

Three Essays on Modeling Consumer Behavior and Its Operations Management Implications

by
Hakjin Chung

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2016

Doctoral Committee:

Professor Hyun-Soo Ahn, Chair
Assistant Professor Eunshin Byon
Assistant Professor Stefanus Jasin
Assistant Professor James Ostler

TABLE OF CONTENTS

LIST OF FIGURES	iv
LIST OF TABLES	vi
LIST OF APPENDICES	vii
ABSTRACT	viii
CHAPTER	
I. Introduction	1
II. Effect of Strategic in a Service System	5
2.1 Introduction and Literature review	5
2.2 Model description	9
2.3 The Customers' Problem	11
2.3.1 Individual Customers	12
2.3.2 Collective Customers	12
2.3.3 Social Customers	13
2.3.4 Threshold Analysis	14
2.4 The Firm's Problem	17
2.4.1 The Profit-Maximizing Firm	18
2.4.2 The Social Firm	22
2.5 Analysis	26
2.5.1 On the Firm's Selfishness	26
2.5.2 On Customers' Selfishness and the Price of Anarchy	30
2.5.3 Price As a Tool to Achieve the Social Optimum . . .	32
2.6 Conclusion	35
III. Dynamic Pricing and Loyalty Programs	37
3.1 Introduction	37

3.1.1	Summary of Contributions	40
3.2	Literature Review	43
3.3	The Model	45
3.3.1	Consumers' Problem	46
3.3.2	The Seller's Problem	52
3.4	The Black-Out Model	62
3.5	Dynamic Adjustment Model	68
3.6	Conclusion	75
 IV. On (Rescaled) Multi-Attempt Consumer Choice Model and Its Implication on Assortment Optimization		 77
4.1	Introduction	77
4.2	Choice Approximation Models	83
4.2.1	Multi-Attempt Model	83
4.2.2	The Rescaled Multi-Attempt Model	85
4.2.3	Numerical Experiments	87
4.3	Assortment Optimization	88
4.3.1	Optimization Formulation	88
4.3.2	Numerical Experiments	90
4.4	Conclusion	93
 V. Conclusions		 94
 APPENDICES		 97
 BIBLIOGRAPHY		 123

LIST OF FIGURES

Figure

2.1	Threshold as a function of service rate at two different prices (p): 4 (left) and 6 (right) when $\lambda = 1$, $R = 10$ and $h = 0.5$	15
2.2	Profit as a function of μ when $p = 5$, $\lambda = 1$, $R = 10$, $h = 1$, and $c(\mu) = 2\mu^3$	19
2.3	Profit as a function of p when $\mu = 0.6$, $\lambda = 1$, $R = 10$, $h = 1$, and $c(\mu) = 2\mu^3$	19
2.4	The original (solid line) and the relaxed (dashed line) profit functions when $\lambda = 1$, $R = 10$, $p = 6$, $h = 1$, and $c(\mu) = 2\mu^3$	20
2.5	Social welfare as a function of the service rate when $\lambda = 1$, $R = 10$, $h = 0.5$, $p = 5$, and $c(\mu) = 2\mu^2$	23
2.6	The service rate choice depending on the firm's objective function and the type of customers. With collective customers (left), the customer utility increases in μ and the social firm always chooses a higher service rate than the selfish firm. However, the individual customer utility is not monotone and in this example, the social firm chooses a lower service rate than the selfish firm. The parameters in both figures are: $R = 10$; $p = 6$; $h = 1$; $c(\mu) = 4.5\mu^4$	28
2.7	Comparison of the profit-maximizing and the welfare-maximizing service rates when customers are social at two different prices (p): 5 (left) and 9 (right). The other parameters in both figures are: $R = 10$; $h = 1$; $c(\mu) = 2\mu^4$. Note that the profit can be higher than welfare because social customers' utility can be negative.	29
2.8	<i>Benefit of Anarchy</i> when $R = 10$, $h = 1$, $p = 5.2$, and $c(\mu) = 3\mu$: The selfish firm sets a lower service rate for collective customers than for individual customers ($\mu^{CP} = 0.72$ and $\mu^{IP} = 1.05$). Consequently, the utility of collective customers is lower than the utility of individual customers ($\mathcal{U}^{CP} = 0.71$ and $\mathcal{U}^{IP} = 1.29$), that is, the PoA = $\frac{1.29}{0.71} = 1.8169$, so individual customers have an 80% higher utility relative to collective customers.	31

2.9	<i>Social Benefit of Customer Selfishness</i> with a selfish firm when $R = 10$, $h = 1$, $p = 6.5$, and $c(\mu) = 3\mu^2$. The selfish firm sets a lower service rate for social customers than for individual customers ($\mu^{IP} = 0.75$ and $\mu^{SP} = 0.86$). As a result, the equilibrium welfare is even lower when customers are social than when selfish ($\mathcal{W}^{IP} = 3.01$ and $\mathcal{W}^{SP} = 2.80$).	32
2.10	Equilibrium Welfare as a function of price when $R = 10$, $h = 1$, $c(\mu) = 2\mu^4$. To achieve the first-best equilibrium, the price should be less than $\bar{p}^{CS} = 2.83$ for collective customers, and between $\bar{p}_L^{IS} = 5.56$ and $\bar{p}_H^{IS} = 6.66$ for individual customers.	33
3.1	Consumer behavior depending on their type (<i>cash-only consumer</i> or <i>loyalty consumer</i>), reservation price (V), and point-worth (Θ) given price p and point requirement q	49
3.2	Examples of expected revenue-to-go functions in price when both $F(\cdot)$ and $G(\cdot)$ are truncated normal (i.e., IGFR): $F \sim \text{truncated } N(30, 10)$, $q = 1$, $t = 1$, and $y = 1$. In figure (a), $\bar{\beta} = 0.6$, $G \sim \text{truncated } N(15, 2)$ and $R = 10$. In figure (b), $\bar{\beta} = 1$, $G \sim \text{truncated } N(28, 1)$ and $R = 30$	54
3.3	The optimal price, $p_t^*(y)$, and the <i>cash-only price</i> , $p_t^c(y)$, depending on the inventory level (left figure; for given $t = 10$) and the remaining time (right figure; for given $y = 8$) when $\lambda = 0.8$, $\bar{\beta} = 0.7$, $F \sim \text{Uniform}[0, 100]$, $G \sim \text{Uniform}[0, 10]$, $q = 10$, and $R = 55$	59
3.4	The optimal strategy when the seller has y units of inventory with t periods to go until the end of the season: $\lambda = 0.8$, $\bar{\beta} = 0.5$, $q = 10$, $R = 55$, $F(\cdot) \sim \text{Uniform}[0, 100]$, $G \sim \text{Uniform}[0, 10]$	64
3.5	Examples of best revenue ($\max_p J_t^Q(p, q, y)$) and best price ($\arg \max_p J_t^Q(p, q, y)$) for each given point requirement when $F \sim N(60, 30)$, $G \sim N(6, 3)$, $q \in \{7.0, 7.2, 7.4, \dots, 12.0\}$, $\bar{\beta}(q) = \frac{1}{\log q}$, $R(q) = 8q$, $t = 1$, and $y = 1$	70
3.6	The optimal price and point requirement depending on the remaining inventory level and period when $\lambda = 0.9$, $Q = \{7, 8, 9, 10, 12\}$, $\bar{\beta}(q) = 0.6 - 0.05q$, $R(q) = 5q$, $F \sim \text{Uniform}[0, 100]$, $G \sim \text{Uniform}[0, 10]$, $t = 9$ (left), and $y = 6$ (right).	72
A.1	Illustration of the LHS and RHS of (A.4) where $\mu = 1$ is the largest solution (left: $b = 5, b_0 = 16$), the middle solution (middle: $b = 3, b_0 = 4$), or the smallest solution (right: $b = 1, b_0 = 1.25$).	100

LIST OF TABLES

Table

2.1	Scenarios based on customers' and firm's objectives	11
3.1	Summary statistics for percentage revenue change from allowing reward sales in every periods (O) and allowing only when it is optimal (B) compared to the bench-mark in which there is no reward sale. The number in the bracket represents the proportion of open decisions for the black-out seller.	67
3.2	Summary statistics for percentage revenue improvement from the seller's discretion at multiple levels.	74
4.1	Comparison of approximation accuracy of various models	87
4.2	Average relative gap in expected revenue for Markovian model and multi-attempt models with its computing time in second.	91
4.3	Average relative gap in expected revenue for constrained (non-scaled and rescaled) 2-attempt models with its computing time in second. .	92

LIST OF APPENDICES

Appendix

A.	Proofs and Technical Details for Chapter II: Effect of Strategic Behavior In a Service System	98
B.	Proofs and Technical Details for Chapter III: Dynamic Pricing and Loyalty Programs	111
C.	Proofs and Technical Details for Chapter IV: On (Rescaled) Multi-Attempt Consumer Choice Model and Its Impli- cation on Assortment Optimization	119

ABSTRACT

Three Essays on Modeling Consumer Behavior and Its OM Implications

by

Hakjin Chung

Chair: Hyun-Soo Ahn

Traditionally, models used in operations management have considered the firm side of the problem by making simplifying assumptions on demand or market: For instance, in many inventory models, demand is simply assumed to be an exogenous random variable. In practice, however, consumers or agents in the market actively make decisions or choices based on self interest. This dissertation aims to analyze how insights and results from traditional models are affected when we account for such active decision making by consumers or the market.

In Chapter II, we study how the customers' decision of joining the queue to receive a service varies by the individual incentive (selfishness) as well as the firm's capacity decision, which also depends on the firm's selfishness. By considering three customer types: individual, collective, and social, and two firm types: profit maximizing and welfare maximizing, we are able to disentangle the effects of selfishness of the customers and the firm, and the interactions between these two in equilibrium. Among other results, we find that there can be a "benefit of selfishness" to consumers and the system, in contrast to the price of anarchy literature. That is, customers acting in

their individual self interest in response to a strategic firm can have a greater overall utility compared to strategically behaving customers.

In Chapter III, we discuss the customers' redemption behavior of loyalty points and its impact on the seller's pricing and inventory rationing strategy. We model the customer choice between cash or loyalty points by characterizing consumers in three dimensions: the amount of points in their accounts, their perceived valuation of points, and their valuation of the product. Applying this choice model into the seller's dynamic pricing model, we characterize the seller's optimal strategy that specifies the optimal cash price, the control of reward sales (black-out decision), and the redemption points. We find that a loyalty program has non-trivial impacts on the seller's strategy, for instance, the seller either adds a premium or offers a discount on price.

In Chapter IV, we study the substitution behavior of customers when their preferred product is not available, and the seller's assortment optimization problem. Motivated by the classic exogenous demand model and the recently developed Markov chain model, we propose a new approximation to the general customer choice model based on random utility called *rescaled multi-attempt* model, in which a customer may consider several substitutes before finally deciding to not purchase anything. The key feature of our proposed approach is that the resulting approximate choice probability can be explicitly written. From a practical perspective, this allows the decision maker to use an off-the-shelf solver, or borrow existing algorithms from literature, to solve a general assortment optimization problem with a variety of real-world constraints.

CHAPTER I

Introduction

Consumers actively make decisions based on their self interest. Taking account of such reactions is very important when the firm makes decisions such as capacity investment, pricing adjustment, and product line selection. While the existing literature of operations management often relies on simplifying assumptions about demand, it is noteworthy that properly incorporating consumers' behavior provides more practical insights. In that vein, this dissertation aims to analyze how insights and results from traditional models are affected, when we account for this active decision making by consumers. Focusing on this goal, we cover a variety of the firm's decisions, such as optimal capacity investment in a service system, revenue management strategies, and assortment planning.

In Chapter II, we analyze how the strategic behavior by consumers affects the firm's capacity decision and the system equilibrium. Some service systems have the benefit of full coordination to maximize total social welfare. For example, highway planners both decide on highway capacities and ration access in congested periods by enforcing arrival thresholds at entrance ramps. Big bandwidth Internet users, such as Netflix or Spotify, can achieve customer coordination, so customers act for their mutual benefit, while the internet access provider maximizes its own profit. We study how the customers' decision of joining the queue to receive a service varies by the

individual incentive (selfishness) as well as the firm’s capacity decision, which also depends on the firm’s selfishness. By considering three customer types: individual, collective, and social, and two firm types: profit maximizing and welfare maximizing, we are able to disentangle the effects of selfishness of the customers and the firm and to answer to the following questions: (i) how the customers decisions to join the system vary by their selfishness, (ii) how the firm’s capacity decision of the system varies by its selfishness, and (iii) what are the interactions between the selfishness of customers’ and the firm’s in equilibrium. Among other results, we find that when the firm endogenously sets the service rate, there can be a “benefit of anarchy,” rather than the usual “price of anarchy.” That is, customers acting in their individual self interest in response to a strategic firm can have a greater overall utility than customers who are acting to maximize their overall utility. Similarly, the selfishness of customers is not always socially costly (in terms of the total welfare - customer utility plus firm profit) when the firm is a profit maximizer. We also find that for selfish customers, higher capacity can lead to lower utility.

In Chapter III, we study a loyalty program and its impact on the seller’s pricing and inventory rationing strategies. Loyalty programs are very popular marketing tools used in many industries, and one of their most significant perks is to allow members to acquire products by redeeming loyalty points (e.g., airline mileage, hotel points, credit card reward points). The amount of goods purchased through loyalty program is quite significant: for example, Hilton issued 4.3 million reward room nights in 2012. In Chapter III, we examine the customers’ redemption behavior of loyalty points and it’s impact on the seller’s strategy. In many loyalty programs, the seller (franchisee) receives monetary compensation from the loyalty point issuer (franchisor or brand headquarters) when consumers use points. We modeled the customer choice between using cash or loyalty points, and incorporated this choice model into the sellers dynamic pricing model where the revenues from both cash sales and reimbursement for

reward sales are embedded in each period. We develop the first mathematical model that characterizes the customers' redemption behavior. Our customer choice model allows consumers to be heterogeneous in three dimensions: the amount of points in their accounts, their perceived valuation of points, and their valuation of the product, and characterizes when consumers use points or cash as a function of the price and the number of redemption points. Applying this choice model into the seller's dynamic pricing model, we characterized the seller's optimal strategy that specifies the optimal cash price, the control of reward sales (black-out decision), and the number of redemption points. We find that with the loyalty program affects the seller's optimal price either lower or higher than without, depending on whether the reimbursement is lower or higher than the optimal price without loyalty program. We also find that, if the seller can discretionally control the point redemption options (e.g., some sellers can allow point redemptions only in certain periods or can dynamically change the redemption requirement points), the optimal strategy is to discourage the reward sales (by blackout or increasing the required points) if the amount of inventory is relatively smaller compared to the remaining periods.

