequalities (4.6) and (4.7) hold for all $s \geqslant t$, which implies $v_s(0) = \xi C + \mu_1 C + v_{s+1}(L)$ and $v_s(m) = v_{s+1}(m-1) + \frac{\mu_1 C}{\beta_1^m}$ for $m \geqslant 1$ and $s \geqslant t$. We want to show the same inequalities hold for $s = t - 1$.

  1. $v_t(0) - v_t(L) = \xi C + \mu_1 C + v_{t+1}(L) - v_t(L)$. When $t \leqslant T - L$, we have

  $$\begin{aligned} v_t(L) - v_{t+1}(L) &= v_{t+1}(L-1) - v_{t+2}(L-1) \\ &= \cdots \\ &= v_{t+L}(0) - v_{t+L+1}(0) \\ &= (\mu_1 + \xi)C + v_{t+L+1}(L) - v_{t+L+1}(0) \\ &\geqslant \mu_1 C, \end{aligned}$$

  where the last inequality follows from the induction hypothesis, i.e., $v_{t+L+1}(0) - v_{t+L+1}(L) \leqslant \xi C$. For $t > T - L$, we have

  $$\begin{aligned} v_t(L) - v_{t+1}(L) &= v_{t+1}(L-1) - v_{t+2}(L-1) \\ &= v_T(L - (T-t)) - v_{T+1}(L - (T-t)) \\ &= \frac{\mu_1 C}{\beta_1^{L-T+t}} + v_{T+1}(L - (T-t) - 1) - v_{T+1}(L - (T-t)) \\ &= \mu_1 C. \end{aligned}$$

  Hence, we conclude that $v_t(0) - v_t(L) = \xi C + \mu_1 C - (v_t(L) - v_{t+1}(L)) \leqslant \xi C$.

2. For $m \geqslant 1$, we have

$$\begin{aligned} v_t(m-1) &= \max\left\{v_{t+1}((m-2)^+), \frac{\xi C}{\beta_1^{m-1}} + v_{t+1}(m+L-1)\right\} + \frac{\mu_1 C}{\beta_1^{m-1}} \\ &\geqslant \frac{\xi C}{\beta_1^{m-1}} + v_{t+1}(m+L-1) + \frac{\mu_1 C}{\beta_1^{m-1}} \end{aligned}$$

and $v_t(m+L) = \frac{\mu_1 C}{\beta_1^{m+L}} + v_{t+1}(m+L-1)$ by induction hypothesis. Hence,

$$v_t(m-1) - v_t(m+L) \geqslant \frac{\xi C}{\beta_1^{m-1}} + v_{t+1}(m+L-1) - v_{s+1}(m+L-1) > \frac{\xi C}{\beta_1^{m}}$$

holds for all $m \geqslant 1$.

Therefore, we showed that inequalities (4.6) and (4.7) hold in each period $t = T, T-1, \ldots, 1$. As a result, the pricing policy

$$p_t(m) = \begin{cases} p^2 & \text{if } m = 0 \\ p^1 & \text{if } m = 1, 2, \ldots \end{cases}$$

is optimal.

Now suppose $\frac{\mu_2}{\mu_1} < L+1-\sum_{i=1}^{L} \beta_1^{-i}$. It suffices to show $v_{t+1}((m-1)^+) \geqslant \frac{\xi C}{\beta_1^m} + v_{t+1}(m+L)$ for all $t = 1, 2, \ldots, T$ and $m = 0, 1, \ldots$ Similarly, we prove by induction on t.

- Base case: When $t = T$, by definition, $v_{T+1}(0) - v_{T+1}(L) = \mu_1 C(L - \sum_{i=1}^{L} \beta_1^{-i}) > \xi C$. For $m \geqslant 1$, $v_{T+1}(m-1) - v_{T+1}(m+L) = \xi C$.

- Induction step: Suppose the statement holds for $s \geqslant t$. Then, for $t-1$,

1. When $m \geqslant 1$, we have

$$\begin{aligned} v_t(m-1) - v_t(m+L) &= v_{t+1}((m-1-1)^+) + \frac{\mu_1 C}{\beta_1^{m-1}} - v_{t+1}(m-1+L) - \frac{\mu_1 C}{\beta_1^{m+L}} \\ &> v_{t+1}((m-1-1)^+) - v_{t+1}(m-1+L) \\ &\geqslant \frac{\xi C}{\beta_1^{m-1}} > \frac{\xi C}{\beta_1^{m}}, \end{aligned}$$

where the middle inequality holds by induction hypothesis.

2. When $m = 0$, for $t \leqslant T-L$, $v_t(L) = \frac{\mu_1 C}{\beta_1^L} + v_{t+1}(L-1) = \cdots = \sum_{i=1}^{L} \frac{\mu_1 C}{\beta_1^i} + v_{t+L}(0)$. Note also that $v_t(0) = \mu_1 C + v_{t+1}(0) = \cdots = L\mu_1 C + v_{t+L}(0)$. Thus, $v_t(0) - v_t(L) = \mu_1 C(L - \sum_{i=1}^{L} \beta_1^{-i}) > \xi C$.

For $t > T - L$,

$$\begin{aligned}
v_t(L) &= \frac{\mu_1 C}{\beta_1^L} + v_{t+1}(L-1) \\
&= \cdots \\
&= \sum_{i=t+L-T}^{L} \frac{\mu_1 C}{\beta_1^i} + v_{T+1}(t+L-T-1) \\
&= \sum_{i=1}^{L} \frac{\mu_1 C}{\beta_1^i} - \mu_1 C(t+L-T-1)
\end{aligned}$$

Thus, $v_t(0) - v_t(L) = (T+1-t)\mu_1 C - v_t(L) = \mu_1 C(L - \sum_{i=1}^{L} \beta_1^{-i}) > \xi C$.

Therefore, we showed that inequality $v_{t+1}((m-1)^+) \geqslant \frac{\xi C}{\beta_1^m} + v_{t+1}(m+L)$ holds for all $t = 1, 2, \ldots, T$ and $m = 0, 1, \ldots$ Consequently, the pricing policy $p_t(m) = p^1$ is optimal for all $m = 0, 1, \ldots$ □

Using this result, we can characterize the optimal pricing policies in the infinite-horizon setting and describe the length of the optimal pricing cycle.

Theorem 4.8. *Suppose $\beta_2 = \beta_1^{-L}$.*
- *If $\frac{\mu_2}{\mu_1} \geqslant L + 1 - \sum_{k=1}^{L} \beta_1^{-k}$, then the optimal policy for the infinite-horizon problem is a cyclic policy $(p^2, p^1, \ldots, p^1)$ with cycle length $L+1$. Moreover, the long-run average revenue is $\frac{C\left(\mu_2 + \mu_1 \sum_{i=1}^{L} \beta_1^{-i}\right)}{L+1}$.*
- *Otherwise, the optimal policy is to always charge p^1. The long-run average revenue is then $C\mu_1$.*

For infinite-horizon problem, the salvage value does not matter and, thus, we may use the specific salvage value specified in Lemma 4.7.

When $\frac{\mu_2}{\mu_1} \geqslant L+1-\sum_{k=1}^{L} \beta_1^{-k}$, applying Lemma 4.7, we conclude that the optimal policy is a cycle of $L+1$ prices: $(p^2, p^1, \ldots p^1)$ (it begins with high price since $n_1 = C$). The long-run revenue is computed by the average revenue of a cycle, namely, $z(n_1) = \frac{C\left(\mu_2 + \mu_1 \sum_{i=1}^{L} \beta_1^{-i}\right)}{L+1}$.

When $\frac{\mu_2}{\mu_1} < L + 1 - \sum_{k=1}^{L} \beta_1^{-k}$, applying Lemma 4.7, we conclude that it is optimal to always use low price p^1, which has an average revenue of $C\mu_1$ per period.

Theorem 4.8 shows that the optimal pricing policy is cyclic with the cycle length $L+1$. The benefit of a cyclic pricing policy is based on the following dynamics. When the number of arrivals does not reach its capacity, a lower price is charged to increase the number of arrivals (hence the potential demand) in the next period. When the number of customers, however, reaches capacity C, choosing a higher price benefits the retailer since he/she will receive a large revenue due to both the large volume of demand and high price. This cyclic policy is optimal only if μ_2/μ_1 is sufficiently large. Specifically, the average revenue of a cycle must compensate for the average long-run revenue of charging the lower price p^1 in the remaining periods of the cycle, i.e., $\frac{C\left(\mu_2+\mu_1\sum_{i=1}^{L}\beta_1^{-i}\right)}{L+1} \geqslant C\mu_1$. Otherwise, it would be optimal to charge low price p^1 in each period.

Case 2: $\beta_1 = \beta_2^{-L}$

To derive the structure of the optimal pricing policy in this case, we employ a different methodology that does not require the dynamic programming formulation. We start with Lemma 4.9, where we show that, after low price, there are at most L consecutive high prices charged.

Lemma 4.9. *Let $L \geqslant 1$ be the integer such that $\frac{1}{\beta_2^{L-1}} < \beta_1 \leqslant \frac{1}{\beta_2^L}$. Then, there exists an optimal pricing policy satisfying the following property: after low price p^1 is charged in a given period, there are at most L consecutive periods in which high price p^2 is charged.*

Proof. First, if there exists $t_0 > 0$ such that $p_t^* = p^2$ for all $t \geqslant t_0$, then since $\beta_2 < 1$, we must have $n_t \to 0$ as $t \to \infty$ and hence, the long-run average revenue is 0 in this case. As this cannot be optimal, p^1 must occur infinitely many times in the optimal pricing policy.

We show the lemma by contradiction. Suppose in an optimal pricing policy $\{p_t\}_{t=1}^{\infty}$, there exists a low price p^1 followed by at least $L+1$ consecutive high price p^2, i.e., there exists $t \geqslant 1$ and some integer $k \geqslant L$ such that $p_{t-1} = p^1$, $p_t = p_{t+1} = \cdots = p_{t+k} = p^2$ and $p_{t+k+1} = p^1$. We consider another pricing policy $\{\tilde{p}_t\}_{t=1}^{\infty}$ where we switch prices in periods $t+k$ and $t+k+1$. That is, $\tilde{p}_s = p_s$ for all $s \neq t+k, t+k+1$ and $\tilde{p}_{t+k} = p^1$, $\tilde{p}_{t+k+1} = p^2$. Then, we have

1. $n_s = \tilde{n}_s$ for all $s = 1, \ldots, t+k$.

2. $n_{t+k+2} = n_t\beta_2^{k+1}\beta_1$ and $\tilde{n}_{t+k+2} = \min\{n_t\beta_2^k\beta_1, C\} \cdot \beta_2 = n_t\beta_2^{k+1}\beta_1$ since $\beta_2^k\beta_1 \leqslant \beta_2^L\beta_1 \leqslant 1$. Thus, $n_s = \tilde{n}_s$ holds for all $s \geqslant t+k+2$.

