CITY UNIVERSITY OF HONG KONG
香港城市大學

Essays on Operations Management: Information Disclosure and Inventory Control
運營管理：信息披露及庫存控制

Submitted to

Department of Management Sciences
管理科學系

in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy
哲學博士學位

by

ZHU Han
朱晗

July 2017
二零一七年七月

Essays on Operations Management: Information Disclosure and Inventory Control

by ZHU Han

Department of Management Sciences

College of Business

City University of Hong Kong

Abstract

This thesis consists of three essays. The first essay is on quality disclosure for experience goods with customer bounded rationality. Deciding whether to disclose quality information is of strategic importance for firms. For experience goods, in practice, customers may not take a firm's nondisclosure strategy as a "signal" to infer low quality. Under a nondisclosure strategy, customers tend to rely on the experiences of those who previously bought the product or experienced the service, and then deduce quality information based on a limited number of these samples (dubbed as "customer bounded rationality"). It remains unclear how customer bounded rationality affects a firm's quality disclosure decision. We build a behavioral model to study firm incentives to disclose quality information of experience goods under customer bounded rationality in the sense of anecdotal reasoning. We find that a firm with a high or low quality level prefers not to disclose information on this quality, whereas a firm with a medium quality level prefers to do so. Our findings are consistent with some recent empirical evidences and provide a new explanation for the incomplete voluntary disclosure observed in many markets, particularly those for experience goods. Ignoring customer bounded rationality can lead to a significant profit loss. When there is congestion in the service context, the demand rate also plays a critical role. We

also provide the managerial implications of our findings.

The second essay establishes performance bounds on the minimum cost of a classic one-warehouse multi-retailer distribution system, in which any inventory replenishment at each location incurs a fixed-plus-variable cost and takes a constant lead time. The optimal policy is unknown and even if it exists, must be extremely complicated. The goal of this essay is to identify an easy-to-compute heuristic policy within the class of modified echelon (r, Q) policies that does not require an integer-ratio property or a synchronized, nested ordering property, yet has certain performance bounds. We first develop a cost upper bound for any given modified echelon (r, Q) policy. Computation of the bound does not require an exact evaluation of the system-wide cost, which is notoriously difficult. We next adopt parameters of the heuristic by minimizing the cost upper bound, which is equivalent to solving a set of independent single-stage (r, Q) systems. With a cost lower bound that has been established in the literature, we then develop easy-to-compute performance bounds for the heuristic policy. Finally, using those bounds, we show that the proposed modified echelon (r, Q) heuristic policy is asymptotically optimal as a pair of system parameters is scaled up, e.g., when the ratios of the fixed cost of the warehouse over those of the retailers become large. Numerical study demonstrates that our proposed heuristic performs well and tends to outperform the echelon-stock (r, nQ) heuristic policy studied in the literature.

The third essay is on inventory control for a single-item periodic-review stochastic inventory system with both minimum order quantity (MOQ) and batch ordering requirements. In each time period, the firm can order either none or at least as much as the MOQ. At the same time, if an order is placed, the order quantity is required to be an integral multiple of a given specific batch size. We first adopt a heuristic policy which is specified by

two parameters (s, k). Applying a discrete time Markov chain approach, we compute the system cost and optimize this (s, k) policy under the long-run average cost criterion. We also consider a simpler one-parameter policy, the so-called S policy, which is a special case of the (s, k) policy. In an intensive numerical study, we find that 1) both policies perform well in comparison with other policies; and 2) the S policy also performs well and is compatible with the (s, k) policy; only in a few cases where demand variation is small, the latter outperforms the former significantly. We also evaluate the effects of some important parameters on system performance.

CITY UNIVERSITY OF HONG KONG
Qualifying Panel and Examination Panel

Surname:	ZHU
First Name:	Han
Degree:	Doctor of Philosophy
College/Department:	Department of Management Sciences

The Qualifying Panel of the above student consists of:

Supervisor(s)

Prof. CHEN Youhua Frank	Department of Management Sciences
	City University of Hong Kong

Qualifying Panel Member(s)

Dr. LI Yanzhi David	Department of Management Sciences
	City University of Hong Kong
Dr. LU Ye	Department of Management Sciences
	City University of Hong Kong

This thesis has been examined and approved by the following examiners:

Dr. LU Ye	Department of Management Sciences
	City University of Hong Kong
Prof. CHEN Youhua Frank	Department of Management Sciences
	City University of Hong Kong
Dr. HU Ming	Rotman School of Management
	University of Toronto
Prof. ZHOU Xiang Sean	Decision Science
	The Chinese University of Hong Kong

Acknowledgments

I would like to express my sincere gratitude to many people who guided, supported and accompanied me during my Ph.D. study at City University of Hong Kong for the past five years. Without their help, it would be impossible for me to go so far.

I would first thank my supervisor, Prof. Frank Chen, for his patient, detailed and expert guidance during my Ph.D. study. He guides me not only in academic research but also in my career and personal life as a great mentor. Without his guidance, support and help, this thesis would be impossible to be accomplished. I am fortunate enough to be a Ph.D. student supervised by him, which I will always be proud of in my life.

Then the sincere thanks go to my coauthors who are Dr. Ming Hu, Dr. Eman Leung, Ms. Xing Liu, Prof. Yi Yang and Dr. Yimin Yu. I have had the privilege of working with and learning from them, and have benefited a lot from the discussions and collaborations with them.

Besides my supervisor and coauthors, I want to express my deeply gratitude to Dr. Yanzhi Li, Dr. Guangwu Liu, Dr. Ye Lu, Dr. Biying Shou and Dr. Stephen Shum for their fundamental courses, constructive guidance and effective assistance during my Ph.D. study.

I would like to give my thanks to Dr. Longmei Chen, Dr. Lijun Gong, Mr. Weixiang Huang, Dr. Huaxiao Shen and Dr. Tingru Zhang, for their company and helps during my Ph.D. life. I wish to thank the core members of Dinner Group as well, including Dr. Ying Mao, Miss Shiyu Wang and Mr. Kun Zhang, for the happy time spent together. It is my great pleasure to be one of them. I also want to thank my Ph.D. fellow students, Dr. Li Chen, Miss Hanqing Liu, Mr. Huiqiang Mao, Dr. Zhenzhen Zhang, to name just a few, for the common struggling experiences.

Besides, special thanks go to my old friends who are Ms. Yin Cai, Mr. Yang Fu, Mr. Chengxi Gao, Mr. Yaoxuan Qin, Ms. Pingping Wang, Miss Teng Wang, Mr. Chi Zhang, and others. Although most of them do not live in Hong Kong, we meet each other at least once each year. Thanks for their comfort, company and encouragements, whenever I suffer from distressful days.

Actually, I am grateful to everyone I met and everything I experienced, even including those disappeared people and those unfortunate things, which make me the one who I am now.

Last but not least, I must express my deepest love and gratitude to my parents for their constant understanding and concerning during my whole life. They have always provided me with self-giving care, unconditional support and endless love, which enable me to focus on my study and research. I sincerely hope that my achievements can make them proud.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Operations management is a field of management that is chiefly concerned with (i) designing and controlling the process of production and (ii) redesigning business operations in the production of goods or services. Both elements of operations management concern a variety of strategic issues, e.g., facility location, inventory control, pricing, process flexibility, and contract design. The goal of operations management is to help managers to make better decisions to minimize costs and/or to maximize profits. This thesis addresses two specific topics: information disclosure and inventory control. The first topic is concerned with maximizing profits, and falls within the category of redesigning business operations. The second is concerned with minimizing an inventory system's total cost, and falls within the category of controlling the process of production. Two separate problems are studied, one for single-echelon inventory systems and the other for multi-echelon systems.

The information disclosure problem is a key component of a business operations strategy. Mandatory disclosure is sometimes required by the government, e.g., it is mandatory for a public listed company to disclose key information in its annual financial reporting. More often, however, a company has wide discretion in deciding whether to disclose information and it is well known that voluntary disclosure is not always appropriate in practice. Decisions about information disclosure have a considerable impact on firm operations, and thus are of great importance.

Inventory control involves stocks of raw materials, work-in-progress, and finished goods. Various approaches and models can be used in developing inventory management systems and practices. Inventory control problems exist in all supply chains, and their strategic importance is fully recognized by top management. As Axsäter [6, page. 2] states: "For those who are working with logistics and supply chains, it is difficult to think of any qualification that is more essential than a thorough understanding of basic inventory models."

1.1 Motivation and Research Problems

The first research problem concerns quality information disclosure in the context of experience goods. Quality is an important factor that greatly influences both customer willingness to pay and firm profits. Firms usually possess more quality information better than their potential customers owing to variety of market research instruments. Potential customers, in contrast, lack the resources and expertise to access reliable information on quality. To alleviate the market inefficiency caused by this kind of information asymmetry, service providers can choose to voluntarily supply verifiable quality information to customers. Grossman [40] and Milgrom [67] demonstrate the classical *unraveling* result: high-quality firms should reveal their quality information because customers will rationally assume low quality if no quality information is disclosed, thus motivating all firms to reveal quality information if there is no direct cost to do so.

However, the unraveling result hinges on a critical assumption, namely, that customers hold a rational expectation about the quality of nondisclosed products. In other words, rational customers will always assume low quality if a firm does not disclose quality information. In practice, however, customers may not be rational enough to realize that a firm's disclosure decision is closely related to its quality information, or is even a direct "signal" of its quality. The salient feature of this work is that it investigates the firm quality disclosure strategy for experience goods by relaxing the aforementioned rational expectation assumption. Faced with unverifiable or no quality information on experience goods, customers will tend to rely on anecdotes from others who have bought or experienced those goods. Then, a natural research question is how customer bounded rationality affects a firm's incentives to disclose quality information.

The second research problem concerns the classical multi-echelon stochastic distribution inventory system, comprising one warehouse and multiple retailers. The retailers replenish their stocks from the warehouse, which in turn replenishes its stock from an external supplier with an unlimited supply. This kind of problem is well studied, but the optimal policy is still unknown. Various heuristic polices have been developed for multi-echelon inventory systems, and extensive numerical experiments have been conducted to show that those policies perform well numerically. Although most of the heuristic policies proposed in the literature make intuitive sense, it is still unclear how large the gap is between a heuristic solution and an optimal one. Therefore, the second problem addresses this issue and fills the gap in the literature on distribution systems by developing a heuristic policy with a performance guarantee.

The third research problem concerns a single-item periodic-review stochastic inventory system with both minimum order quantity (MOQ) and batch ordering requirements. MOQ and batch ordering, applied independently or simultaneously, are two common requirements made by suppliers. Both can help companies take advantage of economies of scale and hence reduce costs. The MOQ requirement means that the order quantity must equal or exceed a specified level if an order is placed. The batch ordering requirement means that the order quantity must be an integral multiple of a specified given batch size.

The coexistence of an MOQ and batch ordering has a two-sided effect. On one hand, requiring a MOQ and batch ordering simultaneously helps suppliers reduce the risk of uncertainty and achieve economies of scale. On the other hand, the two requirements may have a negative effect on buyers' inventory control, especially when MOQs are relatively large compared with their demand, which is not unusual in practice. In such situations, managers need principles or tools to help them control inventory. However, to the best of my knowledge, no research has investigated inventory systems with both MOQ and batch ordering requirements. Therefore, an inventory system with both MOQ and batch ordering requirements are studied in the third problem and several effective heuristic policies are developed.

1.2 Organization of Thesis

The remainder of the thesis is organized as follows. The essay on information disclosure for experience goods with customer bounded rationality is presented in Chapter 2, which comprises an introduction, literature review, model description, analysis, and conclusion. The distribution inventory system problem is studied in Chapter 3, and the single-stage inventory system problem with MOQ and batch ordering is studied in Chapter 4.

Chapter 2

Information Disclosure with Customer Bounded Rationality

2.1 Introduction

Quality is an important factor that largely determines customer willingness to pay and firm profits. However, for experience goods, it is difficult for customers to know the quality without actually consuming them (Nelson [68]). Typically, firms (e.g., retailers or service providers) know their quality better than potential customers do. Firms have a variety of market research instruments (e.g., customer feedback, market surveys, expert evaluations, market data, medical data in the healthcare sector) from which to obtain accurate service quality information. However, potential customers lack the resources and expertise to access such information. To alleviate this issue, service providers can voluntarily supply verifiable quality information to customers via a variety of channels. For example, they can supply such quality information by advertising their products/services in newspapers or on television.

Milgrom [67] and Grossman [40] draw the following classical *unraveling* result: firms should disclose their private quality information if its disclosure and verification are costless. The rationale is as follows. If the disclosure cost is negligible, then customers will rationally assume nondisclosing firms to have the lowest quality, which motivates firms to reveal private quality information. However, in practice there are many markets in which voluntary disclosure is incomplete, especially for experience goods (Spranca et al. [87], Mathios [65], Jin [57], Xiao [93], and Bederson et al. [9]).

As pointed out by Dranove and Jin [26], one of the strong assumptions made in the unraveling result is that customers hold a rational expectation on the nondisclosed quality of products. In practice, however, customers may not be rational enough to realize that a firm's disclosure decision is closely related to its quality information, or is even a direct "signal" of quality. For example, Brown et al. [13] find that some moviegoers do not infer low quality from a cold opening, though rational moviegoers should infer that cold-opened movies are below average in quality.

Customers may not hold rational expectations due to a variety of reasons. Some customers may lack the technical expertise necessary to interpret the disclosed quality information. This lack of expertise may be due to consumers having neither the time nor the education to become knowledgeable enough to understand the information. For example, in the healthcare context, not all customers understand the quality data prepared to assist them in choosing healthcare providers. In particular, for experience goods, even when a firm discloses quality information, customers may unintentionally ignore it because of the complex disclosure format, which can be treated as "nondisclosure". For example, the terms and conditions for credit cards are usually 10 pages in length or even longer. Dranove and Jin [26] state that customers may not pay attention to disclosed information, if they do not understand it or make naive inferences about nondisclosure (see the examples in Fishman and Hagerty [35], Hirshleifer and Teoh [45], Stivers [88], and Schwartz [77]). Moreover, some customers may also think the quality disclosure is mandatory or costless, or even do not believe the information disclosed by the firm itself. In theses examples, customers fail to directly link the disclosure decision with the firm's quality information. Hence, customers may not have rational expectation on the nondisclosed quality of products.

Faced with unverifiable or no quality information on experience goods, customers tend to rely on *anecdotes* from those who have bought the product or have experienced the service previously (Dranove and Jin [26] and Yu et al. [95]). Such anecdotes are easy to obtain via several channels. Customers can collect anecdotal evidences from their friends via word-of-mouth or user-generated content from online channels with information provided by other customers. Customers can also share their usage experiences by posting reviews on independent third-party websites (e.g., epinions.com, consumerreview.com, tripadvisor.com, yelp.com) or social media (e.g., Facebook, Twitter). Empirical studies show that customers' quality perceptions can indeed be affected by other customers' comments and online ratings (see Chen et al. [21] and Lee et al. [61] for related empirical studies). For example, in the healthcare context, people often rely on informal, qualitative information from friends, relatives, and acquaintances in choosing healthcare providers or health plans (Lupton et al. [64], Peters et al. [72], and Huppertz and Carlson [56]).

The salient feature of our paper is to investigate the firm quality disclosure strategy for experience goods by relaxing the rational expectation assumption about nondisclosure for customers. We consider the firm's quality disclosure decision, e.g., whether to disclose verifiable quality information on infrequently purchased experience goods with boundedly rational customers. Such experience goods include electronic gadgets, cars, books, movies, plays, healthcare and financial services whose quality cannot be easily captured without consumption. In the absence of verifiable quality information, other consumers' experiences constitute a key input for those who have not yet experienced the product or service (Hu et al. [49]). Then a natural research question for both academics and practitioners is: What is the impact of customer bounded rationality on a firm's quality disclosure strategy?

To model customer bounded rationality under a nondisclosure strategy, we adopt the *anecdotal reasoning* framework to describe customers' decision-making behavior. Under the anecdotal reasoning framework, customers make decisions based on the sample mean of observed anecdotes or collected samples (Osborne and Rubinstein [70]). Essentially, we assume that because of limitations in customers' ability to gather and process information (Simon [84]). They behave naively, acting as if the mean quality of a small number of samples is perfectly representative of true quality. Anticipating such customer behavior, the firm makes its price and quality disclosure decisions to maximize profits by considering homogeneous customers. If the firm withholds quality information, then potential customers can only infer its quality and make purchase decisions based on other customers' experiences. Otherwise, if the firm discloses quality information (which may involve a quality disclosure cost), then that information becomes public information.

Surprisingly, we find that the optimal disclosure strategy has the following structure: When the firm's quality level is either high or low, it is optimal not to disclose quality information. When it has a medium quality level, however, it may be optimal to disclose quality information. The main managerial insight is that customer bounded rationality is important for firms considering quality disclosure. In particular, the presence of customer bounded rationality negates a high-quality firm's incentive to disclose quality information, but not necessarily that of its medium-quality counterpart. We also show that ignoring customer bounded rationality can lead to a significant profit loss, particularly when the quality disclosure cost is high.

The managerial implications of our results for experience goods are as follows. (1) In contrast to the classic unraveling result, in the presence of customer bounded rationality, high-quality firms do not need to disclose their quality. Social learning or "reputation" under nondisclosure is effective to convey quality information to potential customers. For example, Nike spends minimal effort advertising the quality of its products. Instead, it educates customers through advertising to trust its reputation by reinforcing its brand names. (2) We argue that when the quality disclosure cost is high, firms should invest resources to facilitate customers obtaining samples and the availability of ambient information to reduce customer sample uncertainty, e.g., social media marketing. Indeed, firms are well aware that as social media marketing is increasingly popular (see Tuten and Solomon [90], Hoffman and Fodor [46]).

Finally, we consider congestion in the service context and find that the demand rate also plays a critical role on the quality disclosure decisions. In summary, our findings are complementary to the existing quality disclosure literature by relaxing the rational expectation assumption, and show that under consumer bounded rationality the optimal disclosure strategy is quite different from what the classical unraveling result predicts. Our results also provide a new explanation for the incomplete voluntary disclosure observed in many markets.

The remainder of this chapter proceeds as follows: Section 2.2 reviews the related literature. We present the model in Section 2.3 and analysis and results in Section 2.4. The model is extended to service settings with congestion in Section 2.5. We conclude in Section 2.6.

2.2 Literature Review

Our paper is closely related to the economics and marketing literature on voluntary quality disclosure. Grossman [40] and Milgrom [67] appear to be the earliest papers on voluntary quality disclosure. They draw the classical *unraveling* result: firms should disclose their private quality information if disclosure and quality verification are costless. The rationale is as follows. If the disclosure cost is negligible, then consumers with rational expectations who cannot learn will rationally assume that nondisclosing firms have the lowest quality, which motivates firms to reveal private quality information. Jovanovic [58] argues that a firm should disclose quality information only if its quality level is above a threshold if there is a disclosure cost. Matthews and Postlewaite [66] and Shavell [81] show that mandatory disclosure may motivate sellers to reduce information collection if information acquisition is costly. They assume customers to hold a rational expectation about nondisclosure.

However, some empirical studies, including Spranca et al. [87], Mathios [65], Jin [57], and Xiao [93], show that voluntary disclosure is incomplete in many markets. Board [11] argues that firms may fail to disclose their quality because quality disclosure would intensify price competition with heterogeneous consumers. Guo and Zhao [41] study the relationship between the amount of information disclosed and the timing of disclosure (simultaneous or sequential). In the literature on quality disclosure, a common assumption is that customers have rational expectations in the sense that any nondisclosed product is assumed to be of lowest quality.

As mentioned earlier, the customer rational expectation on nondisclosure is a strong assumption. In practice, customers may not take nondisclosure as a "signal" of low quality (see the examples in Fishman and Hagerty [35], Hirshleifer and Teoh [45], Stivers [88], and Schwartz [77]). Fishman and

Hagerty [35] argue that some consumers lack the technical expertise necessary to interpret disclosed quality information. Brown et al. [13] find that some moviegoers do not expect low quality from a cold opening even though rational moviegoers should infer that cold-opened movies are of poor quality.

As pointed out by Dranove and Jin [26], the unraveling result hinges on the rational expectation assumption that "*Consumers hold a rational expectation on the quality of nondisclosed products*". Our paper fills a gap in the existing literature by relaxing the rational expectation assumption.

Our quality disclosure strategy is similar to the counter-signaling structure in Feltovich et al. [32]. They consider a signaling game with three types of senders. A sender can send costly and noisy information on his or her type to the receiver, and the signal cost is higher for lower types. They show that in equilibrium medium types signal to separate themselves from low types, but high types choose not to signal. However, their setting is different from our quality disclosure setting. Note that in quality disclosure, firms do not send costly noisy quality information with heterogeneous cost functions to customers. Rather, they decide whether to disclose quality information with a fixed cost. Moreover, we assume that customers do not infer quality from a firm's disclosure strategy but infer it from other customers' experiences under nondisclosure.

This work is related to the operations literature on the firm's optimal information disclosure strategy. Hu et al. [48] study a two-period group-buying problem where a firm chooses whether to disclose the number of sign-ups accumulated in the first period. Hu et al. [50] study the impact of real-time delay information on a service system. They find that some amount of information heterogeneity in the population can lead to more efficient outcomes than information homogeneity. In this stream of works, whether the firm discloses sales volume or system state would affect customers' purchase behavior differently.

Our modeling framework of customer bounded rationality follows the recent economics literature on anecdotal reasoning, which is proposed by Osborne and Rubinstein [70]. This modeling framework has been widely applied in a variety of economic settings (see, e.g., Spiegler [85, 86], and Szech [89]). Several recent studies have applied the anecdotal reasoning framework to the marketing and operations management settings. For example, Huang and Yu [53] adopt it to analyze the profitability of opaque selling. Huang and Chen [51] consider the impact of anecdotal reasoning behavior on the pricing and capacity decisions of queueing systems. Huang and Liu [52] examine the impact of anecdotal reasoning behavior on pricing and stocking/capacity decisions. Different from these studies, we consider the impact of customer bounded rationality versus rational expectation on a firm's quality disclosure decision.

The anecdotal reasoning behavior of consumers studied in this work is also related to some recent papers studying operational decisions under social learning. For example, Yu et al. [94] study the impact of consumer-generated quality information (e.g., consumer reviews) on a firm's dynamic pricing strategy in the presence of strategic consumers. Papanastasiou and Savva [71] study dynamic pricing in the presence of social learning and strategic consumers. This essay is different from theses works because this essay studies the impact of customer bounded rationality on a firm's quality disclosure decision.

2.3 Model Setup

We consider a simple model to isolate and demonstrate the impact of customer bounded rationality on a firm's quality disclosure decision. The firm sells an experience good (either a physical product or service) to a population of $\overline{\lambda}$ homogeneous customers. We assume that the product or service is infrequently purchased by individual customers, e.g., durable goods, long-term financial services, healthcare specialist, and high end restaurants. To capture service quality uncertainty of experience goods, we assume that for each individual customer, the associated quality of service ζ is either high, denoted by "H," with probability $\mathbb{P}(\zeta = H) = \alpha$, or low, denoted by "$L$," with probability $\mathbb{P}(\zeta = L) = 1 - \alpha$, where $\alpha \in [0, 1]$. There are two main factors in quality uncertainty for experience goods. (1) Customers are inherently uncertain about consumption even when service quality remains unchanged for individual customers. That uncertainty can be affected by personal factors such as health and mood, which tend to be independent among different customers and unobservable to other customers. For example, customers may have different perceptions of the service quality of the same entertainment-related service. (2) The quality of a service process is uncertain. For example, in hospitals, medical services are uncertain for patients. We assume that the service provider knows the value of α but it is *unknown* to all customers. Throughout the paper, quality level α is fixed.

We assume that customers are homogeneous and risk-neutral. The customer valuations for service quality levels H and L are v_H and v_L, respectively, such that $v_H > v_L$. A customer with perceived valuation V will purchase the service if $V - p \geq 0$, where p is the price. Suppose that the firm decides to disclose its quality information. Then, the customer (expected) valuation of the service is $V = \alpha v_H + (1-\alpha) v_L$, as all customers know α from the firm's disclosure. If the firm decides not to disclose quality information, to model customer bounded rationality, we adopt the anecdotal reasoning framework (see, e.g., Osborne and Rubinstein [70], and Spiegler [85, 86]) and assume that customers have a limited number of samples/anecdotes from which to infer service quality. To isolate the impact of customer bounded rationality on the firm's quality information disclosure decisions, we assume that customers make their purchase decisions based solely on the samples they obtain if the firm does not disclose quality information.

The firm's objective is to decide the price charged for the product/service so as to maximize its expected profit. In addition, the firm has to decide whether (1) to disclose quality information α, i.e., adopt the disclosure strategy, or (2) not to disclose and allow customers to infer the quality information from past customers' experiences, i.e., adopt the nondisclosure strategy. There is a disclosure cost $K \geq 0$. In the basic model, we assume that the firm has sufficient capacity to serve all customers, i.e., congestion is not a major concern. We consider congestion in Section 2.5.

Notably, in the existing literature, it is commonly assumed that customers have no access to samples of past experiences and rationally expect the worst quality level for the firm that does not disclose quality information (Grossman [40], and Milgrom [67]). Our main objective in this paper is to investigate the impact of customer bounded rationality on a firm's quality disclosure decision by relaxing this rational expectation assumption.

2.4 Analysis and Results

In this section, we investigate the firm's optimal pricing strategies under the disclosure and nondisclosure strategies in the presence of customer bounded rationality.

2.4.1 Disclosure Strategy

Under the disclosure strategy, quality information α is known to all customers. Let p_D be the price under the disclosure strategy. Because customers are risk-neutral, the (expected) surplus for each customer is $\alpha v_H+(1-\alpha)v_L-p_D$. A customer will purchase the service only if her surplus is non-negative; otherwise, she will not make the purchase and leave the market. Given that a profit-maximizing firm wants to attract customer population λ, its optimal pricing problem is provided as follows:

$$\begin{aligned} \pi_D^*(\lambda) = & \max_{p_D \geq 0} \; p_D\lambda - K \\ \text{subject to} \;\; & \alpha v_H + (1-\alpha)v_L - p_D \geq 0. \end{aligned} \tag{2.1}$$

The optimal strategy under the disclosure strategy if the firm attracts population λ is given in the following lemma.

Lemma 2.1. *Suppose that a firm attracts customer population λ. Then, under the disclosure strategy, its optimal strategy is as follows.*

$$p_D^* = \alpha v_H + (1-\alpha)v_L. \tag{2.2}$$

In addition, the firm earns a profit $\pi_D^*(\lambda) = [\alpha v_H + (1-\alpha)v_L]\lambda - K$.

The optimal solution for (2.1) can be obtained by letting $\alpha v_H + (1-\alpha)v_L - p_D = 0$. Note that $\pi_D^*(\lambda)$ is an increasing function of λ.

Next, we provide the optimal decision for the firm under the disclosure strategy.

Lemma 2.2. *Let $\tilde{\lambda}_D = \frac{K}{\alpha v_H+(1-\alpha)v_L}$. The firm can earn a profit $\pi_D^*(\overline{\lambda}) = [\alpha v_H + (1-\alpha)v_L]\overline{\lambda} - K$ by adopting the disclosure strategy, which is nonnegative if and only if $\overline{\lambda} \geq \tilde{\lambda}_D$.*

Because customers are homogeneous, if market size $\overline{\lambda}$ is sufficiently large, it is optimal for the firm to enter the market by attracting all customers; otherwise it is not profitable for the firm to enter the market.

2.4.2 Nondisclosure Strategy with Customer Bounded Rationality

In the extant literature, it is typically assumed that rational customers infer low quality in the absence of information disclosure. In practice, however, particularly in the case of experience goods, customers may not take nondisclosure as a signal of low quality. For example, customers may not pay attention to the available information if they do not understand it, or they may make naive inferences about nondisclosure (Dranove and Jin [26]). Customers can also acquire related quality information from other customers' experiences via word-of-mouth or user-generated content and then reason about the firm's quality. For example, Yelp allows users to share their dining experiences in local restaurants. Potential customers may also ask friends who have bought the product before about their experiences. When quality information is not disclosed, potential customers are likely to rely on such anecdotes as "my colleague spent a terrible night at that hotel" or "my friend's Toyota has a smooth ride and is extremely fuel efficient." The salient feature of our model is to incorporate such bounded rationality by relaxing the rational expectation assumption under nondisclosure. We adopt the anecdotal reasoning framework proposed in the recent economics literature (see, e.g., Osborne and Rubinstein [70]; Spiegler [85, 86]) to model customer bounded rationality. Based on the past experiences of others (i.e., word-of-mouth), customers rely on anecdotal reasoning to make their own purchasing decisions.

Remark 2.1. *Note that in the economics and marketing literature, the Bayesian updating rule is sometimes adopted to model customer learning behavior. However, in our context, we consider the experience goods with infrequent purchases, e.g., durable goods, long-term financial services, specialists in* *healthcare, high end restaurants. As we can see below, we assume different generations of customers to enable learning. With the Bayesian updating rule, customers are required to fully rational and update their beliefs in each period, which is not reasonable in our setting. Moreover, there is evidence to show that customers do make decisions based solely on small samples (Tversky and Kahneman [91], Osborne and Rubinstein [70], and Fiedler and Juslin [33]).*

Sequence of Events.

To model customer bounded rationality, we consider a setting in which a firm provides a product/service to different generations of new customers. There are infinitely many discrete "stages" indexed by $k = 1, 2, 3, ...$ Moreover, we also assume each stage to have $\bar{\lambda}$ customers. The firm first commits to using the nondisclosure strategy and to setting price p before time 0. Recall that $\mathbb{P}(\zeta = H) = \alpha \in [0, 1]$ is the probability of obtaining high service quality, where ζ denotes a customer's service quality "realization" from the firm. Customers with bounded rationality do not know the exact value of α. Motivated by Spiegler [85] and Huang and Yu [53], we use the following dynamics to model the anecdotal reasoning process. In Stage 1, generation-1 new customers enter the market and are randomized with equal probabilities (assuming any other strictly positive probability does not affect our results) in making their purchase decisions given that they have *no* information about α. After making their purchasing decisions, the customers obtain their individual service quality realizations, and then they leave the market. In Stage 2, generation-2 customers enter the market. Before making a purchasing decision, each generation-2 customer has an opportunity to communicate with N generation-1 customers, obtaining N samples/anecdotes of the service quality realizations in stage 1. In general, in Stage $k = 1, 2, 3, ...$, each generation-k

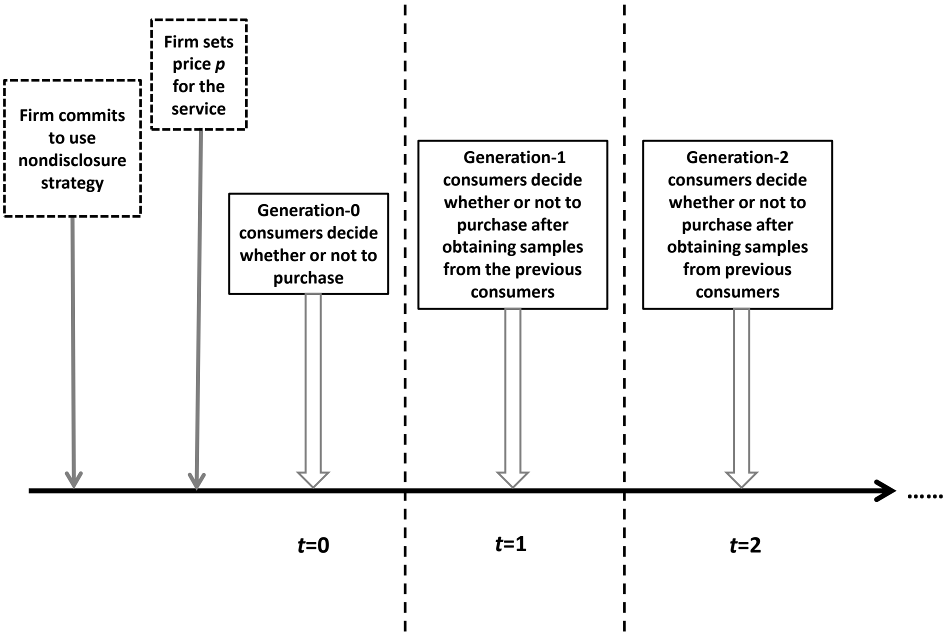

Figure 2.1: Sequence of Events

customer can sample N realizations of generation-(k') customers before making her purchasing decision, for some $k' < k$. Each generation-k customer decides whether to purchase or not based on her samples. However, her *own* service quality offered by the firm at Stage k is an *independent* realization from her samples. Hence, the quality level that she actually obtains probably differs from that in her samples.

$S(N)$ Framework

We now examine the $S(N)$ framework in which each customer is assumed to obtain N samples and "combine" multiple samples by simply taking the sample average. We use the indicator random variable $I_{\{i,j\}}$ to denote the type of sample j that customer i obtains. If the jth sample that customer i obtains is product/service H, then $I_{\{i,j\}} = 1$; otherwise, $I_{\{i,j\}} = 0$. Note that $\mathbb{E}[I_{\{i,j\}}] = \alpha$ and that we focus on the steady-state dynamics of the system described in Figure 2.1. The N samples can be obtained from previous customers, and each customer's samples are independent of those of other customers. Each customer relies on the N samples obtained to make her purchasing decision. Let

$$\alpha_i(N) \equiv \frac{1}{N} \sum_{j=1}^{N} I_{\{i,j\}}$$

be the mean of the samples obtained by customer i. Then, the estimated (expected) valuation of the service for customer i can be expressed as $V = \alpha_i(N)v_H + (1 - \alpha_i(N))v_L$. If $N = \infty$, then all customers learn the exact quality information α. We focus on cases in which $N < \infty$, i.e., each customer can obtain a limited number of samples.

Note that we assume that customers make their purchasing decisions based solely on their obtained samples. Then, customer i purchases from the service provider if and only if

$$\alpha_i(N)v_H + (1 - \alpha_i(N))v_L - p \geq 0. \tag{2.3}$$

Let $\gamma(p)$ be the fraction of customers who purchase the product/service given price p. Based on (2.3), if the customer purchases from the firm, the samples she obtains must satisfy the following condition.

$$\alpha_i(N) \geq \frac{p - v_L}{v_H - v_L}. \tag{2.4}$$

Let $B(n, N, \alpha) = \binom{N}{n}\alpha^n(1 - \alpha)^{N-n}$ be the probability mass function of the binomial distribution with parameters N and α. Then we have the following result.

Lemma 2.3. *Given p, $\gamma(p)$ can be expressed as*

$$\gamma(p) = 1 - \sum_{n=0}^{\lfloor N\frac{p-v_L}{v_H-v_L}-\epsilon\rfloor} B(n, N, \alpha) \tag{2.5}$$

for an arbitrarily small $\epsilon > 0$. In addition, $\gamma(p)$ is decreasing in p.

If the firm decides to adopt the nondisclosure strategy, then it chooses p to maximize its expected profit $\pi(p) = p\gamma(p)\overline{\lambda}$. Solving this optimization is challenging because there is no explicit expression for $\gamma(p)$. However, we have the following result to characterize the optimal strategy.

Lemma 2.4. *The optimal pricing strategy can be described as follows.*

(i) The optimal price p^* satisfies the following condition.

$$p^* = v_L + \frac{j}{N}(v_H - v_L), \textit{for some } j \in \{0, 1, 2, \ldots, N\}.$$

(ii) The optimal price p^* under this nondisclosure strategy can be found by using the following algorithm.

(a) For each $j = 0, 1, \ldots, N$, compute the proportion of customers who purchase, $\gamma_j = 1 - \sum_{n=0}^{j-1} B(n, N, \alpha)$, and the corresponding profit $\pi_j^* = [v_L + \frac{j}{N}(v_H - v_L)]\gamma_j\overline{\lambda}$.

(b) The optimal firm profit is

$$\pi_{ND}^* = \max\{\max_{j=0,1,\cdots,N} \pi_j^*, 0\}.$$

Let $j^* = \arg\max_{j=0,1,\cdots,N} \pi_j^*$ be the optimal price index. The optimal price is $p^* = v_L + \frac{j^*}{N}(v_H - v_L)$ if $\pi_{j^*}^* \geq 0$.

Note that π_{ND}^* is polynomial in α, which is neither convex nor concave. It is unclear how quality level α will affect the optimal price due to the technical challenge imposed by the polynomial function of α. To tackle this problem, we first introduce the notion of log-supermodularity.

Definition 2.4.1. Let $(X, \geq)$ be a lattice. A function $h : X \to \mathbb{R}$ is said to be log-supermodular if it is non-negative and $h(\mathbf{x} \vee \mathbf{y}) \cdot h(\mathbf{x} \wedge \mathbf{y}) \geq h(\mathbf{x}) \cdot h(\mathbf{y})$ for all $\mathbf{x}$, $\mathbf{y} \in X$.