In Chapter IV, we study the substitution behavior of customers and the seller's assortment optimization problem. When their most preferred product is not available, customers may substitute other available products. Understanding such a substitution behavior is mathematically challenging but very important as the seller needs to select an optimal subset of products to offer to maximize revenue. Motivated by the classic exogenous demand model and the recently developed Markov chain model, we propose a new approximation to the general customer choice model based on random utility called the *multi-attempt* model, in which a customer may consider several substitutes before finally deciding to not purchase anything. We show that the approximation error of the multi-attempt model decreases exponentially in the number of attempts. However, despite its strong theoretical performance, the empirical perfor-

mance of the multi-attempt model is not satisfactory. This motivates us to construct a modification of the multi-attempt model called the *rescaled multi-attempt* model. We show that the rescaled 2-attempt model is exact when the underlying true choice model is Multinomial Logit (MNL); if, however, the underlying true choice model is not MNL, we show numerically that the approximation quality of the rescaled 2-attempt model is very close to that of the Markov chain model. The key feature of our proposed approach is that the resulting approximate choice probability can be explicitly written. From a practical perspective, this allows the decision maker to use an off-the-shelf solver, or borrow existing algorithms from the literature, to solve a general assortment optimization problem with a variety of real-world constraints.

CHAPTER II

Effect of Strategic in a Service System

2.1 Introduction and Literature review

We examine the effect of individual incentives (selfishness) in the context of a service system modeled as a single-server queue. Specifically, we consider a system where customers decide to join the queue or not based on their objective, which can range from individual utility maximization to maximization of overall customer utility to social welfare maximization, depending on the level of selfishness. In response to the customers' behavior, the firm decides the rate at which customers are to be served in order to maximize its own objective function, which can be profits (selfish objective) or social welfare (social objective). We examine how the system-equilibrium outcome, characterized by service rate and customer joining behavior, depends on the objectives of the customers and the firm.

Classical economics (*Mas-Colell et al.* 1995) considers the trade-off between consumer surplus and producer surplus in an economic exchange. Recent literature on the price of anarchy (*Roughgarden* 2005 and *Gilboa-Freedman et al.* 2014) is concerned with the differences in policies and outcomes when there are selfish (individually optimizing) customers versus when they are coordinated. The latter research stream takes the infrastructure as given, so does not consider the firm's capacity or infrastructure cost. In contrast, we combine aspects of both streams of research and model the

firm’s infrastructure decisions and costs as well as the customers’ joining decisions, and include the effects of having a profit maximizing firm versus having a firm whose objective is social welfare maximization. By considering three objective functions for the customers: individual, collective (customers as a group), and social (customers plus firm), and two for the firm: selfish and social (customers plus firm), we are able to disentangle the effects of both the price of anarchy and the consumer-producer surplus trade-off. That is, we can look at the benefits of coordination or regulation at multiple levels.

We find, surprisingly, that there can be a “*Benefit of Anarchy*” in contrast to the price of anarchy (PoA) literature. The PoA is defined as the ratio of the individual and collective utilities (the utility of individual customers when they are selfishly trying to maximize their own utility and the utility when they are coordinated to maximize their average utility). We find that, although the PoA must be less than 1 for a fixed service rate (i.e., customer utility must be higher when customers are not selfish), the PoA can be greater than 1 when the firm can endogenously choose the service rate knowing customer behavior.

Some service systems have the benefit of full coordination to maximize total social welfare. For example, highway planners both decide on highway capacities and ration access in congested periods by enforcing arrival thresholds at entrance ramps. Others can achieve customer coordination, so customers act for their mutual benefit, while the firm maximizes its own profit. For example, a big bandwidth user of the internet such as Netflix or Spotify can coordinate the traffic it generates to optimize performance for all its users, taking internet access capacity as fixed. Note that although our model setup is from the perspective of customers and firms choosing actions based on particular objective functions, social or collective customer actions can be enforced through thresholds, without expecting individual customers to choose those actions.

There has been an extensive stream of research that analyzes customers’ strategic

behavior in queueing systems (see the comprehensive review by *Hassin and Haviv* 2003). The first paper in this area is the seminal paper by *Naor* (1969), who studied customers' joining decision after observing the queue length. Naor showed that selfish customers do not consider their negative externalities, thus resulting in suboptimal welfare, and that levying a congestion toll can induce selfish customers to choose the socially optimal decisions. Naor's work has been generalized to multi-server systems (e.g., *Knudsen* 1972 and *Yechiali* 1972), generalized arrival processes (e.g., *Yechiali* 1971), state-dependent pricing (e.g., *Borgs et al.* 2014), and variable service rate (e.g., *Lippman and Stidham* 1977).

While these papers focused on customer joining behavior and the impact of prices, relatively few papers have studied the firm's capacity decision and the interaction of the capacity decision with the customers' joining behavior. *Grassmann* (1979) developed a numerical method to find the welfare-maximizing service rate assuming the customers' behavior is given regardless of the service rate. *De Vany* (1976), *Mendelson* (1985), *Dewan and Mendelson* (1990), *Stidham* (1992), and *Chen and Frank* (2004) studied the short-run pricing and the long-run capacity decisions together. However, this prior work on the capacity problem has not considered state-dependent joining decisions for the customers (i.e., customers do not observe the queue length upon arrival). The only exception is *De Vany* (1976), but he did not explicitly solve the optimization problem. While the assumption of an unobservable queue makes the resulting mathematical model tractable by making the customers' utility function continuous, it assumes customers (who often have bounded rationality) have the ability to correctly estimate the long-run behavior of the system. However, in most practical applications, customers base their decisions on the queue length at the time they arrive, and from this they can reasonably estimate their waiting time. Here we relax the restrictive assumption of earlier work, and study the interaction of the firm's capacity decision with the strategic and dynamic balking behavior of customers.

The remainder of this chapter is organized as follows. Section 2 outlines our queueing model. We study the customers' problem, for a fixed service rate, in section 3, and show that the customer strategy reduces to a threshold such that customers join if and only if the number of customers is less than the threshold. We also show how the threshold changes in the customers' selfishness. In section 4 we consider the firm's problem, given the customer type. The analysis is complicated because of the discrete (threshold) strategy of the customers. Thus, profits and social welfare are discontinuous functions of the service rate. We derive properties of these functions that allow a more complete analysis as well as efficient computation. In section 5, we put what we have learned in sections 3 and 4 together to compare the equilibrium outcomes for different scenarios. All proofs are provided in the Appendix. Our major findings are as follows.

- For a fixed service rate, collective customers (those who maximize their collective utility, as compared with those maximizing their own individual utility or those maximizing social welfare) are least likely to join the queue. Consequently, they minimize the provider's profits, and, of course, they maximize overall customer utility (Proposition II.2(i)).
- For fixed and intermediate service rates (neither very large or very small), individually selfish customers are most likely to join and they maximize the provider's profits and minimize their collective utility. On the other hand, for extreme service rates, social customers are most likely to join, and they maximize provider profits and minimize customer utility (Proposition II.2(ii)).
- When customers are coordinated to maximize their average utility, a profit maximizing firm always under-invests in capacity relative to the socially optimal capacity (Proposition II.11). On the other hand, when customers act to selfishly maximize their individual utility, increasing capacity does not nec-

essarily lead to increasing utility, and, depending on the parameters, a profit maximizing firm may either over-invest or under-invest in capacity relative to the social optimum. When customers are regulated to maximize social welfare, again a profit maximizing firm may either over-invest or under-invest in capacity relative to the social optimum, but now we can characterize the behavior in terms of price. In particular, there exists a cut-off price, below which the profit-maximizing firm under-invests in capacity and over-invests otherwise (Proposition II.13).

- With either a profit maximizing or social welfare maximizing firm, there may be a “*Benefit of Anarchy*” rather than a “*Price of Anarchy*.” That is, the utility of collective customers can be lower than when they are individually selfish (see Section 5.2).
- For a profit-maximizing firm, there can be a “*Social Benefit of Customers’ Selfishness*.” That is, social welfare can be larger for selfish customers (who maximize their individual or collective utility) than for customers who are regulated to maximize social welfare. So, for example, no regulation by a social planner (selfish customers and firm) can generate higher welfare than customer-only regulation (see Section 5.2).
- It is impossible to induce a profit-maximizing firm to choose the socially optimal service rate through price alone, even when customers can be regulated to choose the socially optimal threshold (Proposition II.15).

2.2 Model description

We consider a service system where a firm provides services to delay-sensitive customers. To model this, we consider an $M/M/1$ queue. Without loss of generality, we scale the Poisson arrival rate so that $\lambda = 1$. Upon arrival, a customer observes the

number of customers in the system, the service rate, and the price. She then decides whether or not to join the queue based on the specific objective she wants to maximize (which will be discussed later). We assume that all customers are risk-neutral and identical in terms of willingness-to-pay (i.e., in their benefits of receiving service) and in their disutility for waiting. We denote by R the reward that a customer draws from completing a service, and by p the service fee that a customer pays. If she joins the system, the service is FIFO (First-In-First-Out) and the customer incurs a disutility for waiting which we model as a linear holding cost, h , that is, a cost h per unit time is incurred until she completes the service. If she leaves without being served, there is no gain or loss, thus yielding zero utility.

We assume the firm chooses the service rate μ for the exponential service times at time 0, and the operating cost per unit time, $c(\mu)$, is assumed to be increasing and convex.

To measure the impact of customers' selfishness, we consider three different types of customers : (1) *Individual customers* make decisions that maximize their own individual utility. (2) *Collective customers*, as in the price of anarchy literature, make decisions that maximize the average customer utility. (3) *Social customers* make decisions that maximize the social welfare, which is the sum of average customer utility and the firm's profit. Likewise, to measure the impact of the firm's selfishness, we consider a firm that can be either a profit or a social-welfare maximizer: (1) *The selfish firm* chooses the service rate that maximizes its long-run average profit. (2) *The social firm* chooses the service rate that maximizes the social welfare, or utility plus profit.

In each scenario, we examine the equilibrium behavior of customers and the firm. By comparing equilibrium behavior, utility, profit, and social welfare among different scenarios, we analyze the effect of selfishness and answer the following questions: (i) How do customers behave differently depending on their objective? (ii) How does

	Individual Customers	Collective Customers	Social Customers
Selfish Firm	IP	CP	SP
Social Firm	IS	CS	SS

Table 2.1: Scenarios based on customers' and firm's objectives

the service rate of a firm depend on its objective and the type of customers? (*iii*) Do collective customers achieve a higher customer utility than individual customers? Do social customers generate a higher social welfare than other customer types?

Though the SP combination of social, welfare-maximizing, customers and a selfish, profit-maximizing, firm is unrealistic, it is interesting to note, as we show later, that social welfare can actually be higher under the traditional IP scenario, when both parties are selfish, than under the SP scenario. We also note that a social planner could induce a social optimum (social behavior) through regulation, thresholding/metering, and price.

2.3 The Customers' Problem

In this section, we study how customer behavior (whether to join the queue or not) depends on their selfishness for a given service rate.

Customers decide whether or not to join based on the number of customers in the system when they arrive. Following *Naor* (1969), the joining rule can be specified as a threshold strategy. That is, there exists a threshold n such that a customer joins the system if and only if she observes fewer than n customers in the system upon arrival, and otherwise leaves without being served. We will show that for each of the customer objectives we consider, the optimal decision rule for customers can be always characterized by a threshold, though the particular threshold value depends

on the particular objective, as we detail below.

2.3.1 Individual Customers

If customers are individually selfish, each arriving customer makes her decision based on maximizing her own utility. For a given service rate, μ , the expected utility that an arriving customer can draw when she joins the system with i customers is given by

$$U(i) = R - p - h \frac{i + 1}{\mu}. \quad (2.1)$$

Since the alternative, not joining, gives her zero utility, the individual customer joins the system as long as her expected utility is non-negative. To avoid any ambiguity, we assume that the customer joins the queue if she is indifferent. Then, individual customers join the queue if and only if the number of customers present is less than the *individual threshold*, $n^I(\mu)$, which is given by

$$n^I(\mu) = \max \left\{ n \in \mathbb{N}_0 \mid R - p - h \frac{(n - 1) + 1}{\mu} \geq 0 \right\} = \lfloor b\mu \rfloor, \quad (2.2)$$

where $b = \frac{R-p}{h}$ and $\lfloor \cdot \rfloor$ indicates the largest integer not exceeding the value inside the bracket.

2.3.2 Collective Customers

In the case of collective customers, a customer not only optimizes her own utility but also the utility of the other customers. That is, all customers follow the same decision rule that maximizes the average utility of customers. In our model, the average customer utility depends on how many customers receive service and how long they wait. Given the fact that the individual utility from service, given in (2.1), decreases in the queue length, it can be shown that collective customers always join the queue with i customers if joining the queue with $i + 1$ is optimal, i.e., the collective

customers' optimal pure strategy is also characterized by a threshold.

When the threshold is n , the resulting system is an $M/M/1/n$ queue in which the maximum number of customers in the system is bounded by n . Then the average customers' utility is :

$$\mathcal{U}(\mu, n) = (R - p)\lambda_e(\mu, n) - hL(\mu, n). \quad (2.3)$$

where $\lambda_e(\mu, n) = 1 - \frac{\mu-1}{\mu^{n+1}-1}$ is the effective arrival rate (for those who join), and $L(\mu, n) = \frac{1}{\mu-1} - \frac{n+1}{\mu^{n+1}-1}$ is the average number of customers in the system, $L(\mu, n)$.

Then, the optimal threshold for collective customers (which we call the *collective threshold*), $n^C(\mu)$, is the one that maximizes $\mathcal{U}(\mu, n)$. We can use known results from *Naor* (1969) to find the collective threshold. Letting $b = \frac{R-p}{h}$, we have

$$n^C(\mu) = \arg \max_n \left\{ b\lambda_e(\mu, n) - L(\mu, n) \right\} = \max \left\{ n \in \mathbb{N}_0 \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} \leq b \right\} \quad (2.4)$$

Notice that both the effective arrival rate and the average number of customers in the system are increasing in the threshold. That is, a larger threshold means more customers receive service (increasing the total benefit of service) and more customers in the system on average (increasing the waiting cost). Hence, the collective threshold, $n^C(\mu)$ which maximizes average customer utility, balances the trade-off between the throughput and the waiting cost.