Therefore, it suffices to compare the total revenue in period $t+k$ and $t+k+1$. In fact, the revenue for policy $\{p_t\}_{t=1}^{\infty}$ in these two periods is $n_{t+k}\mu_2 + n_{t+k+1}\mu_1 = n_t\beta_2^k\mu_2 + n_t\beta_2^{k+1}\mu_1$. On the other hand, the revenue for policy $\{\tilde{p}_t\}_{t=1}^{\infty}$ in these two periods is $\tilde{n}_{t+k}\mu_1 + \tilde{n}_{t+k+1}\mu_2 = n_t\beta_2^k\mu_1 + n_t\beta_2^k\beta_1\mu_2 > n_t\beta_2^{k+1}\mu_1 + n_t\beta_2^k\mu_2$ since $\beta_1 > 1 > \beta_2$. Therefore, we achieve a policy with higher total revenue, which contradicts with the optimality of $\{p_t\}_{t=1}^{\infty}$. ☐

Lemma 4.9 has an intuitive explanation: the optimal pricing policy must keep the number of arrivals above a certain level, $\frac{C}{\beta_1}$, such that it can be restored to the maximum capacity C

immediately by charging a low price. This result implies that, if a pricing cycle starts with a low price p^1 and follows by a number of high prices p^2, the cycle length must be less than or equal $L+1$. This will allow us to compare only the average revenues for each cycle with different lengths not exceeding $L+1$, which is demonstrated in the next lemma.

Lemma 4.10. *Let the sequence $\{x_i\}_{i=1}^{\infty}$ satisfy $x_1 > 0, x_2 > 0$, $x_i = x_{i-1}\beta_2$ for $i \geqslant 3$, and let $s_j = \sum_{i=1}^{j} x_i$.*

(a) There exists $k \in \mathbb{N}^+$ such that $\frac{s_1}{1} \leqslant \frac{s_2}{2} \leqslant \cdots \leqslant \frac{s_k}{k} \geqslant \frac{s_{k+1}}{k+1} \geqslant \cdots$, i.e., the sequence $\{\frac{s_j}{j}\}$ increases when $j = 1, 2, \ldots, k$ and decreases when $j = k, k+1, \ldots$.

(b) Define non-negative and monotonic decreasing sequence $\{\xi_i\}$ as $\xi_0 = +\infty$, $\xi_1 = 1$ and $\xi_l = \max\{0, \frac{(l+1)\beta_2^{l-1} - l\beta_2^{l} - 1}{1-\beta_2}\}$ for $l \geqslant 2$. Then $k = \arg\max_j\{\frac{s_j}{j}\}$ whenever $\frac{x_1}{x_2} \in [\xi_k, \xi_{k-1})$.

Proof. (a) If $\frac{s_{i-1}}{i-1} \geqslant \frac{s_i}{i}$ for some $i > 1$, we must have $s_i \geqslant i(s_i - s_{i-1}) = ix_i$. Hence, $s_{i+1} - s_i = x_{i+1} = x_i\beta_2 \leqslant \frac{s_i}{i}$, i.e., $\frac{s_i}{i} \geqslant \frac{s_{i+1}}{i+1}$. Thus, for the sequence $\{\frac{s_j}{j}\}$, once it decreases from k to $k+1$, it will always decrease afterwards.

(b) First, according to part (a), $\arg\max_j\{\frac{s_j}{j}\} = 1$ if and only if $s_1 \geqslant \frac{s_2}{2}$, which is equivalent to $\frac{x_1}{x_2} \geqslant 1$.

Now assume $\arg\max_j\{\frac{s_j}{j}\} \geqslant 2$. According to part (a), $k = \arg\max_j\{\frac{s_j}{j}\}$ if and only if $\frac{s_k}{k} \geqslant \frac{s_{k+1}}{k+1}$ and $\frac{s_k}{k} > \frac{s_{k-1}}{k-1}$. Since $s_{k+1} = s_k + x_{k+1}$, the two above conditions are equivalent to $s_k \geqslant kx_{k+1}$ and $(k-1)x_k > s_{k-1}$.

For $k \geqslant 2$, $s_k = x_1 + x_2\frac{1-\beta_2^{k-1}}{1-\beta_2}$. Thus, when $\frac{x_1}{x_2} \geqslant \xi_k \geqslant \frac{(k+1)\beta_2^{k-1} - k\beta_2^{k} - 1}{1-\beta_2}$, we have $s_k \geqslant kx_2\beta_2^{k-1} = kx_{k+1}$. Similarly, when $\frac{x_1}{x_2} < \xi_{k-1}$ (since $x_1, x_2 > 0$, it implies $\xi_{k-1} > 0$), we have $s_{k-1} < (k-1)x_k$. Therefore, $k = \arg\max_j\{\frac{s_j}{j}\}$ whenever $\frac{x_1}{x_2} \in [\xi_k, \xi_{k-1})$. □

Proof. (a) The proof is by contradiction. Suppose in an optimal pricing policy $\{p_t\}_{t=1}^{\infty}$, the period $s+1$ is the first period satisfying $n_{s+1} \leqslant C\beta_2/\beta_1$. Then, we must have $p_s = p^2$. Let $s+k$ be the first period after period s in which p^1 is charged (otherwise p^2 is charged forever and the long-run average cost equals to zero). We consider the modified policy $\{\tilde{p}_t\}_{t=1}^{\infty}$ with $\tilde{p}_t = p_t$ for all $t \neq s, s+k$. Let $\tilde{p}_s = p^1$, $\tilde{p}_{s+k} = p^2$. Then, we have

- $n_t = \tilde{n}_t$ for all $t = 1, \ldots, s$;
- $\tilde{n}_{s+1} = \tilde{n}_s\beta_1 = n_s\beta_1 = \beta_1 n_{s+1}/\beta_2 \leqslant C$. Moreover, since $p^{s+1} = \cdots, p^{s+k-1} = p^2$, we conclude that $\frac{\tilde{n}_{s+j}}{n_{s+j}} = \frac{\beta_1}{\beta_2}$ for all $j = 1, \ldots, k$.
- $n_t = \tilde{n}_t$ for all $t > s+k$.

Therefore, by changing the policy in periods s and $s+k$, the total revenue can be increased by at least

$$
\begin{aligned}
& n_s(\mu_1-\mu_2)+\sum_{j=1}^{k} n_s\beta_2^j\mu_2(\frac{\beta_1}{\beta_2}-1)+n_s\beta_2^k(\mu_2-\mu_1) \\
= \;& n_s\Big[(\mu_1-\mu_2)+\mu_2\frac{1-\beta_2^k}{1-\beta_2}(\beta_1-\beta_2)+\beta_2^k(\mu_2-\mu_1)\Big] \\
> \;& n_s\Big[(\mu_1-\mu_2)+\mu_2\frac{1-\beta_2^k}{1-\beta_2}(1-\beta_2)+\beta_2^k(\mu_2-\mu_1)\Big] \\
= \;& n_s\Big[\mu_1-\beta_2^k\mu_1\Big]>0,
\end{aligned}
$$

which contradicts the fact that $\{p_t\}_{t=1}^{\infty}$ is an optimal pricing policy.

(b) If there exists two periods, t and $t+j$, with p^2 charged in both periods and $j<L$, then $n_{t+1}\leqslant C\beta_2$. Thus, $n_{t+j+1}\leqslant n_{t+1}\beta_1^{j-1}\beta_2\leqslant C\beta_2^2\beta_1^{L-1}<C\beta_2\beta^{-1}$. This contradicts part (a). □

The following theorem characterizes the optimal policy.

Theorem 4.11. *Suppose $\beta_1=\beta_2^{-L}$ for some $L\in\mathbb{N}^+$. Define $\eta_k=\frac{1}{\xi_k}+(1-\beta_2)$ for $k=1,2,\ldots,L$ and $\eta_{L+1}=+\infty$, where ξ_k is given in Lemma 4.10. If $\frac{\mu_2}{\mu_1}\in[1,\eta_1)$, then the optimal policy is to always charge low price p^1, with average per-period revenue $C\mu_1$. If $\frac{\mu_2}{\mu_1}\in[\eta_{k-1},\eta_k)$ for some $k\in\{2,\ldots,L+1\}$, then the optimal policy is a cyclic policy $(p^2,\ldots,p^2,p^1)$ with cycle of length k. Moreover, the long-run average revenue is $C\frac{s_k}{k}$, where s_k is defined in Lemma 4.10 by the sequence $\{x_i\}_{i=1}^{\infty}$ with the initial two terms $x_1=\mu_1$, $x_2=(\mu_2-\mu_1)+\mu_1\beta_2$.*

Proof. Consider a cycle of length k as a sequence of $k-1$ high prices p^2 followed by a low price p^1. From Lemma 4.9, we know that the length of each cycle in the optimal pricing policy is at most $L+1$. Moreover, since $\beta_1=\beta_2^{-L}$, the initial number of arrivals for each cycle must be C. Thus, we can compute the average revenue for the cycle of length k by $R(k)=\frac{C}{k}(\mu_2+\mu_2\beta_2+\cdots+\mu_2\beta_2^{k-2}+\mu_1\beta_2^{k-1})=C\frac{s_k}{k}$ and $R(1)=C\mu_1=Cs_1$.

For the infinite-horizon problem, assuming $f_k\in[0,1]$ is the frequency of the cycle of length k occurred in the optimal policy. Then, $\sum_{k=1}^{L+1}f_k=1$ and the average revenue is $z(n)=\sum_{k=1}^{L+1}f_kR(k)\leqslant\max_{k=1,2,\ldots,L+1}R(k)$. This means we only need to compare the average revenue of all cycles with length $1,2,\ldots,L+1$ and choose the single cycle length with the highest average revenue.

Thus, using Lemma 4.10, we have

- For $k = 1, 2, \ldots L+1$, when $\frac{x_1}{x_2} \in [\xi_k, \xi_{k-1})$, $\arg\max_{l=1,2,\ldots,L+1} R(l) = k$.
- When $\frac{x_1}{x_2} \in [0, \xi_{L+1})$, $\arg\max_{l=1,2,\ldots,L+1} R(l) = L+1$.