Remark 2.2. *Athey [1] states the following properties for log-supermodularity. (i) If h is positive, then h is log-supermodular if and only if $\log(h(\mathbf{x}))$ is supermodular. (ii) Products of log-supermodular functions are log-supermodular. (iii) A sufficient condition for $\int u(\mathbf{x}, \mathbf{s}) f(\mathbf{s}; \theta) d\mu(\mathbf{s})$ being log-supermodular is that u and f are log-supermodular.*

In particular, log-supermodularity implies nondecreasing strategies.

Lemma 2.5 (Athey [1]). *Suppose that f is non-negative. Then, $\mathbf{x}^*(\theta) \equiv \arg\max_{\mathbf{x} \in A} U(\mathbf{x}, \theta)$ is nondecreasing in θ for all $u : X \times S \to \mathbb{R}_+$ log-supermodular, if and only if U is log-supermodular in $(\mathbf{x}, \theta)$ for all $u : X \times S \to \mathbb{R}_+$ log-supermodular.*

Utilizing the notion of log-supermodularity and Lemma 2.5, we obtain the following result.

Lemma 2.6. *The profit function $\pi_j^*(\alpha)$ is log-supermodular in (j, α). As a result, under the nondisclosure strategy we have (i) the optimal price index $j^*(\alpha)$ is nondecreasing in α; (ii) the optimal price p^* is nondecreasing in α, and the proportion of customers who purchase, γ_{j^*} is nonincreasing in α.*

Despite the noisy sample obtained by each customer, Lemma 2.6 implies that under the nondisclosure strategy, a firm with a high quality level should charge a high price for its service. This is because as α increases, the probability of getting more H samples is higher and the firm is better off by charging a high price.

Lemma 2.6 also implies that, as α increases, the proportion of customer who will buy the product becomes smaller. Although the firm's target customers becomes less, their average valuations are higher. Those high valuations enable the firm to earn a higher marginal profit and ultimately switch its focus from a mass market to this niche segment of customers.

When determining the price under the nondisclosure strategy, a firm makes a trade-off between a higher profit margin and a larger demand quantity. When α is small, setting a low price enables the firm to attract enough customers to make a profit; whereas when α is large, the high profit margin motivates the firm to increase its price and thus brings more profit to the firm than a low price.

Optimal Disclosure Decision

Recall that if the firm discloses quality information, its corresponding optimal profit is $\pi_D^* = [\alpha v_H + (1-\alpha)v_L]\overline{\lambda} - K$, as shown in Section 2.4.1. Let π_{CL} denote the firm's optimal profit in the presence of customer bounded rationality. It follows that $\pi_{CL} = \max\{\pi_{ND}^*, \pi_D^*\}$. Specifically, if $\pi_{ND}^* \geq \pi_D^*$, then the nondisclosure strategy is optimal; otherwise, the disclosure strategy is optimal. It is unclear whether $\pi_{ND}^* - \pi_D^*$ is increasing or decreasing in α. Interestingly, we have the following result.

Proposition 2.1. *There exists* $\underline{\alpha}, \overline{\alpha}$ *such that* $0 < \underline{\alpha} \leq \overline{\alpha} < 1$*, and it is optimal to adopt the nondisclosure strategy for* $\alpha \in [0, \underline{\alpha}] \cup [\overline{\alpha}, 1]$*, whereas it may be optimal to adopt the disclosure strategy for some* $\alpha \in [\underline{\alpha}, \overline{\alpha}]$*.*

Proposition 2.1 states that it is optimal to adopt the nondisclosure strategy when α is either small ($[0, \underline{\alpha}]$) or large ($[\overline{\alpha}, 1]$). If α is medium ($[\underline{\alpha}, \overline{\alpha}]$), then it may be optimal to adopt the disclosure strategy.

One might expect the disclosure strategy to be monotone in the quality level. For example, Lemma 2.6 stipulates that the price is increasing in quality level. However, Proposition 2.1 implies that a firm might be willing to disclose quality information if its quality level is medium, but will adopt the nondisclosure strategy if it is of low or high quality. The key message we want to deliver is that customer bounded rationality has significant implications for a firm's optimal disclosure strategy, which may be much more complicated than what we thought previously. Another important factor in a firm's quality disclosure decision is disclosure cost K. If K is sufficiently large, it is always optimal to adopt the nondisclosure strategy.

Hence, with customer bounded rationality, a firm's quality disclosure decision becomes complex. Correspondingly, the firm should be more careful in deciding whether to disclose quality information in the presence of social media and word-of-mouth. Interestingly, our results are consistent with some recent empirical findings. For example, Luca and Smith [63] show that mid-ranked business schools in the United States are the most likely to disclose their ranking information. Bederson et al. [9] empirically find the quality disclosure strategy is very similar to what our model predicts: when Maricopa County in Arizona adopted voluntary restaurant hygiene grade cards (A, B, C, D), the better A-grade restaurants tended not to disclose their grades, whereas worse As and better Bs disclosed their grades voluntarily.

What is the intuition of our result? There are two forces driving the profitability of nondisclosure: pricing flexibility and the variability of customers' quality perceptions. Note that under disclosure, a firm can charge only the expected valuation. One major advantage of nondisclosure for the firm is pricing flexibility, i.e., it can choose from multiple candidate prices as individual customers may receive different samples. The disadvantage of nondisclosure is that customers may have different perceptions of firm quality (variability), e.g., some customers may underestimate its quality, whereas others may overestimate it. When determining its pricing strategy, a firm makes a trade-off between a higher price and lower demand. When firm quality is high or low, under anecdotal reasoning, there is little variability in customers' quality perceptions (i.e., $N\alpha(1-\alpha)$ for the binomial distribution). Then, although customers' average estimation is consistent with the firm's quality level overall, some customers may overestimate its quality. A high-quality firm can exploit overestimating customers by charging a higher price and capturing most of the customers with a perception of high quality. As a result, the firm has no incentive to disclose quality information. However, a medium-quality firm may anticipate that the samples collected by customers are noisy and that there are a significant proportion of customers who either overestimate or underestimate its quality. As a result, the firm may not be able to utilize the pricing flexibility, as it can charge a lower or higher price only by risking the loss of a significant market share, thereby motivating it to disclose quality information to differentiate itself from low-quality firms.

We now provide a useful refinement of Proposition 2.1. Recall that it is optimal to adopt the nondisclosure strategy in at least two regions: the region close to 0 ($[0, \underline{\alpha}]$) and the region close to 1 ($[\overline{\alpha}, 1]$). More specifically, for the regions closest to 0 and 1, we have the following result.

Corollary 2.3. *There exists $0 < \underline{\alpha}^0 < \underline{\alpha}$ and $\overline{\alpha} < \overline{\alpha}^N < 1$ such that it is optimal to adopt the nondisclosure strategy with $j^* = 0$ for $\alpha \in [0, \underline{\alpha}^0]$, and the nondisclosure strategy with $j^* = N$ for $\alpha \in [\overline{\alpha}^N, 1]$.*

Corollary 2.3 states that for the region closest to 0 ($[0, \underline{\alpha}^0]$), it is always optimal to adopt the nondisclosure strategy with $j^* = 0$, i.e., attracting all customers, and for the region closest to 1 ($[\overline{\alpha}^N, 1]$), it is always optimal to adopt the nondisclosure strategy with $j^* = N$, i.e., attracting customers whose samples are all "H." These results hold because $j^*(\alpha)$ is nondecreasing in α and the nondisclosure strategy is optimal in the regions close to 0 or 1.

To further illustrate the optimal disclosure strategy, Figure 2.2 shows the profits accruing from both the nondisclosure and disclosure strategies. We find that when firm quality level is either high or low ($\alpha \in [0, 0.39]$ or $[0.79, 1]$), it is optimal to adopt the nondisclosure strategy. In addition, when α is at a medium level, the nondisclosure strategy might be optimal. For example, for $\alpha \in [0.42, 0.56], [0.67, 0.76]$, the optimal strategy is the nondisclosure strategy with $j^* = 1$ and $j^* = 2$, respectively. In other regions, disclosure is the optimal strategy.

Note that if each customer is able to obtain enough samples, the firm no longer needs to disclose quality information.

Lemma 2.7. *There exists N_0 such that if $N > N_0$, it is optimal to adopt the nondisclosure strategy.*

Lemma 2.7 shows that when customers have enough samples, the firm should adopt the nondisclosure strategy.

Figure 2.3 shows the impact of N on firm's quality disclosure decision. We can observe that when N is not too large, the firm does not disclose quality information when α is high or low, but does disclose when α is medium (although, in some small regions, it may still be optimal not to disclose quality information). In addition, the disclosure region tends to decrease as N increases: as customer reasoning becomes more effective, the firm has

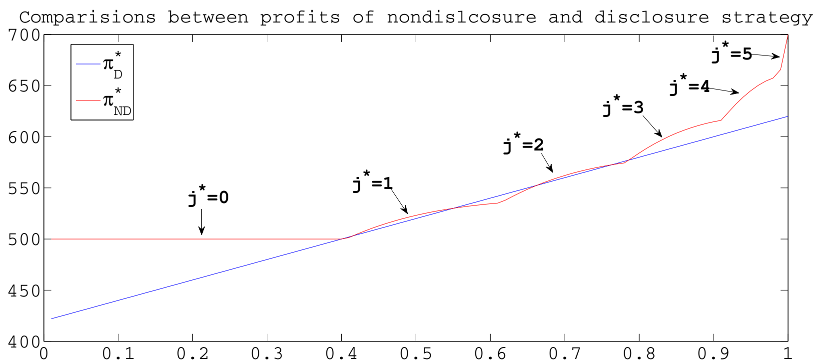

Figure 2.2: Profits for the Nondisclosure and Disclosure Strategies (N=5, $v_H = 7$, $v_L = 5$, $\overline{\lambda}$=100 and K=80)

α	[0,0.40]	[0.41,0.41]	[0.42,0.55]	[0.56,0.61]	[0.62,0.66]
j^*	0	1	1	1	2
OPT	ND	D	ND	D	D
α	[0.67,0.76]	[0.77.0.78]	[0.79,0.91]	[0.92,0.98]	[0.99,1]
j^*	2	2	3	4	5
OPT	ND	D	ND	ND	ND

OPT: Optimal Strategy; ND: Nondisclosure Strategy; D: Disclosure Strategy

Table 2.1: Profits for the Nondisclosure and Disclosure Strategies (N=5, $v_H = 7$, $v_L = 5$, $\overline{\lambda}$=100 and K=80)

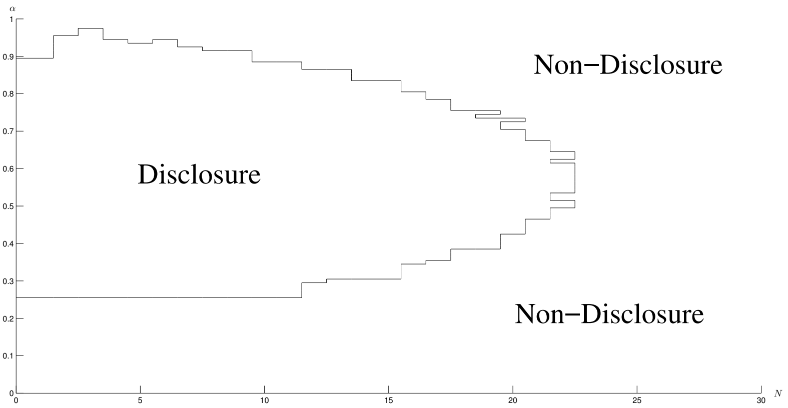

Figure 2.3: Optimal Disclosure Decision under Different N and α ($v_H = 7$, $v_L = 5$, $\overline{\lambda}$=100 and K=80)

fewer incentives to disclose, and when customers have enough samples, the firm should always adopt the nondisclosure strategy.

To better understand how customer bounded rationality affects a firm's disclosure decision, we consider the special case of $N = 1$, i.e., the $S(1)$ framework. In the $S(1)$ framework, each individual customer obtains one anecdote/sample about a service quality realization that occurred in some previous period. The optimal profit π^* is as follows.

$$\pi^* = \{[\alpha v_H + (1-\alpha)v_L]\overline{\lambda} - K, \alpha v_H \overline{\lambda}, v_L \overline{\lambda}\}.$$

Proposition 2.2. *Under the $S(1)$ framework, if the condition that $1 - \frac{K}{v_L \overline{\lambda}} < \frac{v_L}{v_H} < \frac{K}{(v_H - v_L)\overline{\lambda}}$ holds, the firm should always adopt the nondisclosure strategy; otherwise, the optimal strategy can be divided into three segments: the firm should adopt the disclosure strategy when α is at a medium level, and the nondisclosure strategy when α is small or large.*

Proposition 2.2 is consistent with the results in Proposition 2.1; that is, the firm should not disclose quality information when α is relatively large or small, but should disclose it when α is at a medium level.

Our findings underscore the importance of customer bounded rationality and demonstrate that the optimal quality disclosure decision is not straightforward. In addition, our model provides a new explanation for experience goods with incomplete voluntary quality disclosure, e.g., high-quality firms have no incentive to disclose. The managerial implications of our results for experience goods include (1) In contrast to the classic unraveling result, in the presence of customer bounded rationality, high quality firms do not need to disclose their quality. Word-of-mouth or "reputation" under nondisclosure is effective to convey their quality to potential customers. For example, Nike spends minimal effort advertising the quality of its products. High-quality firms should educate customers, e.g., through advertising, that they should trust their "reputation" by reinforcing their brand names. Hence, nondisclosure may not be a "signal" of low quality. (2) We argue that when the quality disclosure cost is high, firms should invest resources to facilitate customers obtaining samples, e.g., in social media marketing. Indeed, firms are well aware of this as social media marketing is increasingly popular (see Tuten and Solomon [90], Hoffman and Fodor [46]).

2.4.3 Impact of Ignoring Customer Bounded Rationality

In this section, we investigate the impact of ignoring customer bounded rationality by using the unraveling result on firm profit (e.g., the firm may simply adopt the unraveling result even in the presence of customer bounded rationality). When consumers hold a rational expectation about nondisclosure, the optimal firm profit under the disclosure strategy is reported in Appendix B.

Now we are ready to measure the value of customer bounded rationality against the result of ignoring it and using the classic unraveling result. We use the profit loss δ between the profit with customer bounded rationality, π_{CL}, and without it, π_{NL}, to measure the impact of ignoring customer bounded rationality:

$$\delta = \frac{\pi_{CL} - \pi_{NL}}{\pi_{NL}} \times 100\%.$$

Lemma 2.8. *The profit loss δ by ignoring customer bounded rationality and using the unraveling result is nondecreasing in K.*

Lemma 2.8 states that if the quality disclosure investment is large, firms should not ignore customer bounded rationality. In this case, firms would be better off by investing resources to facilitate customer obtaining samples.

Figure 2.4 depicts the profit loss by ignoring customer bounded rationality and using the unraveling result under different K and α. It can be seen that the profit loss tends to increase in α, and the profit loss can be up to 30% for a large K.

2.5 Quality Disclosure with Congestion

In practice, customers may have to wait for services if the capacity of a service provider is limited. For example, in hospitals and theme parks, waiting is quite common for patients and visitors, respectively. Customer waiting time may also affect customers' choices. In this section, we analyze the quality disclosure decision in the presence of congestion. We adopt the standard queueing setting; that is, we assume that the firm faces an independent demand stream with customers arriving according to a Poisson process with total rate $\overline{\lambda}$. The firm serves customers one at a time on a first-come, first-served basis. We assume service times to be independent and identically distributed random variables denoted by X. In particular, we assume that X is of the form Y/μ and that Y is an exponentially distributed random variable with mean equal to 1. Hence, service times are exponentially distributed with mean $\mathbb{E}[X] = 1/\mu$. The parameter μ is a scaling parameter that corresponds to the service rate or capacity. There is also a per unit time marginal capacity cost c.

Random variable Y can be viewed as the work content associated with each customer. Given the exponential nature of both customer inter-arrival times and service times, the firm is modeled as an $M/M/1$ queue. There is an extensive body of literature on the economics of queues focusing primarily on the $M/M/1$ queue; see Hassin and Haviv [42] for a review.

Let λ be the effective arrival rate of the firm, which results from customer

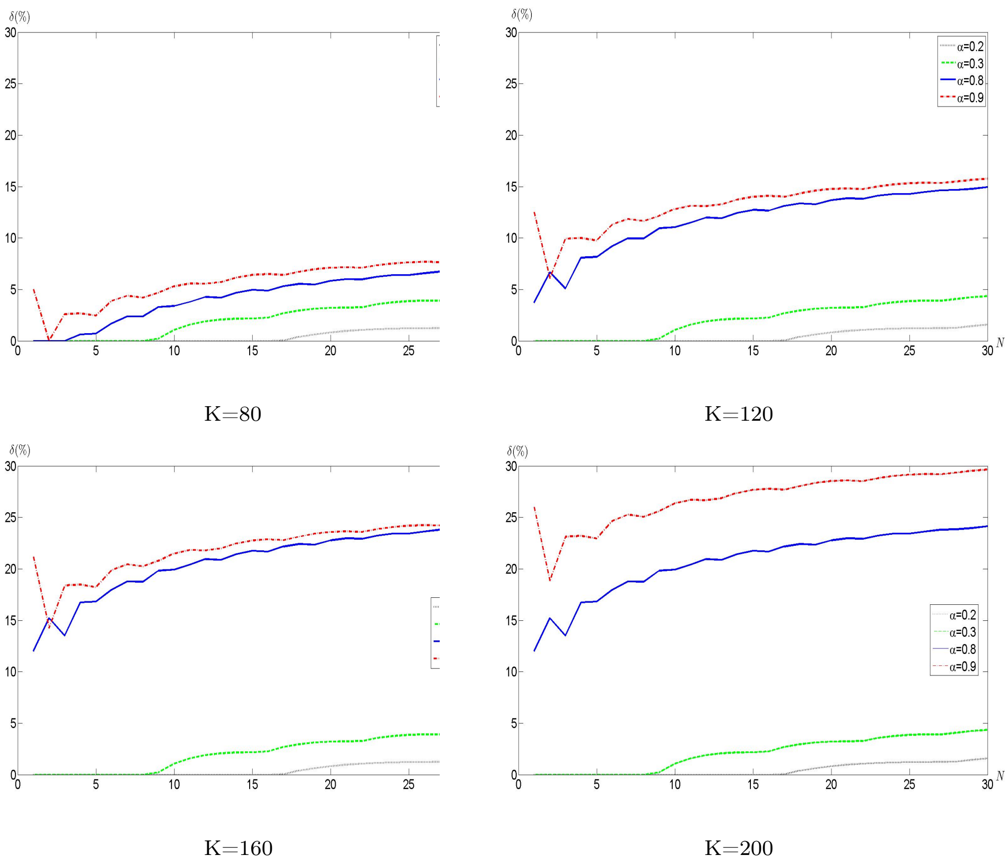

Figure 2.4: Profit Loss δ by Ignoring Customer Bounded Rationality and Using the Unraveling Result ($v_H = 7$,$v_L = 7$, and $\overline{\lambda} = 100$)

choice. Because the capacity is limited, not all customers can join the service provider. Then, the expected delay in equilibrium for customers is $\mathbb{E}[W] = \frac{1}{\mu-\lambda}$. We assume that customers face an invisible queue, i.e., customers cannot observe the exact queue length, demand rate or capacity level. However, they can form rational expectations about the expected delay $\mathbb{E}[W]$. For example, before making a purchase an individual customer may not observe real-time delay information, but its expectation is observable to waiting customers. This is quite different from service quality, which is private information held by individual customers. In some settings, the firm simply announces the expected delay. For example, Disney theme parks list the average waiting times for various activities, and some hospitals also post expected waiting times.

The firm's objective is to decide whether to disclose its quality information and to determine the corresponding price to charge for each service and the capacity level so as to maximize its expected profit.

Given a customer with valuation V, price p for the service, and expected delay $\mathbb{E}[W]$, the customer will purchase the service if $V - p - h\mathbb{E}[W] \geq 0$, where h is the disutility or waiting cost of the delay per unit of time.

Under the voluntary disclosure strategy, quality information α is known to all customers. Let p_D and μ_D be the price and capacity level under the disclosure strategy, respectively. Each customer's surplus is $\alpha v_H + (1-\alpha)v_L - p_D - h\mathbb{E}[W]$, where $\mathbb{E}[W] = \frac{1}{\mu_D - \lambda}$ given arrival rate λ ($\lambda \leq \overline{\lambda}$). Given that a profit-maximizing firm wants to attract customers with demand rate λ, its optimal decision problem is as follows.

$$\pi_D^*(\lambda) = \max_{p_D \geq 0, \mu_D \geq 0} p_D\lambda - c\mu_D - K \tag{2.6}$$
$$\text{subject to } \alpha v_H + (1-\alpha)v_L - p_D - h\frac{1}{\mu_D - \lambda} \geq 0,$$
$$\mu_D > \lambda.$$

Under the nondisclosure strategy, for analytical simplicity and to obtain managerial insights, we focus on the $S(1)$ framework (our results still hold qualitatively under the $S(N)$ framework). More specifically, we consider the $S(1)$ framework in which each customer is assumed to obtain one sample from a previous customer who purchased the service from the firm.

Under the $S(1)$ framework, customer i purchases the service if and only if

$$\alpha_i(1)v_H + (1-\alpha_i(1))v_L \geq p + h\mathbb{E}[W], \tag{2.7}$$

i.e., she makes the purchase if the valuation derived from her sample is no smaller than the selling price plus the expected waiting cost. Under the nondisclosure strategy, the firm can either attract only customers with an "H" sample (H-type customers) or all customers. We provide the optimal selling strategies under the disclosure and nondisclosure strategies in Appendix C.

Define $\lambda_1 = \left(\frac{\sqrt{hc}+\sqrt{(v_L-c)K}}{v_L-c}\right)^2$, $\lambda_2 = \left(\frac{\sqrt{hc}+\sqrt{hc+(v_L-c)K}}{v_L-c}\right)^2$, $\alpha_1 = \left(\frac{\sqrt{hc\overline{\lambda}}-\sqrt{[(v_L-c)\overline{\lambda}-\sqrt{hc\overline{\lambda}}]^2-(v_L-c)K\overline{\lambda}}}{(v_L-c)\overline{\lambda}}\right)^2$, and $\alpha_2 = \left(\frac{\sqrt{hc\overline{\lambda}}+\sqrt{[(v_L-c)\overline{\lambda}-\sqrt{hc\overline{\lambda}}]^2-(v_L-c)K\overline{\lambda}}}{(v_L-c)\overline{\lambda}}\right)^2$. The following proposition characterizes the situations in which the firm should disclose quality information.

Proposition 2.3. *If $\alpha\overline{\lambda}(v_H - v_L) \geq K$, the optimal disclosure strategy is provided in Table 2.2. If $\alpha\overline{\lambda}(v_H - v_L) < K$, the nondisclosure strategy is optimal.*

Conditions		Optimal Strategy
$\overline{\lambda} < \lambda_1$		*Nondisclosure*
$\lambda_1 \leq \overline{\lambda} \leq \lambda_2$	$\alpha < \alpha_1$ *or* $\alpha > \alpha_2$	*Nondisclosure*
	$\alpha_1 \leq \alpha \leq \alpha_2$	*Disclosure*
$\overline{\lambda} > \lambda_2$	$\alpha \leq \alpha_2$	*Disclosure*
	$\alpha > \alpha_2$	*Nondisclosure*

Table 2.2: *Optimal disclosure strategy under the condition* $\alpha\overline{\lambda}(v_H - v_L) \geq K$

Proposition 2.3 shows that the results in Proposition 2.1 still hold in the presence of congestion. With congestion, demand rate $\overline{\lambda}$ also plays a critical role in the quality disclosure decision. When the demand rate is moderate, it is optimal for the firm to disclose quality information when its quality level is medium and not to disclose it when it is either high or low. When the demand rate is high, contrary to the unraveling result, it is optimal for the firm to disclose quality information if its quality level is low but not to disclose it otherwise. In addition, we find that when the demand rate is sufficiently low, the firm has no incentive to disclose quality information. The managerial implication is that in the presence of congestion, the optimal disclosure strategy is jointly determined by firm quality and the demand rate.

Unlike the case without congestion, when the demand rate is high, it is optimal only for a low-quality firm to disclose quality information. Why should the firm disclose its low quality information with the additional quality disclosure cost? Intuitively, with customer bounded rationality, when a firm's quality level is low, it is not easy for customers to receive an H sample. Under the nondisclosure strategy, the firm sets a low price with a capacity investment cost, which adversely affects its profit. Switching to the disclosure strategy might garner the firm more profit in this case.

Figure 2.5 illustrates the optimal disclosure strategy for different combinations of α and $\overline{\lambda}$. If $\overline{\lambda}$ is large, it is optimal to (1) adopt the nondisclosure strategy by attracting all customers when α is small; (2) adopt the disclosure strategy when α is at a medium level; and (3) adopt the nondisclosure strategy by attracting customers who obtained an "H" sample (H-type customers) only when α is large. When $\overline{\lambda}$ is medium, it is optimal not to disclose quality information. Finally, if $\overline{\lambda}$ is small, it is optimal for the firm not to enter the market.

2.6 Concluding Remarks

Deciding whether to disclose quality information on experience goods is of strategic importance for firms. The classical unraveling result stipulates that firms are willing to disclose quality information because rational customers may infer nondisclosure as low quality. In practice, however, in the context of for experience goods, customers may not be rational enough to realize that a firm's disclosure decision is closely related to its quality information, or is even a direct "signal" of quality.

The salient feature of this paper is to relax the rational assumption concerning nondisclosure. Instead, under nondisclosure we assume customers acquire quality information through word-of-mouth or social media on other customers' experiences, i.e., engage in anecdotal reasoning. We examine firm's disclosure of quality information on experience goods under customer bounded rationality. We find that determining the optimal disclosure strategy in this context can be complex. Specifically, our results show that a firm with either a high or low quality level will prefer to hide quality information. A firm with a medium quality level, however, may have an incentive to disclose such information. This counterintuitive result arise with or without the presence of congestion. This is because if the service provider's quality level is either high or low, the information derived from samples by customers is relatively consistent with its actual quality level, and some customers may

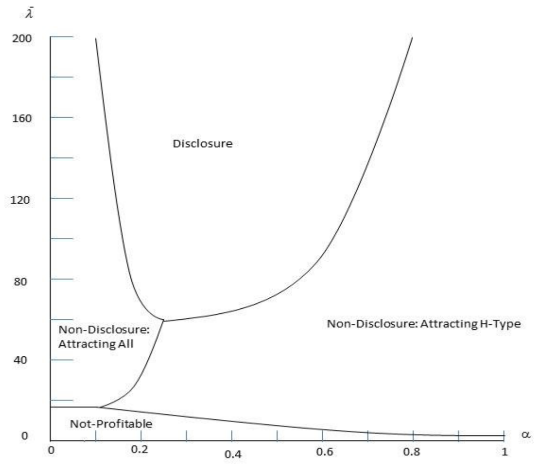

Figure 2.5: Optimal Strategy for Different Combinations of α and $\overline{\lambda}$ ($v_H = 7$, $v_L = 5$, $c = 4$, $K = 30$ and $h = 1$)

overestimate its quality. As customers may obtain different samples, under nondisclosure, particularly if the firm enjoys pricing flexibility: the firm can exploit customer behavior by charging a higher price than its disclosure counterpart. However, when the firm's quality level is medium, the quality information obtained from samples can be very noisy, and many customers may underestimate the firm's quality, thereby motivating the firm to disclose it.

Thus, our results provide a new explanation for the incomplete nature of voluntary disclosure in many markets. In particular, our results imply that word-of-mouth or "reputation" under nondisclosure is effective in conveying firms' high quality information to potential customers. High-quality firms are advised to educate potential customers, e.g., through advertising or social media, to trust their "reputation" by reinforcing their brand names.

Chapter 3

Performance Bounds for Distribution Systems

3.1 Introduction

We consider a continuous-review stochastic distribution inventory system with one warehouse and multiple retailers (OWMR). Customer demands arise only at retailers, and each retailer faces an independent Poisson demand process. The warehouse replenishes the retailers and receives stock from an outside supplier with unlimited capacity. Each shipment, either from the supplier to the warehouse or from the warehouse to a retailer, regardless of its size, takes a positive constant lead time and triggers a positive fixed setup cost. A holding cost is charged for each unit carried at the warehouse and retailers. Excess demand that cannot be satisfied immediately at each retailer is fully backlogged, but incurs a backlogging cost. The objective is to minimize the long-run average system-wide cost.

It is well known that the optimal policy of such a system, even if it exists, must be extremely complicated. For this reason, the research on OWMR models focuses on easy-to-implement heuristic policies. Yet, to the best of our knowledge, no policy has been provided with a performance bound, (See Simchi-Levi and Zhao [83] for a comprehensive review). Recently, Hu and Yang [47] study a serial system with N stages. They introduce a class of so-called modified echelon (r, Q) policies for serial systems (referred to as MERQS): if the echelon inventory position at Stage i is at or below r_i, and Stage $i+1$ at the upper stream has positive on-hand inventory, then a shipment is sent to Stage i to raise its echelon inventory position *as close as possible* to $r_i + Q_i$. For this heuristic, they provide worst-case performance bounds under practical conditions. For example, if the upper stream has a higher fixed cost than the downstream, which is typically the case in practice, their heuristic policy is guaranteed to be 2-optimal. In addition, they identify a set of conditions under which MERQS is asymptotically optimal when certain system parameters scale up. A natural question is whether the notion of modified echelon (r, Q) policies can be extended to analyzing a distribution system so that certain performance bounds can be established. It turns out that the answer is yes, yet the notion and analysis has to be adapted to the distribution system in a nontrivial way.

There exist fundamental differences between the serial and distribution systems. In a serial system, once a shipment arrives at the upstream installation, then sooner or later, all the units of this shipment will be shipped to the subsequent downstream installation, in one or multiple shipments. Whenever a shipment leaves from the upstream installation, it is no doubt that this shipment will be sent to the consecutive downstream installation. However, this is not the case in OWMR models where the units arriving at the warehouse in one shipment may be sent to different retailers. The allocation policy in case of stock shortage at the warehouse is absent in the series system. Specifically, if the on-hand inventory at the warehouse is not enough to satisfy all the ordering requests from retailers, then it is the allocation policy at the warehouse that determines how many units are sent to each requesting retailer. In this way, the allocation policy determines inventory flows and hence affects the system-wide cost directly. In addition, in OWMR models the interaction between the warehouse and one retailer may affect the rest retailers due to the coupling effects among retailers. For example, after satisfying a retailer's order, the warehouse may not have enough stock to fully satisfy the subsequent ordering requests from other retailers. The coupling effects complicate the whole system process and make tracking the system status difficult.

To derive a performance guarantee, it is necessary to identify an upper bound for the proposed policy and a lower bound for the optimal policy. In our case, we take advantage of a lower bound established by Chen and Zheng [19], who decompose the original OWMR system into $N+1$ subsystems with each location/installation corresponding to a single-stage subsystem. This lower bound can be constructed by assuming that each individual subsystem (i.e., installation) runs independently. Moreover, this lower bound is easily computable, because it is the summation of $N+1$ single-stage optimal solutions. Then, we adopt a novel approach to derive an upper bound of the system-wide cost under an arbitrary modified echelon (r, Q) policy without exactly computing it. Based on the upper bound, we identify a heuristic modified echelon (r, Q) policy with the values of r and Q for each installation optimizing a single-stage subsystem. The computational procedure takes the virtue of a standard single-stage (r, Q) system.

For performance bound analysis, we compare the cost upper bound of our heuristic policy with the lower bound of the optimal cost established by Chen and Zheng [19]. The actual performance of our heuristic policy must be better than this provable performance bound. Specifically, we show that the performance of our heuristic policy for an OWMR model is guaranteed to be within $\frac{\sum_{i=1}^{N} C_i^* + \hat{C}_0^*}{\sum_{i=1}^{N} C_i^* + C_0^*}$ times the optimal cost (i.e., $\frac{\sum_{i=1}^{N} C_i^* + \hat{C}_0^*}{\sum_{i=1}^{N} C_i^* + C_0^*}$-optimal hereafter), where C_i^*, C_0^* and $\hat{C}_0^*$ are the optimal costs for single-stage problems and can be easily computed as shown in Zheng [97]. We also provide an alternative performance bound for our heuristic policy, which is determined by the heuristic (r, Q) values of the warehouse and a *specific* retailer. In addition, we provide asymptotic optimality results for our modified echelon (r, Q) heuristic policy when a pair of system parameters, such as fixed setup, holding and shortage costs, are scaled up. It is worthwhile noting that all the theoretical results of our heuristic hold regardless of the sequence in which the warehouse fills orders when there are multiple backlogged retailers (including first-come, first-served as a special case) as long as each retailer's order is being fulfilled as much as possible. Lastly, numerical examples further demonstrate that the proposed (r, Q) policy performs well in most instances and tends to outperform the echelon-stock (r, nQ) heuristic policy studied in Chen and Zheng [20].

Our contribution to the literature on stochastic distribution systems is twofold. First, on the technical side, we fill the gap in the literature on distribution systems by developing a heuristic policy with a performance guarantee. The heuristic policy is easy to compute, as the solutions to several single-stage (r, Q) inventory problems, and moreover, it is asymptotically optimal. Second, on the implication side, the bounds and asymptotic results demonstrate the robustness of single-stage (r, Q) inventory policies. That is, (r, Q) policies based on single-stage systems with some adaption could perform very well even for a distribution system. Moreover, as the existing literature on distribution systems focuses on synchronization between the warehouse and retailers following the lineage of the power-of-two policy, we introduce a class of policies for the distribution system that do not require synchronization, but still perform well with provable performance guarantees.

3.2 Literature Review

The one-warehouse multi-retailer (OWMR) distribution system with stochastic demands has been studied extensively. Various heuristics and corresponding evaluation methods have been developed. For continuous-review OWMR models with one-for-one replenishment policies (or equivalently, order-up-to-S policies), Sherbrooke [82] provides the so-called METRIC (multi-echelon technique for recoverable item control) approximation to evaluate the system-wide cost. This method approximates the retailer lead time as the transportation times plus the average delay at the warehouse due to possible shortage. Graves [39] provides a 2-moment approximation for cost evaluation and numerically shows that this approximation is more accurate than the METRIC which uses only the first moment. Axsäter [3] provides a simple method to evaluate the inventory costs for OWMR systems with independent Poisson demands and one-for-one replenishment policies. This method keeps track of each supply unit as it moves through the system, and then characterizes the time between the placement of an order and the occurrence of its assigned demand unit.

For general batch ordering policies, Duermeyer and Schwarz [27] use the METRIC type approach for batch ordering retailers. Chen and Zheng [20] study an echelon stock (r, Q) policy, and provide exact results and approximations for Poisson and compound Poisson process demands, respectively. Axsäter [4] studies a similar system but uses a very different method to evaluate the system-wide cost exactly with compound Poisson demands. Installation or echelon stock (r, Q) policies with Poisson or compound Poisson demand processes, are also studied in Forsberg [36], Axäster [2], Axsäter [5]. However, all of these papers rely on the key assumption that there exist *no* explicit setup costs at any stage.

For OWMR models, only a few papers take setup costs into consideration.

Clark and Scarf [23] is perhaps the first work dealing with multi-echelon inventory systems with fixed setup costs. The authors provide a cost evaluation method for an OWMR system, under the so-called *balance* assumption. The balance assumption allows *free* lateral transshipments between all retailers, and thus the warehouse stock can be used more efficiently than what is possible in reality. Under this assumption, Federgruen and Zipkin [30] provide a lower bound on the minimum cost for a distribution system where only the warehouse has a setup cost.

Chen and Zheng [19] extend this work to a more general OWMR model with setup costs at all stages. They establish a lower bound on the minimum system-wide cost by decomposing the original system into several subsystems and assuming each subsystem runs independently of each other. Shang et al. [80] consider a distribution inventory system with an (S, T) policy. An (S, T) policy operates as follows: Installation i (either the warehouse or a retailer) reviews its echelon order inventory position every T_i time units and orders up to a base-stock level S_i. The authors develop an evaluation scheme and provide a method to obtain the optimal base-stock level and reorder intervals. For convenience of implementation and coordination across stages, they focus on *synchronized* policies, i.e., when the warehouse receives a shipment at the beginning of a system's order cycle, all retailers place an order. Yet, there remains a lack of rigorous argument showing that the optimal policy for a distribution system with stochastic demands has to satisfy the synchronized property. As observed from Shang et al. [80], even for a synchronized (S, T) policy, the system-wide cost expressions are too complex to be used for further performance bound analysis. In this work, under any given modified echelon (r, Q) policy, we provide an easy-to-compute upper bound on the system-wide cost, which is amenable for performance bound analysis.