2.3.3 Social Customers

Social customers make their decision to maximize the social welfare, which is the sum of the firm's profit and the average customer utility. Note that the operating cost of the firm is given, thus, independent from customers' decisions. Thus, social customers account for the firm's revenue in addition to the customer utility. In other words, if the price is zero, the objectives of social and collective customers are

identical. From this fact, we derive the optimal threshold for social customers (which we call the *social threshold*), which is the same as the collective threshold with $p = 0$, i.e., with $b_0 = \frac{R}{h}$ instead of b .

$$n^S(\mu) = \max \left\{ n \in \mathbb{N}_0 \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} \leq b_0 \right\}. \quad (2.5)$$

2.3.4 Threshold Analysis

The following result summarizes how the thresholds depend on μ .

Lemma II.1. *For all customer types, the optimal policies are characterized by thresholds, which are increasing in the service rate. That is, for customer type α , $\alpha \in \{I, C, S\}$, the threshold changes from $n-1$ to n at $\mu = \underline{\mu}^\alpha(n)$ for $n \in \mathbb{N}$:*

$$\underline{\mu}^I(n) = \min \{ \mu \mid n^I(\mu) = n \} = \frac{n}{b}, \quad (2.6)$$

$$\underline{\mu}^C(n) = \min \{ \mu \mid n^C(\mu) = n \} = \left\{ \mu > 0 \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} = b \right\}, \quad (2.7)$$

$$\underline{\mu}^S(n) = \min \{ \mu \mid n^S(\mu) = n \} = \left\{ \mu > 0 \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} = b_0 \right\}, \quad (2.8)$$

and $\underline{\mu}^\alpha(n)$ is increasing in n .

The intuition behind this lemma is straightforward. The faster the service is, the more utility a customer receives from completing service with the same number of customers upon her arrival to the queue. Thus, for individual and collective customers, if it is optimal to join the queue of length i at rate μ , then it is optimal to join the queue of length i at rate $\mu' > \mu$. Compared to collective customers, social customers additionally consider the firm's revenue, and therefore put more weight on the throughput than the collective customers, and again, if they join at rate μ , they will also join at rate $\mu' > \mu$. For values of μ between the threshold-change points characterized above, the threshold will remain constant: $n^\alpha(\mu) = n$ for $\alpha \in \{I, C, S\}$ and for any $\mu \in [\underline{\mu}^\alpha(n), \underline{\mu}^\alpha(n+1))$ (which we call a *threshold interval*).

The next question is how the thresholds compare across the three customer types. Proposition II.2 tells us that, at the same service rate, collective customers are least likely to join the queue. For intermediate service rates (neither very large nor very small), individual customers are most likely to join, while for extreme service rates, social customers are most likely to join.

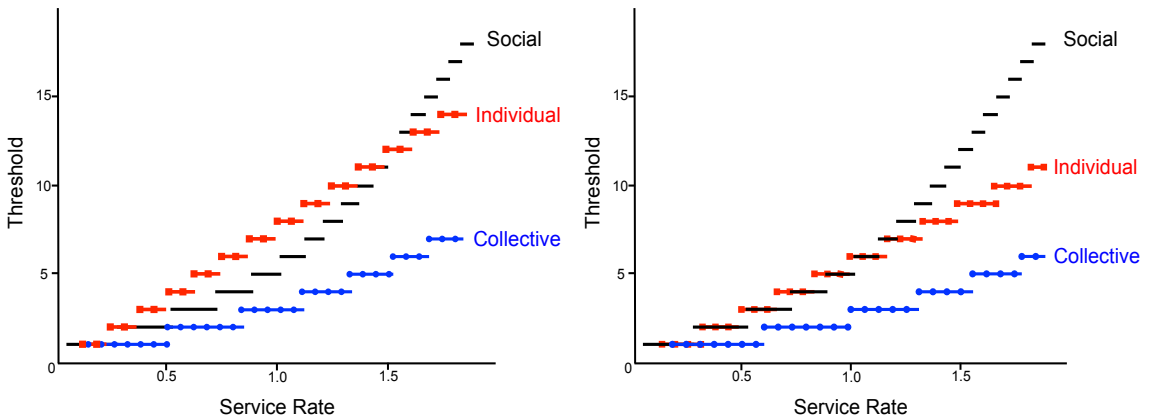
Proposition II.2. *For given service rate μ ,*

(i) *the collective threshold is the smallest, i.e., $n^C(\mu) \leq n^I(\mu)$ and $n^C(\mu) \leq n^S(\mu)$ for any $\mu > 0$.*

(ii) *there exist μ_L and μ_H such that $n^I(\mu) \geq n^S(\mu)$ if $\mu_L \leq \mu \leq \mu_H$, otherwise, $n^I(\mu) \leq n^S(\mu)$.*

(iii) *as the price increases, the individual and collective thresholds decrease while the social threshold remains the same. i.e., $\underline{\mu}^I(n)$ and $\underline{\mu}^C(n)$ increase in p while $\underline{\mu}^S(n)$ does not change.*

Figure 2.1: Threshold as a function of service rate at two different prices (p): 4 (left) and 6 (right) when $\lambda = 1$, $R = 10$ and $h = 0.5$.



Proposition II.2 (i) states that for the same service rate, collective customers increase their average utility by not joining the queue when the other types of customers join. Social customers are more likely to join than collective customers because they include the revenue for joining in their objective function, and individual customers

are also more likely to join than collective customers because they do not consider the negative externality that their joining imposes on other customers.

Proposition II.2 (ii) shows that, depending on μ , social customers can improve social welfare by either joining more or less than individual customers depending on which of the two – increasing revenue or decreasing negative externalities – improves the welfare the most. When the service rate is very low, i.e., $\mu < \mu_L$, the individual threshold is also very small, thus most individual customers do not receive service. In this case, the negative externality caused by an entering customer is small relative to the revenue it brings. When the service rate is very high, $\mu > \mu_H$, the negative externality is again small because very fast service reduces the marginal impact of an extra customer on future customers. Hence, for extreme service rates the revenue benefit dominates the negative externalities, and the social threshold is higher than the individual threshold. On the other hand, for intermediate values of μ , $\mu_L \leq \mu \leq \mu_H$, the negative externality is relatively higher than the boost in revenue. Thus, the social threshold is lower than the individual threshold.

Proposition II.2 (iii) further shows the effect of the service fee, p . First, the social threshold is independent of p , which is just a payment between the customers and the firm. However, individual and collective customers do not consider the firm's revenue, thus, with a higher price they are less likely to join. As a result, the range of μ where the individual threshold is higher than the social threshold shrinks as p increases (see Figure 2.1). If the price is sufficiently high, the individual threshold is always less than the social threshold. If the price is set to be zero, we know that social and collective customers are identical, and the individual threshold is higher than the social threshold.

Given the threshold properties from Proposition II.2, the profit, the customer utility, and the social welfare can be shown to have the following properties. To simplify the notation, in the following corollary we will omit μ (e.g., n^I for $n^I(\mu)$ and

$\Pi(n^I)$ for $\Pi(\mu, n^I(\mu))$, where Π represents the firm's profit function).

Corollary II.3. *For given service rate,*

(i) $\Pi(n)$ is increasing in n .

(ii) $\mathcal{U}(n)$ is increasing in n for $n < n^C$ and decreasing in n for $n \geq n^C$.

(iii) if the service rate is low or high, i.e., $\mu < \mu_L$ or $\mu > \mu_H$,

$\mathcal{W}(n^C) \leq \mathcal{W}(n^I) \leq \mathcal{W}(n^S)$.

(iv) if the service rate is intermediate, i.e., $\mu_L \leq \mu \leq \mu_H$,

$\mathcal{W}(n^C) \leq \mathcal{W}(n^S)$ and $\mathcal{W}(n^I) \leq \mathcal{W}(n^S)$.

First, the firm's profit is increasing in the effective arrival rate and therefore in the threshold. From Proposition II.2, the collective customers are least likely to join so they minimize profits among the three types.

Second, customer utility is decreasing in the threshold given that it is higher than the collective threshold. The collective threshold maximizes customer utility, which is discretely unimodal in the threshold. As more customers join the queue compared to collective customers, the excessive waiting cost diminishes the average customer utility.

As expected, the social welfare is greatest with social customers. But, interestingly, the social welfare with collective customers can be lower than that with individual customers. When this happens, the collective customers' effort to increase the average utility is dominated by the decrease in the firm's revenue (since they increase the average utility by not joining).

2.4 The Firm's Problem

We now consider the firm's problem. Knowing the customers' behavior for each service rate as given in Section 3, the firm chooses the service rate to optimize its objective. We start with the selfish firm that chooses the service rate to maximize

its own profit.

2.4.1 The Profit-Maximizing Firm

From the fact that the optimal strategy of customers is always characterized by a threshold, the firm's average profit when the service rate is μ is given by

$$\Pi(\mu, n^\alpha(\mu)) = p\lambda_e(\mu, n^\alpha(\mu)) - c(\mu) \quad \text{for } \alpha \in \{I, C, S\}.$$

Since the effective arrival rate, $\lambda_e(\mu, n^\alpha(\mu))$, is increasing in μ , both the revenue and the operating cost increase in μ . Thus, the optimal service rate should balance the trade-off between the revenue and the operating cost. However, as the customer threshold is a step function in μ (Lemma II.1), the profit function is not smooth in μ as illustrated in Figure 2.2.

Lemma II.4. *For each customer type $\alpha \in \{I, C, S\}$, the following properties hold for the firm's profit function.*

- (i) $\Pi(\mu, n^\alpha(\mu))$ is strictly concave in μ for μ within each threshold interval where the threshold remains the same, i.e., $\mu \in [\underline{\mu}^\alpha(n), \underline{\mu}^\alpha(n+1))$, for all n .
- (ii) $\Pi(\mu, n^\alpha(\mu))$ is discontinuous and jumps upward at $\underline{\mu}^\alpha(n)$ for $n \in \mathbb{N}$, i.e., $\lim_{\epsilon \rightarrow 0^+} \Pi(\underline{\mu}^\alpha(n) - \epsilon, n-1) < \Pi(\underline{\mu}^\alpha(n), n)$.

At $\underline{\mu}^\alpha(n)$, the profit function jumps upwards because the effective arrival rate increases when a threshold changes from $n-1$ to n . This discontinuity complicates the analysis of the capacity problem, which is in contrast to the price problem for prior literature. In the price-setting problem (Figure 2.3), the profit is also discontinuous in p , that is, as price increases, the customers' threshold decreases and consequently the revenue decreases. However, for values of p inducing the same threshold, the profit function is increasing in p , thus, the profit-maximizing p for given threshold n is always located at the right-end point. Given that only the right-end points of

Figure 2.2: Profit as a function of μ when $p = 5$, $\lambda = 1$, $R = 10$, $h = 1$, and $c(\mu) = 2\mu^3$.

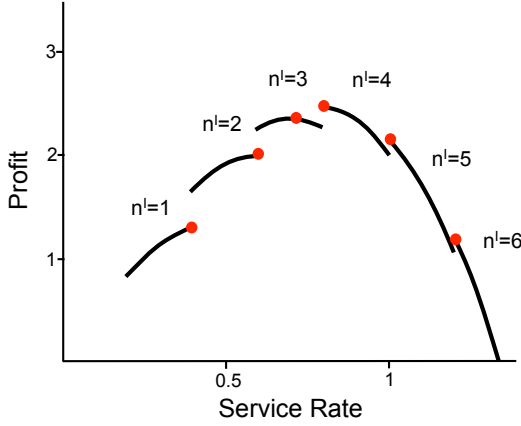
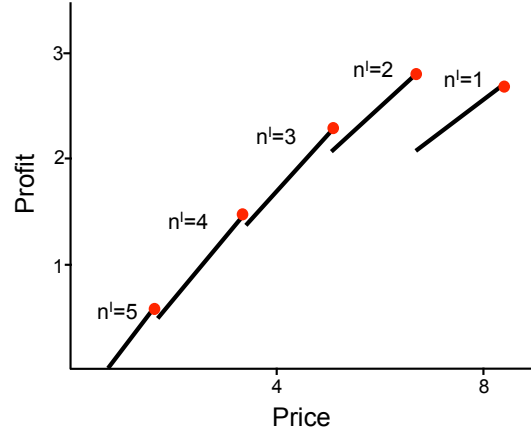


Figure 2.3: Profit as a function of p when $\mu = 0.6$, $\lambda = 1$, $R = 10$, $h = 1$, and $c(\mu) = 2\mu^3$.



each interval need to be considered, the pricing problem can be transformed into the threshold problem, i.e., what threshold the seller needs to induce in order to maximize the profit. On the other hand, in the capacity problem (Figure 2.2), the profit is not monotone within each interval and the optimal μ within this interval can be either the left-end point or an interior point. Thus, in order to solve for the profit-maximizing service rate, one may need to exhaustively search for the optimal μ 's for each threshold interval and then compare them.

Instead, we solve the problem more efficiently by utilizing the upper envelope function. For this, we first relax the constraint that the thresholds be integers and define relaxed thresholds as follows:

$$\tilde{n}^I(\mu) = b\mu, \quad (2.9)$$

$$\tilde{n}^C(\mu) = \left\{ n \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} = b \right\}, \text{ and} \quad (2.10)$$

$$\tilde{n}^S(\mu) = \left\{ n \mid \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{\mu^n(\mu-1)^2} = b_0 \right\}. \quad (2.11)$$

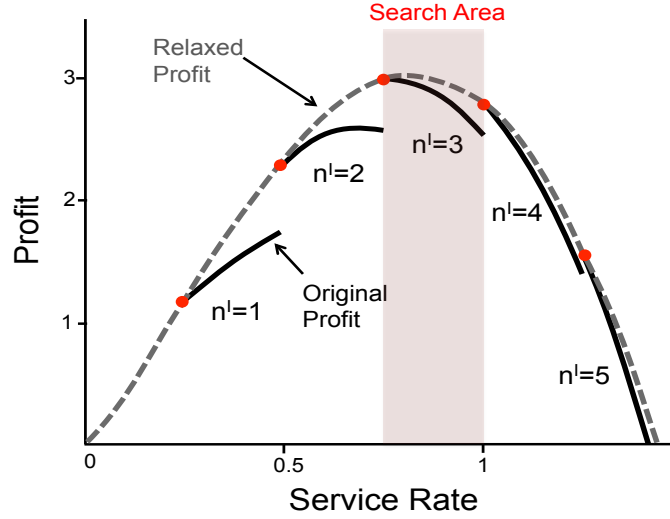
The profit function with the relaxed threshold (which we call the *relaxed profit func-*

tion), $\Pi(\mu, \tilde{n}^\alpha(\mu))$, is given by

$$\Pi(\mu, \tilde{n}^\alpha(\mu)) = p \left(1 - \frac{\mu - 1}{\mu^{\tilde{n}^\alpha(\mu)+1} - 1} \right) - c(\mu),$$

for $\alpha \in \{I, C, S\}$. The relaxed profit function is amenable to analysis and we will use it to find the true optimal service rate more efficiently. The following proposition gives its properties and an efficient means of determining the optimal service rate.