Since $\frac{x_2}{x_1} = \frac{\mu_2}{\mu_1} - (1-\beta_2)$, using the definition of $\{\eta_k\}_{k=1}^{L+1}$, we combine the above two statements: when $\frac{\mu_2}{\mu_1} \in (\eta_{k-1}, \eta_k]$, a cyclic policy $(p^2, \ldots, p^2, p^1)$ with cycle length k is optimal to the infinite-horizon problem. Consequently, the optimal average revenue per period is $R(k) = C\frac{s_k}{k}$. □

Theorem 4.11 indicates that the use of cyclic policy will maximize the average revenue for the retailer. Choosing the optimal cycle length, the retailer can charge once a lower price when the number of customer arrivals is low to increase the demand volume, and then charge a high price for several subsequent periods to increase revenue when the demands are high. We further observe that the cycle length depends on the ratio $\frac{\mu_2}{\mu_1}$, i.e., the larger the ratio, the longer the cycle.

In summary, when charging a lower price has more impact on ranking than charging a higher price, it is beneficial to choose the higher price for multiple periods before switching to a lower price, since charging a lower price for a single period will be enough to bring up the rank of the product.

4.3.4 Main Results for the General Case

In this section, we present the main results for the single-product model with general parameters β_1 and β_2 that satisfy $\beta_2 < 1 < \beta_1$. We divide our results into two cases: (i) $\beta_1\beta_2 < 1$ and (ii) $\beta_1\beta_2 \geqslant 1$. In the first case, $1 < \beta_1 < \beta_2^{-1}$ implies that the high price will decrease the ranking more than the increase in rank by charging a low price. Conversely, in the second case, a low price has more impact on ranking than a high price since $\beta_1 \geqslant \beta_2^{-1} > 1$.

4.3.4.1 Case 1: $1 < \beta_1 < \beta_2^{-1}$

Lemma 4.12. *Let $L \geqslant 1$ be the integer such that $\beta_1^L < \beta_2^{-1} \leqslant \beta_1^{L+1}$. (a) The number of arrivals in an optimal policy will never drop to or below $C\beta_2/\beta_1$. (b) For each high price p^2 charged in an optimal policy, the following (at least) L consecutive periods must have low price p^1 charged.*

Using Lemma 4.12, we can prove the following theorem, which characterizes the structure of the optimal policy for the infinite-horizon model.

Theorem 4.13. *Suppose $1 < \beta_1 < \beta_2^{-1}$ and let $L \geqslant 1$ be the largest integer satisfying $\beta_1^L < \beta_2^{-1}$. Then, the optimal policy for the infinite-horizon problem* (4.1) *must be a cyclic policy. Specifically, an optimal policy either chooses low price p^1 in each period, or it has a cycle which has the form $(p^2, p^1, \ldots, p^1)$ with length $L+1$ or $L+2$.*

Proof. Let $\{p_t\}_{t=1}^{\infty}$ be an optimal pricing policy. We have the following two cases:

- High price p^2 is charged for a finite number of periods. In this case, the long-run average revenue is equivalent to that of the policy of charging low price p^1 at all periods, which equals $C\mu_1$.

- High price p^2 is charged for an infinite number of periods. In this case, we have an infinite number of cycles that start with p^2 followed by a number of low prices p^1. According to Lemma 4.12(b), a cycle must have at least $L+1$ periods with an initial high price followed by at least L low prices.

  Moreover, since high price p^2 is charged for infinite number of periods in the optimal policy, it must be optimal to charge p^2 whenever the number of arrivals reaches C. Therefore, the optimal cycle cannot contain more than $L+1$ low prices.

□

4.3.4.2 Case 2: $\beta_1 \geqslant \beta_2^{-1} > 1$

Based on Lemma 4.9, we have:

Theorem 4.14. *Suppose $\beta_1 \geqslant \beta_2^{-1} > 1$ and let $L \geqslant 1$ be the smallest integer such that $\beta_1 \leqslant \beta_2^{-L}$. Then, the optimal policy for the infinite-horizon problem* (4.1) *must be a cyclic policy with cycles of fixed length in the form of $(p^2, \ldots, p^2, p^1)$, where the length is at most L, or with cycles $(p^2, \ldots, p^2, p^1)$ of possibly varying lengths, but always L or $L+1$.*

Proof. Let $\{p_t\}_{t=1}^{\infty}$ be an optimal pricing policy. A cycle in the pricing policy starts with a number of high price p^2 and ends with a low price p^1. From Lemma 4.9, we know that the length of any cycle is at most $L+1$.

- If the first cycle has a length less than $L+1$, since $\beta_1 > \beta_2^{-(L-1)}$, we conclude that the number of arrivals after the pricing cycle must reach C again. Hence, by the optimality of this cycle, we conclude that optimal policy must repeat this cycle with the same length forever.

- If the first cycle has length $L+1$, then according to Lemma 4.10, we have $\frac{s_1}{1} \leqslant \cdots \leqslant \frac{s_L}{L} \leqslant \frac{s_{L+1}}{L+1}$. Hence, we conclude that the cycle length must be either L or $L+1$.

$\square$

The profitability of cyclic policy is due to charging high price in periods with high demand. This is consistent with what we have already shown in Theorems 4.13 and 4.14. Below we describe these properties in a finite-horizon setting, as a function of *any* starting number of arrivals. We show that there exists a threshold in terms of the number of arrivals in each period that defined when higher price will be charged and when lower price will be charged.

Theorem 4.15. *Consider a finite-horizon problem. In each period t, there exists a threshold value γ_t such that*

$$p_t^* = \begin{cases} p^1 & \textit{if } n_t \leqslant \gamma_t \\ p^2 & \textit{if } n_t > \gamma_t \end{cases}$$

Proof. We show by induction that if $n_t < \frac{C}{\beta_1}$, $p_t^* = p_1$. Suppose this property holds for $t+1$, then for t, we compare the difference between low price and high price, which is defined by

$$f(n_t) = (\mu_1 - \mu_2)n_t + (v_{t+1}(\beta_1 n_t) - v_{t+1}(\beta_2 n_t))$$

Since $n_t < \frac{C}{\beta_1}$, $\beta_2 n_t < \frac{C}{\beta_1}$, we apply the induction hypothesis and obtain

$$\begin{aligned} f(n_t) &= (\mu_1 - \mu_2)n_t + (v_{t+1}(\beta_1 n_t) - v_{t+1}(\beta_2 n_t)) \\ &= (\mu_1 - \mu_2)n_t + (v_{t+1}(\beta_1 n_t) - \mu_1\beta_2 n_t - v_{t+2}(\beta_1\beta_2 n_t)) \\ &\geqslant (\mu_1 - \mu_2)n_t + (\mu_2\beta_1 - \mu_1\beta_2)n_t \geqslant 0. \end{aligned}$$

Now we show that when $n_t \geqslant \frac{C}{\beta_1}$, the difference $f(n_t)$ is strictly decreasing in n_t. To see this, we can rewrite $f(n_t) = (\mu_1 - \mu_2)n_t + (v_{t+1}(C) - v_{t+1}(\min\{\beta_2 n_t, C\})$ since $n_t \geqslant \frac{C}{\beta_1}$. Since $\mu_1 < \mu_2$ and $v_{t+1}(\min\{\beta_2 n_t, C\})$ is non-decreasing in n_t (see Lemma 4.1), we conclude that $f(n_t)$ is decreasing in n_t.

As $f(n_t)$ is positive for $n_t < \frac{C}{\beta_1}$ and decreasing for $n_t \geqslant \frac{C}{\beta_1}$, hence, it has at most one zero point $\gamma \in (\frac{C}{\beta_1})$ and the theorem is proved.[3] $\square$

Theorem 4.15 shows that the optimal pricing policy, given the number of arrivals (or the current rank of the product), must be a threshold policy. This helps to reduce the computational complexity of the problem by using brute-force dynamic program.

[3] Any value of v_{T+1}, which implies that f is non-decreasing will satisfy the initial step of induction.

4.3.5 Numerical Tests

In this section, we present numerical experiments for the single-product model. All the computational tests are conducted using Python 3.6.0 on a 3.40GHz Intel(R) Xeon(R) CPU. Throughout our numerical tests, the total number of periods $T = 1000$ and the maximum number of customer arrivals is $C = 100$.

We first consider a price set $\mathcal{P} = \{p^1, p^2\}$ with a linear relationship between the number of customer arrivals and the past demands, $n_{t+1} = \min\{C, \alpha d_t\}$. As shown in Section 4.3.4, the structure of optimal policies actually depends on the model parameters, namely (β_1, β_2) and (μ_1, μ_2). In addition to specifying available prices and the purchasing probabilities, we compute other model parameters and report them in Tables 4.1 and 4.2. Note that in the optimal cycle column, h represents a high price and l represents a low price.

First, we present the numerical results for the case when $1 < \beta_1 < \beta_2^{-1}$. The linear factor α is set to be 1.2 in the test, i.e., $n_{t+1} = \min\{100, 1.2d_t\}$.

(p^1, p^2)	$(\bar{F}(p^1), \bar{F}(p^2))$	(β_1, β_2)	$L = \lceil -\frac{\log \beta_2}{\log \beta_1} \rceil$	(μ_1, μ_2)	optimal cycle
$(5, 10)$	$(0.95, 0.5)$	$(1.14, 0.6)$	4	$(4.75, 5)$	(l)
$(5, 15)$	$(0.95, 0.5)$	$(1.14, 0.6)$	4	$(4.75, 7.5)$	(l)
$(5, 20)$	$(0.95, 0.5)$	$(1.14, 0.6)$	4	$(4.75, 10)$	(h, l, l, l, l)
$(5, 10)$	$(0.95, 0.6)$	$(1.14, 0.72)$	3	$(4.75, 6)$	(l)
$(5, 15)$	$(0.95, 0.6)$	$(1.14, 0.72)$	3	$(4.75, 9)$	(h, l, l, l)
$(5, 10)$	$(0.9, 0.75)$	$(1.08, 0.9)$	2	$(4.5, 7.5)$	(h, l) or (h, l, l)
$(5, 15)$	$(0.9, 0.6)$	$(1.08, 0.72)$	5	$(4.5, 9)$	(h, l, l, l, l) or (h, l, l, l, l, l)

Table 4.1: Numerical results for $1 < \beta_1 < \beta_2^{-1}$

From Table 4.1, we observe that either a low price is charged in each period or the optimal pricing policy contains cycles that begins with a high price followed by $L - 1$ or L low prices charged in the subsequent periods. The length of the optimal cycle, if exists, is either L or $L + 1$. Moreover, as the ratio $\frac{\mu_2}{\mu_1}$ increases, we observe that cyclic policy outperforms the policy of always charging the lower price. The numerical tests shown in Table 4.1 are consistent with our theoretical results presented in Section 4.3.4.