Although most of the heuristic policies proposed in the literature make intuitive sense, it is still unclear how much the gap is between a heuristic solution and an optimal one. Roundy [76] shows that for the OWMR model with deterministic constant demands, the performance gap between an optimal power-of-two policy and the optimal policy is guaranteed to be within 2%. Other works that analyze performance bounds for deterministic OWMR models include Chen [17], Chan et al. [15] and Levi et al. [62]. Chu and Shen [22] focus on power-of-two ordering policies for a periodic-review distribution inventory system with target service levels under demand uncertainty. The authors show that the proposed heuristic is guaranteed to be 1.26-optimal, compared with the optimal power-of-two policy. Other than this work, we are not aware of other works on performance bound analysis of heuristic policies for a stochastic distribution inventory system.

The paper most closely related to our work is Hu and Yang [47]. The authors propose a heuristic policy in the class of MERQS policies for a continuous-review *serial* inventory system with Poisson demand arrivals. Under their heuristic, the replenishment to the next downstream stage is based on the echelon inventory position of the current stage. If the echelon inventory position is less than a specified level, a shipment is sent to the next downstream to raise its echelon inventory position as close as possible to the order-up-to level. In their analysis, they introduce the definitions of cycles and regular (irregular) shipment periods for each stage. Under their MERQS heuristic for a serial system, the number of shipment periods within one cycle must be an integer. That is, the beginning of one cycle is also the beginning of a shipment period, and the end of this cycle is also the end of the same or another shipment period. This implies that any shipment period associated with a cycle cannot exceed the range of this cycle. However, in an OWMR distribution system, following their definitions, one shipment period of a retailer may exceed a cycle. That is, at the beginning of one cycle, a shipment period belonging to the last cycle might not have been completed, and the last shipment period associated with this cycle may end in one of the subsequent cycles. In addition, as their MERQS heuristic is developed for a serial system, the problem of rationing limited inventory among retailers does not emerge.

In the current work, motivated by MERQS, we develop a class of modified echelon (r, Q) policy for distribution systems (referred to as MERQD). To tackle the aforementioned emerging problems for a distribution system, we treat all retailers as a whole stage, based on which we define *replenishment* and *depletion cycles* with respect to the warehouse. Moreover, we also introduce the definition of *shipment interval* and categorize each shipment interval into two types: *irregular* and *regular*. Based on the inventory position at the beginning of an irregular interval, we further categorize irregular shipments into two types: types I and II. Those definitions are specifically introduced for the distribution system. Unlike counting the number of irregular shipment periods in MERQS for the series system, we focus only on the number of type II irregular shipments within each depletion cycle. Unlike most methods in the literature, we do not compute the exact system-wide cost; instead, we provide an upper bound on the system-wide cost. Compared with the various complicated exact-cost expressions in the literature, the advantage of this cost upper bound is that it is ready for further performance guarantee analysis.

3.3 Model

We formally set up the model and introduce key assumptions. We also review a lower bound and some of its properties established by Chen and Zheng [19], which we use in our subsequent analysis.

3.3.1 Notation and Formulation

We consider a firm that manages a two-echelon distribution inventory system consisting of one warehouse and N retailers (hereafter the OWMR model). For notation convenience, we use Retailer i ($i = 1, 2, \ldots, N$) to denote a specific retailer, and Installation i ($i = 0, 1, \ldots, N$) to denote a specific installation, which can be either the warehouse $i = 0$ or a specific Retailer i. Retailers are replenished from the warehouse, which in turn is replenished from an outside supplier with *unlimited* stock. Retailer i faces a stochastic demand following a Poisson process with stationary rate λ_i. Demands among retailers are assumed to be independent. For Installation i, we denote by $D_i(t, t+\tau]$ the total demand of Installation i over the time interval $(t, t+\tau]$. Specifically, for the warehouse, $D_0(t, t+\tau] = \sum_{i=1}^{N} D_i(t, t+\tau]$ represents the total demand of all retailers over the time interval $(t, t+\tau]$. Let $\lambda_0 \equiv \sum_{i=1}^{N} \lambda_i$. There is a constant lead time $L_i > 0$ for Installation i. In other words, any shipment sent out to Installation i at time t will be received by Installation i at time $t + L_i$. Each *shipment* to Installation i incurs a fixed cost K_i. Without loss of generality, we assume that the variable ordering cost is zero. Let $h_i > 0$ be the echelon holding cost rate at Installation i. Whenever Retailer i runs out of stock, the unmet demand is fully backlogged with a backlog cost rate $p_i > 0$. The firm's objective is to determine a shipment policy that minimizes the long-run average system-wide cost.

As mentioned earlier, we apply a lower bound that has been established in the literature. To proceed, we first review some commonly used concepts. The *echelon inventory* at Installation i is the inventory on hand at Installation i plus the inventories at or in transit to all its downstream stages. (Note that for a retailer, the echelon inventory is merely its on-hand inventory.) The *echelon inventory level* at a retailer is the echelon inventory at that retailer minus the number of customers back-ordered at the retailer.

The echelon inventory level at the warehouse is the echelon inventory at the warehouse minus the total number of customers back-ordered at all retailers. The *echelon inventory position* at Installation i is the sum of the echelon inventory level at Installation i and the inventories in transit to Installation i. For Installation i, define the following inventory variables at time t:

$I_i(t)$ =echelon inventory at Installation i;

$IL_i(t)$ =echelon inventory level at Installation i;

$IP_i^-(t)$ =echelon inventory position before a shipment is sent to Installation i at time t;

$IP_i(t)$ =echelon inventory position after a shipment is sent to Installation i at time t;

$q_i(t)$ =shipping quantity to Installation i;

$B_i(t)$ =backorder inventory level at Installation i, and $B_0(t) = \sum_{i=1}^{N} B_i(t)$;

$OI_i(t)$ =on-hand inventory at Installation i before a shipment is sent to its successor at time t.

The system dynamics are expressed as follows:

$$I_i(t) = IL_i(t) + B_i(t),$$

$$IL_i(t + L_i) = IP_i(t) - D_i(t, t + L_i],$$

$$\sum_{i=1}^{N} IP_i^-(t) = IL_0(t) - OI_0(t), \tag{3.1}$$

$$IP_i(t) = IP_i^-(t) + q_i(t), \tag{3.2}$$

where $i = 0, 1, \ldots, N$. The first equation follows from the definition. The second equation follows because the inventory position of Installation i at time t can be decomposed into two components: the echelon inventory level at Installation i at time $t+L_i$ and the total demand over $(t, t+L_i]$. The third equation follows because the echelon inventory level at the warehouse is the sum of the on-hand inventory at the warehouse and inventory positions at all retailers before any replenishment is sent. The last equation captures the dynamics of inventory positions.

Several remarks are in place. First, at any time t, the shipping quantity to Installation $q_i(t)$ in (3.2) is a decision variable, which in turn can be substituted by the inventory position $IP_i(t)$ as an alternative decision variable. Second, the shipping quantity $q_i(t)$ to Retailer i, $i = 1, 2, \ldots, N$, should be non-negative, i.e., $q_i(t) \geq 0$, and the sum of shipping quantities to retailers should be capped by the on-hand inventory at the warehouse, i.e., $\sum_{i=1}^{N} q_i(t) \leq OI_0(t)$. Third, equations (3.1) and (3.2) jointly imply constraints $\sum_{i=1}^{N} IP_i(t) \leq IL_0(t)$ and $IP_i^-(t) \leq IP_i(t)$, for any $i = 0, 1, 2, \ldots, N$.

There are three cost components: the fixed setup cost (for each shipment), inventory holding cost and backlog cost. For Installation i at time t, the fixed setup cost $K_i\delta(q_i(t) > 0) = K_i\delta(IP_i(t) > IP_i^-(t))$ is incurred, where $\delta(\cdot)$ is an indicator function. The system-wide inventory holding and backlog cost rate at time t can be computed as follows:

$$
\begin{aligned}
& h_0 I_0(t) + \sum_{i=1}^{N} [h_i I_i(t) + p_i B_i(t)] \\
=& h_0[IL_0(t) + \sum_{i=1}^{N} B_i(t)] + \sum_{i=1}^{N} [h_i I_i(t) + p_i B_i(t)] \\
=& h_0 IL_0(t) + \sum_{i=1}^{N} [h_i I_i(t) + (p_i + h_0) B_i(t)]. \qquad (3.3)
\end{aligned}
$$

The original problem with the total controllable long-run average cost can thus be formulated as follows:

$$
\begin{aligned}
(\mathcal{B}) : C^*_{\mathcal{B}} \equiv \min_{\{IP_i(t)\}} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\Big[\int_{t=0}^{T} & \sum_{i=0}^{N} K_i \delta(IP_i(t) > IP_i^-(t)) + h_0 IL_0(t) \\
& + \sum_{i=1}^{N} [h_i I_i(t) + (p_i + h_0) B_i(t)] dt \Big] \\
\text{s.t.} \quad & IP_i^-(t) \le IP_i(t), \ i = 0, 1, \ldots, N; \quad \sum_{i=1}^{N} IP_i(t) \le IL_0(t).
\end{aligned}
$$

Because the optimal policy of such a system, even if it exists, must be extremely complicated, we focus on a class of modified echelon (r, Q) policies as follows.

Definition 3.3.1 (Modified Echelon ($\mathbf{r}$,$\mathbf{Q}$) Policy)**.** The upstream installation ships to a downstream installation on the basis of its observation of the echelon inventory position at the downstream installation. In particular, if the echelon inventory position at Installation i is at or below r_i and the upstream has positive on-hand inventory, then a shipment is sent to Installation i to raise its echelon inventory position *as close as possible* to $r_i + Q_i$.

A critical issue in distribution systems is the *allocation policy* in the case of shortages at the warehouse. Under a modified echelon ($\mathbf{r}$,$\mathbf{Q}$) policy, the allocation policy is consistent with the first-come-first-served rule, i.e., the sequence of the actual shipments to retailers is consistent with that of the retailers' inventory positions reaching their reorder points. In particular, once Retailer i's inventory position is at or below the reorder point r_i, the warehouse must use the on-hand inventory to raise Retailer i's inventory position as close as possible to the order-up-to level $r_i + Q_i$. Therefore, it would never occur that on one hand the warehouse has some on-hand inventory, but on the other hand it keeps Retailer i's inventory position below r_i. Because the external supplier has ample inventory, a modified echelon (**r**,**Q**) policy at the warehouse is in fact a standard echelon (**r**,**Q**) policy with parameters (r_0, Q_0), whereas due to the possible shortage at the warehouse, the modified echelon (**r**,**Q**) policy at retailers behaves differently from the standard (r, Q) policy.

3.3.2 Lower Bound for OWMR Model

We adopt the so-called *induced-penalty bound* on the optimal system-wide cost $C^*_{\mathcal{B}}$, which was established by Chen and Zheng [19]. The authors decompose the stock (each unit) at Retailer i into two components: a common component 0 and a retailer-specific component i. A lower bound of the original system can then be constructed by assuming that the components can be replenished and sold independently. In this sense, each component corresponds to an independent subsystem.

For $i = 1, 2, \ldots, N$, define

$$G_i(y) = \mathbb{E}[h_i(y - D_i(t, t + L_i])^+ + (h_0 + p_i)(y - D_i(t, t + L_i])^-]. \tag{3.4}$$

Note that $G_i(\cdot)$ is convex. Consider a single-stage system with a setup cost K_i, a cost-rate function $G_i(\cdot)$, and a Poisson demand process. Because $G_i(\cdot)$ is convex, the (r, Q) policy is optimal for such a system. Let (r_i^*, Q_i^*) be the optimal (r, Q) policy and C_i the minimum cost. Define

$$G_i^i(y) = \begin{cases} C_i & \text{if } y \le r_i^*, \\ G_i(y) & \text{otherwise,} \end{cases}$$

and

$$G_i^0(y) = G_i(y) - G_i^i(y) = \begin{cases} G_i(y) - C_i & \text{if } y \le r_i^*, \\ 0 & \text{otherwise.} \end{cases}$$

Here $G_i^0(\cdot)$ represents an induced-penalty cost charged to the warehouse, which may be interpreted as the additional cost beyond the absolute minimum at Retailer i due to the warehouse's inability to raise IP_i above r_i^*.

The "system" of component $i = 1, 2, \ldots, N$ is a single-stage system with a setup cost K_i and a loss rate function $G_i^i(\cdot)$, and thus the optimal cost of component i, denoted as C_i^*, can be easily computed. Zheng [97] studies a continuous-review, single-stage inventory model with a constant lead time and a fixed setup cost; see Appendix D. Let L be the shipment lead time and $G(y)$ denote the expected single-stage cost rate incurring at time $t+L$, when the inventory position equals y at time t. They impose the following regularity assumption, and the optimal (r, Q) policy can be efficiently computed.

Assumption 3.3.1 (REGULARITY). (i) $G(y)$ is a convex function and $\lim_{y\to\pm\infty} G(y) = \infty$. (ii) There exist $a > 0$ and $b < 0$, such that $\lim_{y\to+\infty} G'(y) = a$ and $\lim_{y\to-\infty} G'(y) = b$.

The "system" of component 0 is a special OWMR model where only the warehouse has a setup cost K_0 and the loss rate functions at retailers are the induced-penalty cost functions $G_i^0(\cdot)$. Define

$$R(y) \equiv h_0 y + \min_{y_i : \sum_{i=1}^N y_i \le y} \sum_{i=1}^N G_i^0(y_i), \tag{3.5}$$

$$G_0(z) \equiv \mathbb{E}[R(z - D_0(t - L_0, t])], \tag{3.6}$$

$$C_0(r_0, Q_0) \equiv \frac{\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} G_0(y)dy}{Q_0}. \tag{3.7}$$

The minimization in (3.5) is effectively a free inventory position re-balance, which is called a *balancing* assumption in Federgruen and Zipkin [29, 30, 31]. Therefore, given $IL_0(t) = y$, the expected holding and backlogging costs of component 0 at time t are at least $R(y)$. Then, a lower bound on the minimum cost of component 0 can be expressed as (3.7).[1] This lower bound, denoted as C_0^*, can be easily computed; see also Zheng [97]. Specifically, let $(r_0^*, Q_0^*) = \arg\min_{(r_0,Q_0)} C_0(r_0, Q_0)$. Because $G_0(\cdot)$ satisfies Assumption 3.3.1, by Theorem 1 in Zheng [97], C_0^* can be rewritten as

$$\begin{aligned} C_0^* = \min_{(r_0,Q_0)} C_0(r_0, Q_0) &= \min_{(r_0,Q_0)} \frac{\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} G_0(y)dy}{Q_0} \\ &= G_0(r_0^*) = G_0(r_0^* + Q_0^*). \end{aligned} \tag{3.8}$$

Lemma 3.1 (Chen and Zheng [19]). *A lower bound for the original system is* $\underline{C}^* = \sum_{i=0}^N C_i^*$.

Before ending this section, we must point out that C_0^* may take negative values. To see this, note that the holding cost of this single-stage system of component 0 is assessed on the echelon inventory level of the warehouse, other than in a traditional way, assessed on the on-hand inventory; see (3.5)-(3.7). In some extreme cases, e.g., when K_0 and L_i's are quite small, C_0^* can be negative. The following assumption provides a sufficient condition for C_0^* to stay positive.

Assumption 3.3.2. $A \equiv h_0 \sum_{i=1}^N (\lambda_i L_i - \frac{C_i^*}{h_0+p_i}) + \sqrt{\frac{2\lambda_0 K_0 h_0 \underline{p}}{h_0+\underline{p}}} > 0$, where $\underline{p} = \min_{i=1,2,\ldots,N} p_i$.

Lemma 3.2. *Assumption 3.3.2* $\Rightarrow C_0^* > 0$.

The term A in Assumption 3.3.2 is the deterministic counterpart of C_0^* in an inventory system with fluid customer arrivals. Based on Jensen's inequality, it is well known that the inventory costs calculated in the deterministic model underestimate the actual inventory costs with stochastic customer arrivals. Therefore, the optimal cost of the deterministic system provides a

[1] We adopt the convention that discrete units of inventories can be approximated by continuous variables. Such an approximation is widely used in the inventory literature (see Zheng [97] for more discussion).

lower bound on that of the stochastic inventory system with the same system primitives.

3.4 Analysis

Consider that the OWMR system operates under an arbitrarily given modified echelon $(\mathbf{r}, \mathbf{Q})$ policy. Because the size of the shipment to each retailer is physically constrained by the on-hand inventory of the warehouse, Retailer i may not always be able to raise its echelon inventory position to the desired level, $r_i + Q_i$. In that case, the warehouse will feed Retailer i as much as possible. The size of a shipment to Retailer i can be larger or smaller than Q_i. For instance, if the warehouse does not have enough on-hand inventory when Retailer i reaches the reorder point, the shipment can be smaller than Q_i; however, if the warehouse has run out of stock for a while, the size of the shipment sent to Retailer i after the warehouse is back in stock can be larger than Q_i.

Let $C(\mathbf{r}, \mathbf{Q})$ denote the long-run average system-wide cost under the modified echelon $(\mathbf{r}, \mathbf{Q})$ policy. As such a policy is a feasible solution to Problem $(\mathcal{B})$, $C(\mathbf{r}, \mathbf{Q})$ provides a cost upper bound on $C^*_{\mathcal{B}}$.

Observation 3.4.1. For any modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, $C^*_{\mathcal{B}} \leq C(\mathbf{r}, \mathbf{Q})$.

3.4.1 An Upper Bound on $C(\mathbf{r}, \mathbf{Q})$

Given a modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, it is difficult to calculate its long-run average system-wide cost. To overcome this, we propose a novel approach to obtain an upper bound on $C(\mathbf{r}, \mathbf{Q})$. First, we adopt the following cost-accounting scheme (see also Chen and Zheng 19).

Definition 3.4.1 (Cost Accounting Scheme)**.** At time t, we charge the inventory holding cost of the warehouse incurred at time $t + L_0$, and charge the inventory holding and backlog cost of Retailer i incurred at time $t + L_0 + L_i$. In addition, we charge the inventory cost of Retailer i in the form of (3.4) assessed on its inventory position, and charge the warehouse's inventory cost assessed on its echelon inventory level.

As such a cost accounting scheme only shifts costs across time points, the long-run average inventory holding and backlog costs are not affected. The rationale behind it is that an order placed by the warehouse at time t does not affect the inventory holding of the warehouse until time $t + L_0$, and does not affect the inventory holding and backlog costs of Retailer i until time $t + L_0 + L_i$ or later.

Replenishment and Depletion Cycles.

In the subsequent analysis, for notation convenience, we use Installation I to denote the union set of all retailers. A shipment is said to be shipped to Installation I, if it is sent from the warehouse to one of the retailers.

Definition 3.4.2 (Replenishment Cycle and Depletion Cycle)**.** For $j \in \mathbb{N}$, we call $[T_0^j, T_0^{j+1})$ the jth replenishment cycle, where T_0^j is the time epoch of the 1st unit, contained in the jth order of the warehouse, being sent to the warehouse. Similarly, we call $[T_I^j, T_I^{j+1})$ the jth depletion cycle, where T_I^j is the time epoch of the 1st unit, contained in the jth order of the warehouse, being sent to Installation I, i.e., one of the retailers.

Since the depletion cycle is our main focus due to its complexity and is frequently mentioned in the subsequent analysis, we simply use "cycle" to denote a depletion cycle, unless otherwise specified. Note that a depletion cycle $[T_I^j, T_I^{j+1})$ may be an empty set. In that case, all the units contained in the jth order of the warehouse are shipped to a retailer together with one or multiple units in the $(j + 1)$th order of the warehouse. Our analysis focuses on *non-empty* cycles because retailers incur no costs for empty cycles.

Regular and Irregular Shipments.

Over any non-empty cycle $[T_I^j, T_I^{j+1})$, the jth order of the warehouse is shipped to retailers in one or multiple, say $M \in \mathbb{N}$, shipments in total. Of these M shipments, suppose $M_i \in \mathbb{N}$ shipments are sent to Retailer i. Then, $\sum_{i=1}^{N} M_i = M$. Let $T_i^{j,m}$ be the time of the mth shipment sent to Retailer i over cycle $[T_I^j, T_I^{j+1})$, where $m = 1, 2, \ldots, M_i$. By definition, we have $T_I^j \leq T_i^{j,1} \leq \cdots \leq T_i^{j,M_i} < T_I^{j+1}$ for any i. Define $T_i^{j,M_i+1} \equiv T_i^{j+1,1}$. We call $[T_i^{j,m}, T_i^{j,m+1})$ the mth *shipment interval* of Retailer i over the cycle $[T_I^j, T_I^{j+1})$. Note that for the case with $M_i = 0$, i.e., when no shipment is sent to Retailer i, we define $T_i^{j,1} \equiv T_i^{j+1,1}$.

It is easy to see that for $m = 1, \ldots, M_i - 1$, the shipment interval $[T_i^{j,m}, T_i^{j,m+1})$ resides within the cycle $[T_I^j, T_I^{j+1})$. However, the last shipment interval incurred in this cycle, $[T_i^{j,M_i}, T_i^{j,M_i+1})$, may exceed the cycle. Actually, the beginning of this interval must be within this cycle, i.e., $T_i^{j,M_i} \leq T_I^{j+1}$, which is the reason why we associate this interval with the cycle $[T_I^j, T_I^{j+1})$. However, the end of this shipment interval, T_i^{j,M_i+1}, may be *outside* the cycle. This occurs when after the shipment at T_i^{j,M_i}, no shipment is sent to Retailer i over $[T_i^{j,M_i}, T_I^{j+1})$ and the next shipment to Retailer i occurs at $T_i^{j+1,1}$, which is strictly later than T_I^{j+1}. This obviously differs from the notion of the shipment period defined for the serial system by Hu and Yang [47]. There in the series system, the cycle of the upstream stage starts at the beginning of one shipment period of the downstream stage, and ends at the end of another shipment period, i.e., the shipment periods in one cycle cannot exceed the range of its associated cycle.

Depending on the retailer's inventory positions at the beginning and end of a shipment interval, we categorize shipments to retailers and their associated shipment intervals into the following three types.

Definition 3.4.3 (Regular and Irregular Shipment (Interval)).

For a shipment interval $[T_i^{j,m}, T_i^{j,m+1})$, if $IP_i(T_i^{j,m}) = r_i + Q_i$ and $IP_i^-(T_i^{j,m+1}) = r_i$, we call it a *regular* shipment interval of Retailer i; otherwise, we call it an *irregular* shipment interval. In particular, if $IP_i(T_i^{j,m}) = r_i + Q_i$ and $IP_i^-(T_i^{j,m+1}) < r_i$, we call it a *type I irregular* shipment interval; if $IP_i(T_i^{j,m}) < r_i + Q_i$, we call it a *type II irregular* shipment interval. The shipment associated with a regular (type I or II irregular) shipment interval is called a regular (type I or II irregular) shipment.

Note that whether a shipment interval is regular or not depends on the retailer's inventory positions at the beginning and end of this shipment interval. Instead, the type of irregular shipment interval depends only on the inventory position at the beginning of this shipment interval.

To better illustrate this, we plot in Figure 3.1 one scenario of two retailers' inventory positions over the cycle $[T_I^j, T_I^{j+1})$. As shown in Figure 3.1, at the beginning of the cycle $[T_I^j, T_I^{j+1})$, the first shipment over this cycle is sent to Retailer 1 at time epoch $T_I^j = T_1^{j,1}$. The first shipment to Retailer 2 over this cycle occurs at time epoch $T_2^{j,1}$, and the interval $[T_I^j, T_2^{j,1})$ belongs to the last shipment interval associated with the previous cycle. At time epoch T_2^{j,M_i}, the inventory position of Retailer 2 drops to r_2 and thus a shipment from the warehouse is sent to Retailer 2. However, because the warehouse does not have enough on-hand inventory, Retailer 2's inventory position after this shipment is still less than $r_2 + Q_2$. Therefore, the associated shipment interval $[T_2^{j,M_i}, T_2^{j+1,1})$ is a type II irregular shipment interval. Moreover, at the first time after T_1^{j,M_i}, when Retailer 1's inventory position drops to r_1, the warehouse has used up all its on-hand inventory; therefore, Retailer 1's inventory position cannot be raised up to $r_1 + Q_1$ until time epoch $T_1^{j+1,1}$, when the warehouse has enough on-hand inventory due to the arrival of the $(j+1)$th order of the warehouse. Therefore, the last shipment interval $[T_1^{j,M_i}, T_1^{j+1,1})$ of Retailer 1 over this cycle is a type I irregular shipment interval.

A key observation is that for any Retailer i, the first $M_i - 1$ (possibly zero) shipments in any cycle with M_i shipments must be regular. This can be seen by the definition of the modified echelon (r, Q) policy. Imagine two shipments to Retailer i are associated with the same order of the warehouse and that the earlier one does not raise Retailer i's inventory position to the desired level of $r_i + Q_i$. Then, units in the later shipment should be moved to the earlier shipment as much as possible, by the way of how the modified echelon (r, Q) policy works. As a result, either the earlier one becomes a regular shipment or there is only one shipment in that cycle. Therefore, if there are a number of $M_i > 1$ shipments in a cycle, the first $M_i - 1$ (possibly zero) shipments must be regular. However, the last (possibly the only) shipment may not be regular for two possible reasons. First, at the beginning of the last shipment interval, the warehouse may not have enough on-hand inventory to raise Retailer i's inventory position to $r_i + Q_i$, i.e., $IP_i(T_i^{j,M_i}) < r_i + Q_i$. Second, at the end of the last shipment interval, the warehouse may be out of stock and thus Retailer i's inventory position may drop below r_i, i.e., $IP_i^-(T_i^{j,M_i+1}) < r_i$. The two reasons are not exclusive. It is possible that $IP_i(T_i^{j,M_i}) < r_i + Q_i$ and $IP_i^-(T_i^{j,M_i+1}) < r_i$; see Retailer 2 in Figure 3.1.

The following lemma shows that the frequency of both irregular shipments for each retailer and type II irregular shipments for all retailers can be bounded by that of cycle.

Lemma 3.3 (Irregular Shipment Frequency). *(i) For any retailer, there exists at most one irregular shipment interval in each cycle, regardless of whether the cycle is empty or not.*

(ii) Across all retailers, there exists at most one type II irregular shipment interval in each cycle.

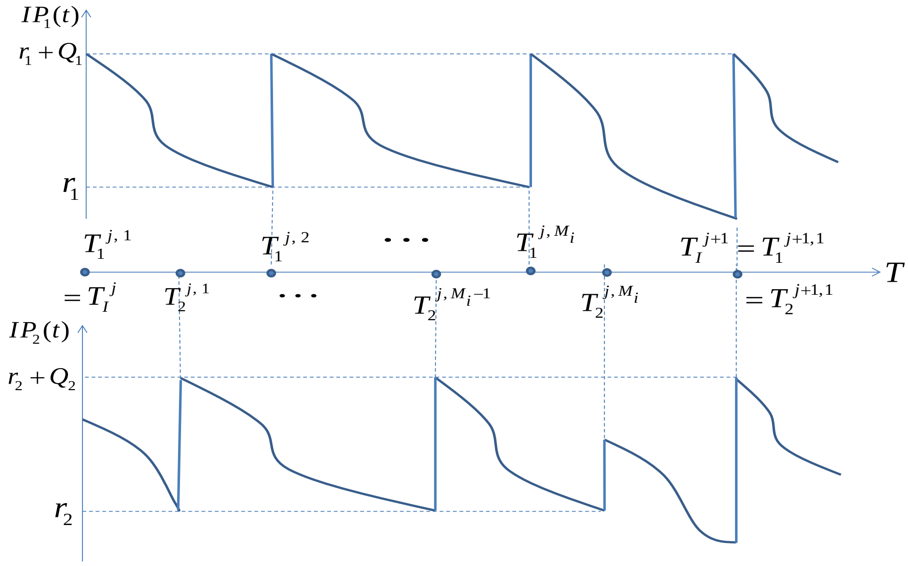

Figure 3.1: Illustration of two retailers' inventory positions over the cycle $[T_I^j, T_I^{j+1})$.

Cost Assessment in Cycles.

Remark 3.1 (Cost in regular and Irregular Shipment Intervals). *We make the following cost assessment.*

(i) The expected cost rate of the fixed, inventory holding and backlog costs at Retailer i in a regular shipment interval is $C_i(r_i, Q_i) = (1/Q_i)[\lambda_i K_i + \int_{r_i}^{r_i+Q_i} G_i(y)dy]$, where $G_i(y)$ is defined in (3.4). The cost rate is the same as that in the single-stage problem with an outside supplier of unlimited supply (Zheng 97).

(ii) For any time t in a non-empty type I irregular shipment interval of Retailer i, the inventory position first drops from $r_i + Q_i$ to r_i and then below r_i. In the former subinterval, the expected cost rate is $C_i(r_i, Q_i)$, whereas in the latter, the expected inventory holding and backlog costs accrue at a rate equal to $G_i(IP_i(t))$.

(iii) For any time t in a non-empty type II irregular shipment interval of Retailer i, the expected inventory holding and backlog costs accrue at a

rate equal to $G_i(IP_i(t))$; in addition, a fixed setup cost, K_i, is incurred for the irregular shipment.

Remark 3.1 states that the total costs of the system consist of four parts: (i) costs in regular shipment intervals, (ii) costs in type I irregular shipment intervals, (iii) inventory and backlog costs in type II irregular shipment intervals and (iv) setup costs in type II irregular shipment intervals. We provide cost upper bounds for the first three parts and the last part separately.

Cost Upper Bound.

We first bound the expected cost rate at all retailers, *excluding* the setup costs for type II irregular shipment intervals, i.e., the first three types of costs discussed in Remark 3.1.

Lemma 3.4. *For any time $t \in [T_I^j, T_I^{j+1}) \neq$, the expected cost rate at Retailer i, denoted by $\hat{\Gamma}_i(IL_0(t))$, which excludes the setup costs in type II irregular shipment intervals, is bounded as follows:*

$$\hat{\Gamma}_i(IL_0(t)) \leq \bar{\Gamma}_i(IL_0(t)) \equiv \begin{cases} \max\{G_i(IL_0(t) - \sum_{j\neq i} r_j + Q_j), G_i(w_i), C_i(r_i, Q_i)\} & \\ \qquad \text{if } IL_0(t) - \sum_{j\neq i} r_j + Q_j \leq r_i, & \\ \max\{G_i(w_i), C_i(r_i, Q_i)\} \qquad \text{otherwise}, & \end{cases}$$

where $w_i \equiv \arg\max_{r_i < z \leq r_i + Q_i}\{G_i(z)\}$.

As an immediate result of Lemma 3.4, we have the following corollary.

Corollary 3.2. *For any time $t \in [T_I^j, T_I^{j+1}) \neq$, the expected cost rate at all retailers, with the setup costs of type II irregular shipments excluded, denoted by $\hat{\Gamma}_I(IL_0(t))$, is bounded as follows:*

$$\hat{\Gamma}_I(IL_0(t)) \leq \bar{\Gamma}_I(IL_0(t)) \equiv \sum_{i=1}^{N} \bar{\Gamma}_i(IL_0(t)). \tag{3.9}$$

Although the upper bound in Corollary 3.2 is provided in terms of $IL_0(t)$, the system-wide cost can be bounded in terms of $IP_0(t)$, based on the relationship $IL_0(t+L_0) = IP_0(t) - D_0(t, t+L_0]$. Because the warehouse replenishes from a supplier with an unlimited supply, the inventory position $IP_0(t)$ at the warehouse in the steady state is uniformly distributed (see Zheng 97), which facilitates our analysis.

We now bound the setup costs associated with type II irregular shipments. By Lemma 3.3(ii), the frequency of incurring irregular shipment intervals is bounded by the frequency of depletion cycles. The following lemma characterizes the expected long-run average length for both the replenishment and depletion cycles.

Lemma 3.5 (Cycle Length). *Under any modified echelon ($\boldsymbol{r}$, $\boldsymbol{Q}$) policy, the long-run average expected cycle length for both replenishment and depletion cycles is Q_0/λ_0, i.e.,* $\lim_{j\to\infty} \mathbb{E}[(T_0^{j+1} - T_0^1)]/j = \lim_{j\to\infty} \mathbb{E}[(T_I^{j+1} - T_I^1)]/j = Q_0/\lambda_0$.

Remark 3.3. *Define*

$$\overline{K} = \max_{i=1,\ldots,N} \{K_i\}.$$

By Lemma 3.3(ii), over any cycle $[T_I^j, T_I^{j+1}) \neq$, the setup cost for type II irregular shipment is incurred once at most and is no more than $\overline{K}$. Consequently, the corresponding setup costs for type II irregular shipments accrue at a rate that is no more than $\overline{K}/(T_I^{j+1} - T_I^j)$. It follows from Lemma 3.5 that the long-run average setup cost for type II irregular shipments has an upper bound $\lambda_0\overline{K}/Q_0$.

For notation convenience, we define $\hat{C}_i(r_i, Q_i) \equiv C_i(r_i, Q_i)$, for $i = 1, 2, \ldots, N$.

Let D_0 denote the total demand over $(0, L_0]$. Moreover, we define

$$\hat{G}_i(y) \equiv \bar{\Gamma}_i(y) - \hat{C}_i(r_i, Q_i), \tag{3.10}$$

$$\Lambda_0(y) \equiv \mathbb{E}[h_0(y - D_0) + \sum_{i=1}^{N} \hat{G}_i(y - D_0)], \tag{3.11}$$

$$\hat{C}_0(r_0, Q_0) \equiv \frac{1}{Q_0}\Big[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\Big]. \tag{3.12}$$

Combining all the cost terms of all the installations, we are ready to present an upper bound on $C(\mathbf{r}, \mathbf{Q})$.

Theorem 3.4 (An Upper Bound). *For any given modified echelon* $(\boldsymbol{r}, \boldsymbol{Q})$ *policy, the long-run average system-wide cost has an upper bound:* $C(\boldsymbol{r}, \boldsymbol{Q}) \leq \sum_{i=0}^{N} \hat{C}_i(r_i, Q_i) + \lambda_0 \overline{K}/Q_0$.

By (3.9)-(3.12), $\sum_{i=0}^{N} \hat{C}_i(r_i, Q_i)$ is an upper bound on the expected warehouse cost and all retailers' costs, with the setup costs of type II irregular shipments excluded. By Remark 3.3, $\lambda_0\overline{K}/Q_0$ is an upper bound on the long-run average setup cost for type II irregular shipments.

3.4.2 Heuristic Policy

We propose the following heuristic modified echelon $(\mathbf{r}, \mathbf{Q})$ policy. In this heuristic policy, for Retailer i, we select $(r_i, Q_i) = (r_i^*, Q_i^*)$, which is the optimal policy for a single-stage system with setup cost K_i and loss rate function $G_i(\cdot)$ as defined in (3.4). With this selection, we have $\hat{C}_i(r_i, Q_i) = C_i(r_i^*, Q_i^*) = C_i^*$ and

$$\mathcal{C}_{\mathcal{B}}^* \leq C(\mathbf{r}, \mathbf{Q})|_{(r_i, Q_i)=(r_i^*, Q_i^*)} \leq \sum_{i=1}^{N} C_i^* + \hat{C}_0(r_0, Q_0) + \frac{\lambda_0 \overline{K}}{Q_0}. \tag{3.13}$$

Plugging (3.12) into (3.13), we have

$$C(\mathbf{r},\mathbf{Q})|_{(r_i,Q_i)=(r_i^*,Q_i^*)} \leq \sum_{i=1}^{N} C_i^* + \frac{1}{Q_0}\Big[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\Big] + \frac{\lambda_0 \overline{K}}{Q_0}. \quad (3.14)$$

We can further tighten the upper bound in (3.14) by minimizing r_0 and Q_0. Specifically, this can be achieved by solving the following optimization problem:

$$\min_{r_0,Q_0} \tilde{C}_0(r_0,Q_0) \equiv \min_{r_0,Q_0} \frac{1}{Q_0}\Big[\lambda_0(K_0+\overline{K}) + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\Big]. \quad (3.15)$$

Note that $\bar{\Gamma}_i(y)$ (see Lemma 3.4) may not be convex for any given r_i and Q_i ($i = 1, 2, \ldots, N$), but it is indeed so for $(r_i, Q_i) = (r_i^*, Q_i^*)$. Consequently, with $(r_i, Q_i) = (r_i^*, Q_i^*)$ for all i, $\hat{G}_i(y)$ in (3.10) and hence $\Lambda_0(y)$ in (3.11) are also convex functions.