Figure 2.4: The original (solid line) and the relaxed (dashed line) profit functions when $\lambda = 1$, $R = 10$, $p = 6$, $h = 1$, and $c(\mu) = 2\mu^3$.



Lemma II.5. For each customer type $\alpha \in \{I, C, S\}$,

(i) $\Pi(\mu, \tilde{n}^\alpha(\mu))$, is differentiable and is an upper envelope of the profit function, that is, $\Pi(\mu, \tilde{n}^\alpha(\mu)) \geq \Pi(\mu, n^\alpha(\mu))$, where equality holds if and only if $\mu = \underline{\mu}^\alpha(n)$ for all n .

(ii) The optimal service rate, $\mu^{\alpha P}$, exists in an interval that contains a maximizer of the relaxed profit function, i.e., $\mu^{\alpha P} \in [\underline{\mu}^\alpha(n_0^\alpha), \underline{\mu}^\alpha(n_0^\alpha + 1)]$ for $n_0^\alpha = n^\alpha(\tilde{\mu})$ where $\tilde{\mu}$ satisfies the first order condition of $\Pi(\mu, \tilde{n}^\alpha(\mu))$.

By using the relaxed profit function, we can specify intervals that contain the profit-maximizing service rate through first order analysis. Then, Lemma II.5 (ii)

implies that we only need to search the intervals, $[\underline{\mu}^\alpha(n_0^\alpha), \underline{\mu}^\alpha(n_0^\alpha + 1)]$, that contain stationary points of the envelope function. For each stationary point, the corresponding interval can be divided into two parts: the interval $[\underline{\mu}^\alpha(n_0^\alpha), \underline{\mu}^\alpha(n_0^\alpha + 1))$ in which the threshold is n_0^α and the end point of $\underline{\mu}^\alpha(n_0^\alpha + 1)$ with threshold $n_0^\alpha + 1$. For the first sub-interval, the original profit function is concave, thus again we can find the maximum by first order analysis.

• **Individual Customers (System IP)**

When customers are individually selfish, the relaxed profit function has the following property.

Lemma II.6. *If customers are individually selfish and the following sufficient condition is satisfied,*

$$c \left(\frac{2}{b} \right) - c \left(\frac{1}{b} \right) \leq \frac{4bp}{(b+2)(b^2+2b+4)}, \quad (2.12)$$

then $n^{IP} > 1$ and the envelope function has a unique optimal solution which is contained in the interval that contains the optimal service rate for the original problem.

From Lemma II.6, we need to search only one interval, identified by the envelope method, if condition (2.12) holds. This is because the relaxed profit function is concave for values of μ such that the induced threshold is greater than 1, and the condition (2.12) implies that the operating cost is small enough relative to the revenue to yield an equilibrium threshold that is greater than 1, $n^{IP} > 1$.

The violation of condition (2.12) implies that the equilibrium threshold can be 1, which is a trivial case where no customer will join the system if a server is busy. Even when condition (2.12) is not satisfied, the envelope method is still efficient because the relaxed profit function is still unimodal for $\mu \geq \frac{2}{b}$. In this case, we need to search at most two intervals, $[\frac{1}{b}, \frac{2}{b}]$ and $[\frac{n_0^\alpha}{b}, \frac{n_0^\alpha+1}{b}]$, for the maximizer of the relaxed profit

function.

• **Collective or Social Customers (System CP and SP)**

When customers are collective or social, the concavity of the relaxed profit function cannot be shown analytically. One difficulty comes from the fact that a closed form for the threshold does not exist (see (2.10) and (2.11)) and we are not able to explicitly write the firm’s profit as a function of μ . Second, the collective and social thresholds follow a convex pattern (see Figure 2.1), which, in turn, can make the relaxed profit function convex. Therefore, analytically there may exist multiple stationary points of the relaxed profit function, so, in principle, we need to examine multiple intervals to find the profit-maximizer. However, in most of our numerical experiments, the profit is unimodal.

2.4.2 The Social Firm

We now consider the optimal service rate of a social firm, e.g., a government agency, that wants to maximize social welfare. The social welfare per unit time is given by

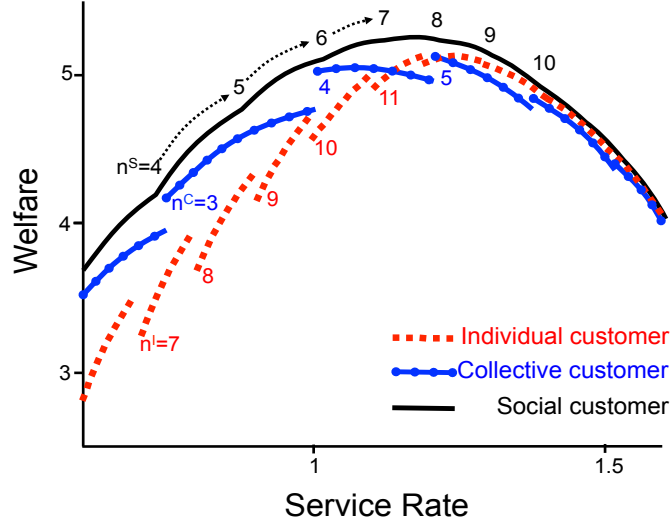
$$\mathcal{W}(\mu, n^\alpha(\mu)) = R\lambda_e(\mu, n^\alpha(\mu)) - hL(\mu, n^\alpha(\mu)) - c(\mu) \quad \text{for } \alpha \in \{I, C, S\}.$$

Now, in addition to the discontinuous effective arrival rate that we had in the profit function, we also have a discontinuous waiting cost, making the analysis even more difficult. The following lemma first characterizes the welfare function within each threshold interval (where the threshold remains the same).

Lemma II.7. *For all customer types, $\alpha \in \{I, C, S\}$,*

(i) $\mathcal{W}(\mu, n^\alpha(\mu))$ is differentiable in μ within each threshold interval, and its slope is strictly greater than the slope of the profit function.

Figure 2.5: Social welfare as a function of the service rate when $\lambda = 1$, $R = 10$, $h = 0.5$, $p = 5$, and $c(\mu) = 2\mu^2$.



(ii) For $\mu > 1$, $\mathcal{W}(\mu, n^\alpha(\mu))$ is concave in μ for μ within each threshold interval.

The first differentiability property is immediate from the differentiability of the effective arrival rate and the average number of customers for a fixed threshold. As long as the threshold remains the same, customer utility, $\mathcal{U}(\mu, n)$, can be shown to be increasing in μ since fast service reduces congestion and, hence, waiting costs. Because the social welfare is the sum of profit and customer utility, this implies that the welfare increases faster in μ than the profit for any given threshold.

The second concavity result follows from the fact that the average number of customers in the system is convex in the service rate μ if the service rate is larger than the arrival rate ($\lambda = 1$). However, when the service rate is below $\lambda = 1$, the average number of customers can be concave in μ . This, in turn, implies that the welfare function within a threshold interval is not necessarily concave if $\mu < 1$. This makes the determination of the optimal service rate difficult.

We now explore how the welfare function (and the optimal service rate) depends on the customer type.

• **Individual Customers (System IS)**

From Lemma II.1, the threshold for individual customers increases by 1 in μ at multiples of $\frac{1}{b}$. When the threshold increases, revenue increases but waiting cost also increases, so social welfare may jump either upwards or downwards. The following lemma characterizes the welfare function with individual customers.

Lemma II.8. *The social welfare with individual customers, $W(\mu, n^I(\mu))$, is discontinuous in μ at the points where the individual threshold changes, i.e., $\mu = \underline{\mu}^I(n)$ for $n \in \mathbb{N}$. Furthermore, there exist μ_L and μ_H , such that if μ is low or high ($\mu < \mu_L$ or $\mu > \mu_H$), the welfare jumps upward when the threshold increases, i.e., $\lim_{\epsilon \rightarrow 0^+} \mathcal{W}(\underline{\mu}^I(n) - \epsilon, n - 1) \leq \mathcal{W}(\underline{\mu}^I(n), n)$. If μ is moderate ($\mu_L \leq \mu \leq \mu_H$), the welfare jumps downwards, i.e., $\lim_{\epsilon \rightarrow 0^+} \mathcal{W}(\underline{\mu}^I(n) - \epsilon, n - 1) \geq \mathcal{W}(\underline{\mu}^I(n), n)$.*

To see the intuition, recall from Proposition II.2 that $n^I(\mu) \leq n^S(\mu)$ if the service rate is either low or high. This implies that if the individual threshold increases (and gets closer to the social threshold), the welfare can be improved. On the other hand, for a moderate value of service rate, when $n^I(\mu) \geq n^S(\mu)$, the negative externality is significant in the individual customers' case. If the threshold increases, the negative externality increases, decreasing social welfare.

Due to the irregular discontinuities, we need to derive all the local maxima for each interval, $[\underline{\mu}^I(n), \underline{\mu}^I(n + 1))$ for $n \in \mathbb{N}$, and compare them to find the socially optimal service rate. Note that the social welfare is negative if $c(\mu) > R$, so we need only consider $\mu < \mu_0$ where μ_0 is such that $c(\mu_0) = R$. Then, we need to find the best service rate within each threshold interval. After comparing all the best service rates within each interval, we determine the optimal service rate of a social firm with individual customers.

• **Collective Customers (System CS)**

The social welfare function with collective customers has the following properties.

Lemma II.9. *The welfare with collective customers, $\mathcal{W}(\mu, n^C(\mu))$, is discontinuous in μ at the points where the collective threshold changes, i.e., $\mu = \underline{\mu}^C(n)$ for $n \in \mathbb{N}$. Furthermore, the welfare jumps upward, i.e., $\lim_{\epsilon \rightarrow 0^+} \mathcal{W}(\underline{\mu}^C(n) - \epsilon, n - 1) < \mathcal{W}(\underline{\mu}^C(n), n)$.*

We only have upward jumps in social welfare because the collective customers' utility is continuous while the profit jumps upward at the points where increasing μ increases the threshold (and hence the effective arrival rate). As in the individual customers' case, we can derive the socially optimal service rate by comparing the local maxima for all of the threshold intervals.

- **Social Customers (System SS)**

The social welfare, when both customers and firm are coordinated to maximize the welfare together, has the following properties.

Lemma II.10. *The social welfare function with social customers $\mathcal{W}(\mu, n^S(\mu))$ is continuous but not differentiable at the points where the social threshold changes, i.e., $\mu = \underline{\mu}^S(n)$ for $n \in \mathbb{N}$.*

Since the social customers change their threshold if and only if it improves the social welfare, the social welfare function is continuous even for the points at which the threshold changes. In particular, if it were discontinuous, say $\mathcal{W}(\mu, n^S(\mu)) < \mathcal{W}(\mu + \epsilon, n^S(\mu) + 1)$ for small enough ϵ , the social welfare at μ could be improved by customers choosing $n^S(\mu) + 1$ instead of $n^S(\mu)$, which contradicts the fact that $n^S(\mu)$ is socially optimal at μ .

Now we can efficiently search for the socially optimal service rate when customers are also social. Since both the firm and customers maximize the same objective, we

can exchange the order of optimizations as follows,

$$\max_{\mu} \mathcal{W}(\mu, n^S(\mu)) = \max_{\mu} \max_n \mathcal{W}(\mu, n) = \max_n \max_{\mu} \mathcal{W}(\mu, n). \quad (2.13)$$

This property can be used to find an optimal service rate more efficiently, as we do not need to find the boundaries of each threshold interval. First, for each n , we find the welfare maximizer over the entire domains given n fixed. Since the welfare function with n fixed is differentiable, it can be found by first- and second-order analysis. Second, by comparing the optimal welfare for each n , we can determine which service rate is the welfare maximizer.

2.5 Analysis

In this section, we study how the equilibrium outcomes change depending on the objectives of the firm and the customers. In particular, we compare the equilibrium service rate and threshold, and examine how the welfare changes as the firm and customers become selfish. We call this the social effect (cost) of selfishness as selfish behavior lowers welfare compared to the all social case (system SS). We also study the effects of selfishness on the profit and the utility components of social welfare.

2.5.1 On the Firm's Selfishness

We first compare the equilibrium outcomes under a profit-maximizing and a social firm for given customer type. This will help us understand the efficiency loss when public service becomes privatized (i.e., service by a government agency vs. service by a private firm).

- **Collective Customers (System CP vs. CS)**

We start with the easiest case, which is when customers are coordinated to maximize

their joint utility, i.e., they are collective customers. From the fact that the social firm generates a higher welfare while the selfish firm earns a higher profit, we know that customer utility is always higher with the social firm, i.e., the social firm generates greater welfare by enhancing customer utility compared to the selfish firm. The following proposition shows how the social firm's service rate differs from the selfish firm's.

Proposition II.11. *With collective customers, the social firm's service rate is at least as large as the selfish firm's, i.e., $\mu^{CP} \leq \mu^{CS}$. As a consequence, the equilibrium threshold with the social firm is also at least as large as with the selfish firm, $n^{CP} \leq n^{CS}$, and hence, more customers are served by the social firm.*

When customers maximize their joint utility, they always achieve a higher utility when they are served faster. That is, the collective customers' utility is strictly increasing in service rate. This implies that the social firm sets a higher service rate than the selfish firm, and the equilibrium threshold is also higher. In other words, the efficiency loss (in terms of a lower welfare) from privatization arises from the selfish firm's underinvestment in capacity.

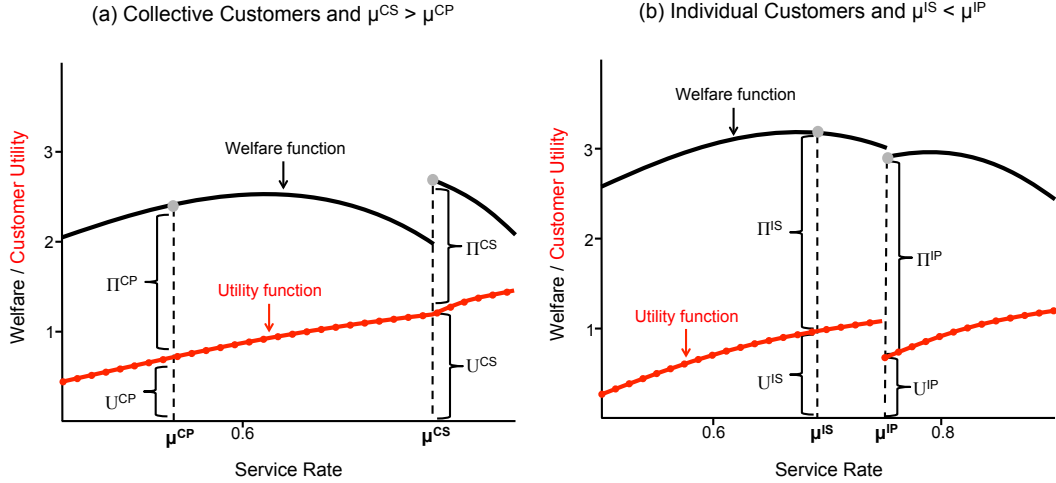
• **Individual Customers (System IP vs. IS)**

Now suppose customers are selfish, i.e., they maximize their individual utility. As with collective customers, one might expect the social firm to set a higher service rate than the selfish firm because it considers average customer utility. However, we find that this is not always the case.