Next, we present our numerical results for the case where $\beta_1 \geqslant \beta_2^{-1} > 1$. We use $\alpha = 1.5$. The results summarized in Table 4.2 show that the optimal pricing policy, for the case $\beta_1 \geqslant \beta_2^{-1} > 1$ must be cyclic. The cycle length can be any integer between 1 and $L + 1$, always starting with a few high prices and ending with a low price. Moreover, we also observe that the cycle length increases when the ratio $\frac{\mu_2}{\mu_1}$ increases. This observation is consistent with our theoretical results presented in Section 4.3.4.

$(p^1,\, p^2)$	$(\bar{F}(p^1),\, \bar{F}(p^2))$	(β_1, β_2)	$L = \lfloor -\frac{\log \beta_1}{\log \beta_2} \rfloor$	(μ_1, μ_2)	optimal cycle
$(5, 10)$	$(0.8, 0.6)$	$(1.2, 0.9)$	2	$(4, 6)$	(h, h, l) or (h, l)
$(5, 50)$	$(0.8, 0.6)$	$(1.2, 0.9)$	2	$(4, 30)$	(h, h, l) or (h, l)
$(5, 10)$	$(1, 0.6)$	$(1.5, 0.9)$	4	$(5, 6)$	(h, l)
$(5, 12.5)$	$(1, 0.6)$	$(1.5, 0.9)$	4	$(5, 7.5)$	(h, h, l)
$(5, 20)$	$(1, 0.6)$	$(1.5, 0.9)$	4	$(5, 12)$	(h, h, h, l)
$(5, 50)$	$(1, 0.6)$	$(1.5, 0.9)$	4	$(5, 30)$	(h, h, h, h, l) or (h, h, h, l)

Table 4.2: Numerical results for $\beta_1 \geqslant \beta_2^{-1} > 1$

In addition to the above tests under the assumption given in Section 4.3.2, we also consider a generalized model where price set contains more than two available prices. For customer valuation distribution, we tested both uniform distribution and exponential distribution. We use function $n_{t+1} = \min\{100, \alpha d_t\}$ to describe the relationship between customers arrivals and previous demand. We also consider multiple choices for α. We report our test results in Table 4.3.

Price set $\mathcal{P}$	$\bar{F}(p)$	α	Optimal pricing policy
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.5	$(10, 7.5, 7.5, 7.5)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.52	$(10, 7.5, 7.5)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.55	$(10, 7.5)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.58	$(10, 10, 7.5)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.6	$(10, 10, 10, 7.5)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	1.8	$(12.5, 12.5, 10)$
$\{5, 7.5, 10, 12.5, 15\}$	$e^{-0.05p}$	2	$(15, 12.5)$
$\{4, 6, \ldots, 20\}$	$(20 - p)/16$	1.25	$(8, 6)$
$\{4, 6, \ldots, 20\}$	$(20 - p)/16$	1.3	$(8, 8, 8, 6)$
$\{4, 6, \ldots, 20\}$	$(20 - p)/16$	1.35	(8)
$\{4, 6, \ldots, 20\}$	$(20 - p)/16$	1.55	$(10, 8)$
$\{4, 6, \ldots, 20\}$	$(20 - p)/16$	1.6	(10)

Table 4.3: Numerical results for generalized model

From Table 4.3, we observe that the optimal pricing policy remain a cyclic policy when multiple prices are considered. Also, as the parameter α increases, higher prices are used more frequently. An intuitive explanation is that larger α reduces the required number of periods (with low price) to recover to the maximum number of arrivals. This encourages the retailer to charge higher price for bigger number of consecutive periods, resulting in higher immediate revenue.

To conclude, we have conducted numerical analysis for a single-product under deterministic demand. Our numerical tests consider discrete price set with different demand functions and different parameters indicating customers' sensitivity to the sales rank. Consistent with our technical results presented in Section 4.3.4, our results show that the optimal pricing policy is cyclical and the cycle length not only depends on the expected revenue for each price candidates but also depends on the customers' sensitivity to the sales rank.

4.4 Single-Product Pricing with Stochastic Arrivals

4.4.1 Model Description

In this section, we generalize the multi-period rank-based dynamic pricing model studied in Section 4.3 by considering a continuous price set $\mathcal{P}$ and stochastic demand. At the beginning of each period t, the retailer first observes the sales rank r_t of the product (larger the rank, the lower is the rank of the product). Then, for given rank r_t, the retailer chooses a price p_t from a continuous price set $\mathcal{P} = [\underline{p}, \bar{p}]$, . The number of customer arriving is a random variable N_t with mean $\mathbb{E}[N_t] = n_t$. The mean value n_t is a non-increasing function of the sales rank index in the current period (i.e., $n_t = g(r_t)$). Let $F(\cdot)$ denote the cumulative distribution of customer's valuation of the product and $\lambda(p) = 1 - F(p)$ denote the proportion of arriving customers who will make the purchase. We will refer to $\lambda(p)$ as the purchasing rate. Thus, the total demand in period t is computed by $D_t = \sum_{i=1}^{N_t} \mathbb{1}\{e_i > p_t\}$, where e_i denotes the valuation for the i-th customer. After demand d_t is realized, the sales rank in the next period (i.e., $r_{t+1} = h(d_t)$) is revealed. The retailer's goal is to set a price $p_t \in \mathcal{P}$ in each period t, such that the total expected revenue is maximized.

The classic demand model using customers' reservation price assumes a stationary customer arrival rate (see, e.g., Gallego and Van Ryzin 1994). Our demand model generalizes this assumption since the number of customer arrivals in each period is non-stationary and depends on the sales rank of the current period, implying that all of the pricing decisions made in the past will affect the rate of arriving customers in all future periods. We use $N_t = \epsilon_1 + \epsilon_2 g(r_t)$ to model the uncertainty of customer arrivals where ϵ_1 is a zero-mean random variable and ϵ_2 is a non-negative random variable with mean 1. This stochastic customer arrivals model combines multiplicative and additive uncertainty models that are most commonly used in literature (see, e.g., Talluri and Van Ryzin 2006).

Note that, given the purchasing rate $\lambda \in [0, 1]$, we can compute the corresponding price $p(\lambda) = F^{-1}(1 - \lambda)$. Thus, one can view the purchasing rate λ as the decision variable, where the price is $p(\lambda)$. The expected revenue per customer per period is then $R(\lambda) := \lambda p(\lambda)$.

Then, the expected revenue in period t is given by $\mathbb{E}[N_t](1 - F(p_t))p_t = n_t R(\lambda_t)$. The following assumption on the revenue function $R(\lambda)$ is imposed in the majority of dynamic pricing literature (see, e.g., Gallego and Van Ryzin 1994).

Assumption 4.16. *The revenue function* $R(\lambda)$ *satisfies* $\lim_{\lambda \to 0} R(\lambda) = 0$ *and it is continuous, bounded, and strictly concave with a bounded largest maximizer defined by* $\lambda^* = \max\{\lambda : R(\lambda) = \max_{\lambda \in [0,1]} R(\lambda)\}$

The concavity of the revenue function comes from the economic assumption that marginal revenue is decreasing and it is used in the operations management literature. Ziya et al. (2004) characterized an equivalent condition on the distribution $F(p)$ and we summarize it in Proposition 4.17.

Proposition 4.17. *(Ziya et al. (2004)) Suppose that the reservation-price cumulative distribution function* $F(p)$ *is twice differentiable and strictly increasing on its domain* $\Omega_p = [p_1, p_2]$ *with* $F(p_1) = 0$*,* $F(p_2) = 1$*. Let* $\rho(p) = f(p)/(1 - F(p))$ *denotes the hazard rate where* $f(p)$ *is the density function. Then,* $R(\lambda)$ *is strictly concave in* λ *if and only if* $2\rho(p) > -f'(p)/f(p)$ *for all* $p \in \Omega_p$*.*

To formally formulate the problem, we use value function $v_{t,T}(r_t)$ to denote the maximum expected revenue from period t to T given that, at the beginning of period t, the number of customer arrivals is n_t. Let $\phi(d_t) = g(h(d_t))$ be the function that maps the current demand to the number of arrivals in the next period (n_{t+1}). Similar to (4.2), we express

$$v_{t,T}(n_t) = \max_{\lambda_t \in [0,1]} \left\{ n_t R(\lambda_t) + \mathbb{E}_{\epsilon_1, \epsilon_2}[v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n_t)))] \right\}, \tag{4.8}$$

with boundary condition $v_{T+1,T}(x) = 0$.

Note that function $\phi(\cdot)$ takes current demand as an input and outputs the number of arriving customers in the next period. Intuitively, higher demand implies higher sales rank in the next period, which then translates into higher number of customer arrivals. Furthermore, intuitively, the gains (from higher demand in the current period on demand in the next period) are decreasing. We reflect these intuitive relationships by assuming:

Assumption 4.18. $\phi(\cdot)$ *is a non-decreasing and concave function.*

This assumption holds, in particular, for the special case analyzed in Section 4.3.1, when the number of customer arrivals has capacity C and is linear in the previous-period demand.

4.4.2 Structural Results

In this section, we derive structural results for the rank-based single-product pricing model with stochastic demand described in the previous section. In each period, after observing the sales rank r_t and its induced mean arrivals, $n_t = g(r_t)$, the retailer decides the optimal purchasing probability $\lambda_t^*(n_t)$ and use the inverse of the reference-price distribution function to compute the optimal price $p_t^* = p(\lambda_t^*) = F^{-1}(1 - \lambda_t^*)$. Since the cumulative distribution function $F(\cdot)$ is non-decreasing, it is clear that $p(\lambda_t)$ is a non-increasing function of λ_t. The following lemma shows the monotonicity of the value function, which will be used in later proofs.

Lemma 4.19. *Under Assumptions 4.16 and 4.18, the value function $v_{t,T}(n)$ defined in* (4.8) *is (a) strictly increasing in the number of arrivals $n > 0$ and (b) concave in the number of arrivals $n > 0$.*

Proof. We show both parts of the lemma by induction.
(a) When $t = T$, $v_{T,T}(n) = nR(\lambda^*)$ for all $n > 0$, which is strictly increasing in n. Suppose $v_{t+1,T}(n)$ is a strictly increasing function of n. Then, for any $n_1 > n_2 > 0$ and any $\lambda \in [0, 1]$, we have

$$\begin{aligned} v_{t,T}(n_1) &= \max_{\lambda_t \in [0,1]} \left\{ n_1 R(\lambda_t) + \mathbb{E}[v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n_1)))] \right\} \\ &> \max_{\lambda_t \in [0,1]} \left\{ n_2 R(\lambda_t) + \mathbb{E}[v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n_2)))] \right\} \\ &= v_{t,T}(n_2), \end{aligned}$$

which completes the proof of part (a)
(b) When $t = T$, $v_{t,T}(n) = nR(\lambda_t)$ is linear (hence concave) in n. Now suppose $v_{t+1,T}(n)$ is concave in n.