Lemma 3.6. *$\Lambda_0(y)$ satisfies Assumption 3.3.1 with $a = h_0$ and $b = -\sum_{i=1}^{N}(p_i + h_0) + h_0$. In addition, $\hat{C}_0(r_0, Q_0)$ is a jointly convex function of r_0 and Q_0, and hence so is $\tilde{C}_0(r_0, Q_0)$.*

Problem (3.15) is a single-stage problem with a fixed setup cost equal to $K_0 + \overline{K}$, and its objective function $\tilde{C}_0(r_0, Q_0)$ is jointly convex in r_0 and Q_0. Therefore, the optimal solution can be efficiently computed (see Federgruen and Zheng 28). Define $(\tilde{r}_0^*, \tilde{Q}_0^*) \equiv \arg\min_{r_0,Q_0} \tilde{C}_0(r_0, Q_0)$. We construct a heuristic modified echelon $(\mathbf{r}, \mathbf{Q})$ policy as follows:

$$(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) = (\hat{r}_0, \hat{Q}_0, \hat{r}_1, \hat{Q}_1, \ldots, \hat{r}_N, \hat{Q}_N) = (\tilde{r}_0^*, \tilde{Q}_0^*, r_1^*, Q_1^*, \ldots, r_N^*, Q_N^*). \quad \text{(MERQD)}$$

As $\Lambda_0(\cdot)$ satisfies Assumption 3.3.1, we have by Theorem 1 in Zheng [97] that $\tilde{C}_0^* \equiv \tilde{C}_0(\tilde{r}_0^*, \tilde{Q}_0^*) = \Lambda_0(\tilde{r}_0^*)$. A tighter upper bound on the system-wide cost can be expressed as follows:

$$C_{\mathcal{B}}^* \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \equiv C(\mathbf{r}, \mathbf{Q})|_{(\mathbf{r},\mathbf{Q})=(\hat{\mathbf{r}},\hat{\mathbf{Q}})} \leq \sum_{i=1}^{N} C_i^* + \tilde{C}_0^*. \tag{3.16}$$

3.4.3 Performance Guarantee and Asymptotic Optimality

Recall that the lower bound of the original system can be expressed as $\underline{C}^* = \sum_{i=0}^{N} C_i^*$. Compared with the upper bound in (3.16), we can derive guaranteed bounds on the effectiveness of our heuristic policy. For notation convenience, let $(\hat{r}_0^*, \hat{Q}_0^*) = \arg\min_{r_0,Q_0} \hat{C}_0(r_0, Q_0)$ and $\hat{C}_0^* = \hat{C}_0(\hat{r}_0^*, \hat{Q}_0^*)$.

Theorem 3.5 (Performance Bounds for General Cases). *(i) The modified echelon* $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ *policy in (MERQD) is at least* $1+(\tilde{C}_0^* - C_0^*)/(\sum_{i=1}^{N} C_i^* + C_0^*)$*-optimal, i.e.,* $\frac{\sum_{i=1}^{N} C_i^* + \tilde{C}_0^*}{\sum_{i=1}^{N} C_i^* + C_0^*}$*-optimal. In addition, if* $C_0^* > 0$*, the modified echelon* $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ *policy in (MERQD) is at least* $\tilde{C}_0^*/C_0^*$*-optimal.*

(ii) Assume Assumption 3.3.2 holds. The modified echelon $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ *policy in (MERQD) is at least* $\max\{\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m} + \frac{1}{4}} + \frac{1}{2}, \frac{1}{\beta_2}\}$*-optimal, where* $m \in \arg\max_{i=1,\ldots,N}\{K_i\}$*,* $\beta_1 \equiv \hat{Q}_0^*/Q_m^*$ *and* $\beta_2 \equiv C_0^*/\hat{C}_0^* \leq 1$*.*[2]

Theorem 3.5(i) directly follows from comparing the induced penalty lower bound and the upper bound established in (3.16). Theorem 3.5(ii) provides a less tight performance bound that depends on the order size ratio $\beta_1 \equiv \hat{Q}_0^*/Q_m^*$, as well as the cost ratio $\beta_2 \equiv C_0^*/\hat{C}_0^*$. Figure 3.2 displays the contour plot of the performance bound as a function of β_1 and β_2. The identified heuristic performs well when the two ratios are large. Denote $f_0 \equiv \lambda_0/\hat{Q}_0^*$

[2]In addition, if $(\beta_2^2 - 1)(\beta_1\lambda_m) + \beta_2\lambda_0 \geq 0$, the heuristic policy in (MERQD) can be alternatively shown to be at least $1 + \lambda_0/(2(\beta_1\beta_2\lambda_m + \sqrt{(\beta_2^2 - 1)(\beta_1\lambda_m)^2 + \beta_1\beta_2\lambda_m\lambda_0}))$-optimal; see the proof of Theorem 3.5. Numerical results show that the bound in Theorem 3.5(ii) tends to perform better but not always.

and $f_i \equiv \lambda_i/Q_i^*$ for $i = 1, 2, \ldots, N$. Note that f_i represents the *replenishment frequency* of Installation $i = 0, 1, \ldots, N$ under our heuristic assuming each installation's replenishment can always be fulfilled. Then, the performance guarantee in Theorem 3.5(ii) can be rewritten as $\max\{1+\sqrt{\frac{\beta_0}{2\beta_2}+\frac{1}{4}}-\frac{1}{2}, \frac{1}{\beta_2}\}$, where $\beta_0 \equiv f_0/f_m$ is the ratio of two replenishment frequencies for single-stage systems. This implies that the theoretical performance bound can be expressed by a replenishment frequency ratio (β_0) and a cost ratio (β_2) for single-stage systems, all of which can be efficiently computed.

Corollary 3.6. *If $2(1/\beta_2 - 1) \geq \lambda_0/(\beta_1\lambda_m)$, then the modified echelon $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ policy in (MERQD) is at least $1/\beta_2$-optimal.*

Corollary 3.6 says that under some condition, the performance bound may depend only on β_2, which is the warehouse's cost ratio between the induced penalty cost and the incurred cost under our heuristic. This can also be observed in Figure 3.2. That is, when β_1 is relatively large such that $2(1/\beta_2 - 1) \geq \lambda_0/(\beta_1\lambda_m)$ holds, there exists a performance bound that becomes independent of β_1; see the contour lines of the performance bound in Figure 3.2 become flat for large values of β_1.

For the case with identical retailers, we can derive a sharper performance bound that does not explicitly depend on the arrival rates.

Theorem 3.7 (Performance Bounds for Identical Retailers)**.** *Supposes all retailers are identical. If Assumption 3.3.2 holds, the modified echelon $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ policy in (MERQD) is at least $\max\{\sqrt{\frac{1}{2\beta_1\beta_2}+\frac{1}{4}}+\frac{1}{2}, \frac{1}{\beta_2}\}$-optimal, where $\beta_1 \equiv \hat{Q}_0^*/Q_i^*$ is the same for any retailer i.*

The following theorem further demonstrates the asymptotic optimality of our heuristic if we take the dominant relationships of cost primitives to the extreme.

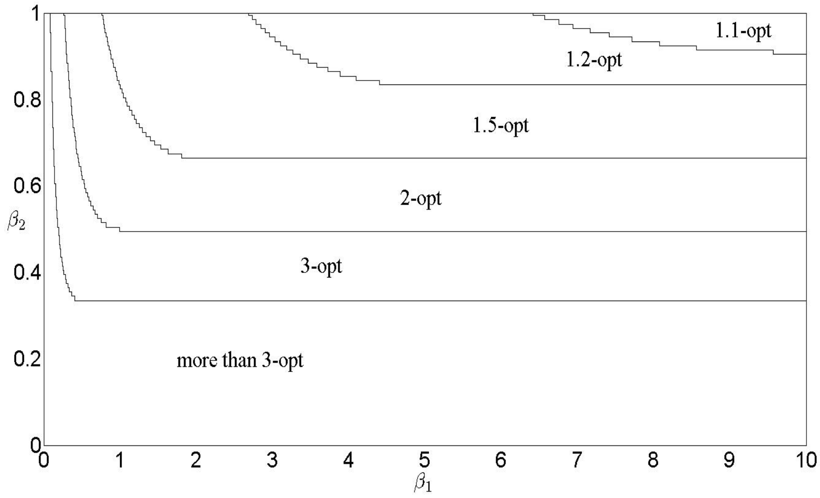

Figure 3.2: Performance bound in Theorem 3.5(ii) as a function of β_1 and β_2.

Theorem 3.8. *The modified echelon* $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{Q}})$ *policy in (MERQD) is asymptotically optimal if for any* $m \in \arg\max_{i=1,\ldots,N}\{K_i\}$*, one of the following conditions holds: (i)* $K_m > 0$ *and* $K_0/K_m \to \infty$*. (ii)* $h_0/h_m \to 0$*. (iii)* $h_0/p_m \to 0$*.*

Theorem 3.8 shows the asymptotic optimality of the heuristic based on the relationship between some system primitives. We make the following remarks on the relationships between system primitives in practical settings, indicating when the heuristic policy is more likely to have a good performance.

Remark 3.9 (Setup Costs)**.** *In an OWMR model where all retailers replenish from a warehouse, the fixed setup cost incurred at the warehouse tends to be much higher than that at each retailer.*

In a distribution system, there are usually larger economies of scales at the warehouse than at retailers. For instance, shipments to the warehouse are usually sent by sea or air cargo from suppliers who may be located far away from the warehouse, such as overseas. However, shipments from the warehouse to nearby retailers are usually sent by truck or van. As a result, shipments to the warehouse tend to incur a much larger cost than shipments to retailers.

Remark 3.10 (Holding Costs). *The echelon inventory holding cost at the warehouse tends to be much lower than that at each retailer.*

The echelon inventory holding usually includes financing and physical handing costs. Because the variable ordering cost at the warehouse is smaller than that at retailers in most cases, the financing cost at the warehouse is proportionally smaller. The physical handing cost also tends to be smaller at the warehouse due to its larger economics of scales. The warehouse is usually located in a suburban area and has a lower out-of-pocket inventory holding cost rate than a downtown retailer.

Remark 3.11 (Shortage Costs). *The inventory shortage cost at each retailer tends to be much larger than its holding cost. By virtue of Remark 3.10, the echelon inventory holding cost at the warehouse therefore tends to be even lower than the shortage cost at each retailer.*

As shown in Huh et al. [55], the ratio between shortage and holding costs is quite large and "[a]t a 25% markup, which is quite common in many retail environments, this ratio is at least 100." In addition to the direct impact of profit loss, shortages at retailers can lead to a loss of customer goodwill and have a long-term negative impact on the retailers' revenue.

Up till now, under the modified echelon (r, Q) policy, the allocation policy at the warehouse in case of shortage (i.e., not enough on-hand inventory at the warehouse to raise multiple retailers' inventory positions to the desired levels) is first-come, first-served. However, all the preceding results still hold for any sequence of serving backlogged retailers if there is more than one. The reason is as follows. Once Retailer i's inventory position drops to r_i

and the warehouse has on-hand inventory, the warehouse should use its on-hand inventory to raise Retailer i's inventory as close as possible to $r_i + Q_i$. However, if Retailer i's inventory position drops to r_i and the warehouse has no on-hand inventory, then Retailer i cannot get replenished and its inventory position will remain r_i or further drop below r_i. As time goes by, the inventory positions of other retailers may also drop to their reorder points, and due to the shortage at the warehouse, their inventory positions cannot be raised either. Suppose at time t a new shipment with Q_0 units arrives at the warehouse. Let J denote the set of retailers whose inventory positions are below their reorder points. In the case $\sum_{i \in J}(r_i + Q_i - IP_i(t)) > Q_0$, all those Q_0 units are immediately sent to retailers. Regardless of the sequence the retailers in set J are served as long as each retailer's order is fulfilled as much as possible, there will be at most one shipment to each retailer and there is at most one type II irregular shipment across all retailers, which implies that Lemma 3.3 still holds. In addition, Lemmas 3.4 and 3.5 are independent of the serving sequence. Therefore, our results are robust to the service sequence at the warehouse in the event of shortage. This property is desirable, as the warehouse may want to prioritize serving a retailer who has a higher backlog penalty cost or a higher expected sales volume.

3.5 Numerical Experiments

We report our numerical study in this section. With a set of comprehensive numerical experiments, we verify the effectiveness of our heuristic policy (including the asymptotic optimality demonstrated in Theorem 3.8) and test its sensitivity to the system parameters. Moreover, we compare our heuristic policy with the standard echelon $(\mathbf{r}, \mathbf{Q})$ policy studied in Chen and Zheng [20].

To test the effectiveness of the heuristic, we compare the upper bound of its system-wide cost, denoted by $\bar{C}$ (see Theorem 3.4), with the induced-penalty lower bound of the optimal cost, denoted by $\underline{C}^*$ (see Lemma 3.1). We define the following performance ratio:

$$\delta_1 \equiv \frac{\bar{C} - \underline{C}^*}{\underline{C}^*} \times 100\%,$$

which is an *upper bound* of the effectiveness of our heuristic policy.

To better ascertain the effectiveness of the heuristic policy, we also report the theoretical performance bounds (see Theorem 3.5 and Footnote 2). We define the following percentages:

$$\delta_2 \equiv \max\{\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m} + \frac{1}{4}} - \frac{1}{2}, \frac{1}{\beta_2} - 1\} \times 100\%.$$

$$\delta_3 \equiv 100\% \times \begin{cases} \min\{\delta_2, \dfrac{\lambda_0}{2(\beta_1\beta_2\lambda_m + \sqrt{(\beta_2^2-1)(\beta_1\lambda_m)^2 + \beta_1\beta_2\lambda_m\lambda_0})}\} & \\ \qquad\qquad \text{if } (\beta_2^2-1)(\beta_1\lambda_m) + \beta_2\lambda_0 \geq 0, & \\ \delta_2 & \text{otherwise.} \end{cases}$$

It follows from the proof of Theorem 3.5 that $\delta_1 \leq \delta_3 \leq \delta_2$ for each instance.

The performance bounds δ_1, δ_2 and δ_3 all use upper bounds of the system-wide cost to replace the real cost of the heuristic policy. It is technically challenging to exactly compute the real cost of the heuristic. To evaluate the exact performance of the heuristic, we use the Monte Carlo simulation method and compute the long-run average cost of the inventory systems under the heuristic policy. We denote by $\tilde{C}$ this real cost obtained by Monte Carlo simulation. We define the following percentage:

$$\delta_4 \equiv \frac{\tilde{C} - \underline{C}^*}{\underline{C}^*} \times 100\%,$$

which is again an upper bound on the performance gap of our heuristic but is presumably tighter than δ_1, δ_2 and δ_3. The complete test set of primitive values is given by $L_0 \in \{0, 1, 2\}$, $L_1 \in \{0, 1, 2\}$, $K_0 \in \{100, 200, 600\}$, $K_1 \in \{10, 20, 40\}$, $h_0 \in \{0.05, 0.1, 0.2\}$, $h_1 \in \{0.3, 0.5, 1\}$, $p_1 \in \{3, 5, 10\}$ and $\lambda_1 \in \{3, 5, 7\}$, with other primitives fixed as $N = 2$, $L_2 = 1$, $K_2 = 20$, $h_2 = 0.5$, $p_2 = 5$ and $\lambda_2 = 5$. By virtue of Huh et al. [55], the holding cost parameters are selected to be much smaller than the shortage costs. All combinations of these primitives provide $3^8 = 6561$ test instances.

The numerical results are summarized in Table 3.1. The average gap δ_1 between our provable cost upper bound in Theorem 3.4 and the induced penalty lower bound is about 9%. Table 3.1 shows that the maximum of δ_1 in our test is about 30%, i.e., the modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy in (MERQD) is provably guaranteed to be at least 1.3-optimal for our test instances. Moreover, the average gap δ_4 between the real cost and induced penalty lower bound is about 6%, with the minimum about 1.39% and the maximum about 20.66%. Axsäter et al. [8] and Gallego et al. [38] point out that the "balance" assumption which is used to derive the induced penalty lower bound may result in large errors from the optimal cost in some cases. In a numerical study, Doğru et al. [25] show that the "balance" assumption can indeed result in a lower bound far way from the optimal cost. This implies that our heuristic may perform much better than the performance bounds reported in Table 3.1, as the optimal policy is unknown and we benchmark the performance of our heuristic against the induced penalty lower bound.

	δ_1	δ_2	δ_3	δ_4
Average (%)	9.19	13.00	12.60	6.11
Standard deviation (%)	4.54	6.75	6.40	3.12
Minimum (%)	2.76	4.06	4.06	1.39
Maximum (%)	29.83	42.95	42.95	20.66

Table 3.1: Overview of performance of modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy.

3.5.1 Sensitivity to System Primitives

We now turn to investigate the impact of system primitives on the performance of the heuristic. To better measure the effect, we consider identical retailers in this subsection. The numerical results are reported in Tables 3.2 and 3.3. Table 3.2 confirms the effectiveness of our heuristic under various h_0 and K_0. It is observed that both β_1 and β_2 decrease as h_0 decreases or K_0 increases, and thus the theoretical performance bound δ_2 becomes smaller and the performance of our heuristic tends to perform better. Specifically, the heuristic has a close-to-optimal performance when $h_0 = 0.05$ and $K_0 = 600$, which is consistent with the asymptotic optimality shown in Theorem 3.8. Table 3.3 demonstrates the effectiveness of the heuristic under various values of p_i and λ_i. As shown in Table 3.3, as p_i increases, both the theoretical performance bounds and the gap between the real cost and the lower bound cost decrease, which means the heuristic performs better. It is also observed that the effectiveness does not seem monotone in λ_i.

h_0	K_0	r_i^*	Q_i^*	r_0^*	Q_0^*	$\hat{r}_0^*$	$\hat{Q}_0^*$	β_1	β_2	δ_1 (%)	δ_2 (%)	δ_3 (%)	δ_4 (%)
$L_0 = 1,\ L_i = 1,\ K_i = 10,\ h_i = 1,\ p_i = 5,\ \lambda_i = 3$													
1	10	1	10	7	13	17	18	1.30	0.58	46.49	75.23	71.78	30.79
	50	1	10	4	28	16	29	2.70	0.71	30.64	41.58	41.58	22.35
	100	1	10	2	39	15	39	3.70	0.76	24.86	31.94	31.94	18.37
	300	1	10	-2	66	13	64	6.30	0.83	17.70	21.07	21.07	13.37
	600	1	10	-7	93	11	89	8.90	0.86	14.38	16.47	16.47	10.71
0.5	10	1	10	8	17	18	24	1.70	0.65	31.36	57.33	53.92	20.87
	50	1	10	5	38	17	40	3.60	0.77	20.14	29.21	29.21	14.48
	100	1	10	4	52	16	53	5.10	0.82	16.17	21.98	21.98	12.36
	300	1	10	0	90	14	89	8.80	0.88	11.25	13.96	13.96	8.45
	600	1	10	-3	126	12	125	12.40	0.90	9.00	10.65	10.65	6.60
0.2	10	1	10	9	26	19	36	2.60	0.74	18.28	37.78	35.72	10.59
	50	1	10	7	57	18	61	5.60	0.84	11.40	18.36	18.36	7.90
	100	1	10	6	80	17	83	7.90	0.88	9.06	13.53	13.53	6.58
	300	1	10	3	138	15	139	13.60	0.92	6.27	8.38	8.38	4.41
	600	1	10	1	194	14	194	19.30	0.94	4.93	6.20	6.20	3.88
0.1	10	1	9	9	37	18	51	4.00	0.81	11.61	24.73	23.58	5.90
	50	1	9	7	80	17	87	8.89	0.89	6.96	11.76	11.76	4.03
	100	1	9	7	111	17	116	12.33	0.92	5.49	8.59	8.59	3.45
	300	1	9	5	192	15	195	21.33	0.95	3.74	5.24	5.24	2.59
	600	1	9	3	272	14	273	30.11	0.96	2.93	3.85	3.85	2.35
0.05	10	1	9	10	50	19	71	5.67	0.85	7.78	17.56	17.06	3.36
	50	1	9	8	112	18	121	12.33	0.92	4.57	8.29	8.29	2.43
	100	1	9	7	157	17	164	17.44	0.94	3.57	6.04	6.04	2.12
	300	1	9	6	271	16	275	30.00	0.97	2.41	3.63	3.63	1.76
	600	1	9	5	382	15	385	42.44	0.97	1.88	2.64	2.64	1.41

Table 3.2: Comparisons between lower bounds and upper bounds under various h_0 and K_0.

p_i	λ_i	r_i^*	Q_i^*	r_0^*	Q_0^*	$\hat{r}_0^*$	$\hat{Q}_0^*$	β_1	β_2	δ_1 (%)	δ_2 (%)	δ_3 (%)	δ_4 (%)
$L_0 = 1,\ L_i = 1,\ K_0 = 100,\ K_i = 10,\ h_0 = 0.1,\ h_i = 1$													
1	3	-3	11	-10	116	6	118	10.27	0.88	8.45	13.17	13.17	6.05
	5	-2	14	-8	149	13	152	10.43	0.88	8.58	13.40	13.40	6.13
	10	0	20	1	210	30	216	10.30	0.88	8.48	13.05	13.05	5.97
	15	3	25	12	258	49	264	10.08	0.88	8.54	13.07	13.07	6.36
3	3	0	10	3	112	15	117	11.10	0.91	6.27	9.86	9.86	4.20
	5	1	13	9	145	23	151	11.08	0.91	6.23	9.75	9.75	4.00
	10	5	18	24	205	45	214	11.28	0.91	6.20	9.57	9.57	3.68
	15	9	22	41	251	67	261	11.32	0.91	6.15	9.40	9.40	3.78
10	3	2	9	10	111	20	116	12.33	0.92	5.09	8.28	8.28	3.25
	5	4	12	17	144	31	150	11.92	0.92	5.23	8.56	8.56	3.20
	10	9	17	37	203	56	212	11.94	0.92	5.26	8.54	8.54	3.36
	15	14	20	56	249	79	260	12.40	0.92	5.13	8.16	8.16	3.09
20	3	3	9	13	111	22	117	12.33	0.92	4.83	8.19	8.19	2.81
	5	5	11	21	143	34	150	11.92	0.92	4.90	8.39	8.33	3.41
	10	11	16	41	203	60	212	12.69	0.93	4.86	8.10	8.10	3.06
	15	16	20	62	248	85	260	12.40	0.93	4.86	8.06	8.00	2.94
50	3	4	9	15	112	25	117	12.33	0.93	4.51	8.10	8.01	2.46
	5	7	11	25	143	37	150	13.09	0.93	4.53	7.97	7.97	2.76
	10	13	16	46	203	65	212	12.69	0.93	4.61	8.05	8.05	2.85
	15	18	20	68	248	91	259	12.35	0.93	4.56	8.08	7.85	2.75

Table 3.3: Comparisons between lower bounds and upper bounds under various p_i and λ_i.

3.5.2 Comparison with Echelon-Stock (r,nQ) Policies

Here we numerically compare the proposed modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy with the so-called echelon-stock (r, nQ) policy that also charges a shipment-based fixed cost. The policy requires all orders in an integer multiple of Q. Chen and Zheng [20] test the performance of the echelon-stock (r, nQ) policy numerically. The policy parameters are obtained through a two-step heuristic algorithm. In the first step, they assume that the demand at each retailer is deterministic with rate λ_i, and then use the algorithm in Roundy [76] to compute power-of-two order quantities. In the second step, given these order quantities, they search for reorder points that minimize the total holding and shortage cost. In particular, the second step can be facilitated because the total holding and shortage cost is convex in reorder points with fixed order quantities, as shown in Chen and Zheng [20]. We denote by r_i' and Q_i' the parameters in the echelon-stock (r, nQ) policy. Recall that the real cost of the modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy obtained by Monte Carlo simulation is denoted by $\tilde{C}$. Similarly, we use $\tilde{C}'$ to denote the real cost of the echelon-stock (r, nQ) policy, which is also obtained by Monte Carlo simulation. Moreover, we define the following percentage:

$$\delta_5 \equiv \frac{\tilde{C}' - \underline{C}^*}{\underline{C}^*} \times 100\%,$$

which measures the performance of the echelon-stock (r, nQ) policy.

We list numerical comparisons between our heuristic and the echelon-stock (r, nQ) policy in Tables 3.4 and 3.5. It can be seen that our proposed policy performs better than the echelon-stock (r, nQ) used in Chen and Zheng [20] for most cases. Compared with their policy, our proposed policy has the following advantages. First, in determining order quantities, our policy takes into account the randomness in customer arrivals, while their policy assumes deterministic demands. Second, we consider setup costs when optimizing the policy parameters, while their policy selects the reorder points to minimize the holding and shortage cost, without setup costs. In other words, the echelon-stock (r, nQ) policy in Chen and Zheng [20] considers demand variability only in the second step and setup costs only in the first step. Instead, we simultaneously consider both in determining the policy parameters; see (3.4) and (3.15). This may explain why our proposed policy performs better than theirs in most cases.

h_0	K_0	r'_i	Q'_i	r'_0	Q'_0	r^*_i	Q^*_i	$\hat{r}^*_0$	$\hat{Q}^*_0$	$\underline{C}^*$	$\tilde{C}'$	$\tilde{C}$	$\delta_5(\%)$	δ_4 (%)
						$L_0 = 1,\ L_i = 1,\ K_i = 10,\ h_i = 2,\ p_i = 5,\ \lambda_i = 3$								
0.5	30	1	6	10	24	1	7	14	33	36.2082	40.5987	40.4156	12.13	11.62
	50	1	6	5	48	1	7	13	40	39.8198	43.0385	43.7741	8.08	9.93
	100	1	6	5	48	1	7	12	54	46.5281	49.2622	50.4100	5.88	8.34
	200	1	6	0	96	1	7	12	74	56.1034	59.3099	60.3576	5.72	7.58
0.2	30	1	6	10	48	0	8	14	51	30.5977	32.6605	32.3352	6.74	5.68
	50	1	6	10	48	0	8	14	62	32.9712	35.1758	34.5400	6.69	4.76
	100	1	6	6	96	0	8	13	83	37.3535	38.9967	38.9470	4.40	4.27
	200	1	6	6	96	0	8	12	115	43.5889	45.2235	45.1626	3.75	3.61
0.1	30	1	6	11	48	0	8	15	71	27.8002	29.6489	28.6153	6.65	2.93
	50	1	6	10	96	0	8	15	86	29.5007	30.8464	30.3451	4.56	2.86
	100	1	6	10	96	0	8	14	117	32.6414	33.9366	33.4587	3.97	2.50
	200	1	6	6	192	0	8	13	161	37.1005	38.2440	37.9214	3.08	2.21
0.05	30	1	6	11	96	0	8	16	99	25.8588	26.8926	26.3359	4.00	1.84
	50	1	6	11	96	0	8	15	122	27.0695	28.1413	27.5359	3.96	1.72
	100	1	6	6	192	0	8	15	164	29.3075	30.2574	29.7592	3.24	1.54
	200	1	6	6	192	0	8	14	226	32.4819	33.3753	32.8961	2.75	1.28
0.01	30	1	6	12	192	0	8	17	220	23.3295	23.9376	23.4170	2.61	0.38
	50	1	6	12	192	0	8	17	269	23.8784	24.5524	23.9658	2.82	0.37
	100	1	6	11	384	0	8	16	365	24.8872	25.4076	24.9653	2.09	0.31
	200	1	6	11	384	0	8	16	503	26.3165	26.9477	26.3970	2.40	0.31

Table 3.4: Comparisons between $(\mathbf{r}, \mathbf{Q})$ and $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policies under various h_0 and K_0.

p_i	λ_i	r'_i	Q'_i	r'_0	Q'_0	r^*_i	Q^*_i	$\hat{r}^*_0$	$\hat{Q}^*_0$	$\underline{C}^*$	$\tilde{C}'$	$\tilde{C}$	$\delta_5(\%)$	δ_4 (%)
$L_0 = 1,\ L_i = 1,\ K_i = 10,\ K_0 = 100,\ h_0 = 0.1,\ h_i = 1$														
5	1	0	4	2	64	0	5	6	68	15.5877	16.3485	16.1299	4.88	3.48
	3	2	6	8	96	1	9	17	116	27.2725	29.7625	28.2124	9.13	3.45
	5	3	10	13	160	3	12	27	151	35.4459	36.8717	36.7580	4.02	3.70
	10	9	10	35	160	7	17	50	213	50.7022	56.8671	52.5554	12.16	3.65
10	1	1	4	4	64	0	6	9	67	16.7343	17.7010	17.3441	5.78	3.64
	3	3	6	14	96	2	9	20	116	29.2143	31.5021	30.1073	7.83	3.06
	5	4	10	22	160	4	12	31	150	37.9151	39.2382	39.1529	3.49	3.26
	10	10	10	44	160	9	17	56	212	54.1700	59.7910	55.9588	10.38	3.30
20	1	1	4	5	64	1	5	10	67	17.8137	18.8559	18.3446	5.85	2.98
	3	4	6	16	96	3	9	22	117	30.9848	33.2126	31.8308	7.19	2.73
	5	5	10	25	160	5	12	34	150	40.1998	41.4787	41.3914	3.18	2.96
	10	12	10	47	160	11	16	60	212	57.2868	62.5457	59.0076	9.18	3.00
50	1	2	4	7	64	2	5	12	67	19.4298	20.0845	20.0069	3.37	2.97
	3	5	6	18	96	4	9	25	117	33.1586	35.3168	34.0333	6.51	2.64
	5	7	10	28	160	7	11	37	150	42.8963	43.9782	44.0404	2.49	2.67
	10	13	10	51	160	13	16	65	212	61.0013	66.9631	62.7135	9.77	2.81

Table 3.5: Comparisons between $(\mathbf{r}, \mathbf{Q})$ and $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policies under various p_i and λ_i.

We also observe that our heuristic is more likely to outperform the echelon-stock (r, nQ) policy in the cases when h_0 is small or K_0 is large. This is consistent with the asymptotic optimality of our proposed policy shown in Theorem 3.8. As h_0 increases or K_0 increases, our heuristic policy seems to converge quickly and become near-optimal, whereas the effectiveness of the echelon-stock (r, nQ) policy does not seem to converge quickly. Moreover, Table 3.5 shows that our proposed policy tends to be more robust as p_i or λ_i varies. For example, when λ_i takes the value of 3 or 10, the gap between the cost of the echelon-stock (r, nQ) policy and the induced penalty lower bound, δ_5, can be more than 10%, whereas the gap for our proposed policy, δ_4, is less than 3.7%. It should also be noted that the echelon-stock (r, nQ) policy can perform better than our proposed policy in some cases, e.g., when h_0 is relatively large; see Table 3.4.

3.6 Conclusion

In this work, we have studied the classic distribution inventory system with setup costs per shipment at each installation. The optimal policy of such a system is unknown, even without setup costs. We have studied a class of modified echelon $(\mathbf{r}, \mathbf{Q})$ policies that do not require an integer-ratio property or a synchronized (nested ordering) property. For a constructed modified echelon $(\mathbf{r}, \mathbf{Q})$ heuristic policy, we have provided 1) worst-case performance guarantees by comparing its performance to a lower bound of the optimal policy and 2) conditions under which the heuristic is asymptotically optimal.

Our work can be easily extended to a distributions system where the warehouse has a minimum order quantity (MOQ) requirement. The MOQ requirement means that whenever an order is placed, the order quantity must be no less than a specified level, say M. Suppose the supplier stipulates this

MOQ requirement to the warehouse, that is, the warehouse, if ordering, must order at least M units each time; otherwise, the warehouse places no order. With some slight adjustment, our heuristic policy in (MERQD) can be modified to accommodate a distribution system with MOQ requirements, that is, $(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) = (\tilde{r}_0^*, \min\{M, \tilde{Q}_0^*\}, r_1^*, Q_1^*, \ldots, r_N^*, Q_N^*)$. Moreover, we can obtain the performance guarantees analogously. For example, the lower bound on the warehouse's cost in (3.8) should be adjusted as $C_0^* = C_0(r_0^*(Q_0^M), Q_0^M)$, where $Q_0^M \equiv \min\{M, Q_0^*\}$.

Finally, we do note that the bounds developed on the proposed modified $(\mathbf{r}, \mathbf{Q})$ policies leave something to be desired, as evidenced by the performance of the numerical experiments. This may be attributable to the complexity of the one-warehouse-multi-retailer system. Nevertheless, our work represents the first attempt in identifying easy-to-compute policies and bounds for distribution systems.

It should be noted the current performance bounds are still dependent on the optimal solutions of single-stage inventory systems. As for future research directions, one may try to derive some closed-form performance bounds which explicitly depend on system primitives. Ideally, a constant performance bound is also expected. Another future direction is to study the relationship between the performance bounds and number of retailers. We can numerically show that the performance bounds converge to a constant as the number of retailer increases. But the closed form of the constant is still unclear.

Chapter 4

Single-Echelon Systems with MOQ and Batch Ordering

4.1 Introduction

In industries, minimum order quantity (MOQ) and batch ordering, applied independently or simultaneously, are two common requirements made by suppliers, both of which can help companies take advantage of economies of scale and hence reduce costs. The MOQ requirement means that the order quantity must equal or exceed a specified level, if an order is placed. The batch ordering requirement means that the order quantity must be an integral multiple of a specified given batch size.

The application of an MOQ is common in practice. With the prevalence of e-commerce, MOQs are becoming more and more common in our lives, especially in online business-to-business sourcing portals such as alibaba.com, where suppliers often set such requirements. MOQs are also applied in manufacturing industries for products that have short lifetimes or long leadtimes. A well-known example is Sport Obermeyer, a fashion sport skiwear manufacturer, which has a minimum production level of 600 garments in Hong Kong and 1200 garments in China per order (Zhao and Katehakis [96]). In fact, MOQ requirements are quite common in China and other low cost manufacturing countries. Low profit margins force manufacturers to pursue large production quantities to break even. On the other hand, batch ordering is also a ubiquitous requirement in industries, because materials often flow in fixed batch sizes in supply chains. For example, raw materials usually arrive at factories in full truckloads, work-in-process is often processed in convenient lot sizes between production stages, and finished goods may be transported in full containers from suppliers to warehouses or distribution centers. Therefore, it is of no surprise that suppliers who apply an MOQ may also require batch ordering. Indeed, our decision to jointly consider both MOQ and batch ordering requirements in this work is largely motivated by our experience with a wholesale company in Hong Kong. For a variety of products, the firm first replenishes its stock from suppliers and then sells to retail customers, and for most of these products, the firm stipulates both MOQ and batch ordering requirements.

The coexistence of an MOQ and batch ordering has a two-sided effect. On the one hand, requiring an MOQ and batch ordering simultaneously helps suppliers reduce the risk of uncertainty and achieve economies of scale. On the other hand, the requirements may have a negative effect on buyers' inventory control, especially when MOQs are relatively large compared with their demand, which is not unusual in practice. Managers in such situations need principles or tools to help control their inventory. However, to the best of our knowledge, no research has investigated inventory systems with both MOQ and batch ordering requirements. Thus, the primary goal of this work is to fill this gap in the literature. In this work, we consider a single product stochastic periodic-review inventory system with both MOQ and batch ordering requirements. The selling firm can make a decision at the beginning of each time period after reviewing the inventory position. When the firm decides to place an order, the order quantity must satisfy both the

MOQ and the batch ordering constraints, where we assume the MOQ is an integral multiple of the batch size. The leftover inventory is carried to the next period and incurs a holding cost, whereas unsatisfied demand is fully backlogged and incurs a backordering cost. The total costs consist of the linear ordering cost, the holding cost, and the backordering cost. The objective is to minimize the long-run average cost of the system.

The optimal policy for the system with only the MOQ requirement, which is partially characterized by Zhao and Katehakis [96], is rather complicated, even without batch ordering. Therefore, for inventory systems with both MOQ and batch ordering requirements, it is necessary to propose some effective heuristic policies, which is the major contribution of our work. Facing the MOQ requirement, many companies apply the (s, S) type policy to control inventories in practice (Zhou et al. [99]). Based on this, we first propose a two-parameter policy with a similar structure, i.e., the (s, k) policy, where $s < k < s + M$ and M represents the MOQ. The (s, k) policy operates as follows: at the beginning of each period, if the inventory position is less than or equal to s, order a quantity that is just sufficient to bring the inventory position to $s+M$ or above (the inventory position after ordering can be larger than $s + M$, because the order quantity must also satisfy the batch ordering requirement); if the inventory position exceeds s but is no more than k, order exactly M; otherwise, order nothing. We identify the bounds for the optimal k, and propose algorithms to find the optimal values of k and s. We also examine a simpler and more easy-to-use policy, i.e., the S policy, which is a special case of the (s, k) policy. The S policy operates in the same way as the (s, k) policy with $s = S - M$ and $k = S - 1$. The numerical study shows that both these polices have close-to-optimal performance in most cases and that there is an overwhelming preponderance to the best (s, S) policy over all examples.

The remainder of this chapter is organized as follows. The literature on MOQ and batch ordering is discussed in Section 4.2. In Section 4.3, the model description and notations are presented. In Section 4.4, we propose a two-parameter (s, k) policy and present algorithms to optimize the policy. A simpler one-parameter policy is introduced in Section 4.5. Numerical examples are conducted in Section 4.6 to measure the effectiveness of these two policies by comparing them with other policies. Finally, Section 4.7 concludes the work by summarizing the findings.