Proposition II.12. *With individual customers, we can have $\mu^{IP} > \mu^{IS}$ or $\mu^{IP} < \mu^{IS}$, depending on the parameters.*

To gain some intuition as to why the social firm might choose a smaller capacity than the selfish firm, first note that, in contrast to collective customers (Proposition

Figure 2.6: The service rate choice depending on the firm’s objective function and the type of customers. With collective customers (left), the customer utility increases in μ and the social firm always chooses a higher service rate than the selfish firm. However, the individual customer utility is not monotone and in this example, the social firm chooses a lower service rate than the selfish firm. The parameters in both figures are: $R = 10$; $p = 6$; $h = 1$; $c(\mu) = 4.5\mu^4$.



II.11), individual customers are not necessarily better off with a higher service rate (a higher threshold), because of their negative externalities. That is, customer utility increases in service rate only within the threshold interval of service rate. Combining this with the fact that the profit might be increasing or decreasing even within the interval, the selfish firm may want to induce a higher threshold (and more revenue) by increasing μ beyond the interval, whereas the social firm may not because of the increased negative externalities and decreased customer utility. Figure 2.6 shows an example where the social firm sets a lower service rate so that the individual threshold is smaller, alleviating excess congestion. Note that the individual customers’ utility decreases when the threshold increases.

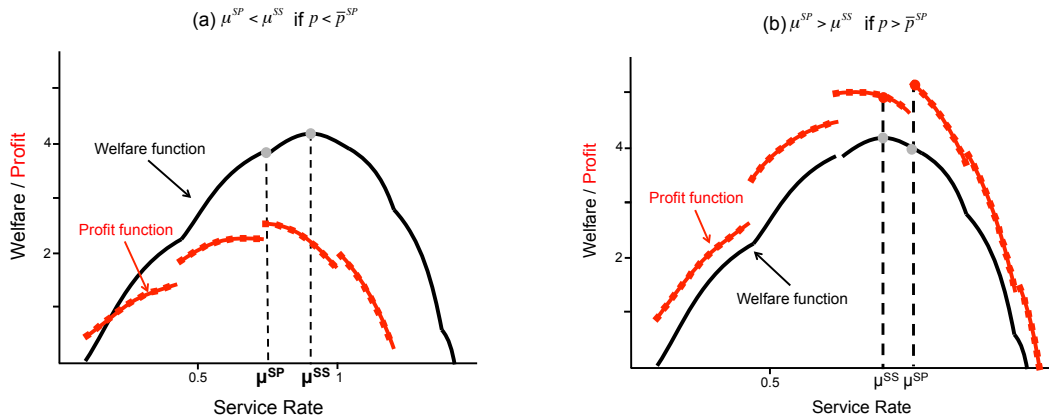
- **Social Customers (System SP vs. SS)**

Lastly, we consider social customers, i.e., those that are regulated by the social planner to maximize welfare. As with individual customers, the utility of social customers decreases when their regulated threshold increases with a higher service rate because

their objective function includes the firm's revenues, so negative externalities for the customers are compensated for by increased profits for the firm. Thus, we can have $\mu^{SP} > \mu^{SS}$ or $\mu^{SP} < \mu^{SS}$ depending on the parameters, as with individual customers (Proposition II.12). However, with social customers we are able to give a complete characterization of their behavior as a function of price, as follows.

Proposition II.13. *With social customers, the selfish firm's service rate is an increasing function of price, while the welfare-maximizing service rate (μ^{SS}) is constant. Hence, there exists a cut-off price (denoted by \bar{p}^{SP}), such that $\mu^{SP} \geq \mu^{SS}$ and $n^{SP} \geq n^{SS}$ if and only if $p > \bar{p}^{SP}$.*

Figure 2.7: Comparison of the profit-maximizing and the welfare-maximizing service rates when customers are social at two different prices (p): 5 (left) and 9 (right). The other parameters in both figures are: $R = 10$; $h = 1$; $c(\mu) = 2\mu^4$. Note that the profit can be higher than welfare because social customers' utility can be negative.



Proposition II.13 implies that the selfish firm sets a higher service rate than the social firm only when the price is high. First, note that μ^{SS} and n^{SS} do not depend on the price because price is just a transfer cost between customers and the firm, and is not included in social welfare. On the other hand, a high price (thus a high margin) implies that the selfish firm can collect high revenue, which encourages the selfish firm to invest in capacity to increase the effective arrival rate. When price is low ($p < \bar{p}^{SP}$), the selfish firm under-invests compared to the socially optimal level

(i.e., $\mu^{SP} < \mu^{SS}$), because the operating cost is relatively higher than the potential revenue.


2.5.2 On Customers' Selfishness and the Price of Anarchy

We now compare the equilibrium outcomes among individual, collective, and social customers holding the firm type fixed. This will help us understand the effect of customers' selfishness (or, equivalently, of customer coordination to maximize their utility or their regulation to maximize welfare) on their utility and overall welfare.

In contrast to the firm's selfishness, customers in our model choose their threshold depending on the service rate chosen by the firm. This means that the effects of customer type on their utility and welfare depend on how the firm sets its service rate for each customer type. We found that the interaction is complicated, and there is no general comparative rule, except for the trivial observation that for the social firm, social welfare is largest when customers are also social. However, we found some surprising behavior in some circumstances.

First, the utility of collective customers can be the lowest among all customer types. Because collective customers are the least likely to join, the (selfish or social) firm may set a significantly lower service rate compared to that for the other customer types and this can cause the collective customers' utility to be lower than for other customer types. This is in contrast to the PoA literature (*Roughgarden* 2005, *Haviv and Roughgarden* 2007, and *Gilboa-Freedman et al.* 2014). PoA is defined as the ratio of selfish customer utility to collective customer utility, i.e., $\frac{U(\mu, n^I(\mu))}{U(\mu, n^C(\mu))}$. Without considering the server's decision (i.e., μ is given), PoA must be less than 1. However, our model incorporates the server's decision, and in this case, PoA can be greater than 1. See Figure 2.8 for an example where the selfish firm under-serves customers and PoA is 1.8169. Summarizing, we can have a “*Benefit of Anarchy*” rather than a “*Price of Anarchy*” when the server under-invests in capacity in response to the

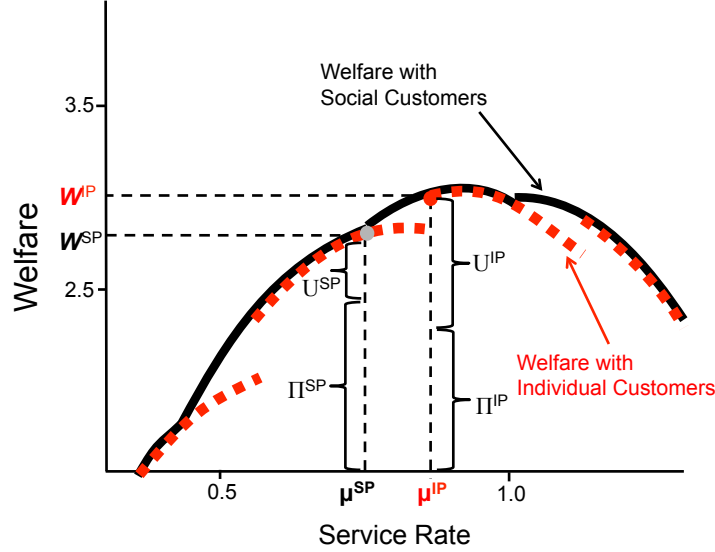
Figure 2.8: *Benefit of Anarchy* when $R = 10$, $h = 1$, $p = 5.2$, and $c(\mu) = 3\mu$: The selfish firm sets a lower service rate for collective customers than for individual customers ($\mu^{CP} = 0.72$ and $\mu^{IP} = 1.05$). Consequently, the utility of collective customers is lower than the utility of individual customers ($U^{CP} = 0.71$ and $U^{IP} = 1.29$), that is, the PoA = $\frac{1.29}{0.71} = 1.8169$, so individual customers have an 80% higher utility relative to collective customers.



demand coordination.

Second, the welfare with social customers can be lower than with the other customer types when the firm is selfish. This implies that the social planner, who wants to maximize social welfare, can be worse off using demand-only regulation (by enforcing thresholds) compared to no regulation at all. Of course, the full coordination of both demand and supply generates the first-best welfare of the system, i.e., \mathcal{W}^{SS} is the maximum possible welfare. And as we observed in Section 5.1, supply regulation (i.e., the social firm) allows the system to achieve higher welfare, i.e., $\mathcal{W}^{\alpha S} \geq \mathcal{W}^{\alpha P}$ for $\alpha \in \{I, C, P\}$. However, demand-only regulation means that social customers may be more likely to join in order to enhance the firm's revenue (thus, improving social welfare) for any given service rate. The selfish firm can take advantage of this behavior and set a lower service rate. Thus, we can have a “*Social Benefit of Customers' Selfishness*” when the firm is also selfish.

Figure 2.9: *Social Benefit of Customer Selfishness* with a selfish firm when $R = 10$, $h = 1$, $p = 6.5$, and $c(\mu) = 3\mu^2$. The selfish firm sets a lower service rate for social customers than for individual customers ($\mu^{IP} = 0.75$ and $\mu^{SP} = 0.86$). As a result, the equilibrium welfare is even lower when customers are social than when selfish ($\mathcal{W}^{IP} = 3.01$ and $\mathcal{W}^{SP} = 2.80$).



2.5.3 Price As a Tool to Achieve the Social Optimum

Of course if either the customers' threshold or the firm's service rate is not regulated to maximize social welfare, the system generates a sub-optimal equilibrium. This social cost of selfish behavior arises because the selfish agent (customers or the firm) maximizes a partial objective (utility or profit). Because the price is a transfer cost between customers and the firm, it only affects these partial objectives, not the overall welfare. Thus, choosing a proper price (which until now has been assumed to be exogeneously given in our model) may compensate for selfishness.

- **Selfish Customers**

First consider the system in which a social planner can regulate the firm, but can only affect customer behavior through the price (i.e., individual/collective customers with the social firm). Then, if customers choose the socially optimal threshold given service rate μ^{SS} , the equilibrium outcome will be the same as the first-best outcome.

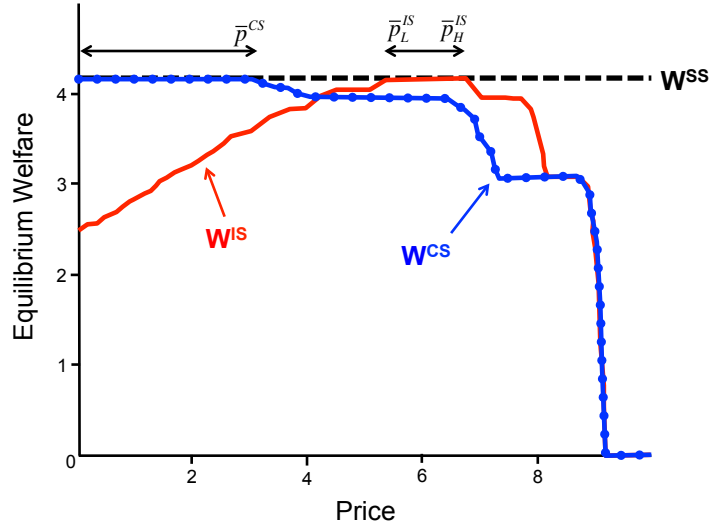
The following proposition shows that there exist intervals for the price such that the first-best outcomes can be obtained with selfish customers.

Proposition II.14. *With the social firm, there exists a price range such that the equilibrium outcome with (individually or collectively) selfish customers equals the first-best equilibrium outcome:*

(i) *For collective customers, there exists a cut-off price \bar{p}^{CS} such that for $p \leq \bar{p}^{CS}$, $\mu^{CS} = \mu^{SS}$, $n^{CS} = n^{SS}$, and $\mathcal{W}^{CS} = \mathcal{W}^{SS}$. Furthermore, the equilibrium welfare decreases in price for $p > \bar{p}^{CS}$.*

(ii) *For individual customers, there exists an interval $(\bar{p}_L^{IS}, \bar{p}_H^{IS})$ such that for $p \in (\bar{p}_L^{IS}, \bar{p}_H^{IS})$, $\mu^{IS} = \mu^{SS}$, $n^{IS} = n^{SS}$, and $\mathcal{W}^{IS} = \mathcal{W}^{SS}$.*

Figure 2.10: Equilibrium Welfare as a function of price when $R = 10$, $h = 1$, $c(\mu) = 2\mu^4$. To achieve the first-best equilibrium, the price should be less than $\bar{p}^{CS} = 2.83$ for collective customers, and between $\bar{p}_L^{IS} = 5.56$ and $\bar{p}_H^{IS} = 6.66$ for individual customers.



Proposition II.14 (i) implies that a social planner, who controls the firm's service rate, can achieve the first-best outcomes with collective customers only if the price is sufficiently low. Collective customers consider the overall customer utility, but act without considering the firm's revenue. However, recall from Section 3 that for a fixed service rate, the social threshold is the special case of the collective threshold

when $p = 0$. That is, if the service rate is fixed at μ^{SS} , $n^C(\mu^{SS}) = n^S(\mu^{SS})$ when $p = 0$. Since the threshold is an integer, the price can be increased from 0, holding μ^{SS} fixed, without affecting social welfare as long as the threshold doesn't increase from $n^S(\mu^{SS})$, giving an upper bound on price, $p \leq \bar{p}^{CS}$. If the price is so high as to decrease the effective arrival rate at μ^{SS} , the service rate has to be adjusted (either upwards to increase the threshold or downwards to save on operating cost), resulting in suboptimal welfare. Since the collective threshold is always lower than the social threshold, increasing the price only makes the gap larger, thus, the resultant equilibrium welfare decreases in price.

Proposition II.14 (ii) shows that there exists a congestion toll (first suggested by Naor, 1969) that induces the socially optimal threshold even with individually selfish customers. The price must be high enough to make them behave as if they considered their negative externality. On the other hand, as with collective customers, it should not be so high as to decrease the effective arrival rate above the socially optimal level.