Since $\phi(\cdot)$ is concave by Assumption 4.18 and $v_{t+1,T}(n)$ is concave and strictly increasing in n, the composite function $v_{t+1,T}(\phi(\cdot))$ is also concave. Thus, for fixed $\lambda_t \in [0, 1]$ and fixed ϵ_1, $\epsilon_2 > 0$, the function $v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n)))$ is concave in n. Taking the expectation with respect to ϵ_1 and ϵ_2, the function $J_{t+1}(\lambda_t, n) = nR(\lambda_t) + \mathbb{E}_{\epsilon_1,\epsilon_2}[v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n)))]$ must also be concave in n. Therefore, using formula (4.8), for any $n_1, n_2 > 0$ and $n_0 =$

$\kappa n_1 + (1-\kappa)n_2$ with any $\kappa \in [0,1]$,

$$\begin{aligned} v_{t,T}(n_0) &= J_{t+1}(\lambda_t^*(n_0), \kappa n_1 + (1-\kappa)n_2) \\ &\leqslant \kappa J_{t+1}(\lambda_t^*(n_0), n_1) + (1-\kappa) J_{t+1}(\lambda_t^*(n_0), n_2) \\ &\leqslant \kappa \max_{\lambda\in[0,1]} \left\{ J_{t+1}(\lambda, n_1) \right\} + (1-\kappa) \max_{\lambda\in[0,1]} \left\{ J_{t+1}(\lambda, n_2) \right\} \\ &= \kappa v_{t,T}(n_1) + (1-\kappa) v_{t,T}(n_2), \end{aligned}$$

which shows the concavity of $v_{t,T}(n)$. □

Our first structural result claims that the optimal purchasing probabilities $\lambda^*(n_t)$ cannot fall below the myopic optimal purchasing probability λ^*, defined in Assumption 4.16.

Proposition 4.20. *Under the Assumption 4.16 and Assumption 4.18, there exists an optimal policy $\{\lambda_t^*(n)\}_{t=1}^T$ such that $\lambda_t^*(n) \geqslant \lambda^*$ for all $n > 0$ and for all $t = 1, 2, \ldots, T$. Consequently, the optimal prices $p_t^*(n) \leqslant p^*$ for all $n > 0$ and for all $t = 1, 2, \ldots, T$.*

Proof. We show the statement by contradiction. Suppose there exists $n > 0$ and $t \in \{1, 2, \ldots, T\}$ such that $\lambda_t^*(n) < \lambda^*$. We will show that we can obtain at least the same revenue with λ^* instead of $\lambda_t^*(n)$.

Since λ^* is the maximizer of single-period revenue function $R(\cdot)$, we have $n_t R(\lambda^*) \geqslant n_t R(\lambda_t^*)$. Moreover, since $\lambda^* > \lambda_t^*$ and both $v_{t+1,T}(\cdot)$ and $\phi(\cdot)$ are non-decreasing functions, we also have $v_{t+1,T}(\phi(\lambda^* n)) \geqslant v_{t+1,T}(\phi(\lambda_t^* n))$ for all $n > 0$. Therefore, by replacing the current policy by $\lambda_t = \lambda^*$, the total expected revenue must satisfy

$$n_t R(\lambda^*) + \mathbb{E}[v_{t+1,T}(\phi(\lambda^*(\epsilon_1 + \epsilon_2 n_t)))] \geqslant n_t R(\lambda_t^*) + \mathbb{E}[v_{t+1,T}(\phi(\lambda_t^*(\epsilon_1 + \epsilon_2 n_t)))],$$

The relationship for prices follows immediately from $p(\lambda) = F^{-1}(1-\lambda)$.

□

Proposition 4.20 can be viewed as a generalization of Proposition 4.2 for the deterministic model. It implies that the optimal pricing policy only focuses on the right-hand-side of single-period revenue curve (as a function of purchasing rate), which is concave and strictly decreasing. Thus, it is sufficient to restrict the purchasing probability λ_t to the region $[\lambda^*, 1]$, i.e., the formulation presented in (4.8) is equivalent to

$$v_{t,T}(n_t) = \max_{\lambda_t \in [\lambda^*, 1]} \left\{ n_t R(\lambda_t) + \mathbb{E}_{\epsilon_1, \epsilon_2}[v_{t+1,T}(\phi(\lambda_t(\epsilon_1 + \epsilon_2 n_t)))] \right\}.$$

Since the function $p(\lambda)$ is non-increasing in λ, we conclude that optimal policy must always choose prices below the myopic optimal price $p^* = F^{-1}(\lambda^*)$. Intuitively speaking, this property holds because any price higher than p^* not only reduces the expected immediate revenue, but also decreases the sales rank in the subsequent periods. Therefore, it would never be optimal to charge a price higher than the myopic optimal price.

While Proposition 4.20 provides a general lower bound for the optimal purchasing probabilities $\lambda_t^*(n)$, the next theorem characterizes the structure of optimal policies by showing the monotonicity of the optimal pricing policies with respect to the sales rank observed by the retailer at the beginning of each period.

Theorem 4.21. *Under Assumptions 4.16 and 4.18, the optimal policy $\lambda_t^*(r)$ must be non-decreasing (equivalently, the optimal price $p_t^*(r)$ must be non-increasing) in the sales rank r.*

Proof. Using the definition of J_t from Lemma 4.19, we observe that $J_{t+1}(\lambda_t, n)$ is also concave in λ_t since $R(\lambda_t)$ is concave. Thus, the optimal solution $\lambda_t^*(n)$ must satisfy the first order condition, i.e.,

$$R'(\lambda_t^*) + \mathbb{E}_{\epsilon_1,\epsilon_2}[(\frac{\epsilon_1}{n} + \epsilon_2) \cdot (v_{t+1,T} \circ \phi)'((\epsilon_1 + \epsilon_2 n)\lambda_t^*))] = 0 \tag{4.9}$$

Suppose there exists $r^1 > r^2 > 0$ such that $\lambda_t^*(r^1) = \lambda_t^1 < \lambda_t^2 = \lambda_t^*(r^2)$. Then, since $g(\cdot)$ is a non-increasing function, we have $n^1 = g(r^1) \leqslant g(r^2) = n^2$. We will use the first order condition (4.9) for (n^1, λ_t^1) and (n^2, λ_t^2). Because the composite function $v_{t+1,T} \circ \phi(\cdot)$ is concave (has a decreasing derivative value) and non-decreasing, the inequality

$$(v_{t+1,T} \circ \phi)'(\lambda_t^1(\epsilon_1 + \epsilon_2 n^1)) \geqslant (v_{t+1,T} \circ \phi)'(\lambda_t^2(\epsilon_1 + \epsilon_2 n^2)) \geqslant 0$$

holds for any ϵ_1 and $\epsilon_2 > 0$. Since $R(\lambda)$ is a strictly concave function, $R'(\lambda)$ must be a strictly decreasing function and hence, $R'(\lambda_t^1) > R'(\lambda_t^2)$. Recall that (n^1, λ_t^1) and (n^2, λ_t^2) satisfy (4.9). By taking expectations and noting that $\frac{1}{n^1} \geqslant \frac{1}{n^2}$, we have a contradiction:

$$\begin{aligned} 0 &= R'(\lambda_t^1) + \mathbb{E}_{\epsilon_1,\epsilon_2}[(\frac{\epsilon_1}{n^1} + \epsilon_2) \cdot (v_{t+1,T} \circ \phi)'((\epsilon_1 + \epsilon_2 n^1)\lambda_t^1))] \\ &> R'(\lambda_t^2) + \mathbb{E}_{\epsilon_1,\epsilon_2}[(\frac{\epsilon_1}{n^2} + \epsilon_2) \cdot (v_{t+1,T} \circ \phi)'((\epsilon_1 + \epsilon_2 n^2)\lambda_t^2))] = 0. \end{aligned}$$

Therefore, for any sales rank $r^1 > r^2$, we must have $\lambda_t^*(r^1) \geqslant \lambda_t^*(r^2)$, which proves that $\lambda_t^*(r)$ is a non-decreasing function in r. As $p(\lambda)$ is a non-decreasing function, we conclude $p_t^*(r)$ is a non-increasing function in r. □

Theorem 4.21 can be viewed as a generalization of Theorem 4.15 which considers a deterministic demand with a two-point price set. This theorem shows that a better sales rank always leads to a higher optimal price. This is intuitive because by raising prices when their products are more visible to customers, retailers can maintain a high demand volume and gain more revenue. Conversely, when the sales rank index is large and the product lacks exposure to the market, dropping prices would be desirable to bring up the sales rank and gain more popularity. This structural result provides a clear direction for how to adjust the retail prices, when observing the sales rank of the products.

It is interesting to compare our structural results (Proposition 4.20 and Theorem 4.21) with those studied by Gallego and Van Ryzin (1994) (Proposition 1 and Theorem 1 in Gallego and Van Ryzin (1994)), in which pricing decisions are made based on the volume of remaining inventory. First, they showed that the optimal prices are lower bounded by the myopic optimal one because the scarcity of the inventory motivates retailer to increase their prices, knowing that not all demand can be satisfied. As a comparison, in our rank-based settings, each additional demand may potentially boost the sales rank in subsequent periods. Hence, there is benefit to decrease the price below the myopic one in order to increase future demand. Second, they showed that optimal price drops when the in-stock inventory is higher. In contrast, with the effect of sales rank, our results show that optimal price raises when the product has better sales rank and more market exposure.

4.4.3 Computational Studies

In this section, we present numerical results we for the stochastic single-product model. All the computational tests are conducted using Python 3.6.0 on a 3.40GHz Intel(R) Xeon(R) CPU.

For the demand function, we consider the following three settings:

1. Linear demand. The reservation-price is uniformly distributed on $[p^l, p^u]$ and $F(p) = \frac{p-p^l}{p^u-p^l}$. Then, given the purchasing probability λ, the corresponding price $p(\lambda) = F^{-1}(1-\lambda) = p^u - (p^u - p^l)\lambda$. Hence, the expected revenue per customer is $R(\lambda) = (p^u - (p^u - p^l)\lambda)\lambda$ with unique maximizer $\lambda^* = \min\{1, \frac{p^u}{2(p^u-p^l)}\}$ and $p^* = \max\{p^l, p^u/2\}$.

2. Log-linear demand. When the reservation-price is exponentially distributed, $F(p) = 1 - e^{-\nu p}$, with $\nu > 0$. The inverse price function is $p(\lambda) = -\frac{\ln \lambda}{\nu}$ and the expected revenue per customer is $R(\lambda) = -\lambda \frac{\ln \lambda}{\nu}$. For this revenue function, the myopic optimal solution is $\lambda^* = e^{-1}$ and $p^* = \frac{1}{\nu}$.