4.2 Literature Review

The existing research on stochastic inventory systems is quite extensive. Here, we mention only a few of the most relevant papers. Many papers focus on problems associated with batch ordering or MOQ separately. The literature related to our work can be divided into two areas: 1) supply chain inventory management with batch ordering; and 2) supply chain inventory management with MOQ.

In the area of batch ordering, Veinott Jr [92] shows the optimality of the (R, Q) policy for a periodic-review inventory system with batch ordering and no fixed ordering cost. This (R, Q) policy operates as follows: at the beginning of each period, if the inventory position is less than the reorder point R, order the smallest integral multiple of the batch size Q that will bring the inventory position to at least R; otherwise order nothing. Chen [18] generalizes Veinott Jr [92]'s result to multi-echelon systems settings and demonstrates the optimality of (R, nQ) policies for multi-stage serial and assembly systems where materials flow in fixed batches and the stochastic demands are stationary over time. Chao and Zhou [16] find the optimal inventory control policy for a multi-echelon serial system with batch ordering and fixed replenishment intervals. They derive a distribution-function solution for its optimal control parameters and design an efficient algorithm for computing those parameters. Huh and Janakiraman [54] extend the work of Veinott Jr [92] and Chen [18] by demonstrating the optimality of echelon (R, nQ) policies for multi-echelon serial systems with nested batch ordering and non-stationary demands. Although it is not optimal in some complex inventory systems with batch ordering, the reorder point, lot-size ordering policy is easy to implement. For this reason, numerous heuristic policies have been proposed, see for example Schwarz and Schrage [78], De Bodt and Graves [24], Gallego [37], Axsäter and Zhang [7], Broekmeulen and van Donselaar [12] and Shang and Zhou [79].

In the area of MOQ, Fisher and Raman [34] consider a two-period model with an MOQ in each period. They work out the optimal order quantity by using stochastic programming methods. Zhao and Katehakis [96] introduce the concept of M-increasing function and first partially characterize the optimal policy for multi-period inventory systems with MOQ. For the uncharacterized part, the authors give easily computable upper bounds and asymptotic lower bounds for these intervals. However, for the characterized part, the optimal policy is complexly structured and difficult to implement in practice. Hellion et al. [43] propose an algorithm with time complexity in $O(T^5)$ for a capacitated lot sizing problem with MOQ and concave costs. For other references on MOQ, the reader is referred to Porteus and Whang [74], Chan and Muckstadt [14], Lee [60], Porras and Dekker [73], and Okhrin and Richter [69].

The most closely related papers to our work are Zhou et al. [99] and Kiesmüller et al. [59]. Zhou et al. [99] propose a two-parameter heuristic policy for a stochastic inventory system with MOQ requirement and demonstrate that the performance of this policy is close to the optimal policy except for a few cases when the coefficient of the demand distribution is very small. Kiesmüller et al. [59] propose a simpler policy, which has only one parameter S. This policy works as follows: no order is placed when the inventory position is not less than the level S; otherwise an order is placed to raise the inventory to S. However, if this order is smaller than the MOQ, the order quantity is increased to the MOQ. The authors show the effectiveness of this policy and develop simple newsvendor inequalities for near-optimal policy parameters. However, both Zhou et al. [99] and Kiesmüller et al. [59] do not consider batch ordering. To the best of our knowledge, our work is the first to study stochastic inventory system with both MOQ and batch ordering requirements. To combine the two requirements, we need to tackle the problem of selecting an order quantity that satisfies both the constraints simultaneously. In a system with only the MOQ constraint, the order quantity can be any integer that is larger than or equal to the MOQ. However, with the addition of batch ordering, the firm has to either round up or round down the order quantity to an integral multiple of the given batch size. Therefore, in our model, the order quantities are subject to two kinds of jumps, which makes the analysis much more difficult.

4.3 Model Description

We consider a periodic-review inventory system for a single item with stochastic demand. The demand D in each period is an independent identically distributed (i.i.d.) random variable. At the beginning of each period, after reviewing the initial inventory position, the retailer decides whether to make a placement. When an order is made, the order quantity is at least M, the minimum order quantity (MOQ). We assume that M is an integral multiple of the batch size Q and $M > Q$. If $M \leq Q$, the problem is reduced to that of Veinott Jr [92], the optimal policy of which can easily be computed. After the order placement, the demand is realized and unsatisfied demand will be backlogged. A penalty cost p per unit will be incurred. Note that the demand is satisfied according to the first-come-first-service rule in our system. At the end of each period, excess inventory will generate an inventory holding cost h per unit per period. In our model, we assume linear variable cost and no fixed ordering cost. Without loss of generality, we assume zero lead-time. Our model can be easily extended to systems with positive lead times using the standard method in Heyman and Sobel [44]. The average cost criterion is used to evaluate the inventory system. Because the linear ordering costs can be ignored under the average cost criterion, the unit ordering cost is set to be 0, e.g., Zheng and Federgruen [98]. The objective is to minimize the long-run average cost of the system.

In the remainder of this chapter, the following notations will be used:

M	Minimum order quantity
Q	Batch size
D_t	Demand in period t
q_t	Order quantity in period t
h	Holding cost per unit per period
p	Penalty cost per unit per period
x_t	The inventory position before ordering in period t
y_t	The inventory position after ordering in period t
Z^+	$\max(0, Z)$
$[a, b]$	The integer numbers between a and b (if a and b are integers, they are included)

The expected cost function $C_t(y_t)$ in period t can be written as

$$C_t(y_t) = h\mathbb{E}[(y_t - D_t)^+] + p\mathbb{E}[(D_t - y_t)^+], \tag{4.1}$$

where y_t and D_t are both integers. We can easily find that $C(y_t)$ is convex and $C(y_t) \to +\infty$ as $|y_t| \to \infty$. In the remainder of this work, we omit the subscript when there is no ambiguity. Let y^* be a minimizer of $C(y)$. We also assume that x, M, and Q are all integers.

4.4 The Two-Parameter Heuristic Policy

As above mentioned, Zhao and Katehakis [96] find the structure of the optimal policy for the system with an MOQ to be rather complex and conclude that such an optimal policy is not practically implementable. The presence of batch ordering makes the problem even more complicated. Therefore, it is necessary to develop some easily implementable polices that have good performance. Based on the analysis of multi-period stochastic inventory system with an MOQ (see Zhou et al. [99]) and the optimal policy structure for batch ordering (see Veinott Jr [92]), we propose a modified (s, k) policy: given an initial inventory position x_t and two integer parameters s and k, where $s < k < s + M$, the order quantity q_t is

$$q_t = y_t - x_t = \begin{cases} M + mQ, & \text{if } x_t \le s; \\ M, & \text{if } s < x_t \le k; \\ 0, & \text{if } x_t > k. \end{cases} \tag{4.2}$$

where $m \ge 1$, and m is the unique integer such that $0 < y_t - (s+M) \le Q$. That is, when x_t is not larger than s, order up to y_t, such that $s + M < y_t \le s + M + Q$; when x_t is larger than s but does not exceed k, order exactly M; and when x_t is above k, do not make an order. This policy is an extension of the (s, k) policy proposed in Zhou et al. [99], where no batch ordering exists. For this reason, we simply call our modified (s, k) policy the (s, k) policy in the remainder of this work. To identify the optimal policy parameters s and

k that minimize the long-run average cost, we use a discrete time Markov chain with transition matrix $\mathbf{P}$ and analyze the system under two cases: $\Delta \geq Q$ and $\Delta < Q$, where Δ is defined as the difference between k and s, i.e., $\Delta = k - s$.

4.4.1 Case 1: $\Delta \geq Q$

Under the condition $\Delta \geq Q$, Δ has a finite state space $[Q, M-1]$. In our (s,k) policy, the inventory position after order placement in the $(t+1)^{th}$ period is $y_{t+1} = x_{t+1} + q_{t+1} = y_t - D_t + q_{t+1}$, so

$$y_{t+1} = \begin{cases} y_t - D_t, & \text{if } y_t - D_t > k; \\ y_t - D_t + M, & \text{if } s < y_t - D_t \leq k; \\ y_t - D_t + M + mQ, & \text{if } y_t - D_t \leq s. \end{cases} \tag{4.3}$$

We can see that $\{y_i\}$ is a discrete time Markov chain (DTMC) and has the finite state space $[k+1, k+M]$. The state space can be split into three segments:

1. $[k+1, s+M]$: If $y_{t+1} \in [k+1, s+M]$, it means that $q_{t+1} = 0$, and $D_t = y_t - y_{t+1}$.

2. $[s+M+1, s+M+Q]$: If $y_{t+1} \in [s+M+1, s+M+Q]$, there are three possibilities: $q_{t+1} = 0$, $q_{t+1} = M$, or $q_{t+1} = M + mQ$. The three terms correspond to $D_t = y_t - y_{t+1}$, $D_t = y_t - y_{t+1} + M$, and $D_t = y_t - y_{t+1} + M + mQ$, respectively.

3. $[s+M+Q+1, k+M]$: If $y_{t+1} \in [s+M+Q+1, k+M]$, then $q_{t+1} = 0$ or $q_{t+1} = M$, and $D_t = y_t - y_{t+1}$, or $D_t = y_t - y_{t+1} + M$.

It is easy to compute the transition probabilities $P_{i,j} = Prob(y_{t+1} = j | y_t = i)$, and hence the transition matrix $\mathbf{P}$:

$$P_{i,j} = \begin{cases} p_{(i-j)^+}, & \text{for } j \in [k+1, s+M] \\ & \forall i \in [k+1, k+M], \\ \sum_{m=0}^{\infty} p_{(i+M+mQ-j)} + p_{(i-j)^+}, & \text{for } j \in [s+M+1, s+M+Q] \\ & \forall i \in [k+1, k+M], \\ p_{(i+M-j)} + p_{(i-j)^+}, & \text{for } j \in [s+M+Q+1, k+M] \\ & \forall i \in [k+1, k+M], \end{cases} \tag{4.4}$$

where $p_k = Prob(D_t = k)$ and $p_{(i-j)^+}$ equals to p_{i-j} if $i \geq j$, and zero otherwise. For convenience of notation, m is allowed to take the value of 0 in the expression.

Because the Markov chain is irreducible and positive recurrent, the unique steady state probabilities $\vec{\pi} = \{\pi_1, \pi_2, ..., \pi_M\}$ exist, where π_i denotes the long-run average proportion of time in which the inventory position y is $k+i$. Because the Markov chain is also aperiodic, π_i is also the limiting probability that the chain is in state i.

Let $t \to \infty$, we can have

$$\begin{cases} \sum_{i=1}^{M} \pi_i = 1 \\ \vec{\pi} P = \vec{\pi} \end{cases} \tag{4.5}$$

Therefore, we can calculate the stationary probabilities by solving the linear equations (4.5). Before exploring the properties of the (s, k) policy, we must point out that given the MOQ M, the batch size Q, and the demand distribution, the transition matrix $\mathbf{P}$ and stationary probabilities $\vec{\pi}$ depend only on Δ, and are independent of s. Note that a given Δ has and only has one corresponding $\mathbf{P}$ and $\vec{\pi}$. Now we can calculate the long-run average cost for this case:

$$L(\Delta, k) = \sum_{i=1}^{M} \pi_i C(k+i). \tag{4.6}$$

We can derive the following proposition:

Proposition 4.1. *For a given $\Delta \geq Q$, $L(\Delta, k)$ is convex in k.*

Let k^{*1} be the value at which L reaches its minimum for a given $\Delta \geq Q$. Note that k^{*1} is in fact a function of Δ. For convenience of notation, we use k^{*1} to denote the corresponding optimal k for a given $\Delta \geq Q$.

Proposition 4.2. *Given $\Delta \geq Q$, k^{*1} satisfies $y^* - M \leq k^{*1} < y^* \leq k^{*1} + M$.*

For a given Δ, the preceding propositions help us narrow the search space of k^{*1}. Based on the propositions, we design the following algorithm to compute Δ^* and the corresponding k^{*1} that minimize the long-run average cost.

Algorithm 4.1 Policy Optimization for $\Delta \geq Q$

Set $L^{*1} \leftarrow$Inf;
for $\Delta = Q$ **to** $M - 1$
  calculate $\mathbf{P}$ by (4.4);
  calculate $\vec{\pi}$ by (4.5);
  for $k = y^* - M$ **to** $y^* - 1$
    calculate $L(\Delta, k)$ by (4.6);
    if $L(\Delta, k) < L^{*1}$
      $L^{*1} \leftarrow L(\Delta, k)$, $k^{*1} \leftarrow k$, $\Delta^{*1} \leftarrow \Delta$;

When computing the transition matrix $\mathbf{P}$, we find there are many repeated calculations. To avoid calculating the same probability repeatedly, we also provide a recursive method to get $\mathbf{P}$ for this case. For a given Δ, the corresponding transition matrix $\mathbf{P}$ can be divided into three parts by column, and each part is a submatrix. The first part is the first $(M - \Delta)$ columns of $\mathbf{P}$ with $P_{i,j} = p_{(i-j)^+}$ in this part. The second part is a submatrix with Q columns consecutive to the first part. In the second part,

$P_{i,j} = p_{(i-j)^+} + \sum_{m=0}^{\infty} p_{(i+M+mQ-j)}$. The third part is the last $\Delta - Q$ columns with $P_{i,j} = p_{(i+M-j)} + p_{(i-j)^+}$. With an abuse of notation, let P^{Δ} denote the transition matrix for a given Δ. If $\Delta = Q$, we directly calculate P^{Δ} by (4.4). Otherwise, we can calculate P^{Δ} recursively. Now, assume we already know $P^{\Delta-1}$. The following algorithm enables us to calculate P^{Δ} recursively. The

Algorithm 4.2 Calculating $\mathbf{P}$ for $\Delta \geq Q$

if $\Delta = Q$
    calculate $\mathbf{P}$ by (4.4)
else
    for $i = 1$ **to** M
        for $j = 1$ **to** $M - \Delta$
            $P_{i,j}^{\Delta} \leftarrow P_{i,j}^{\Delta-1}$;
    $j \leftarrow M - \Delta + 1$;
    for $i = 1$ **to** $M - 1$
        $P_{i,j}^{\Delta} \leftarrow P_{i+1,j+1}^{\Delta-1}$;
    $i \leftarrow M$;
    $P_{i,j}^{\Delta} \leftarrow P_{i,j}^{\Delta-1} + \sum_{m=0}^{\infty} p_{(i+M+mQ-j)}$;
    for $i = 1$ **to** M
        for $j = M - \Delta + 2$ **to** $M - \Delta + Q$
            $P_{i,j}^{\Delta} \leftarrow P_{i,j}^{\Delta-1}$;
    if $\Delta = Q + 1$
        for $i = 1$ **to** $M - 1$
            $P_{i,j}^{\Delta} \leftarrow Prob(D = i)$;
        $i \leftarrow M$;
        $P_{i,j}^{\Delta} \leftarrow Prob(D = M) + Prob(D = 0)$;
    else
        $j \leftarrow M - \Delta + Q + 1$;
        for $i = 1$ to $M - 1$
        $P_{i,j}^{\Delta} \leftarrow P_{i+1,j+1}^{\Delta-1}$;
    $i \leftarrow M$;
    $P_{i,j}^{\Delta} \leftarrow Prob(D = \Delta - Q - 1) + Prob(D = M + \Delta - Q - 1)$;
    for $i = 1$ **to** M
        for $j = M - \Delta + Q + 2$ **to** M
            $P_{i,j}^{\Delta} \leftarrow P_{i,j}^{\Delta-1}$;

motivation for Algorithm 4.2 is quite simple: $P_{i,j}$ depends only on the value of $i - j$. If $P_{i,j}^{\Delta}$ and $P_{i,j}^{\Delta-1}$ are in the same part, then they have an identical expression and hence $P_{i,j}^{\Delta} = P_{i,j}^{\Delta-1}$. If $P_{i,j}^{\Delta}$ and $P_{i,j}^{\Delta-1}$ are not in the same part, then $P_{i,j}^{\Delta}$ and $P_{i,j+1}^{\Delta-1}$ (if exists) must be in the same part. With $i - j$ being a constant, $P^{\Delta}_{i,j} = P^{\Delta-1}_{i+1,j+1}$ for $i \in [1, M-1]$, $j \in [1, M-1]$. For the case $i = M$ or $j = M$, we calculate $P^{\Delta}_{i,j}$ separately.

4.4.2 Case 2: $\Delta < Q$

At the beginning of this subsection, we must point out that as we show in the numerical study, Case 2 does arise. Under the condition $\Delta < Q$, Δ has a finite state space $[1, Q-1]$. Similar to the case $\Delta \geq Q$, the inventory position after a possible ordering is still a discrete time Markov chain, and the Markov chain has the finite state space $[k+1, s+M+Q]$ in this case. The state space can also be split into three segments:

1. $y_{t+1} \in [k+1, s+M]$: If $y_{t+1} \in [k+1, s+M]$, it means that $q_{t+1} = 0$, hence $D_t = y_t - y_{t+1}$.

2. $y_{t+1} \in [s+M+1, k+M]$: If $y_{t+1} \in [s+M+1, k+M]$, it means that $q_{t+1} = 0, M$, or $M + mQ$. The three terms correspond to three possibilities $D_t = y_t - y_{t+1}, y_t - y_{t+1} + M$, and $y_t - y_{t+1} + M + mQ$, respectively. Recall that m is the largest integer such that $0 < y_{t+1} - (s+M) \leq Q$ and $m \geq 1$.

3. $y_{t+1} \in [k+M+1, s+M+Q]$: If $y_{t+1} \in [k+M+1, s+M+Q]$, it means that $q_{t+1} = 0$ or M, then $D_t = y_t - y_{t+1}$, or $D_t = y_t - y_{t+1} + M + mQ$.

The transition probabilities $P_{i,j} = Prob(y_{t+1} = j|y_t = i)$ can be calculated easily and the transition matrix $\mathbf{P}$ is

$$P_{i,j} = \begin{cases} p_{(i-j)^+}, & \text{for} j \in [k+1, s+M] \\ & \forall i \in [k+1, s+M+Q], \\ \sum_{m=0}^{\infty} p_{(i+M+mQ-j)} + p_{(i-j)^+}, & \text{for} j \in [s+M+1, k+M] \\ & \forall i \in [k+1, s+M+Q], \\ \sum_{m=1}^{\infty} p_{(i+M+mQ-j)} + p_{(i-j)^+}, & \text{for} j \in [k+M+1, s+M+Q] \\ & \forall i \in [k+1, s+M+Q], \end{cases} \tag{4.7}$$

Again, let $t \to \infty$, the limiting probabilities are the steady state probabilities $\vec{\pi} = \{\pi_1, \pi_2 \ldots \pi_{s+M+Q-k}\}$. We can calculate the stationary probabilities by solving the linear equations

$$\begin{cases} \Sigma_{i=1}^{s+M+Q-k} \pi_i = 1 \\ \vec{\pi} P = \vec{\pi} \end{cases} \tag{4.8}$$

The transition matrix $\mathbf{P}$ and stationary probabilities $\vec{\pi}$ still only depend on Δ, and are independent of s. Now we have the long-run average cost for this case:

$$L(\Delta, k) = \sum_{i=1}^{M+Q-(k-s)} \pi_i C(k+i) = \sum_{i=1}^{M+Q-\Delta} \pi_i C(k+i). \tag{4.9}$$

The following propositions are useful in searching for the optimal parameters.

Proposition 4.3. *For a given $\Delta < Q$, $L(\Delta, k)$ is convex in k.*

For a given $\Delta < Q$, let k^{*2} be the value of k at which $L(\Delta, k)$ reaches minimum.

Proposition 4.4. *Given $\Delta < Q$, k^{*2} satisfies $y^* - M - Q + \Delta \leq k^{*2} < y^* \leq k^{*2} + M + Q - \Delta$.*

Based on these propositions, we provide the following algorithm to compute Δ^* and the corresponding k^{*2} for this case.

Algorithm 4.3 Policy Optimization for $\Delta < Q$

Set $L^{*2} \leftarrow$Inf;
for $\Delta = 1$ **to** $Q - 1$
    calculate $\mathbf{P}$ by (4.7);
    calculate $\vec{\pi}$ by (4.8);
    for $k = y^* - M - Q + \Delta$ **to** $y^* - 1$
        calculate $L(\Delta, k)$ by (4.9);
        if $L(\Delta, k) < L^{*2}$
            $L^{*2} \leftarrow L(\Delta, k)$, $k^{*2} \leftarrow k$, $\Delta^{*2} \leftarrow \Delta$;

We also propose a recursive method for calculating $\mathbf{P}$ for this case. In this case, $\mathbf{P}$ is a $(M - \Delta + Q) \times (M - \Delta + Q)$ matrix. Obviously, the size of $\mathbf{P}$ will decrease as Δ increases. $\mathbf{P}$ can also be divided into three parts in this case. The first part is the first $M - \Delta$ columns, with $P_{i,j} = p_{(i-j)^+}$ in this part. The second part is a submatrix with Δ columns consecutive to the first part. In this part, $P_{i,j} = p_{(i-j)^+} + \sum_{m=0}^{\infty} p_{(i+M+mQ-j)}$. The third part is the last $Q - \Delta$ columns, with $P_{i,j} = p_{(i-j)^+} + \sum_{m=1}^{\infty} p_{(i+M+mQ-j)}$. Similar to the case of $\Delta \geq Q$, we also provide a recursive algorithm to calculate $\mathbf{P}$. Again with an abuse of notation, we write P^{Δ} to denote the corresponding $\mathbf{P}$ for a given Δ.

Algorithm 4.4 Calculating $\mathbf{P}$ for $\Delta < Q$

if $Q = 1$
    calculate $\mathbf{P}$ by (4.7)
else
    for $i = 1$ **to** $M - \Delta + Q$
        for $j = 1$ **to** $M - \Delta + Q$
            if $j = M - \Delta + 1$
                $P_{i,j}^{\Delta} \leftarrow P_{i+1,j+1}^{\Delta-1}$;
            else
                $P_{i,j}^{\Delta} \leftarrow P_{i,j}^{\Delta-1}$;

Now, we can get two pairs of solutions (Δ^{*1}, k^{*1}) and (Δ^{*2}, k^{*2}) for two cases separately. To get the global optimal solution, we only need to compare the two pairs and find the optimal pair of (Δ^*, k^*) that minimizes $L(\Delta, k)$, i.e., $(\Delta^*, k^*) = \arg\min\{L(\Delta^{*1}, k^{*1}), L(\Delta^{*2}, k^{*2})\}$. The following algorithm describes the whole (s, k) policy optimization.

Algorithm 4.5 Policy Optimization for the (s, k) policy

Set $L^{*1} \leftarrow$Inf, $L^{*2} \leftarrow$Inf;
for $\Delta = Q$ **to** $M - 1$
    calculate $\mathbf{P}$ by Algorithm 4.2;
    calculate $\vec{\pi}$ by (4.5);
    for $k = y^* - M$ **to** $y^* - 1$
        calculate $L(\Delta, k)$ by (4.6);
        if $L(\Delta, k) < L^{*1}$
            $L^{*1} \leftarrow L(\Delta, k)$, $k^{*1} \leftarrow k$, $\Delta^{*1} \leftarrow \Delta$;
for $\Delta = 1$ **to** $Q - 1$
    calculate $\mathbf{P}$ by Algorithm 4.3;
    calculate $\vec{\pi}$ by (4.8);
    for $k = y^* - M - Q + \Delta$ **to** $y^* - 1$
        calculate $L(\Delta, k)$ by (4.9);
        if $L(\Delta, k) < L^{*2}$
            $L^{*2} \leftarrow L(\Delta, k)$, $k^{*2} \leftarrow k$, $\Delta^{*2} \leftarrow \Delta$;
if $L^{*1} > L^{*2}$
    $L^* \leftarrow L^{*2}$, $k^* \leftarrow k^{*2}$, $\Delta^* \leftarrow \Delta^{*2}$;
else
    $L^* \leftarrow L^{*1}$, $k^* \leftarrow k^{*1}$, $\Delta^* \leftarrow \Delta^{*1}$;

4.5 The One-Parameter Heuristic Policy

In this section, we develop a simpler heuristic policy: when the initial inventory position is not less than S, no order is placed; when the initial inventory position is less than S but exceeds $S - M$, order exactly M; otherwise, order $M + mQ$, where m is the smallest integer such that $x_t + M + mQ > S$. The motivation for this policy is quite intuitive. Veinott Jr [92] has shown the optimality of the (R, Q) policy for batch ordering. The key principle of the (R, Q) policy is to order the smallest integral multiple of Q to bring the inventory position above the reorder point R. In our model, there is an MOQ constraint, so the order quantity cannot be less than M. Based on this, the order quantity in period t is q_t, such that

$$q_t = y_t - x_t = \begin{cases} M + mQ, & \text{if } x_t \leq S - M; \\ M, & \text{if } S - M < x_t < S; \\ 0, & \text{if } x_t \geq S. \end{cases} \tag{4.10}$$

We need to point out that when there is no MOQ constraint, i.e., $M = 0$, the modified S policy reduces to the simple (R, Q) policy. Or when $Q = 1$, the modified S policy reduces to the S policy of Kiesmüller et al. [59]. For this reason, we simply call the modified S policy the S policy. In our model, because $M > Q$, the state space of the inventory level after ordering in the S policy is $[S, S + M - 1]$. The inventory position after ordering is a discrete time Markov chain, and we can calculate the transition probability $P_{i,j} = Prob(y_{t+1} = j | y_t = i)$ and transition matrix $\mathbf{P}$, where

$$P_{i,j} = \begin{cases} p_{(i-j)^+}, & \text{for } j = S, \\ & \forall i \in [S, S + M - 1], \\ \sum_{m=0}^{\infty} p_{(i+M+mQ-j)} + p_{(i-j)^+}, & \text{for } j \in [S + 1, S + Q], \\ & \forall i \in [S, S + M - 1], \\ p_{(i+M-j)} + p_{(i-j)^+}, & \text{for } j \in [S + Q + 1, S + M - 1], \\ & \forall i \in [S, S + M - 1]. \end{cases} \tag{4.11}$$

We can also use the same method as that of the (s, k) policy to calculate the stationary probabilities $\vec{\pi}$ by (4.5) and hence the long-run average cost can be calculated by $L(S) = \sum_{i=1}^{M} \pi_i C(S + i - 1)$.

Note that the S policy is in fact a special case of the (s, k) policy, where $s = S - M$ and $k = S - 1$. Therefore, Proposition 1 and Proposition 2 still hold for the S policy and the policy optimization is simpler: because this is a special case of the (s,k) policy where $\Delta = k-s = M-1$, we can calculate the transition matrix and stationary probabilities, and then identify the optimal S^* that minimizes $L(S)$ from the set $[y^* - M + 1, y^*]$ as we can do for the general case $\Delta \geq Q$. The complexity of solving the linear equation (4.5) is $O(M^3)$, if a Gaussian elimination is used. Because most of the calculation lies in this part, the total complexity of the S policy is $O(M^3)$, while the complexity of the (s,k) policy is $O(M^4)$, because Δ can take M different values.

4.6 Numerical Study

In this section, we conduct numerical experiments to test the performance of the (s,k) policy and the S policy. We consider two different discrete demand distributions:

- Normal distribution

  To ensure nonnegative demand, we truncate the normal distribution at zero to avoid negative demand.

- Poisson distribution

We conduct numerical studies with respect to the following parameters: the batch size Q, the expected demand per period $\mathbb{E}(D)$, the critical ratio $p/(p+h)$, and the demand coefficient of variation (we can omit this factor in the case of a Poisson distribution). We assume $M = 30$ unless otherwise specified. Other parameters are chosen as follows. The holding cost $h = 1$ is fixed. Batch size Q varies as 3, 5, 6, 10, and 15. $\mathbb{E}(D)$ takes the values of 10, 15, 20, 30, and 40; $p/(h+p)$ varies as 0.80, 0.85, 0.90, and 0.95, and the demand coefficient of variation (c.v.) of the normal distribution takes the values of 0.1, 0.2, 0.3, and 0.4. The complete set of parameter values is given in Table 4.1. All possible combinations of the parameters give us 400 instances for normally distributed demand, and 100 instances for Poisson distributed demand. We must point that we can get the same results as Zhou et al. [99] and Kiesmüller et al. [59], if the same parameter values are selected and Q is fixed to be 1.

h=1;	$M = 30$;	$Q \in \{3, 5, 6, 10, 15\}$;
$\mathbb{E}(D) \in \{10, 15, 20, 30, 40\}$; $p/(h+p) \in \{0, 80, 0.85, 0.90, 0.95\}$;		
$c.v. \in \{0.1, 0.2, 0.3, 0.4\}$ for normal distribution		

Table 4.1: Base parameter values for the numerical experiments

To better illustrate the performance of the (s, k) policy and the S policy, we compare them with two other policies. The first policy is "the optimal policy" that achieves the minimal average cost among all admissible policies. We use value iteration to compute the optimal long-run average cost and the optimal policy is computed as follows.

We initially compute the minimal average cost of a certain number of periods. Then we keep increasing a fixed number of periods, computing the minimal average cost of these periods, and comparing the deviation of the two costs. The iteration does not end until the deviation is insensitive to the increments of periods. It can shown that the long-run average cost converges to a constant as the number of periods increases. We take this constant as the optimal cost, which is the minimal cost among all admissible polices. This kind of method is known as value-iteration method; see Bertsekas et al. [10] for more details.

We compare the long-run average costs of the (s, k) policy and the S policy to the optimal cost. Denote the average cost of the (s, k) policy by $C_{s,k}$, and C_S for the S policy, and the optimal cost by C_{OPT}. For each instance, we use G_1 and G_3 to denote the gaps between the costs of these policies, as follows:

$$G_1 = \frac{C_{s,k} - C_{OPT}}{C_{OPT}} * 100\%$$

and

$$G_3 = \frac{C_S - C_{OPT}}{C_{OPT}} * 100\%$$

Another alternative is the (s, S) policy with $S - s = M$ due to the minimum order quantity constraint. This policy is often used in practice to control inventories when there is a minimum order quantity constraint, see Zhou et al. [99] and Robb and Silver [75]. Kiesmüller et al. [59] also compare their proposed heuristic policy to this (s, S) policy. Due to the batch size constraint, in our (s, S) policy, when the inventory position is less than or equal to s, an order is placed to bring the inventory position above S but no more than $S + Q$; otherwise do not order. If x_t denotes the inventory position in period n before ordering, the order quantity q_t can be described as follows:

$$q_t = \begin{cases} 0, & x_t > s; \\ M + m \times Q, & x_t \leq s. \end{cases}$$

where m is the smallest integer subject to $x_t + q_t > S$. Note that this (s, S) policy is in fact a special case of our (s, k) policy where $\Delta = k - s = 0$ (assume Δ is allowed to be 0).

For each instance, we also test the performance of the (s, k) policy and the S policy by comparing them to the best (s, S) policy. Denote the average cost for the best (s, S) policy by $C_{s,S}$, and use G_2 and G_4 to denote the gaps between the costs of these policies, as follows:

$$G_2 = \frac{C_{s,S} - C_{s,k}}{C_{s,k}} * 100\%$$

and

$$G_4 = \frac{C_{s,S} - C_S}{C_S} * 100\%$$

We calculate the average gap, the maximal gap, and the minimal gap, which are denoted by avgG_j, minG_j, and maxG_j ($j = 1, 2, 3, 4$), respectively.

4.6.1 Performance of the (s, k) Policy

The numerical results of the examples of the (s, k) policy for normally distributed demand are given in Table 4.2. Table 4.2 summarizes the gaps G_1 between the optimal policy and the (s, k) policy, and the gaps G_2 between the (s, k) policy and the best (s, S) policy. In each row, only one parameter is fixed while all others vary. For example, for $Q = 6$, all the other parameters in this row vary and result in 80 instances of normal demand.

Factor	Value	avgG_1	maxG_1	avgG_2	minG_2	maxG_2
Q	3	1.43	25.11	37.15	1.11	150.98
	5	1.32	24.32	41.76	3.75	182.92
	6	1.29	23.52	44.68	7.23	235.18
	10	0.91	20.33	54.72	25.65	266.56
	15	0.71	19.03	59.45	30.18	228.50
$\mathbb{E}(D)$	10	0.10	2.12	29.56	16.50	49.49
	15	0.52	5.63	34.56	17.83	53.79
	20	0.71	6.05	27.82	1.11	70.01
	30	4.29	25.11	66.52	29.27	150.98
	40	0.04	0.20	79.28	18.40	266.56
c.v.	0.1	4.41	25.11	72.31	1.11	266.56
	0.2	0.08	1.12	53.40	18.26	112.41
	0.3	0.01	0.01	36.72	17.95	58.03
	0.4	0.03	0.20	27.77	16.50	42.91
p/(h+p)	0.8	0.70	13.59	49.23	4.20	228.50
	0.85	0.96	16.87	48.51	2.19	266.56
	0.9	1.28	20.86	46.94	1.11	255.01
	0.95	1.59	25.11	45.52	1.56	235.18

Table 4.2: Performance of the (s, k) policy (normal distribution)

It can be seen that the performance of the (s, k) policy is quite close to that of the optimal polices in most cases. For example, at c.v. $= 0.2$, the maximum G_1 is 1.12%. G_1 tends to decrease as Q increases. When c.v. takes the value of 0.2, 0.3, or 0.4, the average G_1 is less than 0.1%. These results indicate that the (s, k) policy performs better when the c.v. values are relatively large.

Table 4.2 shows that the performance of the (s, k) policy is close to the optimal performance on average for normally distributed demand. However, there are still a few cases in which the (s, k) policy performs much worse than the optimal policy. When the coefficient variation of demand is small, e.g., c.v. $= 0.1$, the maximum gap G_1 between the optimal policy and the (s, k) policy is 25.11% when $Q = 3$, $\mathbb{E}(D) = 30$, and $p/(h + p) = 0.95$. This is the worst case over all 400 instances.

Table 4.2 also demonstrates that the (s, k) policy performs much better than the best (s, S) policy over all examples. The gap, G_2, between the (s, k) policy and the best (s, S) policy can be quite substantial. For example, the maximum G_2 is 266.56% at $Q = 10$, $\mathbb{E}(D) = 40$, c.v. $= 0.1$, and $p/(h + p) = 0.85$. The average G_2 with respect to c.v.=0.4 is 27.77%, while the average G_2 with respect to the other parameters are all larger than this value. These results indicate that the (s, k) policy always outperforms the best (s, S) policy.

These observations are consistent with those of Zhou et al. [99], who do not consider batch ordering. By conducting numerical examples, Zhou et al. [99] show that their two-parameter policy always outperforms the best (s,S) policy, and the performance of their policy improves for normally distributed demand as c.v. increases.

We next present the results of the examples for Poisson distributed demand as shown in Table 4.3. The gap, G_1, between the optimal policy and

Factor	Value	avgG_1	maxG_1	avgG_2	minG_2	maxG_2
Q	3	0.34	2.60	47.64	17.39	122.72
	5	0.31	2.47	53.91	20.44	134.40
	6	0.29	2.41	56.76	22.04	141.50
	10	0.26	2.09	64.46	28.83	156.88
	15	0.07	0.94	67.41	39.35	150.11
$\mathbb{E}(D)$	10	0.00	0.00	27.59	17.39	44.98
	15	0.05	0.21	31.55	22.11	46.03
	20	0.01	0.01	30.88	19.66	49.73
	30	1.21	2.60	75.15	59.32	90.01
	40	0.00	0.01	125.00	71.37	156.88
p/(h+p)	0.8	0.08	0.49	59.82	19.32	156.88
	0.85	0.19	1.18	59.45	19.06	154.80
	0.9	0.31	1.94	58.48	18.32	145.42
	0.95	0.44	2.60	54.39	17.39	133.18

Table 4.3: Performance of the (s, k) policy (Poisson distribution)

the (s, k) policy is not larger than 2.60% over all instances. As Q increases, G_1 tends to decrease. The results in Table 4.3 indicate that the (s, k) policy achieves nearly minimal cost for Poisson distributed demand.

Table 4.3 also demonstrates that the (s, k) policy outperforms the best (s, S) policy for Poisson distributed demand. The maximum gap G_2 ranges from 44.98% to 156.88%. As Q increases, G_2 tends to increase. Even though G_2 tends to decrease as $p/(h+p)$ increases, the minimum G_2 is still as large as 17.39%. These results show that the (s, k) policy performs much better than the best (s, S) policy for Poisson distributed demand.

4.6.2 Performance of the S Policy

The numerical results of the examples of the S policy for normally distributed demand are given in Table 4.4. Table 4.4 shows that the performance of the S policy is close to that of the optimal policy except for a few cases at c.v. = 0.1. For example, the maximum G_3 are 9.51%, 6.61%, and 5.19%, at c.v.