• **Selfish Firm**

Now suppose that a social planner can only regulate the customers' behavior while the firm is selfish. In particular, suppose the social planner can set the threshold to $n^S(\mu)$ for any μ , and tries to set the price to induce a selfish firm to choose the socially optimal service rate, μ^{SS} . From Proposition II.13, we observed that the selfish firm's service rate increases in price when customers are social. This is because increasing the price increases the slope of the profit function, while the welfare function (its slope and threshold intervals) does not change in price. So, one might expect that there exists a price such that $\mu^{SP} = \mu^{SS}$. However, this is not the case. For any given threshold interval (where the social threshold remains the same), the welfare function is strictly steeper than the profit function in μ because the difference (customer utility) is strictly increasing in μ for any given n (see Lemma II.7 (i)). Therefore, the first-order conditions for μ^{SP} and μ^{SS} are always different, and the service rates that

maximize welfare and profit must be different. Thus we have the following.

Proposition II.15. *A social planner cannot regulate the selfish firm through price.*

2.6 Conclusion

We consider a queueing system in which a firm sets the service rate and each customer decides whether to join or not depending on the expected waiting cost. We particularly study how the performance (i.e., customer utility, firm profit, and overall social welfare) depends on the selfishness in the objectives that customers and the firm use to make decisions.

(i) For a given service rate, collective customers are least likely to join, and have lower waiting cost. (ii) In addition, individually selfish customers maximize the provider's profits and minimize their collective utility when service rate is moderate (neither very large nor very small). On the other hand, for extreme service rates, social customers maximize profit and minimize customer utility. (iii) We also show that, as long as the threshold remains the same, the social firm always provides service at a higher level than the selfish firm. These results are intuitive as we focus on either the customers' decision or the firm's decision by fixing the action of the counter-party.

However, when the interplay between the customers and the firm is considered (i.e., choosing the best response), the equilibrium outcome exhibits complicated and counter-intuitive behavior.

(i) *Benefit of Anarchy*: In contrast to the case with a fixed service rate in which individual customers always draw lower utility than collective customers (“*Price of Anarchy*”), the utility of collective customers can be lower than that of selfish customers when either a selfish firm or a social firm chooses the service rate. This happens particularly when the (selfish or social) firm sets a lower service rate for

collective customers because of their lower joining rate.

(ii) Social Benefit of Customers' Selfishness: We also find that either individual or collective customers can generate greater overall social welfare than social customers when the firm sets the service rate to maximize its profit.

(iii) Social Benefit of Lower Capacity: When the firm chooses the service rate, it changes the threshold and effective arrival rate. Thus, the result that the social firm provides a higher service rate is no longer true when there is an interplay between customers and the firm. In fact, when customers are either individual or social, customer utility decreases whenever the threshold increases as a response to a higher service rate. We find that the social firm can choose a lower service rate than the selfish firm and increase utility by reducing congestion.

(iv) Price control is insufficient to regulate the system: In the queueing literature, it is well known that a price (i.e., a congestion toll) is a mechanism that induces otherwise selfish customers to behave in a socially efficient way (*Naor 1969*). We show that, in contrast, it is impossible to induce a profit-maximizing firm to choose the socially optimal service rate through price alone, even when customers can be regulated to choose the socially optimal threshold.

CHAPTER III

Dynamic Pricing and Loyalty Programs

3.1 Introduction

Loyalty or reward programs are widely used in a number of industries including hospitality, financial service, transportation, and retailing. Many loyalty programs offer points (e.g. Delta's Skymiles, Hilton's HHonor points, and Member's reward points by American express) to consumers for their purchases of products or services, which can be later redeemed to buy perks, products, or services. Therefore, loyalty points are considered as an alternative form of currency. The option of purchasing with points has significant impacts both on consumers' purchase behavior and the seller's revenue. Indeed, both membership size and activities of many loyalty programs are astronomical. Between 2000 and 2014, loyalty program memberships in the U.S. tripled from 1.1 to 3.3 billions, that is an average of 10 memberships per person or 29 per household (*Colloquy* 2015). With the growth in membership, a significant number of points are issued and redeemed. In each year, newly issued points in the U.S. have a monetary value of \$48 billion (*Gordon and Hlavinka* 2011). Frequent flyer miles are considered as one of the world's most valuable currency, with an estimated 14 trillion miles worth more than \$700 billion (*Economist* 2005a). Marriott's program alone had more than 40 million loyalty members with 225 billion points issued to its members in 2013 (*Marriott* 2013). A recent survey further shows that consumers ac-

tively manage and redeem these loyalty points, and these redemptions account for a substantial portion of goods and services purchased (*NerdWallet* 2015). Passengers of Southwest Airline, for instance, redeemed 6.2 million award flights, which represented 11% of revenue passenger miles flown in 2014, up from 9% in 2012. Hilton HHonor program awarded 5.4 million reward nights and more than 95,800 items through the HHonors Global Shopping Mall (*HiltonWorldwide* 2015a).

Depending on how the points are issued and redeemed, the impact of point redemption to the seller is quite different. In a stand-alone loyalty program, such as restaurant loyalty programs and airlines, the seller is the point issuer and, at the same time, the provider for goods or services. In this case, the seller bears the cost of redeemed goods or services, making a redemption a cost activity without revenue. Since it is a pure cost activity, its implication on the revenue (at least from a short-run perspective) is quite straight-forward. It should be noted that, due to its magnitude and monetary value (not to mention, its value in marketing and consumer loyalty), accumulated points are important in the firm's accounting statement. In accounting, accumulated points are treated as a liability either through deferred revenue or incremental cost (*Chapple et al.* 2010).

On the other hand, there are many loyalty programs where point issuers are different from firms that actually provide a service or product. For instance, about 85 percent of Hilton's 4300 properties are franchised and operated by property management companies (*HiltonWorldwide* 2015b). In such case, when a consumer uses accumulated points to stay in the franchisee property, the property management company (franchisee) collects compensatory revenue from a hotel headquarter who manages the loyalty program. When consumers use their credit-card points to buy an airline ticket or to buy frequent mileages from an airline (transfer points), the seller (the airline) gets compensatory revenue from the credit card company who is the original point issuer. In fact, this type of transactions (called by *the sales of miles*) is

the largest contributor of airlines' ancillary revenue. For instance, US major airlines generate \$10.7 billion of ancillary revenue through frequent flier program (*IdeaWorks* 2015).

In this paper, we study the implications of a loyalty program on the seller's pricing and inventory rationing policies regarding reward sales. In particular, we focus on a situation where the revenue implication of loyalty program is less obvious, that is, when the seller gets reimbursed for a reward sale from a point issuer. It should be noted that, even among loyalty programs where sellers are different from issuers, the terms and conditions of reward sales that apply to sellers and consumers vary significantly within and across industry. In some hotel loyalty programs, point requirements for a reward stay are often fixed and determined by the hotel headquarter. Consequently, an individual seller (e.g., a property management company) cannot easily change these policies, and have a little control over reward sales. For instance, one night stay at Hyatt Regency San Francisco is fixed to 15,000 points, which contrasts to the cash price that is dynamically changing from \$209 to \$524 (*BoardingArea* 2015). In some cases, the issuer prevents even the use of operational tools such as rationing or black-out (e.g., Hilton properties with *no black-out* policy for standard rooms). On the other hand, even in hotel reward stays, the seller can change point requirement quite frequently depending on availability and posted price for suites. For instance, one night stay of *grand king corner suite* at Conrad Seoul (a premier Hilton property) on June 3, 2016 can be booked using either 435,000 Won or 102,542 HHonors points. However, one night stay of the same room on June 5, 2016 requires either 450,000 Won or 106,078 HHonors points (about 4.24 Won per 1 Hhonor point in both cases).

Given these various settings of loyalty program, our central question is how the terms of reward sales (e.g., reimbursement rate, operational discretion over reward sales) affect the seller's pricing and rationing decisions. Specifically, (*i*) how does the

reward sale affect the seller’s posted price?, (ii) if the seller can dynamically ration inventory for reward sales, when should the seller block reward sales, and when should not?, and (iii) if the seller can also set point requirement, how does the seller change the point requirement along with cash price?

In order to address our research questions, we first model a consumer’s decision: how he/she compares purchase options and decides whether to buy or not, and, if buying, how to pay for the product (either in cash or with loyalty points). We reflect the fact that consumers are heterogeneous in their willingness to pay, point balance, and perceived value of a point (i.e., how much does 1 loyalty point worth to a consumer?). In particular, we characterize a consumer’s decision as a function of price (cash price and point requirement) and rationing decisions. We then incorporate this into a dynamic pricing (and rationing) model for the seller who wants to maximize the total expected revenue – the sum of the revenue from cash sales and the revenue from reward sales. To account for the fact that sellers have different levels of discretion over reward sales, we consider several settings in which the level of operational discretion over reward sales differs. We first examine the setting where the seller can only change cash price and has no discretion over the terms of reward sales. In this case, no rationing (e.g., black out) is allowed, and both point requirement and reimbursement are fixed. We then consider the case in which the seller can dynamically ration (*black out*) inventory for reward sales in addition to choosing the price. Lastly, we generalize our analysis to the case where the seller can change both price and point requirement, and the reimbursement rate depends on the point requirement.

3.1.1 Summary of Contributions

Our paper is the first paper to examine the dynamic pricing problem of a seller whose product can be purchased with cash or loyalty points. Specifically, we model how consumers decide whether to buy or not, and if buying, whether to pay in cash

or points depending on their reservation price, point balance, and perceived value of a point. Incorporating this consumer model into the seller's problem, we show that how the presence of loyalty program influences the seller's pricing and inventory rationing decision.

We first show that a consumer will view two different purchase options (cash vs. point) as imperfect substitutes. As a result, as cash price (or point requirement) increases, some consumers switch from paying cash to redeeming points (or *vice versa*), but the overall demand decreases (as a result of being imperfect substitutes). We find that the impact of reward sales on the seller's price is non-trivial. One may think that, allowing reward purchase will increase the marginal value of the inventory (since the seller can sell a product in more ways), and, as a result, increase the price. We find that this is not the case: In fact, the seller's price can be either higher or lower than the seller's price when a reward sale is not available (we call it *cash-only price*) depending on the gap between the reimbursement rate (which is the revenue a seller earns from a reward sale) and the *cash-only price*. For instance, when the reimbursement rate is low relative to the *cash-only price*, the seller offers a discount in order to induce more cash sales. In the opposite case, the seller increases cash price (adds a premium) in order to induce more consumers to buy with points while charging a higher price to cash consumers.

We examine how discretion over the terms of a reward sale affects the seller's decision and resultant revenue. We find that it is optimal to block out a reward sale (i.e., disallow point redemption) when inventory level (relative to the remaining selling season) or a reimbursement rate is very low. In both cases, the *cash-only price* will be significantly higher than a reimbursement rate. Thus, allowing a reward sale only decreases the seller's revenue as some consumers choose to buy with points instead of cash. To avoid this, the seller blocks the reward sale completely and sell products by cash only. However, unless the gap between the *cash-only price* and

the reimbursement rate is significant, allowing reward sales is generally better for the seller as the total sales increase. That is, even when the seller’s revenue from a reward sale (the reimbursement rate) is lower than the *cash-only price*, allowing a reward sale can be still optimal as long as the benefit from increasing the likelihood of a sale is significant. We also analyze the case where the seller can dynamically change both price (for cash sales) and point requirement (for reward sales), and characterize the optimal policy. We find that further segmenting reward sales with multiple tiers of point requirements only marginally increases the revenue from the seller with the option to black-out reward sales.

Since consumers in our model are heterogeneous in multiple dimensions, the resulting seller’s problem is not necessarily tractable under the existing assumption used in the classic pricing literature (e.g., increasing generalized failure rate). However, we show that, under a set of reasonable assumptions that are common in many practical scenarios, the revenue function is indeed well-behaved (e.g., the revenue function is unimodal in price, the price-point pair is monotone in inventory and time). In fact, we believe that our paper is one of the first papers in dynamic pricing literature that study the case when consumers are heterogeneous in several attributes and have multiple options for purchase.

The remainder of this paper is organized as follows. Section 2 provides a survey of relevant literature. Section 3 outlines our base model, in which a consumer choice model is embedded into a dynamic pricing problem, and describes its analytical results regarding the effect of reward sales. In sections 4 and 5, we provide extensions of our base model to the cases in which the seller further decides whether to allow or block reward sales and the seller dynamically changes point requirement along with price.

3.2 Literature Review

Our paper is related to mainly two streams of research – loyalty program and dynamic pricing. Impact of loyalty programs has been extensively studied in marketing, both from consumers’ and firm’s perspective. In particular, a substantial body of marketing literature focuses on the long-term effect of loyalty program. *Bijmolt et al.* (2010), *Dorotic et al.* (2012), and *Breugelmans et al.* (2014) provide comprehensive reviews. Many of these studies attempt to establish empirical evidences on positive impact of loyalty programs or suggest several underlying mechanism leading to positive impacts. For instance, a loyalty program generally improves a consumer’s evaluation of the good and induces repeat-purchasing behavior. The reasons studied include switching cost (*Carlsson and Löfgren* 2006), psychological barrier (*Bolton et al.* 2000 and *Hallberg* 2004), and strategic behavior (*Lewis* 2004). There are also a few studies that find the limited impact of loyalty program (*Uncles et al.* 2003, *Sharp and Sharp* 1997). While there are conflicting results on the effect of loyalty program on customers’ retention, several papers explain how the seller can improve their profit through a loyalty program. *Kim et al.* (2001) show that a loyalty program creates a switching cost, which then allows the seller to charge a higher price than the seller without a loyalty program. In another study, *Kim et al.* (2004) show that the seller can use reward sales to reduce excess capacity in a low-demand season. In a competitive setting, *Shin and Sudhir* (2010) demonstrates how the seller can strategically manage loyalty program to gain a competitive advantage (e.g., offering a discount to its own or a competitor’s loyalty members). While our paper also focuses on the seller’s profit, we zoom in the seller’s pricing and rationing decisions when some customers are trying to buy the product with reward points (which may not be always highly profitable).