3. Logit demand. This demand model is based on multinomial-logit (MNL). For single-product model, the purchasing probability is $1 - F(p) = \frac{e^{-bp}}{\tau + e^{-bp}}$, where parameter b denotes the price sensitivity. Under this model, the price can be computed as $p(\lambda) = \frac{1}{b}\ln(\frac{1-\lambda}{\tau\lambda})$ and the expected revenue per customer is $R(\lambda) = \frac{\lambda}{b}\ln(\frac{1-\lambda}{\tau\lambda})$. The myopic optimal purchasing probability in this case is $\lambda^* \in [0,1]$, which solves $(1-\lambda^*)\ln\left(\frac{1-\lambda^*}{\tau\lambda^*}\right) = 1$ and the optimal price is $p^* = \frac{1}{b}\ln(\frac{1-\lambda^*}{\tau\lambda^*})$.

To have closely comparable range of prices choices, we appropriately parametrize all three demand functions. For linear demand, we use $p^l = 2$ and $p^u = 10$, for log-linear demand, we use $\nu = 0.5$, and for logit demand, we use $b = 0.25$ and $\tau = 0.1$.

To obtain the number of arrival customers in the next period, we use a concave function $\phi(d_t) = \min\{C, \alpha_1 d_t^{\alpha_2}\}$ where $\alpha_2 \in (0,1]$, $\alpha_1 > 0$ and C is the upper bound on the number of arrivals. We use $C = 100$, $\alpha_1 = 3.5$, and $\alpha_2 = 0.8$.

For stochastic customer arrivals N_t, we consider both Poisson distribution with mean $n_t = g(r_t)$ and binomial distribution with mean $n_t = g(r_t)$ and the total number of trials C.

4.4.3.1 Impact of Sales Rank and Time Horizon

To test the impact of time horizon on the optimal pricing policies, we vary the total number of periods T from 1 to 100 and compare both optimal policies and optimal average expected revenue per-period. We draw the optimal price curve as a function of initial number of arrivals, for different time horizons, $T = 1, 2, 3, 100$. Also, we compare the average revenue curve for optimal policy versus myopic policy, as a function of number of periods. For the last comparison, the initial number of arrivals is the maximum possible $n_0 = C = 100$. We present our numerical test results in Figures 4.3–4.8 and summarize our observations below.

1. For each demand type and customer arrivals distribution, we observe that the price is non-decreasing in the number of customer arrivals. With low customer arrivals, the retailer has an incentive to reduce price to increase customer arrivals in the following periods. On the other hand, a better sales rank already attracts more customers and allows retailer to increase the price (although the price will still be lower than myopic). This observation is consistent with Proposition 4.20 and Theorem 4.21 in Section 4.4.2.

2. The fewer time-periods remaining, the higher the price. An intuitive explanation is that retailers are more worried about their reputation when they need to stay in the market for a longer time. Hence, the longer the time horizon, the more important it is to improve products' popularity in the market.

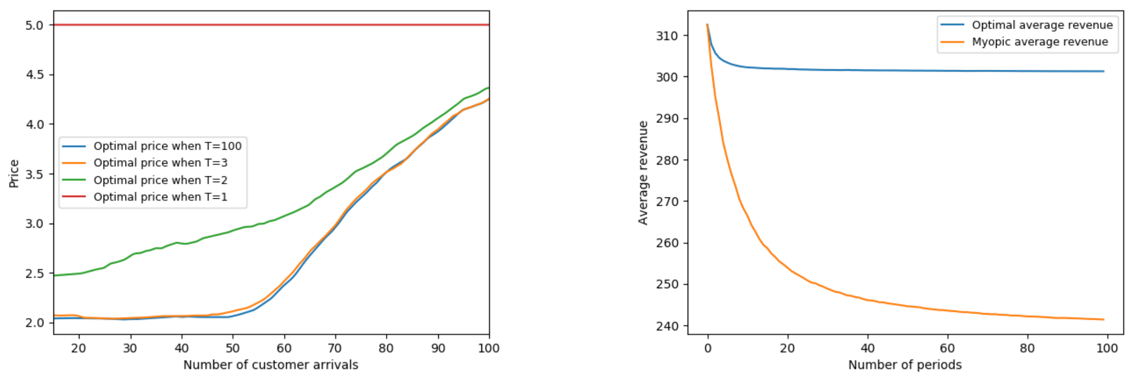

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.3: Poisson arrivals and linear demand with $\bar{F}(p) = \frac{10-p}{8}$

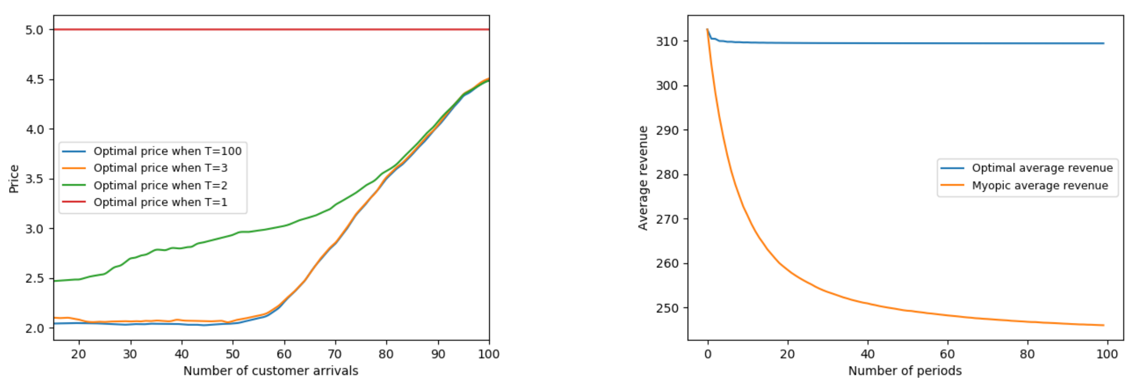

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.4: Binomial arrivals and linear demand with $\bar{F}(p) = \frac{10-p}{8}$

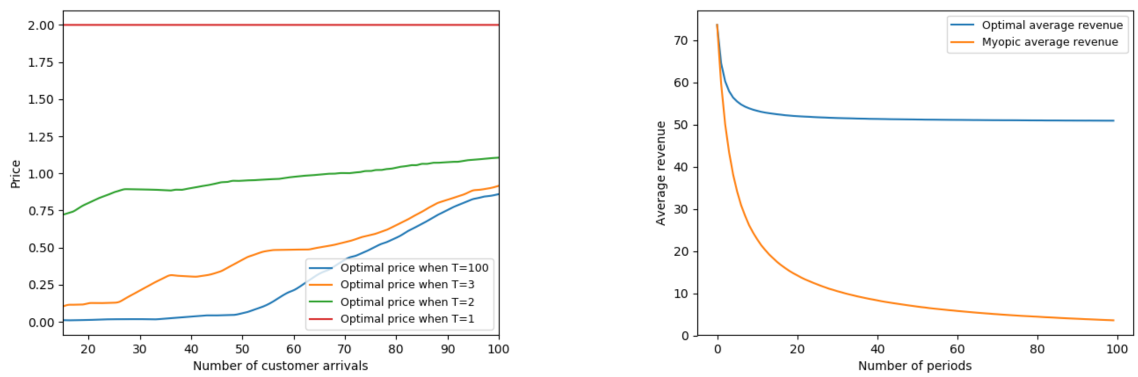

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.5: Poisson arrivals and loglinear demand with $\bar{F}(p) = e^{-0.5p}$

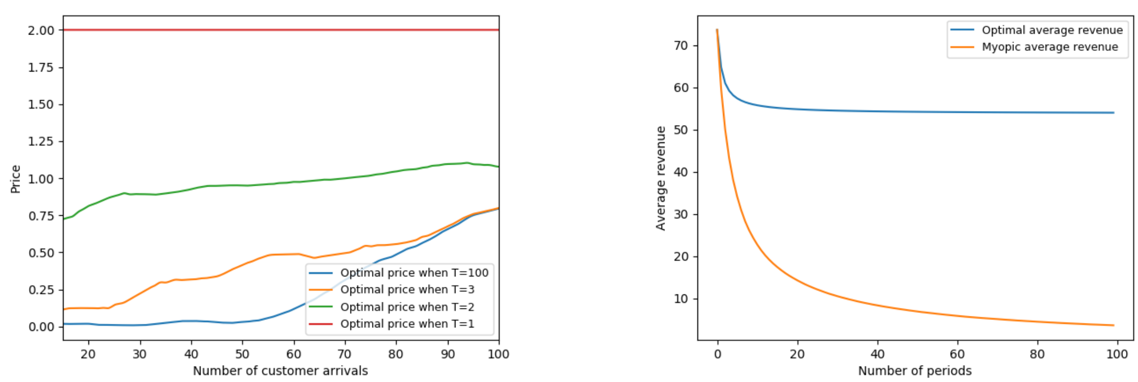

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.6: Binomial arrivals and loglinear demand with $\bar{F}(p) = e^{-0.5p}$

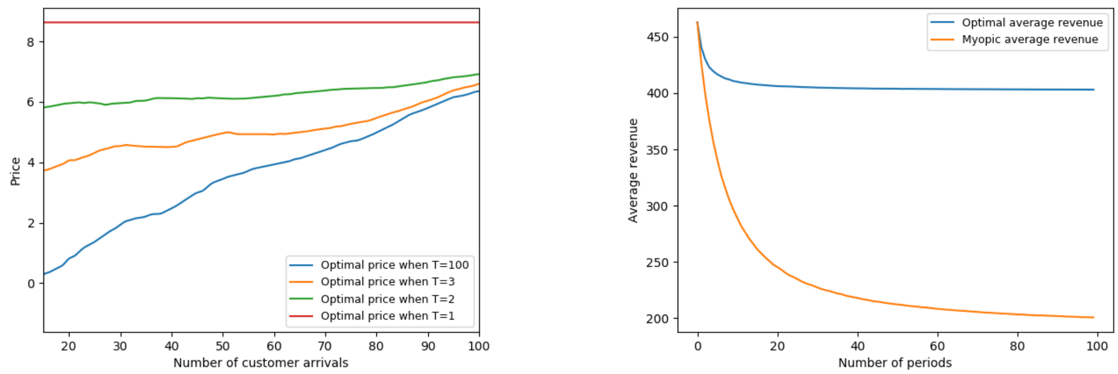

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.7: Poisson arrivals and logit demand with $\bar{F}(p) = \frac{e^{-0.25p}}{0.1+e^{-0.25p}}$