= 0.2, 0.3, and 0.4, respectively. The corresponding average G_3 are all at the 1% level. G_3 tends to decreases, as Q increases. This is identical with our former analysis for the (s, k) policy, recalling that the S policy is a special case. Table 4.4 indicates that the S policy performs well in the cases with relatively large values of demand coefficient variation.

Factor	Value	avgG_3	maxG_3	avgG_4	minG_4	maxG_4
Q	3	4.38	40.93	32.91	-0.73	102.28
	5	3.64	35.28	38.36	3.75	182.92
	6	3.34	33.33	41.64	7.23	235.18
	10	1.96	23.67	53.09	24.18	266.56
	15	1.16	19.47	58.70	29.34	228.50
$\mathbb{E}(D)$	10	0.48	5.50	29.09	16.49	49.49
	15	2.51	16.08	31.99	17.77	50.28
	20	1.05	7.75	27.40	-0.73	70.01
	30	9.88	40.93	57.64	23.91	102.28
	40	0.56	2.21	78.59	15.94	266.56
c.v.	0.1	8.53	40.93	66.09	-0.73	266.56
	0.2	1.15	9.51	51.69	18.21	112.24
	0.3	0.90	6.61	35.45	17.94	55.94
	0.4	1.01	5.19	26.54	15.94	42.91
p/(h+p)	0.8	3.03	40.93	45.55	4.20	228.50
	0.85	2.97	38.86	45.50	2.19	266.56
	0.9	2.89	36.15	44.67	0.34	255.01
	0.95	2.70	33.82	44.06	-0.73	235.18

Table 4.4: Performance of the S policy (normal distribution)

However, in the cases with small values of demand coefficient variation, e.g. at c.v. = 0.1, the S policy may not perform very well. For example, the average G_3 is 8.53%, and the maximum G_3 is 40.93% at c.v. = 0.1.

Table 4.4 also demonstrates that the S policy outperforms the best (s, S) policy for normally distributed demand in most instances. Over all 400 instances, there is only one exception, where G_4 = -0.73% at Q = 3, $\mathbb{E}(D)$ = 20, c.v. = 0.1, and $p/(h+p)$ = 0.95. In other instances, G_4 can be substantial. For example, the average G_4 with respect to Q varies between

32.91% and 58.70%. As Q increases, G_4 tends to increase. Recalling that G_3 tends to decrease as Q increases, this may lead to the conclusion that the S policy performs better when Q is relatively large.

We next test the performance of the S policy for Poisson distributed demand as shown in Table 4.5. We find that the average G_3 is not larger than 5.50% over all instances. The maximum G_3 is 10.22%, which is attained at $Q = 3$, $\mathbb{E}(D) = 30$, and $p/(h + p) = 0.8$.

Factor	Value	avgG_3	maxG_3	avgG_4	minG_4	maxG_4
Q	3	1.83	10.22	45.19	17.39	122.72
	5	1.48	8.36	51.96	20.44	134.40
	6	1.32	7.45	55.04	22.04	141.50
	10	0.81	4.65	63.53	28.83	156.88
	15	0.27	1.86	67.06	39.35	150.11
$\mathbb{E}(D)$	10	0.00	0.01	27.59	17.39	44.98
	15	0.05	0.25	31.54	22.07	46.03
	20	0.16	0.38	30.69	19.48	49.72
	30	5.50	10.22	67.96	57.28	78.31
	40	0.00	0.01	125.00	71.36	156.88
p/(h+p)	0.8	1.34	10.22	57.62	19.31	156.88
	0.85	1.20	9.21	57.76	19.05	154.80
	0.9	1.08	8.35	57.24	18.32	145.42
	0.95	0.95	7.18	53.61	17.39	133.18

Table 4.5: Performance of the S policy (Poisson distribution)

Table 4.5 also shows that the S policy performs better than the best (s, S) policy for Poisson distributed demand in all instances. The gap, G_4, between the S policy and the best feasible (s, S) policy is substantial. Over all 400 instances, G_4 varies between 17.39% and 156.88%.

These observations are consistent with Kiesmüller et al. [59], recalling that when Q=1, our S policy reduces to the one-parameter policy of Kiesmüller et al. [59], who also show, by the way of numerical study, the superiority of the one-parameter policy to the (s, S) policy.

The above-mentioned numerical examples are all taken when $M = 30$.

We next present the results of the (s, k) policy and the S policy for different values of M. In the remainder of this subsection, Q varies as 3, 5, and 6 and M can take different values. Because M is an integral multiple of Q, we let M take the values of 10, 15, and 20 when $Q = 5$, otherwise M varies as 12, 18, and 24. $\mathbb{E}(D)$ varies as 5, 10, 15, 20, and 25. Other parameters take the same values as shown in Table 4.1, resulting in a total of 720 instances for normally distributed demand. Table 4.6 shows the performance of the (s, k) policy and the S policy with different values of M. From Table 4.6, we can draw several similar conclusions to the case with M=30. First, both the (s, k) policy and the S policy have good performance close the the optimal policy in most cases. For example, the average G_1 and G_3 with respect to different M are all at the 1% level. Second, the (s, k) policy and the S policy perform much worse than the optimal policy when the demand variation is small, i.e. c.v. $= 0.1$. The average G_1 and G_3 are 2.29% and 2.31%, respectively, at c.v. $= 0.1$, while they are all less than the 0.1% at other c.v. levels. The maximum G_1 and G_3 are 21.77% and 29.22%, which are also much larger than the values at higher c.v. levels. We should also note that for a fixed Q, G_1 and G_3 tend to increase as M increases. For example, when Q=5, the average G_1 increases from 0.73% to 1.15% as M increases from 10 to 20.

4.6.3 Sensitivity Analysis

In this section, we first compare the performances of the (s, k) policy and the S policy. In fact, in some instances, they have the same performance, because the S policy is a special case of the (s, k) policy. However, we are still interested in the differences between the policies over all instances. We use G_5 to denote the gap between the (s, k) policy and the S policy as follows:

$$G_5 = \frac{C_S - C_{s,k}}{C_{s,k}} * 100\%$$

Factor	Value	avgG_1	maxG_1	avgG_3	maxG_3
Q	3	0.51	9.00	0.52	15.53
	5	0.90	21.77	0.91	29.22
	6	0.37	6.83	0.38	6.83
$\mathbb{E}(D)$	5	0.01	0.33	0.01	0.33
	10	0.78	18.38	0.78	19.66
	15	1.13	20.73	1.13	22.73
	20	0.91	21.77	0.92	29.22
	25	0.15	3.24	0.16	6.95
c.v.	0.1	2.29	21.77	2.31	29.22
	0.2	0.08	1.08	0.09	6.95
	0.3	0.01	0.03	0.02	5.37
	0.4	0.01	0.02	0.02	4.33
p	0.8	0.45	11.91	0.47	29.22
	0.85	0.54	14.24	0.56	26.93
	0.9	0.67	18.79	0.68	26.34
	0.95	0.72	21.77	0.73	25.01
M	10	0.73	18.38	0.73	19.66
	15	0.83	20.73	0.83	22.73
	20	1.15	21.77	1.16	29.22
	12	0.24	7.01	0.24	7.01
	18	0.40	8.31	0.40	15.53
	24	0.69	9.00	0.70	9.04

Table 4.6: Performance of the two policies for different values of M

To better illustrate the differences, we take all the instances at $Q = 5$ as an example, and enumerate the results in Table 4.7. It can be seen that the (s, k) policy does not perform significantly better than the S policy except for a few cases with small values of demand coefficient variation. For example, when c.v. $= 0.1$, the maximum G_5 is 20.42% at $\mathbb{E}(D) = 30$. However, when c.v. is relatively large, the S policy performs nearly as well as the (s, k) policy. The maximum G_5 is less than 8%, and in many instances these two policies have the same performance. These results indicate that the S policy is a recommendable substitute for the (s, k) policy at a high c.v. level. This is also consistent with the finding of Kiesmüller et al. [59] that the performance of the one-parameter policy is close to the performance of the two-parameter policy when there is no batch ordering constraint.

G_5	p/(h+p)	E(D)=10	E(D)=15	E(D)=20	E(D)=30	E(D)=40
Normal	0.80	2.07	10.82	0.00	20.42	0.00
c.v.=0.1	0.85	2.73	10.76	0.00	15.35	0.00
	0.90	2.88	10.25	0.00	10.05	0.00
	0.95	1.36	7.84	0.36	4.66	0.00
Normal	0.80	0.00	1.25	0.13	7.75	0.09
c.v.=0.2	0.85	0.00	1.53	0.15	6.34	0.07
	0.90	0.00	1.60	0.09	4.87	0.05
	0.95	0.00	1.64	0.02	3.40	0.03
Normal	0.80	0.00	0.02	0.69	5.35	1.14
c.v.=0.3	0.85	0.00	0.02	0.53	4.64	1.06
	0.90	0.00	0.01	0.49	4.00	0.90
	0.95	0.00	0.01	0.31	3.05	0.68
Normal	0.80	0.00	0.16	1.37	4.33	1.86
c.v.=0.4	0.85	0.00	0.12	1.19	3.86	1.67
	0.90	0.00	0.10	0.88	3.36	1.48
	0.95	0.00	0.04	0.64	2.70	1.22
Poisson	0.80	0.00	0.01	0.26	7.86	0.00
	0.85	0.00	0.01	0.29	6.23	0.00
	0.90	0.00	0.01	0.21	5.01	0.00
	0.95	0.00	0.00	0.10	3.22	0.00

Table 4.7: Comparison of the (s, k) policy and the S policy at $Q = 5$

We next study the effect of the $\mathbb{E}(D)$. To better illustrate the effect of $\mathbb{E}(D)$, we plot in Figure 4.1 the long-run average costs (normalized by the global minimum of $C(y)$) of the (s, k) policy as a function of $\mathbb{E}(D)$ under different levels of c.v. with a fixed $M = 30$. Figure 4.1 illustrates that for a given M,

- when $\mathbb{E}(D)$ is relatively small, e.g., less than M, a lower level of demand variation leads to a higher decreasing rate of average costs as $\mathbb{E}(D)$ increases.

- when $\mathbb{E}(D)$ is relatively large, e.g., larger than M, the average cost is not affected significantly by the different demand variation levels as $\mathbb{E}(D)$ increases.

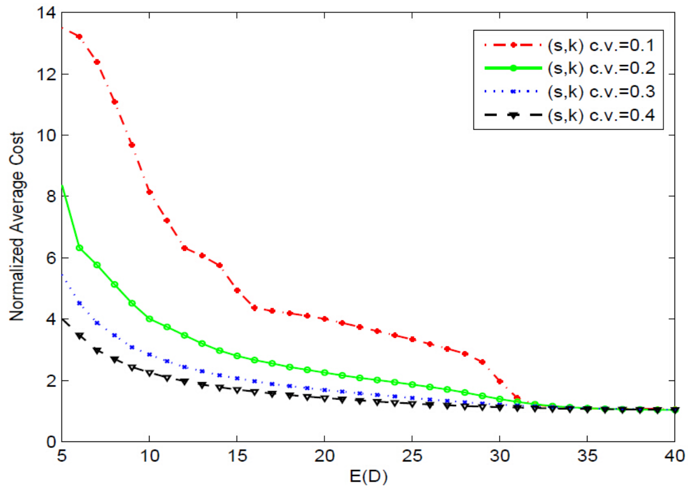

Figure 4.1: The impact of $\mathbb{E}(D)$ in the (s,k) policy with $p/(h+p) = 0.9$ and $Q = 5$

The reason for this is that as $\mathbb{E}(D)$ increases, the constraint of the minimum order quantity becomes looser and looser. Similar results can be drawn for the S policy. We plot the results for the S policy in Figure 4.2.

Before ending this section, we must point out that the case of $\Delta < Q$ can appear, although $\Delta \geq Q$ in most of our numerical examples. For example, for normally distributed demand with $\mathbb{E}(D) = 100$ and $c.v. = 1$, we can get the optimal $\Delta = 1$ and the optimal $k = 183$, when $M = 8$, $Q = 4$, and $p/(h+p) = 0.8$.

4.7 Conclusion

In this work, we design an algorithm for a heuristic two-parameter policy, the (s,k) policy, to control stochastic inventories with minimum order quantity and batch size constraints. Applying a Markov chain approach, we compute

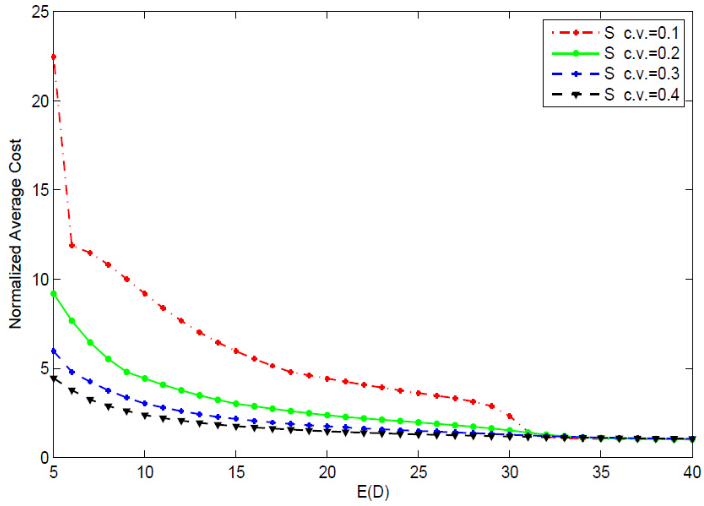

Figure 4.2: The impact of $\mathbb{E}(D)$ in the S policy with $p/(h+p) = 0.8$ and $Q = 3$

the system costs and provide recursive algorithms to optimize the policy under the long-run average cost criterion. We also develop the computational procedure for a simpler policy, the S policy, which is motivated by the (R, Q) policy in batch ordering and is a special case of the (s, k) policy. Numerical studies are conducted to demonstrate the effectiveness of these two policies with respect to the optimal policy and the (s, S) policy for both normally and Poisson distributed demand. Overall, the S policy has a good performance close to that of the (s, k) policy; only in a few cases with small demand variation, the latter outperforms the former significantly.

For future directions, one may try to develop some heuristic policies with easily computable performance bounds. These performance bounds provide a worst-case optimal gap between the optimal policy and the heuristics, and thus are of great importance. Moreover, the performance bounds are allowed to depend on the optimal solutions of single-stage models, which can be easily computed as shown in Zheng [97]. Second, one try to study multi-echelon inventory systems with a MOQ constraint. A possible start-up model is a two-stage serial inventory system with a MOQ requirement at the upstream installation, which is the basis of distribution or even more general systems.

Bibliography

[1] S. Athey. Monotone comparative statics under uncertainty. *Quarterly Journal of Economics*, pages 187–223, 2002.

[2] S. Axäster. Evaluation of installation stock based (r, q)-policies for two-level inventory systems with poisson demand. *Operations Research*, 46 (3), 1998.

[3] S. Axsäter. Simple solution procedures for a class of two-echelon inventory problems. *Operations Research*, 38(1):64–69, 1990.

[4] S. Axsäter. Simple evaluation of echelon stock (r, q) policies for two-level inventory systems. *IIE transactions*, 29(8):661–669, 1997.

[5] S. Axsäter. Exact analysis of continuous review (r, q) policies in two-echelon inventory systems with compound poisson demand. *Operations Research*, 48(5):686–696, 2000.

[6] S. Axsäter. *Inventory control*, volume 90. Springer, 2006.

[7] S. Axsäter and W.-F. Zhang. A joint replenishment policy for multi-echelon inventory control. *International Journal of Production Economics*, 59(1):243–250, 1999.

[8] S. Axsäter, J. Marklund, and E. A. Silver. Heuristic methods for centralized control of one-warehouse, n-retailer inventory systems. *Manufacturing & Service Operations Management*, 4(1):75–97, 2002.

[9] B. B. Bederson, G. Z. Jin, P. Leslie, A. J. Quinn, and B. Zou. Incomplete disclosure: Evidence of signaling and countersignaling. *The American Economic Journal: Microeconomics*, forthcoming, 2017.

[10] D. P. Bertsekas, D. P. Bertsekas, D. P. Bertsekas, and D. P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific Belmont, 1995.

[11] O. Board. Competition and disclosure. *The Journal of Industrial Economics*, 57(1):197–213, 2009.

[12] R. A. Broekmeulen and K. H. van Donselaar. A heuristic to manage perishable inventory with batch ordering, positive lead-times, and time-varying demand. *Computers & Operations Research*, 36(11):3013–3018, 2009.

[13] A. L. Brown, C. F. Camerer, and D. Lovallo. To review or not to review? limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, 4(2):1–26, 2012.

[14] E. Chan and J. Muckstadt. The effects of load smoothing on inventory levels in a capacitated production and inventory system. Technical report, Cornell University Operations Research and Industrial Engineering, 1999.

[15] L. M. A. Chan, A. Muriel, Z.-J. M. Shen, D. Simchi-Levi, and C.-P. Teo. Effective zero-inventory-ordering policies for the single-warehouse multiretailer problem with piecewise linear cost structures. *Management Science*, 48(11):1446–1460, 2002.

[16] X. Chao and S. X. Zhou. Optimal policy for a multiechelon inventory system with batch ordering and fixed replenishment intervals. *Operations Research*, 57(2):377–390, 2009.

[17] F. Chen. Effectiveness of (r, q) policies in one-warehouse multiretailer systems with deterministic demand and backlogging. *Naval Research Logistics (NRL)*, 47(5):422–439, 2000.

[18] F. Chen. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations Research*, 48(3):376–389, 2000.

[19] F. Chen and Y.-S. Zheng. Lower bounds for multi-echelon stochastic inventory systems. *Management Science*, 40(11):1426–1443, 1994.

[20] F. Chen and Y.-S. Zheng. One-warehouse multiretailer systems with centralized stock information. *Operations Research*, 45(2):275–287, 1997.

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

[22] L. Y. Chu and Z.-J. M. Shen. A power-of-two ordering policy for one-warehouse multiretailer systems with stochastic demand. *Operations Research*, 58(2):492–502, 2010.

[23] A. J. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science*, 6(4):475–490, 1960.

[24] M. A. De Bodt and S. C. Graves. Continuous-review policies for a multi-echelon inventory problem with stochastic demand. *Management Science*, 31(10):1286–1299, 1985.

[25] M. K. Doğru, A. De Kok, and G. Van Houtum. A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems. *Flexible Services and Manufacturing Journal*, 21(3-4):114–147, 2009.

[26] D. Dranove and G. Z. Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–63, 2010.

[27] B. Duermeyer and L. Schwarz. A model for the analysis of system service level in warehouse retailer distribution systems: the identical retailer case. *Multi-level Production/Inventory Control Systems: Theory and Practice*, 16:163–193, 1981.

[28] A. Federgruen and Y.-S. Zheng. An efficient algorithm for computing an optimal (r, q) policy in continuous review stochastic inventory systems. *Operations Research*, 40(4):808–813, 1992.

[29] A. Federgruen and P. Zipkin. Allocation policies and cost approximations for multilocation inventory systems. *Naval Research Logistics Quarterly*, 31(1):97–129, 1984.

[30] A. Federgruen and P. Zipkin. Approximations of dynamic, multilocation production and inventory problems. *Management Science*, 30(1):69–84, 1984.

[31] A. Federgruen and P. Zipkin. Computational issues in an infinite-horizon, multiechelon inventory model. *Operations Research*, 32(4):818–836, 1984.

[32] N. Feltovich, R. Harbaugh, and T. To. Too cool for school? signalling and countersignalling. *RAND Journal of Economics*, 33(4):630–650, 2002.

[33] K. Fiedler and P. Juslin. Taking the interface between mind and environment seriously. *Information Sampling and Adaptive Cognition*, pages 3–29, 2006.

[34] M. Fisher and A. Raman. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research*, 44(1): 87–99, 1996.

[35] M. J. Fishman and K. M. Hagerty. Mandatory versus voluntary disclosure in markets with informed and uninformed customers. *Journal of Law, Economics, and Organization*, 19(1):45–63, 2003.

[36] R. Forsberg. Exact evaluation of (r, q)-policies for two-level inventory systems with poisson demand. *European Journal of Operational Research*, 96(1):130–138, 1997.

[37] G. Gallego. New bounds and heuristics for (q, r) policies. *Management Science*, 44(2):219–233, 1998.

[38] G. Gallego, Ö. Özer, and P. Zipkin. Bounds, heuristics, and approximations for distribution systems. *Operations Research*, 55(3):503–517, 2007.

[39] S. C. Graves. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science*, 31(10):1247–1256, 1985.

[40] S. J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24 (3):461–483, 1981.

[41] L. Guo and Y. Zhao. Voluntary quality disclosure and market interaction. *Marketing Science*, 28(3):488–501, 2009.

[42] R. Hassin and M. Haviv. *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59. Springer Science & Business Media, 2003.

[43] B. Hellion, F. Mangione, and B. Penz. A polynomial time algorithm to solve the single-item capacitated lot sizing problem with minimum order quantities and concave costs. *European Journal of Operational Research*, 222(1):10–16, 2012.

[44] D. Heyman and M. Sobel. Stochastic models in operations research. vol. ii: Stochastic optimization. 1984.

[45] D. Hirshleifer and S. H. Teoh. Limited attention, information disclosure, and financial reporting. *Journal of accounting and economics*, 36(1): 337–386, 2003.

[46] D. L. Hoffman and M. Fodor. Can you measure the roi of your social media marketing? *MIT Sloan Management Review*, 52(1):41, 2010.

[47] M. Hu and Y. Yang. Modified echelon (r, q) policies with guaranteed performance bounds for stochastic serial inventory systems. *Operations Research*, 62(4):812–828, 2014.

[48] M. Hu, M. Shi, and J. Wu. Simultaneous vs. sequential group-buying mechanisms. *Management Science*, 59(12):2805–2822, 2013.

[49] M. Hu, J. Milner, and J. Wu. Liking and following and the newsvendor: Operations and marketing policies under social influence. *Management Science*, 62(3):867–879, 2015.

[50] M. Hu, Y. Li, and J. Wang. Efficient ignorance: Information heterogeneity in a queue. *Management Science*, 2017.

[51] T. Huang and Y.-J. Chen. Service systems with experience-based anecdotal reasoning customers. *Production and Operations Management*, 24 (5):778–790, 2015.

[52] T. Huang and Q. Liu. Strategic capacity management when customers have boundedly rational expectations. *Production and Operations Management*, 24(12):1852–1869, 2015.

[53] T. Huang and Y. Yu. Sell probabilistic goods? a behavioral explanation for opaque selling. *Marketing Science*, 33(5):743–759, 2014.

[54] W. T. Huh and G. Janakiraman. Technical noteąłon optimal policies for inventory systems with batch ordering. *Operations Research*, 60(4): 797–802, 2012.

[55] W. T. Huh, G. Janakiraman, J. A. Muckstadt, and P. Rusmevichientong. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*, 55(3):404–420, 2009.

[56] J. W. Huppertz and J. P. Carlson. Consumers' use of hcahps ratings and word-of-mouth in hospital choice. *Health Service Research*, 45(6p1): 1602–1613, 2010.

[57] G. Z. Jin. Competition and disclosure incentives: an empirical study of hmos. *The Rand Journal of Economics*, 36(1):93, 2005.

[58] B. Jovanovic. Truthful disclosure of information. *The Bell Journal of Economics*, pages 36–44, 1982.

[59] G. Kiesmüller, A. De Kok, and S. Dabia. Single item inventory control under periodic review and a minimum order quantity. *International Journal of Production Economics*, 133(1):280–285, 2011.

[60] C.-Y. Lee. Inventory replenishment model: lot sizing versus just-in-time delivery. *Operations Research Letters*, 32(6):581–590, 2004.

[61] Y.-J. Lee, K. Hosanagar, and Y. Tan. Do i follow my friends or the

crowd? information cascades in online movie ratings. *Management Science*, 61(9):2241–2258, 2015.

[62] R. Levi, R. Roundy, D. Shmoys, and M. Sviridenko. A constant approximation algorithm for the one-warehouse multiretailer problem. *Management Science*, 54(4):763–776, 2008.

[63] M. Luca and J. Smith. Strategic disclosure: The case of business school rankings. *Journal of Economic Behavior & Organization*, 112:17–25, 2015.

[64] D. Lupton, C. Donaldson, and P. Lloyd. Caveat emptor or blissful ignorance? patients and the consumerist ethos. *Social Science & Medicine*, 33(5):559–568, 1991.

[65] A. D. Mathios. The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market. *Journal of Law and Economics*, 43(2):651–677, 2000.

[66] S. Matthews and A. Postlewaite. Quality testing and disclosure. *The RAND Journal of Economics*, 16(3):328–340, 1985.

[67] P. R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.

[68] P. Nelson. Information and consumer behavior. *Journal of Political Economy*, 78(2):311–329, 1970.

[69] I. Okhrin and K. Richter. An o (t 3) algorithm for the capacitated lot sizing problem with minimum order quantities. *European Journal of Operational Research*, 211(3):507–514, 2011.

[70] M. J. Osborne and A. Rubinstein. Games with procedurally rational players. *American Economic Review*, pages 834–847, 1998.

[71] Y. Papanastasiou and N. Savva. Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*, 63(4): 919–939, 2016.

[72] E. Peters, J. Hibbard, P. Slovic, and N. Dieckmann. Numeracy skill and the communication, comprehension, and use of risk-benefit information. *Health Affairs*, 26(3):741–748, 2007.

[73] E. Porras and R. Dekker. An efficient optimal solution method for the joint replenishment problem with minimum order quantities. *European Journal of Operational Research*, 174(3):1595–1615, 2006.

[74] E. L. Porteus and S. Whang. *Supply chain contracting: non-recurring engineering charge, minimum order quantity, and boilerplate contracts.* Graduate School of Business, Stanford University, 1999.

[75] D. J. Robb and E. A. Silver. Inventory management with periodic ordering and minimum order quantities. *Journal of the Operational Research Society*, pages 1085–1094, 1998.

[76] R. Roundy. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Science*, 31(11):1416–1430, 1985.

[77] A. Schwartz. How much irrationality does the market permit? *The Journal of Legal Studies*, 37(1):131–159, 2008.

[78] L. B. Schwarz and L. Schrage. Optimal and system myopic policies for multi-echelon production/inventory assembly systems. *Management Science*, 21(11):1285–1294, 1975.

[79] K. H. Shang and S. X. Zhou. Optimal and heuristic echelon (r, nq, t) policies in serial inventory systems with fixed costs. *Operations Research*, 58(2):414–427, 2010.

[80] K. H. Shang, Z. Tao, and S. X. Zhou. Optimizing reorder intervals for two-echelon distribution systems with stochastic demand. *Operations Research*, 63(2):458–475, 2015.

[81] S. Shavell. Acquisition and disclosure of information prior to sale. *RAND Journal of Economics*, 25(1):20–36, 1994.

[82] C. C. Sherbrooke. Metric: A multi-echelon technique for recoverable item control. *Operations Research*, 16(1):122–141, 1968.

[83] D. Simchi-Levi and Y. Zhao. Performance evaluation of stochastic multi-echelon inventory systems: A survey. *Advances in Operations Research*, 2012, 2011.

[84] H. A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118, 1955.

[85] R. Spiegler. Competition over agents with boundedly rational expectations. *Theoretical Economics*, 1(2):207–231, 2006.

[86] R. Spiegler. The market for quacks. *The Review of Economic Studies*, 73(4):1113–1131, 2006.

[87] M. Spranca, D. E. Kanouse, M. Elliott, P. F. Short, D. O. Farley, and R. D. Hays. Do consumer reports of health plan quality affect health plan selection? *Health Service Research*, 35(5 Pt 1):933, 2000.

[88] A. Stivers. Unraveling of information: Competition and uncertainty. *Topics in Theoretical Economics*, 4(1):1–16, 2004.

[89] N. Szech. Becoming a bad doctor. *Journal of Economic Behavior & Organization*, 80(1):244–257, 2011.

[90] T. L. Tuten and M. R. Solomon. *Social Media Marketing*. SAGE, 2014.

[91] A. Tversky and D. Kahneman. Belief in the law of small numbers. *Psychological Bulletin*, 76(2):105, 1971.

[92] A. Veinott Jr. The optimal inventory policy for batch ordering. *Operations Research*, 13(3):424–432, 1965.

[93] M. Xiao. Is quality accreditation effective? evidence from the childcare market. *International Journal of Industrial Organization*, 28(6):708–721, 2010.

[94] M. Yu, L. Debo, and R. Kapuscinski. Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *Management Science*, 62(2):410–435, 2015.

[95] Y. Yu, S. Benjaafar, and Y. Gerchak. Capacity sharing and cost allocation among independent firms with congestion. *Production and Operations Management*, 24(8):1285–1310, 2015.

[96] Y. Zhao and M. Katehakis. On the structure of optimal ordering policies for stochastic inventory systems with minimum order quantity. *Probability in the Engineering and Informational Sciences*, 20(2):257, 2006.

[97] Y.-S. Zheng. On properties of stochastic inventory systems. *Management Science*, 38(1):87–103, 1992.

[98] Y.-S. Zheng and A. Federgruen. Finding optimal (s, s) policies is about as simple as evaluating a single policy. *Operations Research*, 39(4):654–665, 1991.

[99] B. Zhou, Y. Zhao, and M. Katehakis. Effective control policies for stochastic inventory systems with a minimum order quantity and linear costs. *International Journal of Production Economics*, 106(2):523–531, 2007.

Appendix A

Proofs

Proof of Lemma 2.1. Note that it is optimal to let

$$\alpha v_H + (1-\alpha)v_L - p_D = 0,$$

i.e., to provide zero surplus for customers. This is because the objective function is increasing in p_D. Then we have

$$p_D = \alpha v_H + (1-\alpha)v_L.$$

Since $\pi_D = p_D\lambda - K$, we have

$$\pi_D = (\alpha v_H + (1-\alpha)v_L)\lambda - K.$$

Therefore, the optimal price is $p_D^* = \alpha v_H + (1-\alpha)v_L$ and the optimal profit is $\pi_D^*(\lambda) = [\alpha v_H + (1-\alpha)v_L]\lambda - K$.▮

Proof of Lemma 2.2. The firm can make a nonnegative profit by adopting the disclosure strategy, if and only if $\pi_D^* = \pi_D^*(\overline{\lambda}) = (\alpha v_H + (1-\alpha)v_L)\overline{\lambda} - K \geq 0$, i.e., $\overline{\lambda} \geq \tilde{\lambda}_D$, where $\tilde{\lambda}_D$ is defined as $\tilde{\lambda}_D = \frac{K}{\alpha v_H + (1-\alpha)v_L}$; otherwise when $\overline{\lambda} < \tilde{\lambda}_D$, $\pi_D^* < 0$, and the firm should not enter the market.▮

Proof of Lemma 2.3. Recall that if the customer purchases the service, the samples he obtained must satisfy $\alpha_i(N) \geq \frac{p-v_L}{v_H-v_L}$ Therefore, if a customer obtains N samples and decides to purchase the service, the proportion of H samples must be at least $\frac{p-v_L}{v_H-v_L}$. In other words, among the N samples the customer obtained, the number of H samples should be at least $N\frac{p-v_L}{v_H-v_L}$. Because the number is an integer, this means the number of H samples should be larger than $\lfloor N\frac{p-v_L}{v_H-v_L} - \epsilon \rfloor$. Therefore, if the number of H samples is larger than $\lfloor N\frac{p-v_L}{v_H-v_L} - \epsilon \rfloor$, the customer purchases the service; otherwise, the customer would not purchase the service. ∎

Proof of Lemma 2.4. (i) Observing

$$\gamma(p) = 1 - \sum_{n=0}^{\lfloor N\frac{p-v_L}{v_H-v_L} - \epsilon \rfloor} B(n, N, \alpha),$$

it follows that it is optimal for the firm to set the price p such that $N\frac{p-v_L}{v_H-v_L}$ is an integer due to the floor function. This is because if the term is not an integer, the firm can make more profits by increasing p until the term is an integer with γ unchanged. Therefore, $N\frac{p-v_L}{v_H-v_L}$ must be an integer. Denote this integer as j, i.e., $j = N\frac{p-v_L}{v_H-v_L}$, and it follows that $p^* = v_L + \frac{j}{N}(v_H - v_L), j = 0, 1, 2, \ldots, N$.

(ii) For any given j, we know that the optimal price p_j^* for this j must satisfy $p_j^* = v_L + \frac{j}{N}(v_H - v_L)$. By substituting the preceding formula in $\pi(p) = \gamma p\overline{\lambda}$, then the profit function for this particular j, π_j^*, can be expressed as $\pi_j^* = [v_L + \frac{j}{N}(v_H - v_L)]\gamma_j\overline{\lambda}$. We do the same computation for each $j = 0, 1, 2, \ldots, N$, and find the largest one (denoted as $\pi_{j^*}^*$) among all π_j^*. If $\pi_{j^*}^* \geq 0$, then the optimal price p^* is $p_{j^*}^*$. ∎

Proof of Proposition 2.1. We consider two special cases of the nondisclosure strategy.

Case 1: suppose $j = 0$. In this case, $\lim_{\alpha \to 0} \pi_j^* - \pi_D^* = K \geq 0$, i.e., the firm is better off by adopting the nondisclosure strategy. Because of the continuity $\pi_0^* - \pi_D^*$, it follows that there exists a positive $\underline{\alpha}$ such that it is optimal to adopt the nondisclosure strategy for $\alpha \in [0, \underline{\alpha}]$.

Case 2: suppose $j = N$. In this case, we have $\lim_{\alpha \to 1} \pi_N^* - \pi_D^* = K \geq 0$. Because of the continuity $\pi_N^* - \pi_D^*$, it follows that there exists a positive $\overline{\alpha} < 1$ such that it is optimal to adopt the nondisclosure strategy for $\alpha \in [\overline{\alpha}, 1]$. This completes the proof. ∎

Proof of Corollary 2.3. This result directly follows from Proposition 2.1 and Lemma 2.6. Recall that Proposition 2.1 states that it is optimal to adopt the nondisclosure strategy when α is small ($[0, \underline{\alpha}]$) or large ($[\overline{\alpha}, 1]$), and Lemma 2.6 states that $j^*(\alpha)$ is nondecreasing in α. Therefore, it is optimal to adopt the nondisclosure strategy with $j^* = 0$ for $\alpha \in [0, \underline{\alpha}^0]$, and the nondisclosure strategy with $j^* = N$ for $\alpha \in [\overline{\alpha}^N, 1]$.∎

Proof of Lemma 2.7. This is due to the fact $\alpha_i(N) \to \alpha$ as $N \to \infty$ based on the law of large numbers.∎

Proof of Proposition 2.2. $\pi_H^* \geq \pi_L^*$, if and only if $\alpha \geq \frac{v_L}{v_H}$. Based on this, we next discuss the optimal profits under two cases: $\alpha \geq \frac{v_L}{v_H}$ and $\alpha < \frac{v_L}{v_H}$.

(1) When $\alpha \geq \frac{v_L}{v_L}$, we only need to compare π_H^* and π_D^*. $\pi_D^* - \pi_H^* = (1 - \alpha)v_L\overline{\lambda} - K$. If $\alpha \leq 1 - \frac{K}{v_L\overline{\lambda}}$, we have $\pi_D^* \geq \pi_H^*$ and the disclosure strategy is optimal; otherwise, the nondisclosure strategy (attracting H-type customers only) is optimal.

(2) When $\alpha < \frac{v_L}{v_H}$, we only need to compare π_L^* and π_D^*. $\pi_D^* - \pi_L^* = \alpha(v_H - v_L)\overline{\lambda} - K$. If $\alpha \geq \frac{K}{(v_H - v_L)\overline{\lambda}}$, we have $\pi_D^* \geq \pi_L^*$ and the disclosure strategy is optimal; otherwise, the nondisclosure strategy (attracting all customers) is optimal.

Therefore, we have

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \min\{\frac{K}{(v_H - v_L)\overline{\lambda}}, \frac{v_L}{v_H}\}; \\ \pi_D^*, & \text{if } \frac{K}{(v_H - v_L)\overline{\lambda}} < \alpha < \frac{v_L}{v_H}; \\ \pi_D^*, & \text{if } \frac{v_L}{v_H} < \alpha < 1 - \frac{K}{v_L\overline{\lambda}}; \\ \pi_H^*, & \text{if } \alpha > \max\{\frac{v_L}{v_H}, 1 - \frac{K}{v_L\overline{\lambda}}\}. \end{cases}$$

Because the relationships among $\frac{K}{(v_H - v_L)\overline{\lambda}}$, $\frac{v_L}{v_H}$ and $1 - \frac{K}{v_L\overline{\lambda}}$ are uncertain, we need to consider each possible scenario.