Specifically, instead of focusing on the impact of lingering behavioral changes induced by a loyalty program such as retention or repeating purchases (which most

marketing literature does), we focus on the redemption of points and how it inflicts on the seller profit maximizing price. In that vein, our paper is closely related to a handful of papers that study how consumers assess and use loyalty points (*Liston-Heyes* 2002, *Kivetz and Simonson* 2002a, *Liu* 2007 and *Basumallick et al.* 2013). For instance, *Kivetz and Simonson* (2002a) show how a consumer evaluates loyalty points depending on effort to acquire loyalty points (e.g., accumulating points through credit card spending vs. accumulating points by flying). *Dreze and Nunes* (2004) propose mental accounting model to explain how a consumer evaluates loyalty points and show that using loyalty points may lower psychological cost for purchasing goods and services. In a similar vein, *Kivetz and Simonson* (2002b) find that some consumers use points to indulge in luxury goods without feeling guilty. While these papers focus on how they use points, *Stourm et al.* (2015) explain why many consumers stockpile loyalty points using cognitive and psychological incentives, both of which enable consumers to value points differently than cash. Since we consider a consumer comparing two purchase methods (cash vs. point), the endowed point balance (i.e., how many points do you currently have?) and the perceived value of points (i.e., how much is 1 point worth to you?) are equally important determinants of whether a consumer buys the product or not and using what method. In this regard, *Kopalle et al.* (2012), *Dorotic et al.* (2014), and *Chun et al.* (2015) are closely relevant to our paper. These papers show that allowing point redemptions may cannibalize cash sales and revenue. Departing from these, our focus is on the seller who has to operationalize pricing and rationing decisions when facing consumers who can purchase with cash or points.

Our paper is also closely related to papers in dynamic pricing. *Bitran and Caldentey* (2003) and *Elmaghraby and Keskinocak* (2003) provide extensive reviews on earlier works in this area. These area focuses on how the seller should adjust the price based on remaining time and inventory throughout the selling horizon in different settings:

single product (*Gallego and Van Ryzin 1994* and *Bitran and Mondschein 1997*), multiple products (*Zhang and Cooper 2005* and *Maglaras and Meissner 2006*), impact of strategic consumers (*Aviv and Pazgal 2008*, *Ahn et al. 2007*), and negotiation (*Kuo et al. 2011*). Up to our best knowledge, this is the first paper that examines the interaction between price and reward purchase (point redemption) in a dynamic pricing setting.

Like the majority of pricing papers, we derive analytic results and extract insights therein. To gain analytic tractability of optimal price, we introduce technical condition that is analogous to *Ziya et al. (2004)* and *Lariviere (2006)*). However, in our model, a consumer has to compare two options (three if we include “do not purchase”) depending on realization of three idiosyncratic attributes – reservation price, point balance, and point worth. The details will be discussed later. As we will see, simply extending IGFR to multiple dimensions does not result in a well-behaved profit function. We provide a sufficient condition and explain why it is not a big concern in many practical cases.

3.3 The Model

We consider the seller’s problem in the context of a dynamic pricing model in *Gallego and Van Ryzin (1994)* and *Bitran and Mondschein (1997)*. This enables us to highlight how reward sales through point redemption affects the seller’s pricing (and rationing) policy compared to the seller that does not offer a loyalty program. Following the convention, we consider a seller with limited inventory of a product over a predetermined selling horizon, which is divided into T periods, indexed by $t \in \{1, 2, \dots, T\}$. Each period is short enough so that at most one consumer arrives in a given period with probability $\lambda \in (0, 1)$. In each period, the seller can sell its product either in cash (*cash sale*) or with points (*reward sale*). That is, a consumer

can purchase a product either by paying p in cash or redeeming q loyalty points (we call a number of redeemed points for purchasing as a *point requirement*). After seeing both price and point requirement, a consumer decides whether to buy or not, and, if buying, chooses the method of payment. The seller receives the revenue p from each cash sale. On the other hand, the seller receives the reimbursement for each reward sale. In many situation, the amount of reimbursement depends on point requirement, q , thus it is reasonable to assume that the reimbursement rate, $R(\cdot)$, is a non-decreasing function of q . The seller's goal is to maximize its total expected revenue (the revenue from cash sales plus the revenue from reimbursements) throughout the selling season. In particular, we consider several types of the seller depending on the discretionary control over the terms of reward sales (e.g., whether a seller can block reward sales, whether a seller can change point requirement), which will be discussed later in detail.

3.3.1 Consumers' Problem

An arriving consumer observes cash price (p) and point requirement (q), and decides whether to buy or not, and how to pay for a product (either paying p in cash or redeeming q points). We assume that a consumer maximizes her utility, which depends on the following three attributes: (i) reservation price, i.e., how much she values the product (ii) point balance which determines whether a consumer is able to buy with points or not, and (iii) perceived value of a point, i.e., how much she values points (in terms of monetary denomination). We assume these attributes are consumer idiosyncratic, thus heterogeneous.

Specifically, we model the consumer's reservation price by a random variable, V , which follows distribution $F(\cdot)$ and density $f(\cdot)$. If cash is the only available form of payment, consumers with reservation price higher than the current price will purchase the product. Thus, $\bar{F}(p) := 1 - F(p)$ represents the probability that a consumer

prefers buying with cash to not buying.

We assume that point balance is also a heterogeneous attribute. This reflects the fact that not all consumers belong to a loyalty program (i.e., they have zero point balance) and, among members, only some of them have accumulated more than q points. We model the point balance as a non-negative random variable, W , with distribution $\beta(\cdot)$. Then, $\bar{\beta}(q) := 1 - \beta(q)$ represents the proportion of consumers with sufficient points: We call them *loyalty consumers* and we call consumers without enough points (either a non-member or having insufficient points) as *cash-only consumers*.

In addition to reservation price and point balance, the perceived value of a point, which is how one point is worth to a consumer, is also an idiosyncratic attribute. This reflects the fact that not all consumers with the same reservation price and point balance will behave the same. The point-worth is highly subjective and depend on many factors (*Liston-Heyes* 2002, *Liu* 2007, and *Basumallick et al.* 2013). In fact, loyalty points do not have the universally accepted conversion rate. In addition, how each consumer assesses loyalty points varies depending on how she acquires points and how she intends to use. For instance, depending on how a consumer has accrued points (e.g., credit card usage, paying for points, and promotion gift), the perceived value of points widely varies (*Kivetz* 2003). In addition, the value of points also depends on how she intends to use. A consumer can use 80,000 Chase Ultimate Rewards points to buy a business-class round-trip ticket from US to Europe (which has the cash-value of \$6,000) or to buy a \$800 Amazon gift card (*PointsGuy* 2015). Although it is obvious that using points for a flight ticket saves a lot more cash, some consumers will opt in for a gift card (which has less cash value) as the value of a business-class ticket to Europe differs significantly from one person to another. This observation implies that the value of 80,000 points can be quite different depending on the intended use. We model this attribute using a non-negative random variable Θ with distribution $G(\cdot)$ and density $g(\cdot)$. That is, a point-worth of $\Theta = \theta$ (i.e., 1


point is equal to $\$ \theta$) means that q points has the equivalent monetary value of $\$ \theta q$ to that consumer, making her indifferent between paying $\$ \theta q$ in cash and redeeming q points.

We assume that all three attributes – V , W , and Θ are independent. Not only this assumption gives us a traction for analysis, it also reflects the fact that there is no direct evidence or underlying cognitive mechanism to relate one attribute to another. For instance, as mentioned earlier, the point worth has many compounding factors including the method of acquisition and the intended use. Even among elite members (e.g., an Hhonor diamond member who accumulates more than 120,000 points per year), some members choose to redeem points frequently while other consumers choose to stockpile points for a grand adventure. Hence, the amount of points a consumer accumulates does not necessarily dictate her point-worth (as a result, her redeeming choice). In addition, even for the same person, the perceived value may change per circumstance. For instance, the perceived value of points when an elite member makes a hotel reservation for a personal vacation will be significantly different from that when the same person makes a reservation for a business trip. Consequently, there is no obvious correlation among these attributes.

For given p and q , the utility from each purchase option depends on these three attributes – v , w , and θ (realizations of V , W , and Θ). If she obtains a product by paying p in cash, the corresponding utility is $v - p$. If she purchases by redeeming q points, her utility is $v - \theta q$. If she does not purchase, she earns the reservation utility of the outside option, which is set to be zero. In case of a *cash-only consumer* who does not have enough points (i.e., $w < q$), her decision is reduced to either buying with cash or not buying, i.e., $\max \{v - p, 0\}$. On the other hand, a *loyalty consumer* (who has enough points) compares three options (cash purchase, point purchase and no purchase) and makes a choice to maximize her utility, i.e., $\max \{v - p, v - \theta q, 0\}$.

Given this consumer model, the following Lemma and Figure 1 characterizes the

Figure 3.1: Consumer behavior depending on their type (*cash-only consumer* or *loyalty consumer*), reservation price (V), and point-worth (Θ) given price p and point requirement q .



likelihood of each behavior by a random consumer (or depending on the type of a consumer).

Lemma III.1. *Suppose that price p and point requirement q are given.*

(a) *Consumer decision falls into three cases:*

- *With probability $P(\text{cash purchase}) = \beta(q)\bar{F}(p) + \bar{\beta}(q)\bar{F}(p)\bar{G}(p/q)$, a consumer will buy with cash.*
- *With probability $P(\text{point purchase}) = \bar{\beta}(q) \int_0^{p/q} \bar{F}(qx)dG(x)$, a consumer will buy with points.*
- *With probability $1 - P(\text{cash purchase}) - P(\text{point purchase})$, a consumer will not purchase.*

(b) *Compared to a cash-only consumer, a loyalty consumer will buy with cash less likely, but in overall (using cash and points), buy a product more likely.*

Part (a) describes who buys with each payment method. With probability $\beta(q)$, a consumer is a *cash-only consumer*, thus, she purchases if and only if $v \geq p$ (as shown in Figure 1-a). With probability $\bar{\beta}(q)$, a consumer has enough points and her choice depends on her v and θ . Cash purchase is optimal if and only if $v \geq p$ and $\theta q \geq p$,

point redemption is optimal if and only if $v \geq \theta q$ and $p > \theta q$. When her reservation price is low and point worth is high, then she does not purchase (as shown in Figure 1-b). Note again that all three attributes – reservation price, point balance, and point worth – affect a consumer’s choice.

Part (b) shows that a *loyalty consumer* is more likely to purchase (combining cash and point purchase) than a *cash-only consumer* (who can only purchase with cash). This implies that allowing point redemption increases the overall probability of a purchase, but decreases the probability of a cash purchase. The decrease in cash sales implies that some consumers (those whose point balance is high and point-worth is low) will switch the payment from cash to points (*Point B* in Figure 1-b) given that these consumers would have paid in cash if there is no option to buy with points. On the other hand, allowing point redemption creates a new stream of demand from *loyalty consumers* who find the price p is too high but redeeming q points is not too costly (*Point A* in Figure 1-b). Adding these consumers to a mix, point redemption increases the likelihood of total sales while substituting the likelihood of cash sales.

The next result further characterizes how consumers respond to a change in price or point requirement. We note that the results hold regardless of distributions.

Lemma III.2. (a) *If price p increases given point requirement q ,*

$$\frac{dP(\text{cash purchase})}{dp} \leq 0, \quad \frac{dP(\text{point purchase})}{dp} \geq 0, \quad \frac{dP(\text{total purchase})}{dp} \leq 0.$$

(b) *If point requirement q increases given price p ,*

$$\frac{dP(\text{cash purchase})}{dq} \geq 0, \quad \frac{dP(\text{point purchase})}{dq} \leq 0, \quad \frac{dP(\text{total purchase})}{dq} \leq 0.$$

Lemma III.2 shows that cash and point purchases are indeed imperfect substitutes as the increase of price (or point requirement) results in some consumers switching from cash to point purchase (or vice versa) and others not buying. If price increases, a

cash purchase becomes less attractive while making a point purchase relatively more attractive. The result shows that a decrease in cash sales is always bigger than an increase in reward sales. As a result, the probability of a purchase (by cash and points) decreases. One may think this imperfect substitution is an obvious result: *cash-only consumers* (who cannot use points) simply get priced-out as p increases. This argument does not tell the whole story. To see this, suppose that all consumers are *loyalty consumers* (they have enough points). Even in this case, the probability of a purchase still decreases, thus, two options still remain to be imperfect substitutes. This is because consumers whose perceived value of a loyalty point is high enough ($\theta q > v$) will not consider point purchase and drop out if price increases (i.e., *Cash A* in Figure 1-b).

On the other hand, when a point requirement increases, some consumers who buy with points will switch to either pay with cash or not purchase at all. This is derived from the two facts: As q increases, point-purchase becomes less attractive and fewer people become eligible to use points (i.e., $\bar{\beta}(q)$ decreases in q).

Remark 1. There are cases where consumers receive a slightly different product when purchasing with points, making the product purchased with cash more valuable. For instance, many loyalty programs issue points only for a cash purchase, but not for point-redeeming consumers. In other example, passengers with award flight tickets will be placed to a different priority for a seat upgrade. Our model can accommodate such scenario (with a slight modification), and thus, adding so does not alter the results or ensuing insights of our model. To see this, suppose that a consumer accrues \bar{q} points only if she purchases with cash. To capture this, we add a term that represents the difference in form of a disutility: $v - \theta(q + \bar{q})$ becomes the utility from a point purchase. Other than this minor change, everything else remains the same.

3.3.2 The Seller's Problem

We now consider the seller's pricing problem in each period. As mentioned in the introduction, one of the main questions is to examine how the seller's discretion over the terms of reward sales (e.g., open-close point redemption or changing point requirement) affects the seller's (pricing) decision and revenue. To do this, we first start with a simple model where the seller always allows point redemptions, and the point requirement and reimbursement rate are fixed throughout a selling season. Later, we will gradually relax the restrictions in two ways: In section 4, we consider a seller who can block the reward sales (*black-out model*). In section 5, we consider a seller who chooses not only price but also point requirement in each period (*dynamic adjustment model*).