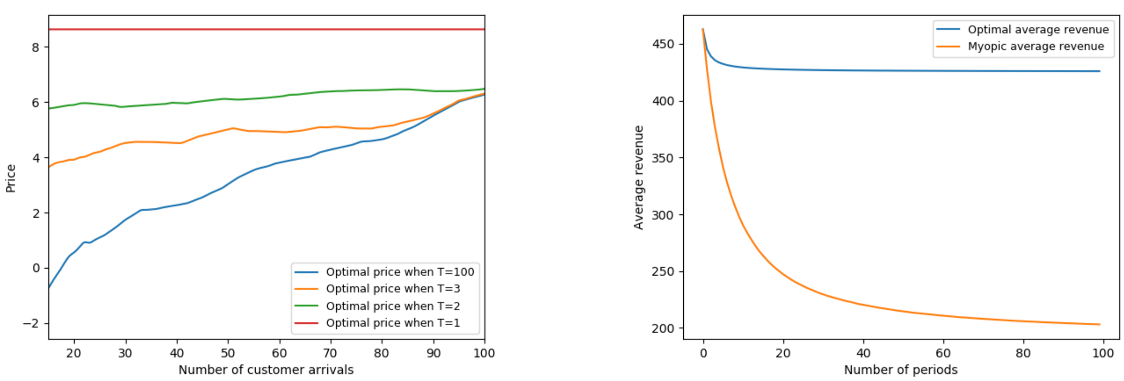

(a) Optimal pricing policy (b) Expected revenue per period

Figure 4.8: Binomial arrivals and logit demand with $\bar{F}(p) = \frac{e^{-0.25p}}{0.1+e^{-0.25p}}$

3. When looking at Figures 4.3–4.8, we observe that the myopic price and optimal price for long horizons are significantly different. Similarly, myopic revenue per period and optimal revenue per period are very different for long horizons. The revenue for short horizons is obviously highly influenced by the number of initial arrivals, but this dependency disappears for longer horizons.

4. Surprisingly, for some cases (e.g., linear demand), a relatively short time horizon ($T = 3$) is sufficient for the pricing policy to converge to the optimal pricing policy for very long time horizons (say, $T = 100$). This can be seen by comparing Figures 4.3–4.4 with Figures 4.5–4.8. When the range of potential revenues is relatively small (Figures 4.3–4.4), it is worthwhile to aggressively move towards the best operating policy. However, when the effect of current pricing on revenue is high (Figures 4.5–4.8), it is not worthwhile to "sacrifice" the whole period's revenue for the speed of convergence.

5. We also observe that the pricing policy seems to have two types of behavior: In Figures 4.3–4.4, we have flat pricing of approximately $p = 2$ for $n < 50$ and, then, price (approximately) linearly increasing in n. This can be intuitively explained as a smooth trade-off between price and number of customer arrivals. The fewer arrivals implies more aggressive pricing policy. However, when we approach the minimal price of $p = 2$, we cannot accelerate that process further, as explained in the point above. Similarly, in Figures 4.5–4.6, price $p = 0$ allows us to reach the maximum purchasing rate $\lambda = 1$, this establishing a threshold below which pricing cannot be used to increase the number of customer arrivals. Note that a different observation is drawn from Figures 4.7–4.8, where negative price may be used as an incentive and which rapidly increases customer arrivals.

   Note that the average revenue as a function of time horizon is decreasing. This is because we initiate the study with the most optimistic case, where number of arrivals is $C = 100$.

4.5 Concluding Remarks

In this chapter, we investigated the dynamic pricing problem with the effect of sales rank. Specifically, we showed the optimality of cyclic pricing policy for the model with deterministic customer arrivals and discrete price set. For stochastic customer arrivals with continuous price set, we derived an upper bound for the optimal pricing policies and showed the monotonicity of the optimal pricing policies. Based on both theoretical and numerical results, we found that the optimal pricing policy tends to increase prices when better sales rank is observed. Our technical results provide guidelines for online retailers to dynamically adjust their pricing policies when sales rank is observed.

One of the future research directions is to study a dynamic pricing problem with both sales rank information and inventory considerations. It would be interesting to characterize the optimal pricing policies as well as providing efficient heuristics that can achieve near-optimal performance. Another direction is to study pricing algorithms based on demand learning where the demand function is unknown and based on features of available products.

CHAPTER V

Conclusions

This dissertation focuses on the emerging features of the e-commerce supply chain management and revenue management. Due to the significance of service quality in today's customer-driven business environment, we studied service-level constrained inventory control system in the first two essays. Such inventory systems can help retailers reduce stockouts probability to satisfy customers' demand and maintain a good reputation. The first essay focused on a static multi-period planning problem and formulated data-driven models to optimize the best inventory decisions with minimum cost and the best pricing decisions with maximum revenue. The second essay focused on developing efficient algorithms for dynamic inventory control problems in which inventory decisions can be altered in real time after observing the realized demand. In addition to the traditional backlogging model, the essay also studied remanufacturing inventory system with stochastic returned products. The third essay considered the dynamic pricing problem under the sales rank information. By studying several mathematical formulations, we found out the impact of ranking information on the optimal pricing policy and on the total revenue. The three essays presented in previous chapters provided methods and insights in today's growing online retail marketplace.

There are several directions for future studies. First, due to the availability of past sales data and customer's browsing data, demand learning and prediction can be important in optimizing inventory and pricing decisions. Thus, an interesting direction for future research is to study non-parametric data-driven models by providing efficient learning algorithms and deriving asymptotic performance bounds. Second, since retailers usually need to make inventory decisions together with their pricing decisions, it is worthwhile to investigate a dynamic joint pricing and inventory problem under sales rank information.

BIBLIOGRAPHY

BIBLIOGRAPHY

Absi, N., S. Kedad-Sidhoum, S. Dauzère-Pérès. 2011. Uncapacitated lot-sizing problem with production time windows, early productions, backlogs and lost sales. *International Journal of Production Research* **49**(9) 2551–2566.

Ahn, H.-S., M. Gümü, P. Kaminsky. 2007. Pricing and manufacturing decisions when demand is a function of prices in multiple periods. *Operations Research* **55**(6) 1039–1057.

Anderson, E. T., G. J. Fitzsimons, D. Simester. 2006. Measuring and mitigating the costs of stockouts. *Management Science* **52**(11) 1751–1763.

ARI. 2015. Lead times, downtime and production delays. `http://www.arifleet.com/publications/lead_times_delays/#item1648`. Accessed August 2015.

Aviv, Y., Y. Levin, M. Nediak. 2009. Counteracting strategic consumer behavior in dynamic pricing systems. *Consumer-Driven Demand and Operations Management Models*. Springer, New York, 323–352.

Bass, F. M. 1969. A new product growth for model consumer durables. *Management Science* **15**(5) 215–227.

Bertsimas, D., I. C. Paschalidis. 2001. Probabilistic service level guarantees in make-to-stock manufacturing systems. *Operations Research* **49**(1) 119–133.

Besbes, O., I. Lobel. 2015. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science* **61**(1) 92–110.

Blogger, Rack. 2016. No time like the present: Solve seasonal ecommerce demand challenges now. `https://blog.rackspace.com/solve-seasonal-demand-website-now`. Accessed Jan, 2018.

Bookbinder, J. H., J.-Y. Tan. 1988. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science* **34**(9) 1096–1108.

Boyaci, T., G. Gallego. 2001. Serial production/distribution systems under service constraints. *Manufacturing & Service Operations Management* **3**(1) 43–50.

Brynjolfsson, E., Y. Hu, D. Simester. 2011. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science* **57**(8) 1373–1386.

Brynjolfsson, E., Y. Hu, M. D. Smith. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science* **49**(11) 1580–1596.

Brynjolfsson, E., Y. Hu, M. D. Smith. 2010. Research commentarylong tails vs. superstars: The effect of information technology on product variety and sales concentration patterns. *Information Systems Research* **21**(4) 736–747.

Brynjolfsson, E., M. D. Smith. 2000. Frictionless commerce? a comparison of internet and conventional retailers. *Management Science* **46**(4) 563–585.

CEMEX. 2015. Annual report 2014. `http://www.cemex.com/CEMEX_AR2014/pdf/download/00_CX14OK.pdf`. Accessed August 2015.

Chao, X., X. Gong, C. Shi, C. Yang, H. Zhang, S. X. Zhou. 2017. Approximation algorithms for capacitated perishable inventory systems with positive lead times. Forthcoming at Management Science.

Chao, X., X. Gong, C. Shi, H. Zhang. 2015. Approximation algorithms for perishable inventory systems. *Operations Research* **63**(3) 585–601.

Chen, F., J.-S. Song. 2001. Optimal policies for multiechelon inventory problems with markov-modulated demand. *Operations Research* **49**(2) 226–234.

Chen, F. Y., D. Krass. 2001. Inventory models with minimal service level constraints. *European Journal of Operational Research* **134**(1) 120–140.

Chen, X., P. Hu, S. He. 2013. Technical notepreservation of supermodularity in parametric optimization problems with nonlattice structures. *Operations Research* **61**(5) 1166–1173.

Chen, X., D. Simchi-Levi. 2012. *Pricing and inventory management, R. Philips and O. Ozer, eds.*. Oxford University Press, Oxford.

Chen, Y., V. F. Farias. 2018. Robust dynamic pricing with strategic customers. Forthcoming in Mathematics of Operations Research.

Chen, Y., V. F. Farias, N. K. Trichakis. 2018. On the efficacy of static prices for revenue management in the face of strategic customers. Working Paper, MIT.

Chen, Y., Q. Wang, J. Xie. 2011. Online social interactions: A natural experiment on word of mouth versus observational learning. *Journal of Marketing Research* **48**(2) 238–254.

Cheung, M., A. Elmachtoub, R. Levi, D. B. Shmoys. 2016. The submodular joint replenishment problem. *Mathematical Programming* **158**(1-2) 207–233.

Chevalier, J. A., D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* **43**(3) 345–354.

Cohen, M., I. Lobel, R. Paes-Leme. 2016. Feature-based dynamic pricing. Working Paper, New York University.

Cohen, M. A., N. Agrawal, V. Agrawal. 2006. Winning in the aftermarket. *Harvard Business Review* **84**(5) 129–138.

Conlisk, John, Eitan Gerstner, Joel Sobel. 1984. Cyclic pricing by a durable goods monopolist. *The Quarterly Journal of Economics* **99**(3) 489–505.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* **49**(10) 1287–1309.

Fine, C. H., J. C. Lafrance, D. Hillebrand. 1996. The us automobile manufacturing industry. *Meeting the Challenge: US Industry Faces the 21st Century*. Citeseer, 35.

Fitzsimons, G. J. 2000. Consumer response to stockouts. *Journal of Consumer Research* **27**(2) 249–266.

Fleder, D., K. Hosanagar. 2009. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management Science* **55**(5) 697–712.