Case 1: $\frac{K}{(v_H - v_L)\overline{\lambda}} < \frac{v_L}{v_H} < 1 - \frac{K}{v_L\overline{\lambda}}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{K}{(v_H - v_L)\overline{\lambda}}; \\ \pi_D^*, & \text{if } \frac{K}{(v_H - v_L)\overline{\lambda}} < \alpha < 1 - \frac{K}{v_L\overline{\lambda}}; \\ \pi_H^*, & \text{if } \alpha > 1 - \frac{K}{v_L\overline{\lambda}}. \end{cases}$$

Case 2: $\frac{v_L}{v_H} < \frac{K}{(v_H - v_L)\overline{\lambda}} < 1 - \frac{K}{v_L\overline{\lambda}}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{v_L}{v_H}; \\ \pi_D^*, & \text{if } \frac{v_L}{v_H} < \alpha < 1 - \frac{K}{v_L\overline{\lambda}}; \\ \pi_H^*, & \text{if } \alpha > 1 - \frac{K}{v_L\overline{\lambda}}. \end{cases}$$

Case 3: $\frac{K}{(v_H - v_L)\overline{\lambda}} < 1 - \frac{K}{v_L\overline{\lambda}} < \frac{v_L}{v_H}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{K}{(v_H - v_L)\overline{\lambda}}; \\ \pi_D^*, & \text{if } \frac{K}{(v_H - v_L)\overline{\lambda}} < \alpha < \frac{v_L}{v_H}; \\ \pi_H^*, & \text{if } \alpha > \frac{v_L}{v_H}. \end{cases}$$

Case 4: $\frac{v_L}{v_H} < 1 - \frac{K}{v_L\overline{\lambda}} < \frac{K}{(v_H - v_L)\overline{\lambda}}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{v_L}{v_H}; \\ \pi_D^*, & \text{if } \frac{v_L}{v_H} < \alpha < 1 - \frac{K}{v_L\overline{\lambda}}; \\ \pi_H^*, & \text{if } \alpha > 1 - \frac{K}{v_L\overline{\lambda}}. \end{cases}$$

Case 5: $1 - \frac{K}{v_L\overline{\lambda}} < \frac{v_L}{v_H} < \frac{K}{(v_H - v_L)\overline{\lambda}}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{v_L}{v_H}; \\ \pi_H^*, & \text{if } \alpha > \frac{v_L}{v_H}. \end{cases}$$

Case 6: $1 - \frac{K}{v_L\overline{\lambda}} < \frac{K}{(v_H - v_L)\overline{\lambda}} < \frac{v_L}{v_H}$

$$\pi^* = \begin{cases} \pi_L^*, & \text{if } \alpha < \frac{K}{(v_H - v_L)\overline{\lambda}}; \\ \pi_D^*, & \text{if } \frac{K}{(v_H - v_L)\overline{\lambda}} < \alpha < \frac{v_L}{v_H}; \\ \pi_H^*, & \text{if } \alpha > \frac{v_L}{v_H}. \end{cases}$$

We can find that in most cases (except case 5), α is divided into three segments, and the firm should adopt the disclosure strategy when α is at a medium level. In all cases, even in case 5, when α is small (large), the firm should adopt the nondisclosure strategy to attract all customers (H-type customers only). ∎

Proof of Lemma 2.8. Recall that when the customers do have learning opportunities, the optimal profit under the nondisclosure strategy is $\pi_{ND}^* = \max\{\max_{j=0,1,2,\ldots,N} \pi_j^*, 0\}$ where $\pi_j = [v_L + \frac{j}{N}(v_H - v_L)]\gamma_j\overline{\lambda}$ for any $j = 0, 1, 2, \ldots, N$. In particular, $\pi_0 = v_L\overline{\lambda}$. Therefore, we always have $\pi_{ND}^* \geq v_L\overline{\lambda}$.

Based on the relationship of π_D^*, π_{ND}^* and $v_L\overline{\lambda}$, we consider the following possible scenarios. If $\pi_D^* \geq \pi_{ND}^*$ and $\pi_D^* \geq v_L\overline{\lambda}$, then we have $\pi_{CL} = \pi_{NL} = \pi_D^*$ and $\delta = 0$. If $\pi_{ND}^* \geq \pi_D^* \geq v_L\overline{\lambda}$, then $\delta = (\pi_{ND}^* - \pi_D^*)/\pi_D^*$. Because

π_D^* is decreasing in K, δ is increasing in K. If $\pi_{ND}^* \geq v_L\overline{\lambda} \geq \pi_D^*$, then $\delta = (\pi_{ND}^* - v_L\overline{\lambda})/(v_L\overline{\lambda})$, which is independent of K.∎

Proof of Proposition 2.3. For the new expressions of some notations, e.g., π_D^*, π_H^*,π_L^*, see Appendix B. Under the condition $\alpha\overline{\lambda}(v_H - v_L) \geq K$, since $\pi_L^* \leq \pi_D^*$, we only need to compare π_H^* and π_D^*. $\pi_H^* - \pi_D^* = (v_L - c)\overline{\lambda}\alpha - 2\sqrt{ch\overline{\lambda}}\sqrt{\alpha} + (c - v_L)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}} + K$. Note that this a quadratic function of $\sqrt{\alpha}$. The discriminant Δ for the equation

$$\pi_H^* - \pi_D^* = (v_L - c)\overline{\lambda}\alpha - 2\sqrt{ch\overline{\lambda}}\sqrt{\alpha} + (c - v_L)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}} + K = 0 \quad \text{(A.1)}$$

is $\Delta = (-2\sqrt{hc\overline{\lambda}})^2 - 4(v_L - c)\overline{\lambda}[-(v_L - c)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}} + K] = 4\{[v_L - c)\overline{\lambda} - \sqrt{hc\overline{\lambda}}]^2 - (v_L - c)\overline{\lambda}K\}$.

Based on the sign of Δ and the value of $\sqrt{\overline{\lambda}}$, we next divide $\overline{\lambda}$ into four sets: S1=$\{\overline{\lambda}|\frac{\sqrt{hc}-\sqrt{(v_L-c)K}}{v_L-c} < \sqrt{\overline{\lambda}} < \sqrt{\overline{\lambda}_1}\}$, S2=$\{\overline{\lambda}|\sqrt{\overline{\lambda}} \leq \frac{\sqrt{hc}-\sqrt{(v_L-c)K}}{v_L-c}\}$, S3=$\{\overline{\lambda}|\sqrt{\overline{\lambda}_1} < \sqrt{\overline{\lambda}} < \sqrt{\overline{\lambda}_2}\}$, and S4=$\{\overline{\lambda}|\sqrt{\overline{\lambda}} \geq \sqrt{\overline{\lambda}_2}\}$.

Note that $\overline{\lambda} \in S1$, if and only if $\Delta < 0$; otherwise if $\overline{\lambda} \in \{S2 \cup S3 \cup S4\}$, $\Delta \geq 0$.

Case 1: $\overline{\lambda} \in S1$. In this case, we have $\Delta < 0$, and hence equation (A.1) has no real solution. It follows that $\pi_H^* \geq \pi_D^*$ in this case, and the firm should adopt nondisclosure strategy.

In the following three cases, $\overline{\lambda} \in \{S2 \cup S3 \cup S4\}$ and we have $\Delta \geq 0$. Equation (A.1) has two roots$x_1 = \frac{\sqrt{hc\overline{\lambda}}-\sqrt{[(v_L-c)\overline{\lambda}-\sqrt{hc\overline{\lambda}}]^2-(v_L-c)K\overline{\lambda}}}{(v_L-c)\overline{\lambda}}$ and $x_2 = \frac{\sqrt{hc\overline{\lambda}}+\sqrt{[(v_L-c)\overline{\lambda}-\sqrt{hc\overline{\lambda}}]^2-(v_L-c)K\overline{\lambda}}}{(v_L-c)\overline{\lambda}}$. Note that $x_1 \leq x_2$.

Case 2: $\overline{\lambda} \in S2$. If $hc < (v_L - c)K$, then it is trivial to consider this case. We next consider the scenario $hc \geq (v_L - c)K$. Note that in this case, we have

$$
\begin{aligned}
&\sqrt{\bar{\lambda}} \leq \frac{\sqrt{hc}-\sqrt{(v_L-c)K}}{v_L-c} \\
&\Rightarrow (v_L-c)\bar{\lambda} < \sqrt{hc\bar{\lambda}} \\
&\Rightarrow x_1 \geq \frac{\sqrt{hc\bar{\lambda}}-\sqrt{[(v_L-c)\bar{\lambda}-\sqrt{hc\bar{\lambda}}]^2}}{(v_L-c)\bar{\lambda}} = 1
\end{aligned}
$$

Hence, Equation (A.1) has no solutions located in the interval (0,1). In other words, $\pi_H^* \geq \pi_D^*$ for $\alpha \in (0,1)$ and hence the firm should adopt the nondisclosure strategy.

In the following two cases, we have $\sqrt{\bar{\lambda}} \geq \sqrt{\lambda_1} \geq \frac{\sqrt{hc}}{v_L-c}$, and thus $x_2 \leq \frac{\sqrt{hc\bar{\lambda}}+\sqrt{[(v_L-c)\bar{\lambda}-\sqrt{hc\bar{\lambda}}]^2}}{(v_L-c)\bar{\lambda}} = 1$. We proceed to discuss the sign of x_1.
Case 3: $\bar{\lambda} \in S3$. In this case, we have $\sqrt{\lambda_1} \leq \sqrt{\bar{\lambda}} < \sqrt{\lambda_2}$, and it follows that $0 < x_1 \leq x_2 \leq 1$. Therefore, if $x_1^2 < \alpha < x_2^2$, we have $\pi_H^* < \pi_D^*$, and the voluntary disclosure strategy is better; otherwise the nondisclosure strategy is better.
Case 4: $\bar{\lambda} \in S4$. In this case, we have $x_1 \leq 0 < x_2 \leq 1$. Therefore, if $\alpha < x_2^2$, we have $\pi_H^* < \pi_D^*$, and the firm should adopt the voluntary disclosure strategy; otherwise, the firm should adopt the nondisclosure strategy. ∎

Proof of Lemma 3.2. Recall that $G_i(y) = \mathbb{E}[h_i(y - D_i(t, t+L_i])^+ + (h_0+p_i)(y - D_i(t, t+L_i])^-]$. We define

$$
G_i^d(y) = h_i(y - \lambda_i L_i)^+ + (h_0 + p_i)(y - \lambda_i L_i)^-,
$$

for any $i = 1, 2, \ldots, N$. Note that $G_i^d(y)$ is the single-stage cost of Retailer i assuming deterministic demands. It follows form Jensen's inequality that for any y,

$$
G_i^d(y) \leq G_i(y),
$$

which means that the inventory costs calculated in the deterministic model underevaluate the actual inventory cost with stochastic demands.

We define

$$G_i^{id}(y) = \begin{cases} C_i^* & \text{if } y \leq \lambda_i L_i - C_i^*/(h_0 + p_i), \\ G_i^d(y) & \text{otherwise,} \end{cases}$$

and

$$\begin{aligned} G_i^{0d}(y) &= G_i^d(y) - G_i^{id}(y) = (h_0 + p_i)(\lambda_i L_i - C_i^*/(h_0 + p_i) - y)^+ \\ &= \begin{cases} G_i^d(y) - C_i^* & \text{if } y \leq \lambda_i L_i - C_i^*/(h_0 + p_i), \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Note that $G_i^d(\lambda_i L_i - C_i^*/(h_0 + p_i)) = C_i^* = G_i(r_i^*)$. It follows that $\lambda_i L_i - C_i^*/(h_0 + p_j) \leq r_i^*$. Therefore, we have that for any y,

$$G_i^{0d}(y) \leq G_i^0(y).$$

Let D_0 denote the total demand over $(0, L_0)$. It follows that $\mathbb{E}[D_0] = \lambda_0 L_0$. We have

$$\begin{aligned} G_0(y) =& \mathbb{E}[h_0(y - D_0) + \min_{y_i : \sum_{i=1}^N y_i \leq y - D_0} \sum_{i=1}^N G_i^0(y_i)] \\ \geq& h_0(y - \lambda_0 L_0) + \min_{y_i : \sum_{i=1}^N y_i \leq y - \lambda_0 L_0} \sum_{i=1}^N G_i^0(y_i) \qquad (A.2) \\ \geq& h_0(y - \lambda_0 L_0) + \min_{y_i : \sum_{i=1}^N y_i \leq y - \lambda_0 L_0} \sum_{i=1}^N G_i^{0d}(y_i), \end{aligned}$$

where the first inequality follows from Jensen Inequality and the second inequality follows from $G_i^{0d}(y) \leq G_i^0(y)$. Let $j = \arg\min_{i=1,2,\ldots,N} p_i$ and y_i^*, $i = 1, 2, \ldots, N$, denote the value at which $\min_{y_i : \sum_{i=1}^N y_i \leq y - \lambda_0 L_0} \sum_{i=1}^N G_i^{0d}(y_i)$ attains its optimal value. Note that $G_i^{0d}(y) = (h_0 + p_i)(\lambda_i L_i - C_i^*/(h_0 + p_i) - y)^+$

is strictly decreasing in $y \in (-\infty, \lambda_i L_i - C_i^*/(h_0 + p_i)]$ and equals to 0 in $y \in [\lambda_i L_i - C_i^*/(h_0 + p_i), \infty)$. It follows that $y_i^* = \lambda_i L_i - C_i^*/(h_0 + p_i)$ for any $i \neq j$ and $y_j^* = y - \lambda_0 L_0 - \sum_{i \neq j} y_i^*$. Therefore,

$$\begin{aligned} G_0(y) \geq & h_0(y - \lambda_0 L_0) + \min_{y_i : \sum_{i=1}^N y_i \leq y - \lambda_0 L_0} \sum_{i=1}^N G_i^{0d}(y_i) \\ = & h_0(y - \lambda_0 L_0) + G_j^{0d}(y_j^*) \\ = & h_0(y - \lambda_0 L_0) + (h_0 + p_j)(\sum_{i=1}^N \lambda_i L_i - C_i^*/(h_0 + p_i) - (y - \lambda_0 L_0))^+ \\ \equiv & G_0^d(y). \end{aligned} \tag{A.3}$$

Note that $G_0^d(y)$ is a convex function with minimal value $h_0 \sum_{i=1}^N (\lambda_i L_i - \frac{C_i^*}{h_0 + p_i})$. We next consider a single-stage inventory system with setup cost K_0 and inventory cost rate $G_0^d(y)$. It can easily be seen that the optimal cost of this system, denoted by $G_0^d(r_0^{d*})$, is not greater than $G_0(r_0^*)$, i.e., $G_0(r_0^*) \geq G_0^d(r_0^{d*})$. It can be further verified that $G_0^d(r_0^{d*}) = h_0 \sum_{i=1}^N (\lambda_i L_i - \frac{C_i^*}{h_0 + p_i}) + \sqrt{\frac{2\lambda_0 K_0 h_0 p_j}{h_0 + p_j}}$. Therefore, under the assumption $h_0 \sum_{i=1}^N (\lambda_i L_i - \frac{C_i^*}{h_0 + p_i}) + \sqrt{\frac{2\lambda_0 K_0 h_0 p_j}{h_0 + p_j}} > 0$, we have $C_0^* = G_0(r_0^*) > 0$.∎

Proof of Lemma 3.3. (i) We prove the lemma by contradiction. Suppose there exists two irregular intervals for Retailer i over the cycle $[T_I^j, T_I^{j+1})$. We denote by $[\hat{T}_i^{j,1}, \tilde{T}_i^{j,1})$ the first irregular shipment interval. As there exists at least one shipment after the first irregular shipment interval, the warehouse must have some on-hand inventory after the first irregular shipment interval. It implies that under the modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, the warehouse has enough inventory to raise Retailer i's inventory position to $r_i + Q_i$ at time $\hat{T}_i^{j,1}$, that is, $IP_i(\hat{T}_i^{j,1}) = r_i + Q_i$. On the other hand, at the end of the first irregular interval (once Retailer i's inventory position reaches r_i), the warehouse also has on-hand inventory to raise the inventory position of Retailer i above r_i. That is, $IP_i^-(\tilde{T}_i^{j,1}) = r_i$. It follows from the definition of regular shipment interval that the first irregular shipment interval is indeed a regular shipment interval. Therefore, there exists at most one irregular shipment for Retailer i in each cycle, whether the cycle is empty or not.

(ii) Consider the last shipment from warehouse to retailers in a specified cycle. That is, the warehouse has enough inventory to fulfill retailers' orders before the last shipment which in turn implies that retailers' inventory position after these shipments should be $r_i + Q_i$. In short, the shipments before the last shipment are either regular or type I irregular shipments. However, for the last shipment, the warehouse may not have enough inventory to raise retailer's inventory position to the target one. Then, that retailer will experience a type II irregular shipment interval. In summary, there exists at most one type II irregular shipment interval over one cycle. ∎

Proof of Lemma 3.4. By Remark 3.1, we first construct a cost bound, excluding the setup costs in type II irregular shipment intervals, in terms of $IP_i(t)$ and $C_i(r_i, Q_i)$ as follows. If Retailer i is in a regular shipment interval, then we charge $C_i(r_i, Q_i)$; if Retailer i is in a type I irregular shipment interval, then we charge $\max\{G_i(IP_i(t)), C_i(r_i, Q_i)\}$; if Retailer i is in a type II of irregular shipment interval, then we charge $G_i(IP_i(t))$.

Let $OI_0^+(t)$ denote the on-hand inventory at the warehouse at time t. By the definition, $IL_0(t) = OI_0^+(t) + \sum_{j=1}^{N} IP_j(t)$. Under the modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, we have $IP_j(t) \leq r_j + Q_j$ for any $j = 1, 2, \ldots, N$. It follows that

$$OI_0^+(t) + IP_i(t) = IL_0(t) - \sum_{j \neq i} IP_j(t) \geq IL_0(t) - \sum_{j \neq i} (r_j + Q_j). \qquad \text{(A.4)}$$

We prove the result by considering the ensuing two scenarios.

Scenario I: $IL_0(t) - \sum_{j \neq i}(r_j + Q_j) > r_i$. By (A.4), we have $r_i < OI_0^+(t) + IP_i(t)$. If $OI_0^+(t) = 0$, then it follows that $r_i < IP_i(t) \leq r_i + Q_i$. Now consider the case with $OI_0^+(t) > 0$. In this case, the warehouse has excess inventory after shipment to retailers, which implies that the inventory position of Retailer i must be above r_i, i.e., $r_i < IP_i(t) \leq r_i + Q_i$. In each case, we have $r_i < IP_i(t) \leq r_i + Q_i$, but Retailer i could be in either a regular or an irregular shipment interval. Therefore, to obtain an upper bound on $\hat{\Gamma}_i(IL_0(t))$, we charge the larger one between expected cost rates of the regular and irregular shipment intervals. That is, the cost of Retailer i excluding the setup cost in a type II irregular shipment interval, is no more than $\max\{G_i(IP_i(t)), C_i(r_i, Q_i)\}$. Moreover, by the definition of w_i, it follows that $\max\{G_i(IP_i(t)), C_i(r_i, Q_i)\} \leq \max\{G_i(w_i), C_i(r_i, Q_i)\}$. Therefore, in this scenario, $\hat{\Gamma}_i(IL_0(t)) \leq \max\{G_i(w_i), C_i(r_i, Q_i)\}$.

Scenario II: $IL_0(t) - \sum_{j\neq i}(r_j + Q_j) \leq r_i$. If $r_i < IP_i(t) \leq r_i + Q_i$, then as mentioned above, the cost of Retailer i excluding the setup cost in a type II irregular shipment interval, should be bounded by $\hat{\Gamma}_i(IL_0(t)) \leq \max\{G_i(w_i), C_i(r_i, Q_i)\}$. Otherwise, if $IP_i(t) \leq r_i$, then we have $OI_0^+(t) = 0$ and Retailer i must be in an irregular shipment interval, but could be in either a type I or a type II irregular shipment interval. As a result, $\hat{\Gamma}_i(IL_0(t)) \leq \max\{G_i(IP_i(t)), C_i(r_i, Q_i)\}$. Recall that $IL_0(t) - \sum_{j\neq i}(r_j + Q_j) \leq OI_0^+(t) + IP_i(t)$, thus, $IL_0(t) - \sum_{j\neq i}(r_j + Q_j) \leq IP_i(t) \leq r_i$. Due to the convexity of $G_i(\cdot)$, we can obtain that $G_i(IP_i(t)) \leq \max\{G_i(IL_0(t) - \sum_{j\neq i}(r_j + Q_j)), G_i(r_i)\} \leq \max\{G_i(IL_0(t) - \sum_{j\neq i}(r_j + Q_j)), G_i(w_i)\}$. Therefore, if $IL_0(t) - \sum_{j\neq i}(r_j + Q_j) \leq r_i$, we have $\hat{\Gamma}_i(IL_0(t)) \leq \max\{G_i(IL_0(t) - \sum_{j\neq i}(r_j + Q_j)), G_i(w_i), C_i(r_i, Q_i)\}$.$\blacksquare$

Proof of Lemma 3.5. We first consider the replenishment cycle. Because the warehouse faces an ample supply, there are Q_0 units of demand in total over any cycle $[T_0^j, T_0^{j+1})$. As the demand arrives following a homogeneous Poisson process with rate λ_0, the expectation of the length of any replenishment cycle j at the warehouse is $\mathbb{E}[T_0^{j+1} - T_0^j] = Q_0/\lambda_0$. Therefore, the result holds for the warehouse.

Now consider the depletion cycle. By the definition of T_0^j, the order in the jth cycle will arrive at the warehouse at time $T_0^j + L_0$. Denote by Δt_I^j the time lag between the arrival time at the warehouse and the receiving time of Installation I, i.e., $\Delta_I^j \equiv T_I^j - (T_0^j + L_0)$. Δt_I^j can be understood as the sojourn time of the 1st unit in the jth order at the warehouse. Clearly, $\Delta_I^j \geq 0$. By the definition of modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, we know that $IP_0^-(T_0^j) = r_0$ and $IP_0(T_0^j) = r_0 + Q_0$. It follows that $IL_0(T_0^j + L_0) = IP_0(T_0^j) - D_0(T_0^j, T_0^j + L_0] = r_0 + Q_0 - D_0(T_0^j, T_0^j + L_0]$. At time $T_0^j + L_0$ and before units contained in the warehouse's jth order are sent to Installation I, there are at most $(r_0 - D_0(T_0^j, T_0^j + L_0] - \sum_{i=1}^N r_i)^+$ units in addition to r_i units in Retailer i's inventory position. Let $t(x) \equiv \max_{i=1,2,\ldots,N} \inf\{t \geq 0 | D_i(0,t] \geq x\}$. Thus, for all j, we have

$$0 \leq \Delta_I^j \leq t(r_0 - D_0(T_0^j . T_0^j + L_0] - \sum_{i=1}^{N} r_i)^+ \leq t(r_0 - \sum_{i=1}^{N} r_i)^+. \tag{A.5}$$

Note that the lower and upper bound on Δ_I^j depend only on the policy parameters. Therefore,

$$\begin{aligned}
&\lim_{j\to\infty} \mathbb{E}[T_I^{j+1} - T_I^1]/j \\
=&\lim_{j\to\infty} \mathbb{E}[(T_0^{j+1} + L_0 + \Delta_I^{j+1}) - (T_0^1 + L_0 + \Delta_I^1)]/j \\
=&\lim_{j\to\infty} \mathbb{E}[T_0^{j+1} - T_0^1]/j + \mathbb{E}[\Delta_I^{j+1} - \Delta_I^1]/j \\
=&\lim_{j\to\infty} \mathbb{E}[T_0^{j+1} - T_0^1]/j \\
=&Q_0/\lambda_0,
\end{aligned} \tag{A.6}$$

where the third equality is a direct result of (A.5) and the last is true because of $\mathbb{E}[T_0^{j+1} - T_0^j] = Q_0/\lambda_0$ for all j. Therefore, the result also holds for depletion cycle.$\blacksquare$

Proof of Theorem 3.4. By (3.9) and the definition of $\hat{G}_i(y)$ in (3.10), we obtain

$$\hat{\Gamma}_I(IL_0(t)) \leq \bar{\Gamma}_I(IL_0(t)) = \sum_{i=1}^{N} \hat{G}_i(IL_0(t)) + \sum_{i=1}^{N} \hat{C}_i(r_i, Q_i). \tag{A.7}$$

We denote by $\Gamma_0(IP_0(t))$ the total expected cost rate of all installations at time t when the inventory position of the warehouse is $IP_0(t)$, where we exclude the setup costs incurred at the warehouse and the setup costs of type II irregular shipment intervals incurred at retailers. By such a definition, $\Gamma_0(IP_0(t))$ constitutes two parts: (i) the inventory holding cost at the warehouse, and (ii) the total costs at all retailers excluding the setup costs in type II irregular shipment intervals. Then, according to the cost accounting scheme (see Definition 3.4.1), we have

$$\begin{aligned}
\Gamma_0(IP_0(t)) =& \mathbb{E}[h_0(IP_0(t) - D_0)] + \mathbb{E}[\hat{\Gamma}_I(IP_0(t) - D_0)] \\
\leq& \mathbb{E}[h_0(IP_0(t) - D_0)] + \mathbb{E}[\sum_{i=1}^{N} \hat{G}_i(IP_0(t) - D_0)] + \sum_{i=1}^{N} \hat{C}_i(r_i, Q_i) \\
=& \Lambda_0(IP_0(t)) + \sum_{i=1}^{N} \hat{C}_i(r_i, Q_i),
\end{aligned} \tag{A.8}$$

where the inequality follows from (A.7), and the last equality from (3.11).

Because the warehouse has an unlimited supply form the external supplier, under the modified echelon $(\mathbf{r}, \mathbf{Q})$ policy, the inventory position of the warehouse, $IP_0(t)$, is uniformly distributed on $\{r_0+1, \ldots, r_0+Q_0\}$. Therefore, by the definition of $\Gamma_0(IP_0(t))$, the long-run average system-wide cost, with the setup costs of type II irregular shipment intervals incurred at retailers being excluded, can be bounded as follows

$$
\begin{aligned}
&\frac{1}{Q_0}\Big[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \Gamma_0(y)dy\Big] \\
\leq &\frac{1}{Q_0}\Big[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} [\Lambda_0(y) + \sum_{i=1}^{N} \hat{C}_i(r_i, Q_i)]dy\Big] \\
= &\sum_{i=0}^{N} \hat{C}_i(r_i, Q_i), \qquad\qquad (A.9)
\end{aligned}
$$

where the inequality is due to (A.8) and the equality holds true due to (3.12).

Finally, combining (A.9) and Remark 3.3, we can obtain that the long-run average system-wide cost can be bounded as: $C(\mathbf{r}, \mathbf{Q}) \leq \sum_{i=0}^{N} \hat{C}_i(r_i, Q_i) + \lambda_0 \overline{K}/Q_0$.∎

Proof of Lemma 3.6. As $(r_i, Q_i) = (r_i^*, Q_i^*)$, $\hat{G}_i(y)$ is convex and thus, so is $\Lambda_0(y)$. Since $\lim_{y\to+\infty} \hat{G}_i(y) = 0$, it follows that $\lim_{y\to+\infty} \Lambda_0(y) = \infty$. Moreover, since $\lim_{y\to+\infty} \hat{G}_i'(y) = 0$ and $\lim_{y\to-\infty} \hat{G}_i'(y) = -(h_0+p_i)$, we have $\lim_{y\to+\infty} \Lambda_0'(y) = h_0 > 0$ and $\lim_{y\to-\infty} \Lambda_0'(y) = -\sum_{i=1}^{N}(h_0 + p_i) + h_0 < 0$, and thus $\lim_{y\to-\infty} \Lambda_0(y) = \infty$. Therefore, $\Lambda_0(y)$ satisfies Assumption 3.3.1. The convexity of $\hat{C}_0(r_0, Q_0)$ follows directly from the result of the single-stage model in Zheng [97]. ∎

Proof of Theorem 3.5. (i) As shown in Lemma 3.1, the lower bound cost is $\sum_{i=0}^{N} C_i^*$. As shown in (3.16), the upper bound cost is $\sum_{i=1}^{N} C_i^* + \tilde{C}_0^*$. It follows that the modified echelon $(\hat{\mathrm{r}}, \hat{\mathbf{Q}})$ policy in (MERQD) is at least $1+(\tilde{C}_0^* - C_0^*)/(\sum_{i=1}^{N} C_i^* + C_0^*)$-optimal. The last result can be obtained by showing $(\sum_{i=1}^{N} C_i^* + \tilde{C}_0^*)/(\sum_{i=1}^{N} C_i^* + C_0^*) \leq \tilde{C}_0^*/C_0^*$. It suffices to show $(\sum_{i=1}^{N} C_i^* + \tilde{C}_0^*)C_0^* \leq (\sum_{i=1}^{N} C_i^* + C_0^*)\tilde{C}_0^*$. Then, the desired result directly holds because $C_0^* \leq \tilde{C}_0^*$.

(ii) By (3.13) and the definition of $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$, we have that for any (r_0, Q_0), $\mathcal{C}_{\mathcal{B}}^* \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \leq \sum_{i=1}^{N} C_i^* + \hat{C}_0(r_0, Q_0) + \lambda_0 \overline{K}/Q_0$. By Lemma D.1(iii) and

Lemma D.2(ii), we can obtain that for any (r_0, Q_0),

$$\begin{aligned}\mathcal{C}_{\mathcal{B}}^* \le C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) &\le \sum_{i=1}^{N} C_i^* + \hat{C}_0(r_0, Q_0) + \frac{\lambda_0 \overline{K}}{Q_0} \\ &\le \sum_{i=1}^{N} C_i^* + \epsilon(\frac{Q_0}{\hat{Q}_0^*})\hat{C}_0^* + \frac{\lambda_0 C_m^* Q_m^*}{2\lambda_m Q_0}. \end{aligned} \tag{A.10}$$

To facilitate the comparison, we select Q_0 in (A.10) as follows:

$$\begin{aligned}\check{Q}_0 &\equiv \arg\min_{Q_0} \Big\{ \epsilon(\frac{Q_0}{\hat{Q}_0^*})\hat{C}_0^* + \frac{\lambda_0 C_m^* Q_m^*}{2\lambda_m Q_0} \Big\} \\ &= \sqrt{\frac{(\hat{Q}_0^*)^2 \hat{C}_0^* + (\lambda_0 C_m^* Q_m^*)\hat{Q}_0^*/\lambda_m}{\hat{C}_0^*}} \\ &= \hat{Q}_0^* \sqrt{1 + \frac{\lambda_0 C_m^* Q_m^*}{\lambda_m \hat{C}_0^* \hat{Q}_0^*}}. \end{aligned}$$

Replacing Q_0 in (A.10) with $\check{Q}_0$, we can obtain a new upper bound.

$$\mathcal{C}_{\mathcal{B}}^* \le C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \le \sum_{i=1}^{N} C_i^* + \hat{C}_0^* \sqrt{1 + \frac{\lambda_0 C_m^* Q_m^*}{\lambda_m \hat{C}_0^* \hat{Q}_0^*}}. \tag{A.11}$$

By (A.11), the relative gap between $\mathcal{C}_{\mathcal{B}}^*$ and $C(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ is bounded as follows:

$$\begin{aligned}(C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) - \mathcal{C}_{\mathcal{B}}^*)/\mathcal{C}_{\mathcal{B}}^* &\le \hat{C}_0^* \Big(\sqrt{1 + \frac{\lambda_0 C_m^* Q_m^*}{\lambda_m \hat{C}_0^* \hat{Q}_0^*}} - \beta_2 \Big) / \mathcal{C}_{\mathcal{B}}^* \\ &\le \hat{C}_0^* \Big(\sqrt{1 + \frac{\lambda_0 C_m^* Q_m^*}{\lambda_m \hat{C}_0^* \hat{Q}_0^*}} - \beta_2 \Big) / (\beta_2 \hat{C}_0^* + \sum_{i=1}^{N} C_i^*) \\ &\le \hat{C}_0^* \Big(\sqrt{1 + \frac{\lambda_0 C_m^* Q_m^*}{\lambda_m \hat{C}_0^* \hat{Q}_0^*}} - \beta_2 \Big) / (\beta_2 \hat{C}_0^* + C_m^*). \end{aligned}$$

To obtain the desire result, it is sufficient to show the following stronger statement: for any $x_1, x_2 > 0$, $x_2(\sqrt{1 + \lambda_0 x_1/(\lambda_m x_2 \beta_1)} - \beta_2)/(\beta_2 x_2 + x_1) \le \alpha \equiv \max\{\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m} + \frac{1}{4}} - \frac{1}{2}, \frac{1}{\beta_2} - 1\}$, which is equivalent to $\alpha^2 x^2 + [2\beta_2\alpha(1+\alpha) - \lambda_0/(\lambda_m\beta_1)]x + \beta_2^2\alpha(\alpha+2) + \beta_2^2 - 1 \ge 0$ for any $x > 0$. By verifying that (1) $2\beta_2\alpha(1+\alpha) - \lambda_0/(\lambda_m\beta_1) \ge 0$ when $\alpha \ge \sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m} + \frac{1}{4}} - \frac{1}{2}$; and (2)

$\beta_2^2\alpha(\alpha+2)+\beta_2^2-1\geq 0$ when $\alpha\geq\frac{1}{\beta_2}-1$, the desired result ensues.

We prove the alternative bound mentioned in Footnote 2. Following part (ii), it suffices to show that the quadratic function $\alpha^2x^2+[2\beta_2\alpha(1+\alpha)-\lambda_0/(\lambda_m\beta_1)]x+\beta_2^2\alpha(\alpha+2)+\beta_2^2-1\geq 0$ when $\alpha\geq\lambda_0/(2(\beta_1\beta_2\lambda_m+\sqrt{(\beta_2^2-1)(\beta_1\lambda_m)^2+\beta_1\beta_2\lambda_m\lambda_0}))$ under the condition $(\beta_2^2-1)(\beta_1\lambda_m)+\beta_2\lambda_0\geq 0$. The bound is then established by verifying that the quadratic function $f(x)=\alpha^2x^2+[2\beta_2\alpha(1+\alpha)-\lambda_0/(\lambda_m\beta_1)]x+\beta_2^2\alpha(\alpha+2)+\beta_2^2-1$ has a zero discriminant, i.e., $\Delta=4\alpha^2(1-\lambda_0\beta_2/(\lambda_m\beta_1))-4\alpha\beta_2\lambda_0/(\lambda_m\beta_1)+(\lambda_0/(\lambda_m\beta_1))^2=0$. ∎

Proof of Corollary 3.6. It is easy to verify if $2(1/\beta_2-1)\geq\lambda_0/(\beta_1\lambda_m)$, then we have $\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m}+\frac{1}{4}}+\frac{1}{2}\leq\frac{1}{\beta_2}$, i.e., $\max\{\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m}+\frac{1}{4}}+\frac{1}{2},\frac{1}{\beta_2}\}=\frac{1}{\beta_2}$. ∎

Proof of Theorem 3.7. Because all retailers are identical, we have $Q_m^*=Q_i^*$ and $\lambda_m=\lambda_i$ in this case. By (A.11), the relative gap between $\mathcal{C}_{\mathcal{B}}^*$ and $C(\hat{\mathbf{r}},\hat{\mathbf{Q}})$ is bounded as follows:

$$\begin{aligned}(C(\hat{\mathbf{r}},\hat{\mathbf{Q}})-\mathcal{C}_{\mathcal{B}}^*)/\mathcal{C}_{\mathcal{B}}^* &\leq \hat{C}_0^*\Big(\sqrt{1+\frac{\lambda_0 Q_i^* C_i^*}{\lambda_i\hat{Q}_0^*\hat{C}_0^*}}-\beta_2\Big)/\mathcal{C}_{\mathcal{B}}^*\\ &=\hat{C}_0^*\Big(\sqrt{1+\frac{NQ_i^*C_i^*}{\hat{Q}_0^*\hat{C}_0^*}}-\beta_2\Big)/(\beta_2\hat{C}_0^*+NC_i^*).\end{aligned}$$

To obtain the desire result, it is sufficient to show the following stronger statement: for any $x_1,x_2>0$, $x_2(\sqrt{1+x_1/(x_2\beta_1)}-\beta_2)/(\beta_2x_2+x_1)\leq\alpha\equiv\max\{\sqrt{\frac{1}{2\beta_1\beta_2}+\frac{1}{4}}-\frac{1}{2},\frac{1}{\beta_2}-1\}$, which is equivalent to $\alpha^2x^2+[2\beta_2\alpha(1+\alpha)-1/(\beta_1)]x+\beta_2^2\alpha(\alpha+2)+\beta_2^2-1\geq 0$ for any $x>0$. By verifying that (1) $2\beta_2\alpha(1+\alpha)-1/(\beta_1)\geq 0$ when $\alpha\geq\sqrt{\frac{1}{2\beta_1\beta_2}+\frac{1}{4}}-\frac{1}{2}$; and (2) $\beta_2^2\alpha(\alpha+2)+\beta_2^2-1\geq 0$ when $\alpha\geq\frac{1}{\beta_2}-1$, the desired result ensues. ∎

Proof of Theorem 3.8. It follows from Theorem 3.5 that to prove asymptotic optimality of the modified echelon $(\mathbf{r},\mathbf{Q})$ policy, it is sufficient to first show the following statements: (i) $\lim_{K_0/K_m\to\infty}\beta_1=\infty$ and $\lim_{K_0/K_m\to\infty}\beta_2=$

1; (ii) $\lim_{h_0/h_m \to 0} \beta_1 = \infty$ and $\lim_{h_0/h_m \to 0} \beta_2 = 1$.