Florian, M., J. K. Lenstra, A. H. G. Rinnooy Kan. 1980. Deterministic production planning: Algorithms and complexity. *Management Science* **26**(7) 669–679.

Gallego, G., Ö. Özer. 2001. Integrating replenishment decisions with advance demand information. *Management Science* **47**(10) 1344–1360.

Gallego, G., G. J. Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40**(8) 999–1020.

Gallego, G., G. J. Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* **45**(1) 24–41.

Gallo, A. 2014. The value of keeping the right customers. `https://hbr.org/2014/10/the-value-of-keeping-the-right-customers`. Accessed Jan, 2018.

Gallo, C. 2012. *The Apple Experience: Secrets to Building Insanely Great Customer Loyalty (ENHANCED EBOOK)*. McGraw Hill Professional.

Ghiani, G., G. Laporte, R. Musmanno. 2005a. Introducing logistics systems. *Introduction to Logistics Systems Planning and Control* 1–23.

Ghiani, G., G. Laporte, R. Musmanno. 2005b. *Introducing Logistics Systems*. John Wiley & Sons, Ltd.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*, vol. 53. Springer.

Goh, J., M. Sim. 2011. Robust optimization made easy with rome. *Operations Research* **59**(4) 973–985.

Gong, X., X. Chao. 2013. Technical note-optimal control policy for capacitated inventory systems with remanufacturing. *Operations Research* **61**(3) 603–611.

González-Ramírez, R. G., N. R. Smith, R. G. Askin. 2011. A heuristic approach for a multi-product capacitated lot-sizing problem with pricing. *International Journal of Production Research* **49**(4) 1173–1196.

Goolsbee, A., J. Chevalier. 2002. Measuring prices and price competition online: Amazon and barnes and noble. Tech. rep., National Bureau of Economic Research.

Graves, S., H. Meal, S. Dasu, Y. Qin. 1986. Two-stage production planning in a dynamic environment. *S. Axsater, C. Schneeweiss, E. Silver, eds. Multi-Stage Production Planning and Control. Lecture Notes in Economics and Mathematical Systems. Springer-Verlag, Berlin, Germany* 9–43.

Halman, N., D. Klabjan, M. Mostagir, J. Orlin, D. Simichi-Levi. 2009. A fully polynomial time approximation scheme for single-item stochastic lot-sizing problems with discrete demand. *Mathematics of Operations Research* **34**(3) 674–685.

Hamilton, J. D. 1994. *Time series analysis*, vol. 2. Princeton University Press.

Heath, D. C., P. L. Jackson. 1994a. Modeling the evolution of demand forecasts ith application to safety stock analysis in production/distribution systems. *IIE Transactions* **26**(3) 17–30.

Heath, D. C., P. L. Jackson. 1994b. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution system. *IIE Transactions* **26**(3) 17–30.

Ho-Dac, N. N., S. J. Carson, W. L. Moore. 2013. The effects of positive and negative online customer reviews: do brand strength and category maturity matter? *Journal of Marketing* **77**(6) 37–53.

Hsu, T. 2017. Alibabas singles day sales hit new record of 25.3 billion. `https://www.nytimes.com/2017/11/10/business/alibaba-singles-day.html`. Accessed Jan, 2018.

Huang, K., S. Ahmed. 2005. The value of multi-stage stochastic programming in capacity planning under uncertainty. *Stochastic Programming E-Print Series* **15**.

Huang, K., S. Ahmed. 2008. On a multi-stage stochastic programming model for inventory planning. *INFOR: Information Systems and Operational Research* **46**(3) 155–163. doi:10.3138/infor.46.3.155.

Jing, X., M. Lewis. 2011. Stockouts in online retailing. *Journal of Marketing Research* **48**(2) 342–354.

Kannan, P. K. 2017. Digital marketing: A framework, review and research agenda. *International Journal of Research in Marketing* **34**(1) 22–45.

Katok, E., D. Thomas, A. Davis. 2008. Inventory service-level agreements as coordination mechanisms: The effect of review periods. *Manufacturing & Service Operations Management* **10**(4) 609–624.

Kleywegt, A. J., A. Shapiro, T. Homem-de Mello. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* **12**(2) 479–502.

Levi, R., G. Janakiraman, M. Nagarajan. 2008a. A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research* **33**(2) 351–374.

Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(2) 284–302.

Levi, R., R. O. Roundy, D. B. Shmoys, V. A. Truong. 2008b. Approximation algorithms for capacitated stochastic inventory control models. *Operations Research* **56**(5) 1184–1199.

Levi, R., C. Shi. 2013. Approximation algorithms for the stochastic lot-sizing problem with order lead times. *Operations Research* **61**(3) 593–602.

Liu, Y., W. L. Cooper. 2015. Optimal dynamic pricing with patient customers. *Operations Research* **63**(6) 1307–1319.

Lobel, I. 2016. Dynamic pricing with heterogeneous patience levels. Working Paper, New York University.

Luedtke, J., S. Ahmed. 2008. A sample approximation approach for optimization with probabilistic constraints. *SIAM Journal on Optimization* **19**(2) 674–699.

Mancilla, C., R. Storer. 2012. A sample average approximation approach to stochastic appointment sequencing and scheduling. *IIE Transactions* **44**(8) 655–670.

McCormick, G.P. 1976. Computability of global solutions to factorable nonconvex programs: Part i–convex underestimating problems. *Mathematical Programming* **10**(1) 147–175.

Mills, E. S. 1959. Uncertainty and price theory. *The Quarterly Journal of Economics* 116–130.

Mills, T. C. 1991. *Time Series Techniques for Economists*. Cambridge University Press.

Mudambi, S. M., D. Schuff. 2010. Research note: What makes a helpful online review? a study of customer reviews on amazon. com. *MIS Quarterly* 185–200.

Nagarajan, V., C. Shi. 2016. Approximation algorithms for inventory problems with submodular or routing costs. *Mathematical Programming* **160**(1) 225–244.

Nahmias, S. 1993. *Production and operations analysis*. Irwin, Homewood, IL.

Nielsen. 2010. Online shopping trends. *USA: The Nielsen Company* .

Noe, T., G. Parker. 2005. Winner take all: Competition, strategy, and the structure of returns in the internet economy. *Journal of Economics & Management Strategy* **14**(1) 141–164.

Özer, Ö., R. Phillips. 2012. *The Oxford handbook of pricing management*. Oxford University Press, Oxford, UK.

Pagnoncelli, B. K., S. Ahmed, A. Shapiro. 2009. Sample average approximation method for chance constrained programming: theory and applications. *Journal of Optimization Theory and Applications* **142**(2) 399–416.

Pan, X., B. T. Ratchford, V. Shankar. 2004. Price dispersion on the internet: A review and directions for future research. *Journal of Interactive Marketing* **18**(4) 116–135.

Pochet, Y., L. A. Wolsey. 1988. Lot-size models with backlogging: Strong reformulations and cutting planes. *Mathematical Programming* **40**(1) 317–335.

Qiang, S., M. Bayati. 2016. Dynamic pricing with demand covariates. Working Paper, Stanford University.

Remy, S., L. Guerard, R. Dirlea. 2010. The old psychology of internet pricing. `http://www.atkearney.com.au/documents/10192/522404/The_Old_Psychology_of_Internet_Pricing.pdf/df0b3f95-f9e9-4a71-98e7-07e5f99f2c7a`. Accessed Jan, 2018.

Robert, K. 1998. The net: A market too perfect for profits. *BusinessWeek* .

Rosen, S. 1981. The economics of superstars. *The American Economic Review* **71**(5) 845–858.

Sethi, S. P., F. Cheng. 1997. Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research* **45**(6) 931–939.

Shang, K. H., J.-S. Song. 2006. A closed-form approximation for serial inventory systems and its application to system design. *Manufacturing & Service Operations Management* **8**(4) 394–406.

Shen, Z.-J. M., X. Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production and Operations Management* **16**(6) 713–728.

Simchi-Levi, D., X. Chen, J. Bramel. 2014. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*. Springer Series in Operations Research and Financial Engineering, Springer New York.

Snyder, L. V., Z.-J. M. Shen. 2011. *Fundamentals of supply chain theory*. John Wiley & Sons.

Talluri, K. T., G. J. Van Ryzin. 2006. *The Theory and Practice of Revenue Management*, vol. 68. Springer Science & Business Media.

Tao, Z., S. X. Zhou. 2014. Approximation balancing policies for inventory systems with remanufacturing. *Mathematics of Operations Research* **39**(4) 1179–1197.

Thomas, J. 1970. Price-production decisions with deterministic demand. *Management Science* **16**(11) 747–750.

Truong, V.-A. 2014. Approximation algorithm for the stochastic multiperiod inventory problem via a look-ahead optimization approach. *Mathematics of Operations Research* **39**(4) 1039–1056.

Tucker, C., J. Zhang. 2011. How does popularity information affect choices? a field experiment. *Management Science* **57**(5) 828–842.

Verweij, B., S. Ahmed, A. J. Kleywegt, G. Nemhauser, A. Shapiro. 2003. The sample average approximation method applied to stochastic routing problems: a computational study. *Computational Optimization and Applications* **24**(2-3) 289–333.

Wei, L., S. Jasin, L. Xin. 2017. On the deterministic approximation of lost-sales inventory system with sequential probabilistic service level constraint. Working Paper, University of Michigan, Ann Arbor, MI, USA.

Yoo, B., K. Kim. 2012. Does popularity decide rankings or do rankings decide popularity? an investigation of ranking mechanism design. *Electronic Commerce Research and Applications* **11**(2) 180–191.

Zangwill, W. I. 1969. A backlogging model and a multi-echelon model of a dynamic economic lot size production system – a network approach. *Management Science* **15**(9) 506–527.

Zhang, H., S. Jasin. 2018. Online learning and optimization of (some) cyclic pricing policies for revenue management with patient customers. Working Paper, University of Michigan.

Zhang, H., C. Shi, X. Chao. 2016. Approximation algorithms for perishable inventory systems with setup costs. *Operations Research* **64**(2) 432–440.

Zhang, M., S. Küçükyavuz, S. Goel. 2014. A branch-and-cut method for dynamic decision making under joint chance constraints. *Management Science* **60**(5) 1317–1333.

Zhou, S. X., Z. Tao, X. Chao. 2011. Optimal control of inventory systems with multiple types of remanufacturable products. *Manufacturing & Service Operations Management* **13**(1) 20–34.

Zipkin, P. 2000. *Foundations of Inventory Management*, vol. 2. McGraw-Hill, New York.

Ziya, S., H. Ayhan, R. D. Foley. 2004. Relationships among three assumptions in revenue management. *Operations Research* **52**(5) 804–809.