(i) Let $r_i(Q_i) = \arg\min_{r_i} C_i(r_i, Q_i)$ for $i = 0, 1, 2, \ldots, N$. Define $A_i(Q_i) \equiv Q_i G_i(r_i(Q_i)) - \int_0^{Q_i} G_i(r_i(y))dy$. By Lemma D.1, $A_i(Q_i)$ is a continuous and strictly increasing function such that $A_i(Q_i^*) = \lambda_i K_i$ and $A_i(0) = 0$. Let $A_i^{-1}(x)$ be the inverse function of $A_i(Q_i)$. Then, $Q_i^* = A_i^{-1}(\lambda_i K_i)$ and $A_i^{-1}(0) = 0$. Similarly, let $\hat{r}_0(Q_0) = \arg\min_{r_0} \hat{C}_0(r_0, Q_0)$ and define $\hat{A}_0(Q_0) \equiv Q_0 \Lambda_0(\hat{r}_0(Q_0)) - \int_0^{Q_0} \Lambda_0(\hat{r}_0(y))dy$. Let $\hat{A}_0^{-1}(x)$ be the inverse function of $\hat{A}_0(Q_0)$. Then, $\hat{Q}_0^* = \hat{A}_0^{-1}(\lambda_0 K_0)$ and $\hat{A}_0^{-1}(0) = 0$. Let $\gamma = K_m/K_0$. We have

$$\begin{aligned}
\lim_{K_0/K_m \to \infty} \beta_1 &= \lim_{K_0/K_m \to \infty} \frac{\hat{A}_0^{-1}(\lambda_0 K_0)}{A_m^{-1}(\lambda_m K_m)} = \lim_{\gamma \to 0} \frac{\hat{A}_0^{-1}(\lambda_0 K_0)}{A_m^{-1}(\lambda_m \gamma K_0)} \\
&= \frac{\hat{A}_0^{-1}(\lambda_0 K_0)}{A_m^{-1}(\lambda_m K_0)} \lim_{\gamma \to 0} \frac{A_m^{-1}(\lambda_m K_0)}{A_m^{-1}(\lambda_m \gamma K_0)} = \infty,
\end{aligned}$$

where the last equality follows form the fact that $A_m^{-1}(x)$ is continuous and $A_m^{-1}(0) = 0$.

It remains to show $\lim_{K_0/K_m \to \infty} \beta_2 = 1$. Recall that $A_0(Q)$ and $\hat{A}_0(Q)$ are both increasing convex functions. Let $r(Q_0) = \arg\min_{r_0} C_0(r_0, Q_0)$. Define $H_0(Q) \equiv G_0(r(Q))$ and $\hat{H}_0(Q) \equiv \Lambda_0(\hat{r}_0(Q))$. It is easy to see that $A_0'(Q) = QH_0(Q)$ and $\hat{A}_0' = Q\hat{H}_0(Q)$. Moreover, because $G_0(y)$ and $\Lambda(y)$ satisfy Assumption 3.3.1, by Zheng [97], $H_0(Q)$ and $\hat{H}_0(Q)$ are increasing convex functions with asymptotic slope $-ab/(a-b)$ as $Q \to \infty$, where parameters a and b are defined in Assumption 3.3.1. That is, we have

$$\lim_{Q \to \infty} H_0'(Q) = \frac{h_0 \underline{p}}{h_0 + \underline{p}}, \tag{A.12}$$

$$\lim_{Q \to \infty} \hat{H}_0'(Q) = \frac{h_0[\sum_{i=1}^N (p_i + h_0) - h_0]}{\sum_{i=1}^N (p_i + h_0)}, \tag{A.13}$$

where $\underline{p} = \min_{i=1,\ldots,N} p_i$. With the above analysis, under the condition $K_m >$

0, we can obtain

$$
\begin{aligned}
\lim_{K_0/K_m \to \infty} \beta_2 &= \lim_{\gamma \to 0} \frac{H_0[A_0^{-1}(\lambda_0 K_m/\gamma)]}{\hat{H}_0[\hat{A}_0^{-1}(\lambda_0 K_m/\gamma)]} \\
&= \lim_{\gamma \to 0} \frac{H_0'[A_0^{-1}(\lambda_0 K_m/\gamma)] \frac{dA_0^{-1}(\lambda_0 K_m/\gamma)}{d\gamma}}{\hat{H}_0'[\hat{A}_0^{-1}(\lambda_0 K_m/\gamma)] \frac{d\hat{A}_0^{-1}(\lambda_0 K_m/\gamma)}{d\gamma}} \\
&= \lim_{\gamma \to 0} \frac{H_0'[A_0^{-1}(\lambda_0 K_m/\gamma)]}{\hat{H}_0'[\hat{A}_0^{-1}(\lambda_0 K_m/\gamma)]} \lim_{Q \to \infty} \frac{\hat{A}_0'(Q)}{A_0'(Q)} \\
&= \lim_{\gamma \to 0} \frac{H_0'[A_0^{-1}(\lambda_0 K_m/\gamma)]}{\hat{H}_0'[\hat{A}_0^{-1}(\lambda_0 K_m/\gamma)]} \lim_{Q \to \infty} \frac{Q\hat{H}_0'(Q)}{QH_0'(Q)} \\
&= 1,
\end{aligned}
$$

where the second equality results from the L'Hospital's Rule, the third from the rules for derivative of inverse functions, and the last from (A.12) and (A.13).

(ii) Let $\xi \equiv h_0/h_m$. Then $G_m(y)$ can be written as $G_m(y) = \mathbb{E}[h_m(y - D_m(t, t+L_m]) + (h_m + \xi h_m + p_m)(y - D_m(t, t+L_m])^-]$. It follows that for any $p_m > 0$ and $h_m > 0$, Q_m^* converges to a constant when $\xi \to 0$. Then we proceed to show that $\lim_{\xi \to 0} \hat{Q}_0^* = \infty$. Let $\bar{H}_0(Q) \equiv \frac{h_0[\sum_{i=1}^N (p_i + h_0) - h_0]}{\sum_{i=1}^N (p_i + h_0)} Q$, $\bar{A}_0(Q) \equiv \frac{h_0[\sum_{i=1}^N (p_i + h_0) - h_0]}{2\sum_{i=1}^N (p_i + h_0)} Q^2$, $\bar{Q}_0 \equiv \sqrt{\frac{2\lambda_0 K_0 \sum_{i=1}^N (p_i + h_0)}{h_0[\sum_{i=1}^N (p_i + h_0) - h_0]}}$. Then, it is easy to check that $\bar{A}_0' = Q\bar{H}_0'(Q)$ and $\bar{A}_0(\bar{Q}_0) = \lambda_0 K_0$. Because $\hat{H}_0(Q)$ is an increasing convex functions, by (A.13), we have $\hat{H}_0'(Q) \leq \frac{h_0[\sum_{i=1}^N (p_i + h_0) - h_0]}{\sum_{i=1}^N (p_i + h_0)} = \bar{H}_0'(Q)$, which implies that $\hat{A}_0'(Q) \leq \bar{A}_0'(Q)$. Note that $\hat{A}_0(0) = \bar{A}_0(0) = 0$. Therefore, we must have $\hat{A}_0(Q) \leq \bar{A}_0(Q)$.We then have $\hat{A}_0(\bar{Q}_0) \leq \bar{A}_0(\bar{Q}_0) = \lambda_0 K_0$. Because $\hat{A}_0(\hat{Q}_0^*) = \lambda_0 K_0$ and $\hat{A}_0(Q)$ is an increasing function, we have $\hat{Q}_0^* \geq \bar{Q}_0$. On the other hand, because $K_0 > 0$ and

$$
\lim_{\xi \to 0} \bar{Q}_0 = \sqrt{\frac{2\lambda_0 K_0 \sum_{i=1}^N (p_i + \xi h_m)}{\xi h_m [\sum_{i=1}^N (p_i + \xi h_m) - \xi h_m]}} = \infty, \tag{A.14}
$$

we have $\lim_{\xi \to 0} \hat{Q}_0^* = \infty$. Therefore, we have $\lim_{\xi \to 0} \beta_1 = \infty$.

It remains to show $\lim_{\xi\to 0}\beta_2 = 1$. Recall that we always have $C_0^* \le \hat{C}_0^* \le \hat{C}_0^*(r_0^*, Q_0^*)$, where the second equality follows from the definition of $\hat{C}_0^*$. Therefore, it suffices to show $\lim_{\xi\to 0}\hat{C}_0^*(r_0^*, Q_0^*) = \lim_{\xi\to 0} C_0^*$. Following the same logic of showing $\lim_{\xi\to 0}\hat{Q}_0^* = \infty$, one can easily prove $\lim_{\xi\to 0} Q_0^* = \infty$. Then we have

$$\lim_{\xi\to 0}\frac{C_0^*}{\hat{C}_0^*(r_0^*, Q_0^*)} = \lim_{\xi\to 0}\frac{H_0(Q_0^*)}{\hat{H}_0(Q_0^*)} = \lim_{\xi\to 0}\frac{H_0'(Q_0^*)}{\hat{H}_0'(Q_0^*)} = \lim_{\xi\to 0}\frac{\frac{\xi h_m \underline{p}}{\xi h_m + \underline{p}}}{\frac{\xi h_m[\sum_{i=1}^N (p_i + \xi h_m) - \xi h_m]}{\sum_{i=1}^N (p_i+\xi h_m)}} = 1, \tag{A.15}$$

where the second equality holds by the L'Hospital's Rule, and the third holds due to (A.12) and (A.13).

(iii) The proof emulates that of (ii). Let $\mu \equiv h_0/p_m$. Then $G_m(y)$ can be written as $G_m(y) = \mathbb{E}[h_m(y - D_m(t, t+L_m]) + (h_m + \mu p_m + p_m)(y - D_m(t, t+L_m])^-]$. It follows that for any $p_m > 0$ and $h_m > 0$, Q_m^* converges to a constant when $\mu \to 0$. In addition, (A.14) in this case is expressed as $\lim_{\mu\to 0}\bar{Q}_0 = \sqrt{\frac{2\lambda_0 K_0 \sum_{i=1}^N (p_i+\mu p_m)}{\mu p_m[\sum_{i=1}^N (p_i + \mu p_m) - \mu p_m]}} = \infty$. Therefore, we have $\lim_{\mu\to 0}\hat{Q}_0^* = \infty$ and $\lim_{\mu\to 0}\beta_1 = \infty$. To see $\lim_{\mu\to 0}\beta_2 = 1$, note that (A.15) in this case is expressed as

$$\lim_{\mu\to 0}\frac{C_0^*}{\hat{C}_0^*(r_0^*, Q_0^*)} = \lim_{\mu\to 0}\frac{\frac{\mu p_m \underline{p}}{\mu p_m + \underline{p}}}{\frac{\mu p_m[\sum_{i=1}^N (p_i+\mu p_m) - \mu p_m]}{\sum_{i=1}^N (p_i + \mu p_m)}} = 1.$$

∎

Proof of Proposition 4.1. $C(\cdot)$ is convex in k for each i, so is their sum. Therefore, $L(\Delta, k)$ is convex in k for a given Δ.∎

Proof of Proposition 4.2. Based on the definitions of k^{*1} and y^*, we prove the proposition by contradiction. First, assume $k^{*1} > y^*$. Because $C(y)$ is non-decreasing when $y > y^*$, $C(k^{*1}) > C(y^*)$, and hence $C(k^{*1}+i) > C(y^*+i)$, $\forall i \in [1, M]$. Therefore, we should have $L(\Delta, k^{*1}) = \sum_{i=1}^M \pi_i C(k^{*1}+$

$i) > \sum_{i=1}^{M} \pi_i C(y^* + i) = L(\Delta, y^*)$. This contradicts the definition of k^{*1}, and hence $k^{*1} \leq y^*$. Second, assume $k^{*1} = y^*$. Then $L(\Delta, k^{*1}) = \sum_{i=1}^{M} \pi_i C(k^{*1} + i) = \sum_{i=1}^{M} \pi_i C(y^* + i) > \sum_{i=1}^{M} \pi_i C(y^* + i - 1) = \sum_{i=1}^{M} \pi_i C(k^{*1} + i - 1) = L(\Delta, k^{*1} - 1)$. This also contradicts the definition of k^{*1}. We thus now have $k^{*1} < y^*$. Similarly, we can prove that $y^* - M \leq k^{*1}$ again by contradiction. Assume that $k^{*1} < y^* - M$. Because $C(y)$ is convex and y^* is the minimizer of $C(y)$, it can be easily seen that $C(k^{*1}) > C(y^* - M) > C(y^*)$. Therefore, $C(k^{*1}+i) > C(y^*-M+i), \forall i \in [1, M]$, and hence $L(\Delta, k^{*1}) > L(\Delta, y^*-M)$. This means that k^{*1} is not optimal, which contradicts the definition of k^{*1}. Therefore, we have $y^* \leq k^{*1} + M$. ∎

Proof of Proposition 4.3. Given $\Delta < Q$, $M + Q - \Delta$ is fixed, and $L(\Delta, k)$ is convex, because it is the summation of several convex functions $\pi_i C(k + i)$. ∎

Proof of Proposition 4.4. By contradiction, assume $k^{*2} \geq y^*$, then $C(k^{*2}+i) > C(k^{*2}+i-1), \forall i \in [1, M+Q-\Delta]$, and $L(\Delta, k^{*2}) = \sum_{i=1}^{M+Q-\Delta} \pi_i C(k^{*2}+i) > \sum_{i=1}^{M+Q-\Delta} \pi_i C(k^{*2} + i - 1) = L(\Delta, k^{*2} - 1)$. This contradicts the definition of k^{*2}, so $k^{*2} < y^*$. Now, assume $y^* - M - Q + \Delta > k^{*2}$, then $C(k^{*2} + i) > C(y^* - M - Q + \Delta + i), \forall i \in [1, M + Q - \Delta]$. Then we have $L(\Delta, k^{*2}) > L(\Delta, y^* - M - Q + \Delta)$, which contradicts the definition of k^{*2}. Therefore, $y^* - M - Q + \Delta \leq k^{*2} < y^*$.∎

Appendix B

The Unraveling Result

Recall that if the firm adopts the disclosure strategy, it can earn a profit $\pi_D^*(\overline{\lambda})$. Suppose $\pi_D^*(\overline{\lambda}) \geq 0$, i.e., $\overline{\lambda}$ is large enough that the firm can make nonnegative profits under the disclosure strategy.

We next focus on the optimal firm profit under the nondisclosure strategy. Grossman [40] shows that "*consumers with rational expectations will assume that the monopolist is of the worst possible quality consistent with his disclosure when he makes less than a full disclosure*." Following the existing literature, if the firm does not disclose its quality information, customers rationally assume the lowest valuation v_L for the service. Given price p, each customer's surplus is $v_L - p$. Then, for the profit maximizing firm, the optimal pricing problem is provided as follows:

$$\pi^*(\lambda) = \max_{p \geq 0} p\lambda \tag{B.1}$$
$$\text{subject to } \; v_L - p \geq 0,$$

The optimal strategy under the nondisclosure strategy for customers without learning opportunities is provided in the following lemma.

Lemma B.1. *Given that the firm attracts a population λ of customers without learning opportunities, the optimal price is $p^* = v_L$ and the firm earns a profit $\pi^*(\lambda) = v_L\lambda$.*

Proof. The result can be obtained by letting $v_L - p = 0$. The rest of the proof mimics that in Lemma 2.1.∎

By comparing the profits under disclosure and nondisclosure strategy, we characterize the optimal disclosure decision for the firm in the following lemma.

Proposition B.1 (Threshold Policy)**.** *Let* $\tilde{\alpha} = \frac{K}{(v_H - v_L)\overline{\lambda}}$. *When customers have no learning opportunities, the optimal quality disclosure decision for the firm is a threshold policy: it is optimal for the firm to disclose quality information if* $\alpha \geq \tilde{\alpha}$*; otherwise, the firm should not disclose. The optimal profit for the firm is* $\pi_{NL} = \max\{\pi_D^*, v_L\overline{\lambda}\}$. *In addition, if* $\overline{\lambda}$ *is larger, then the firm is more likely to disclose its quality information.*

Proof. It is optimal for the firm to adopt the disclosure strategy, if and only if $\pi_D^*(\overline{\lambda}) = [\alpha v_H + (1-\alpha)v_L]\overline{\lambda} - K \geq \pi^*(\overline{\lambda}) = v_L\overline{\lambda}$, i.e., $\alpha \geq \tilde{\alpha}$, where $\tilde{\alpha}$ is defined as $\tilde{\alpha} = \frac{K}{(v_H - v_L)\overline{\lambda}}$; otherwise when $\alpha < \tilde{\alpha}$, we have $\pi_D^*(\overline{\lambda}) < \pi^*(\overline{\lambda})$, and the firm should adopt the nondisclosure strategy.∎

This result is consistent with that of Jovanovic [58]. Note that without customer learning, $\overline{\lambda}$ only affects $\tilde{\alpha}$ but not the structure of the optimal disclosure decision.

Appendix C

Quality Disclosure with Congestion

Under the disclosure strategy, the quality information α is known to all customers. Let p_D, μ_D be the price and capacity level under the disclosure strategy, respectively. For each customer, her surplus is

$$\alpha v_H + (1-\alpha)v_L - p_D - h\mathbb{E}[W],$$

where $\mathbb{E}[W] = \frac{1}{\mu_D - \lambda}$ given the arrival rate λ ($\lambda \leq \overline{\lambda}$). Given that the profit maximizing firm would like to attract customers with demand rate λ, its optimal decision problem is provided as follows:

$$\begin{aligned}\pi_D^*(\lambda) = \max_{p_D \geq 0, \mu_D \geq 0} \; & p_D\lambda - c\mu_D - K \\ \text{subject to} \;\; & \alpha v_H + (1-\alpha)v_L - p_D - h\frac{1}{\mu_D - \lambda} \geq 0, \\ & \mu_D > \lambda.\end{aligned} \tag{C.1}$$

Then, given that the firm attract customers with demand rate λ, the optimal strategy under the disclosure strategy is provided in the following proposition.

Lemma C.1. *Suppose the firm attracts customers with demand rate λ, then* *the optimal strategy is as follows:*

$$\mu_D^*(\lambda) = \lambda + \sqrt{\frac{h\lambda}{c}}, \tag{C.2}$$

$$p_D^*(\lambda) = \alpha v_H + (1-\alpha)v_L - \sqrt{\frac{hc}{\lambda}}. \tag{C.3}$$

In addition, the firm earns a profit $\pi_D^*(\lambda) = [\alpha v_H + (1-\alpha)v_L - c]\lambda - 2\sqrt{hc\lambda} - K$.

Proof. Note that it is optimal to let

$$\alpha v_H + (1-\alpha)v_L - p_D - h\mathbb{E}[W] = 0,$$

i.e., to provide zero surplus for customers. Then we have

$$p_D = \alpha v_H + (1-\alpha)v_L - h\mathbb{E}[W].$$

Since $\pi_D = p_D\lambda - c\mu_D - K$, we have

$$\pi_D = (\alpha v_H + (1-\alpha)v_L - h\frac{1}{\mu - \lambda})\lambda - c\mu_D - K.$$

One can verify that π_D is a concave function in μ. Based on the first-order optimality condition, we can obtain the optimal solution $\mu_D^*(\lambda) = \lambda + \sqrt{\frac{h\lambda}{c}}$, $p_D^*(\lambda) = \alpha v_H + (1-\alpha)v_L - \sqrt{\frac{hc}{\lambda}}$ and $\pi_D^*(\lambda) = [\alpha v_H + (1-\alpha)v_L - c]\lambda - 2\sqrt{hc\lambda} - K$. ∎

The optimal solution for (C.1) can be easily obtained by letting $\alpha v_H + (1-\alpha)v_L - p_D - h\mathbb{E}[W] = 0$. It is intuitive that as λ increases, then the firm should charge a higher price and invest in a large capacity. Finally, note that $\pi_D^*(\lambda)$ is a convex function of λ.

Next, we provide the optimal decision for the firm under the disclosure strategy.

Lemma C.2. *Let* $\tilde{\lambda}_D = \left(\frac{\sqrt{hc}+\sqrt{hc+K[\alpha v_H+(1-\alpha)v_L-c]}}{\alpha v_H+(1-\alpha)v_L-c}\right)^2$. *If* $\overline{\lambda} \geq \tilde{\lambda}_D$, *the firm can obtain an optimal profit* $\pi_D^* = \pi_D^*(\overline{\lambda})$ *by adopting the disclosure strategy; otherwise if* $\overline{\lambda} < \tilde{\lambda}_D$, *the firm should not enter the market.*

Proof. We know that $\pi_D^*(\lambda) = [\alpha v_H + (1-\alpha)v_L - c]\lambda - 2\sqrt{hc\lambda} - K$. The only positive root in term of $\sqrt{\lambda}$ for the equation $[\alpha v_H + (1-\alpha)v_L - c]\lambda - 2\sqrt{hc\lambda} - K = 0$ is $\sqrt{\lambda} = \sqrt{\tilde{\lambda}_D}$, or equivalently, $\lambda = \tilde{\lambda}_D$, where $\tilde{\lambda}_D$ is defined as

$$\tilde{\lambda}_D = \left(\frac{\sqrt{hc} + \sqrt{hc + K[\alpha v_H + (1-\alpha)v_L - c]}}{\alpha v_H + (1-\alpha)v_L - c}\right)^2.$$

Note that π_D^* is a quadratic function in term of $\sqrt{\lambda}$. Therefore, $\pi_D^* = \pi_D^*(\overline{\lambda}) \geq 0$ when $\overline{\lambda} \geq \tilde{\lambda}_D$ and the firm can make a nonnegative profit by adopting the disclosure strategy; otherwise when $\overline{\lambda} < \tilde{\lambda}$, $\pi_D^* < 0$, and the firm should not enter the market. ∎

Since customers are homogeneous, as long as the total arrival rate $\overline{\lambda}$ is large enough, it is optimal for the firm to enter the market by adopting the disclosure strategy, and otherwise it is not profitable for the firm to enter the market.

For analytical simplicity and to obtain managerial insights, we focus on the $S(1)$ framework (our results still hold qualitatively under the $S(N)$ framework). Specifically, we study the $S(1)$ framework where each arrival in the Poisson arrival process is assumed to obtain one sample from a previous arrival who purchased the service from the firm.

Under the $S(1)$ framework, customer i purchases if and only if

$$\alpha_i(1)v_H + (1-\alpha_i(1))v_L \geq p + h\mathbb{E}[W],$$

i.e., she purchases the service if the valuation derived from her sample is no smaller than the selling price plus the expected waiting cost.

Then we have the following lemma.

Lemma C.3. *Under the nondisclosure strategy, (i) suppose the firm commits to attract customers who obtained an "H" sample only: if* $\overline{\lambda} \geq \frac{4ch}{\alpha(v_H-c)^2}$*, the optimal strategy is* $p_H^* = v_H - \sqrt{\frac{hc}{\alpha\overline{\lambda}}}$*,* $\mu_H^* = \alpha\overline{\lambda} + \sqrt{\frac{\alpha\overline{\lambda}h}{c}}$ *with the optimal profit* $\pi_H^* = \alpha(v_H - c)\overline{\lambda} - 2\sqrt{\alpha hc\overline{\lambda}}$*; otherwise if* $\overline{\lambda} < \frac{4ch}{\alpha(v_H-c)^2}$*, the firm should not enter the market; (ii) suppose the firm commits to attract all customers: if* $\overline{\lambda} \geq \frac{4ch}{(v_L-c)^2}$*, the optimal strategy is* $p_L^* = v_L - \sqrt{\frac{hc}{\overline{\lambda}}}$*,* $\mu_L^* = \overline{\lambda} + \sqrt{\frac{\overline{\lambda}h}{c}}$ *with the optimal profit* $\pi_L^* = (v_L - c)\overline{\lambda} - 2\sqrt{hc\overline{\lambda}}$*; otherwise if* $\overline{\lambda} < \frac{4ch}{(v_L-c)^2}$*, the firm should not enter the market.*

Proof. (i): The firm adopts the nondisclosure strategy to attract H-type customers only. In this case, the proportion of H-type customers is α. Because only customers who obtained an H sample would purchase this service, given the total population λ (including customers who obtained an L sample), the population of customers who actually purchase is $\alpha\lambda$, and hence the expected delay is $\mathbb{E}[W] = \frac{1}{\mu-\alpha\lambda}$. Note that under the nondisclosure strategy, customers do not know the exact value of α. If the firm aims to attract H-type customers only, then the surplus for H-type customers must be nonnegative, and the utility for L-type customers must be negative. Let μ_H, p_H and π_H denote the capacity, price and profit in this case, respectively. Given the total population λ, then the firm's profit maximization problem is provided as follows:

$$
\begin{aligned}
\pi_H^*(\lambda) = \max_{p_H \geq 0, \mu_H \geq 0} \ & \alpha p_H \lambda - c\mu_H \\
\text{subject to } \ & v_H - p_H - h\frac{1}{\mu_H - \alpha\lambda} \geq 0, \\
& v_L - p_H - h\frac{1}{\mu_H - \alpha\lambda} < 0, \\
& \mu_H > \alpha\lambda,
\end{aligned}
$$

The optimal solution can be obtained by letting $v_H - p_H - h\mathbb{E}[W] = 0$, where $\mathbb{E}[W] = \frac{1}{\mu_H - \alpha\lambda}$, and hence $p_H = v_H - \frac{h}{\mu_H - \alpha\lambda}$. By substituting the preceding formula in $\pi_H = \alpha p_H \lambda - c\mu_H$, we obtain $\pi_H = (v_H - \frac{h}{\mu - \alpha\lambda})\alpha\lambda - c\mu_H$. Note that π_H is concave in μ_H. Based on the first-order optimality condition, we can obtain the optimal solution $\mu_H^*(\lambda) = \alpha\lambda + \sqrt{\frac{\alpha\lambda h}{c}}$ and $p_H^*(\lambda) = v_H - \sqrt{\frac{hc}{\alpha\lambda}}$. The firm earns a profit $\pi_H^*(\lambda) = \alpha(v_H - c)\lambda - 2\sqrt{\alpha hc\lambda}$ under the optimal solution. Note that $\pi_H^*(\lambda) \geq 0$, if and only if $\overline{\lambda} \geq \frac{4ch}{\alpha(v_H - c)^2}$. Therefore, the firm can make nonnegative profits if $\overline{\lambda} \geq \frac{4ch}{\alpha(v_H - c)^2}$; otherwise the firm should not enter the market.

(ii): The firm adopts the nondisclosure strategy to attract all customers. Notice that if the firm is able to attract L-type customers, it is able to attract H-type customers as well since $v_H > v_L$ (keep in mind all customers are subject to the same expected delay). Because all customers are willing to purchase, then the expected delay is $\mathbb{E}[W] = \frac{1}{\mu_L - \lambda}$ in this case. Let μ_L, p_L and π_L denote the capacity, price and profit in this case, respectively. The firm's profit maximization problem is given by

$$\begin{aligned} \pi_L^*(\lambda) = & \max_{p_L \geq 0, \mu_L \geq 0} \; p_L\lambda - c\mu_L \\ \text{subject to } & \; v_L - p_L - h\frac{1}{\mu_L - \lambda} \geq 0, \\ & \mu_L > \lambda. \end{aligned}$$

The optimal solution can be obtained by letting $v_L - p_L - h\mathbb{E}[W] = 0$ where $\mathbb{E}[W] = \frac{1}{\mu_L - \lambda}$.Then $p_L = v_L - \frac{h}{\mu_L - \lambda}$. The rest of the proof is similar to case (i). ∎

The profit for the optimal nondisclosure strategy can be described as

$$\pi_{ND}^* = \max\{\pi_H^*, \pi_L^*, 0\} = \max\{\alpha(v_H - c)\overline{\lambda} - 2\sqrt{\alpha hc\overline{\lambda}}, (v_L - c)\overline{\lambda} - 2\sqrt{hc\overline{\lambda}}, 0\}.$$

We first focus on the case in which the firm can make nonnegative profits under nondisclosure strategy by attracting either all customers or H-type customers only. The following proposition provides the optimal nondisclosure strategy.

Lemma C.4. *When $\overline{\lambda} \geq \max\{\frac{4hc}{(v_L-c)^2}, \frac{4hc}{\alpha(v_H-c)^2}\}$, i.e., $\overline{\lambda}$ is large enough, there exists $\tilde{\alpha} = \left(\frac{2\sqrt{hc\overline{\lambda}}+\sqrt{4hc\overline{\lambda}-4(v_H-c)\overline{\lambda}(-(v_L-c)\overline{\lambda}+2\sqrt{hc\overline{\lambda}})}}{2(v_H-c)\overline{\lambda}}\right)^2$ such that it is optimal to attract customers who obtained an "H" sample only if $\alpha > \tilde{\alpha}$; otherwise it is optimal to attract all customers.*

Proof. When $\overline{\lambda}$ is large enough, i.e., $\overline{\lambda} \geq \max\{\frac{4hc}{(v_L-c)^2}, \frac{4hc}{\alpha(v_H-c)^2}\}$, we have $\pi_H^* \geq 0$ and $\pi_L^* \geq 0$. It is optimal to attract H-type customers only if $\pi_H^* \geq \pi_L^*$. $\pi_H^* - \pi_L^* = (v_H - c)\overline{\lambda}\alpha - 2\sqrt{hc\overline{\lambda}\alpha} - (v_L - c)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}}$. This is a quadratic function of $\sqrt{\alpha}$. Since $(v_L - c)\overline{\lambda} - 2\sqrt{hc\overline{\lambda}} > 0$, the only positive root of in term of $\sqrt{\alpha}$ for the equality

$$(v_H - c)\overline{\lambda}\alpha - 2\sqrt{hc\overline{\lambda}\alpha} - (v_L - c)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}} = 0,$$

is

$$\sqrt{\alpha} = \frac{2\sqrt{hc\overline{\lambda}} + \sqrt{4hc\overline{\lambda} - 4(v_H - c)\overline{\lambda}(-(v_L - c)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}})}}{2(v_H - c)\overline{\lambda}}.$$

In addition, this root is less than 1, which can be easily seen from $(v_H - c)\overline{\lambda}\alpha - 2\sqrt{hc\overline{\lambda}\alpha} - (v_L - c)\overline{\lambda} + 2\sqrt{hc\overline{\lambda}} > 0$ when $\alpha = 1$. ∎

Intuitively, if the expected quality level is high enough, then it is more profitable by attracting H-type customers only; otherwise, it is optimal to attract all customers.

Corollary C.1. *The optimal capacity under the nondisclosure strategy is discontinuous in α. Moreover, it is non-monotonic at $\tilde{\alpha}$.*

Proof. This is implied by Lemma C.4. If $\alpha < \tilde{\alpha}$, it is optimal to attract all customers, and corresponding capacity is $\mu_L^* = \overline{\lambda}+\sqrt{\frac{\overline{\lambda}h}{c}}$, which is independent of α. Then $\lim_{\alpha\to\tilde{\alpha}^-} = \overline{\lambda}+\sqrt{\frac{\overline{\lambda}h}{c}}$. If $\alpha > \tilde{\alpha}$, then it is optimal to attract H-type customers only, and the corresponding optimal capacity is $\mu_H^* = \alpha\overline{\lambda}+\sqrt{\frac{\alpha\overline{\lambda}h}{c}}$. Then $\lim_{\alpha\to\tilde{\alpha}^+} = \tilde{\alpha}\overline{\lambda}+\sqrt{\frac{\tilde{\alpha}\overline{\lambda}h}{c}}$. Hence the optimal capacity is discontinuous at $\tilde{\alpha}$. It is non-monotonic because $\alpha\overline{\lambda}+\sqrt{\frac{\alpha\overline{\lambda}h}{c}}$ is nondecreasing in α and $\tilde{\alpha}\overline{\lambda}+\sqrt{\frac{\tilde{\alpha}\overline{\lambda}h}{c}} < \overline{\lambda}+\sqrt{\frac{\overline{\lambda}h}{c}}$.∎

Corollary C.1 shows that under the nondisclosure strategy, the capacity decision should be made cautiously, i.e., a higher service quality not necessarily implies the need for a larger capacity. This is due to the fact that when quality level is large enough, then the firm is more profitable by only attracting customers who obtained an "H" sample.

The following proposition provides the optimal nondisclosure strategy under the condition $\overline{\lambda} < \max\{\frac{4hc}{(v_L-c)^2}, \frac{4hc}{\alpha(v_H-c)^2}\}$.

Proposition C.1. *Under the condition $\overline{\lambda} < \max\{\frac{4hc}{(v_L-c)^2}, \frac{4hc}{\alpha(v_H-c)^2}\}$, the optimal nondisclosure strategy can be described as follows: (1) If $\frac{4hc}{\alpha(v_H-c)^2} \leq \overline{\lambda} < \frac{4hc}{(v_L-c)^2}$, the firm should adopt nondisclosure strategy to attract customers who obtained an "H" sample only. (2) If $\frac{4hc}{(v_L-c)^2} \leq \overline{\lambda} < \frac{4hc}{\alpha(v_H-c)^2}$, the firm should adopt nondisclosure strategy to attract all customers. (3) Otherwise the firm should not enter the market.*

Proof. If the firm aims to attract all customers, it can make non-negative profit, if and only if $\overline{\lambda} \geq \frac{4hc}{(v_L-c)^2}$. If the firm aims to attract H-type customers only, it can make non-negative profit, if and only if $\overline{\lambda} \geq \frac{4hc}{\alpha(v_H-c)^2}$. If $\frac{4hc}{(v_L-c)^2} \leq \overline{\lambda} < \frac{4hc}{\alpha(v_H-c)^2}$, the firm cannot make non-negative profit by attracting H customers only, but the firm can make non-negative profits by attracting attracting all customers. Therefore, attracting all customers is the optimal strategy in this case. Similarly, if $\frac{4hc}{\alpha(v_H-c)^2} \leq \overline{\lambda} < \frac{4hc}{(v_L-c)^2}$, then the firm cannot make non-negative profit by attracting all customers, but it can make non-negative profit by attracting H-type customers only. Therefore, attracting H-type customers only is the optimal strategy in this case. If $\overline{\lambda} < \min\{\frac{4hc}{(v_L-c)^2}, \frac{4hc}{\alpha(v_H-c)^2}\}$, $\pi_H^* < 0$ and $\pi_L^* < 0$, and the firm should not enter the market.∎

Appendix D

Single-Stage Inventory Model

Assume that each shipment incurs a fixed setup cost K. We adopt the convention that discrete units of inventories can be approximated by continuous variables. We will follow that convention throughout. Consequently, the long-run average cost of the single-stage problem can be approximated as follows:

$$C(r,Q) = \frac{\lambda K + \int_r^{r+Q} G(y)dy}{Q}. \tag{D.1}$$

Clearly, the approximation is adequate when Q is sufficiently large. The objective is to determine the values of r and Q such that minimize $C(r,Q)$. For any fixed Q, define $r(Q) \equiv \arg\min_r C(r,Q)$. If the optimal solution is not unique, we choose the largest one; this convention will be used throughout the section. Also, define $C(Q) \equiv C(r(Q),Q)$, $(r^*,Q^*) \equiv \arg\min_{r,Q} C(r,Q)$ and $C^* \equiv C(r^*,Q^*)$. Finally, let $H(Q) \equiv G(r(Q))$ and $A(Q) \equiv QH(Q) - \int_0^Q H(y)dy$.

Lemma D.1 (Zheng 97). *Under Assumption 3.3.1, the following results hold: (i) $C(r,Q)$ is jointly convex in r and Q. (ii) $G(r(Q)) = G(r(Q)+Q)$ and $G(r^*) = G(r^*+Q^*) = C^*$. (iii) $C(Q)/C^* \le \epsilon(Q/Q^*)$, where $\epsilon(q) = (q + q^{-1})/2$. (iv) $H(Q)$ is an increasing convex function, $\int_0^Q H(y)dy \ge \frac{1}{2}QH(Q)$ and $H(Q^*) = C^*$. (v) $A(Q)$ is an increasing function and $A(Q^*) = \lambda K$.*

By Lemma D.1, we further obtain the following properties.

Lemma D.2. *(i) For any* $y \in [r^*, r^* + Q^*], G(y) \leq C^*$. *(ii)* $\lambda K \leq \frac{1}{2}C^*Q^*$.