Since q and R are the fixed, we use, for ease of notation, β for $\beta(q)$, $\bar{\beta}$ for $\bar{\beta}(q)$, and R for $R(q)$ throughout this section. We incorporate Lemma 1 and embed consumer choice for given (p, q) in a seller's dynamic pricing problem. We let $J_t(p, y)$ represents the revenue-to-go function when the seller charges price p when it has y units in inventory with t -periods to go and then follows the optimal policy starting from $t - 1$ and onwards:

$$\begin{aligned}
 V_t(y) &= \max_{p \geq 0} J_t(p, y) \quad \text{for } y > 0, t = 1, 2, \dots, T, \text{ and} & (3.1) \\
 V_0(y) &= 0 \text{ for } y \geq 0, \quad \text{and} \quad V_t(0) = 0 \text{ for } t = 1, 2, \dots, T,
 \end{aligned}$$

where

$$\begin{aligned}
J_t(p, y) = & \lambda \bar{\beta} \left[\bar{F}(p) \bar{G}(p/q) (p + V_{t-1}(y-1)) + \int_0^{p/q} \bar{F}(qx) dG(x) (R + V_{t-1}(y-1)) \right. \\
& \left. + \left(1 - \bar{F}(p) \bar{G}(p/q) - \int_0^{p/q} \bar{F}(qx) dG(x) \right) V_{t-1}(y) \right] \\
& + \lambda \beta \left[\bar{F}(p) (p + V_{t-1}(y-1)) + F(p) V_{t-1}(y) \right] + (1 - \lambda) V_{t-1}(y).
\end{aligned}$$

We define $\Delta_t(y) = V_t(y) - V_t(y-1)$, which represents the marginal value of inventory, and rewrite $J_t(p, y)$ as

$$\begin{aligned}
J_t(p, y) = & \lambda \bar{\beta} \left[\bar{F}(p) \bar{G}(p/q) (p - \Delta_{t-1}(y)) + \int_0^{p/q} \bar{F}(qx) dG(x) (R - \Delta_{t-1}(y)) \right] \\
& + \lambda \beta \left[\bar{F}(p) (p - \Delta_{t-1}(y)) \right] + V_{t-1}(y). \tag{3.2}
\end{aligned}$$

We observe from (3.2) that $J_t(p, y)$ has three sources of revenue.

1. the revenue of cash sales to a *cash-only consumer* :

$$\lambda \beta \bar{F}(p) (p - \Delta_{t-1}(y)).$$

2. the revenue of cash sales to a *loyalty consumer* :

$$\lambda \bar{\beta} \bar{F}(p) \bar{G}(p/q) (p - \Delta_{t-1}(y)).$$

3. the revenue of reward sales to a *loyalty consumer* :

$$\lambda \bar{\beta} \int_0^{p/q} \bar{F}(qx) dG(x) (R - \Delta_{t-1}(y)).$$

Note that the first term $\lambda \beta \bar{F}(p) (p - \Delta_{t-1}(y))$, has been extensively studied in the literature, as this is the revenue term in class dynamic pricing problems. It has been shown that this term is unimodal in p if $F(\cdot)$ has an increasing generalized failure rate (IGFR) – *Ziya et al.* (2004) and *Lariviere* (2006). Following similar derivation, the cash-sale revenue to a *loyalty consumer* is also unimodal in p if both $F(\cdot)$ and

$G(\cdot)$ are IGFR. Lastly, from Lemma III.2(a), more consumers will choose to use points as price increases. Thus, the reward-sale revenue increases in p . Although each portion is unimodal under simple condition (IGFR), note that $J_t(p, y)$ is not necessarily unimodal in p as it is the sum of these three functions. This imposes a challenge in characterizing optimal price and its properties.

Figure 3.2: Examples of expected revenue-to-go functions in price when both $F(\cdot)$ and $G(\cdot)$ are truncated normal (i.e., IGFR): $F \sim \text{truncated } N(30, 10)$, $q = 1$, $t = 1$, and $y = 1$. In figure (a), $\bar{\beta} = 0.6$, $G \sim \text{truncated } N(15, 2)$ and $R = 10$. In figure (b), $\bar{\beta} = 1$, $G \sim \text{truncated } N(28, 1)$ and $R = 30$.

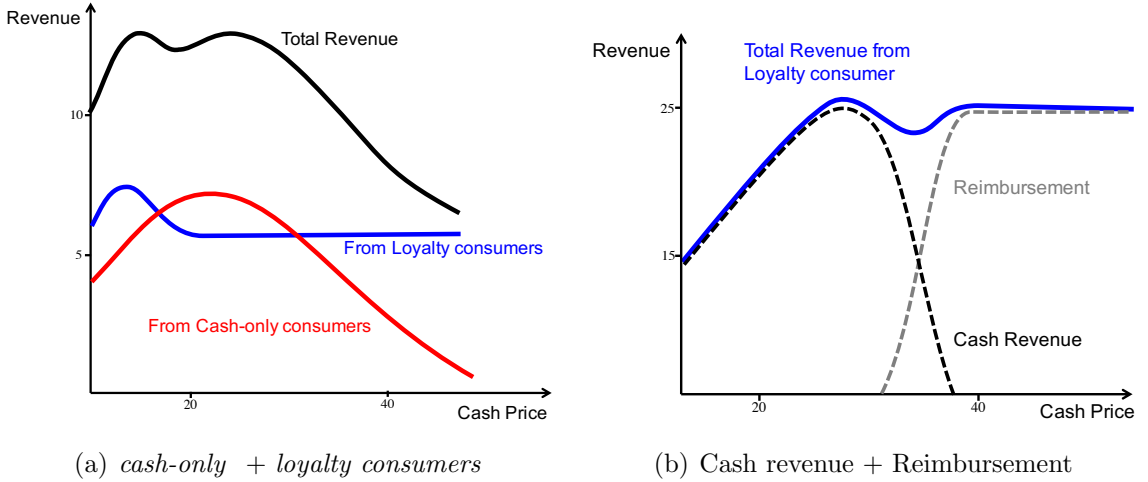


Figure 2-a and b illustrate two examples that show how the seller's revenue changes in price. One may think that the lack of unimodality comes from the fact that there are two different subgroups, *cash-only consumers* and *loyalty consumers*, as illustrated in Figure 2-a. This explanation, however, is not always true. The expected revenue function is not unimodal even when all consumers are *loyalty consumers* (Figure 2-b). To see why this is the case, note that *loyalty consumers* always compare two options. Depending on how *loyalty consumers* buy, the seller receives two different revenues, price and reimbursement rate. As price increases, marginal consumers whose valuations are close to p drop out from cash purchase, but only some of them substitute to buy with points. Thus, the change of revenue in price not only depends

on $F(\cdot)$ but also on $G(\cdot)$, and the interaction between two. Because of this, the expected revenue-to-go function is not unimodal even when $F(\cdot)$ and $G(\cdot)$ are IGFR.

Nonetheless, we identify a sufficient condition under which the revenue function is unimodal in price, allowing us to characterize the optimal price and comparative statics.

Theorem III.3. *If the following three conditions are satisfied for a given q :*

A1. $F(\cdot)$ and $G(\cdot)$ are with an increasing failure rate (IFR),

A2. $\frac{f(p)}{\bar{F}(p)} \geq \frac{g(p/q)}{\bar{G}(p/q)}$ for any p , and

A3. $\frac{f(p)}{\bar{F}(p)} / \frac{\beta g(p/q)}{1 - \beta \bar{G}(p/q)}$ is non-decreasing in p ,

there exists a unique optimal price, $p_t^(y)$, that satisfies the first-order condition for $J_t(p, y)$.*

Condition *A1* is common in the pricing literature (see *Ziya et al. 2004* and *Lariviere 2006*) and is satisfied by a large range of probability distributions including uniform, normal, and exponential and their truncated versions. This condition guarantees that cash-sale revenues from *cash-only* and *loyalty consumers* are unimodal, respectively. As mentioned above, because of substitution between cash purchase and point redemption, IGFR of F and G does not guarantee the unimodality. Two conditions, *A2* and *A3*, address this. Recall from Lemma 2-(a) that an increase in price induces some cash-paying consumers either to not buying (*pricing-out effect*) or switching to point-purchase (*substitution effect*). Condition *A2* implies that, as p increases, the number of consumers who are priced out increases faster than the number of consumers who switch to point purchase. In other words, the *pricing-out effect* dominates the *substitution effect*. Condition *A3* further implies that the *pricing-out effect* grows faster than the *substitution effect* as price increases. Both conditions imply that an increase in price always has a greater impact on the pricing-out effect than the substitution effect.

It can be shown that conditions $A2$ and $A3$ can be satisfied if the condition in the following corollary is satisfied. If this is the case, the revenue function is unimodal even when all consumers are *loyalty consumers* ($\beta = 1$).

Corollary III.4. *The sufficient conditions for $A2$ and $A3$ are as follows: for a given q , the hazard ratio of the reservation price distribution to the q point-worth distribution, $\frac{f(p)}{F(p)} \bigg/ \frac{g(p/q)}{G(p/q)}$, is greater than 1 and increasing in price.*

It turns out these two conditions are not restrictive from practical perspective. *Stourm et al.* (2015) find that consumers tend to accumulate and stockpile points (for distant future use) rather than using immediately. For instance, only about a quarter of newly issued frequent-flyer miles (or equivalently 6.3% of the accumulated miles) is redeemed in 2004, while the rest are accumulated or expired (*Economist* 2005b). Moreover, at least one third of issued points (\$15 billion worth) in the United States each year are never redeemed and go expired (*Colloquy* 2011). There are many plausible reasons for this, including cognitive cost (e.g., redeeming process is costly) and psychological motivation (e.g., a high value on the possession of points itself) as discussed by *Stourm et al.* (2015). Furthermore, points required for redemption is often set quite high, thus, a consumer needs to have a sufficiently large balance of points. All of these suggest that consumers do not immediately use up points as soon as they are qualified. In these situations, $A2$ and $A3$ are not restrictive assumptions. As explained above, conditions $A2$ and $A3$ imply that the primary impact of increase in price is the *pricing-out effect*, not the *substitution effect* (from cash to point use). When consumers are hesitant to use points (instead, they rather hold and stockpile points for future use), the conditions are easily satisfied. Thus, Theorem III.3 shows that, although the revenue function is not unimodal in general, such a pathological behavior does not arise in many situations where loyalty points are accumulated and redeemed.

Building on this, we now examine the property associated with the optimal price.

For this, we first define an auxiliary optimization problem, which represents a pricing problem of the seller who does not offer reward sales in period t , but follows the optimal policy from period $t - 1$ and onward. That is, the revenue-to-go from period $t - 1$ onward is still $V_{t-1}(\cdot)$ as defined by (3.1).

$$V_t^c(y) = \max_p \left\{ \lambda \bar{F}(p) (p + V_{t-1}(y - 1)) + (1 - \lambda \bar{F}(p)) V_{t-1}(y) \right\}, \quad y > 0, \quad t = 1, \dots, T. \quad (3.3)$$

Let $p_t^c(y)$ be the solution to (3.3), which we call as *cash-only price*, the best price that the seller can charge if only the cash sale is available in period t . We also denote the optimal price that the seller will charge considering both cash and reward sales (the solution to (3.2)) by $p_t^*(y)$. As the seller has an additional channel to sell using points, one can suggest that the marginal value of inventory increases. This logic is further supported by the fact that, for a given price, the likelihood of total sales is always higher when the reward sales are available (Lemma III.1-b). If this intuition holds, the optimal price must be always higher than the *cash-only price*: i.e., $p_t^*(y) > p_t^c(y)$ for any t and y . However, the next result shows that this is not the case.

Proposition III.5. *Given y units of inventory and t periods to go until the end of the season:*

(a) *If the reimbursement rate is higher than the cash-only price price, $R > p_t^c(y)$, the reward sales induce the seller to add a premium on the cash-only price, i.e., $p_t^c(y) < p_t^*(y) < R$.*

(b) *If the reimbursement rate is lower than the cash-only price, $R < p_t^c(y)$, the reward sales induce the seller to offer a discount on the cash-only price, i.e., $R < p_t^*(y) < p_t^c(y)$.*

This proposition indicates that the optimal price, $p_t^*(y)$, can be higher or lower than the *cash-only price* $p_t^c(y)$ depending on whether $p_t^c(y)$ is higher or lower than R

(Note that $p_t^c(y)$ monotonically changes in y and t). To elaborate further, consider the case when $R > p_t^c(y)$. Suppose that the seller charges the *cash-only price*, $p_t^c(y)$, but allows a reward purchase. Since $R > p_t^c(y)$, the seller gains more from a reward sale than a cash sale. From Lemma 2-a, the seller can increase reward sales by increasing price; in other words, the seller adds a premium: $p_t^c(y) < p_t^*(y)$. By converting more consumers to buy with points, the seller's revenue increases. As stated, the amount of a premium is chosen so that $p_t^*(y)$ does not exceed R . To understand why, the primary reason for a premium is to induce some consumers (with a relatively lower willingness to pay in cash) to buy with points while earning a good portion of revenue from cash sales. This is especially beneficial to the seller as the revenue from consumers who buy with points is higher than their willingness to pay in cash. However, when this premium becomes too high ($p > R$), then a very few will buy in cash in the first place, thus the cash-sale alone is suboptimal. Furthermore, a good portion of consumers who should have bought with cash now buy with points. Thus, the seller is no longer able to price discriminate effectively if a premium is too high. Hence, $p_t^*(y)$ should always lie between $p_t^c(y)$ and R .

Similarly, when $R < p_t^c(y)$, the seller will prefer cash sales to reward sales. If the seller charges $p_t^c(y)$ while allowing reward sales, the seller will get the lower reimbursement rate from any consumer who will buy with points. In anticipation of this, the seller offers a discount to induce more consumers to buy with cash. Once again, $p_t^*(y)$ should always lie between R and $p_t^c(y)$ because, if price falls below R , the seller does not earn much from both cash and reward sales.

The following result further shows that the region that the seller offers a discount or premium is monotone in inventory level and remaining periods.

Corollary III.6. *Given y units of inventory and t periods to go until the end of the season:*


(a) *If it is optimal to offer a discount from the cash-only price, then it is also optimal*

to offer a discount with $y - 1$ units of inventory or/and $t + 1$ periods to go.

(b) If it is optimal to add a premium to the cash-only price, then it is also optimal to add a premium with $y + 1$ units of inventory or/and $t - 1$ periods to go.

This result identifies when the reward sales will reduce a mark-up or mark-down that is applied in a dynamic pricing policy. In the standard dynamic pricing model (Gallego and Van Ryzin 1994 and Bitran and Mondschein 1997), there exists a mark-up when the inventory level changes from y to $y - 1$ given t periods to go. The proposition says that in regions where the *cash-only price* is sufficiently high (compared to R), the seller must reduce such a mark-up (and keep the price not too high), otherwise, a sizable portion of consumers will buy with points. Likewise, in regions where the *cash-only price* is very low (lower than R), the seller needs to reduce a mark-down (and keep the price not too low) in order to induce consumers to purchase with points. Thus, the reward sales have non-trivial effects on the seller's price. In fact, the reward sales attenuate the price fluctuation of dynamic pricing as it reduces the extent of a mark-up or mark-down. Figure 3 illustrates this result.

Figure 3.3: The optimal price, $p_t^*(y)$, and the *cash-only price*, $p_t^c(y)$, depending on the inventory level (left figure; for given $t = 10$) and the remaining time (right figure; for given $y = 8$) when $\lambda = 0.8$, $\bar{\beta} = 0.7$, $F \sim \text{Uniform}[0, 100]$, $G \sim \text{Uniform}[0, 10]$, $q = 10$, and $R = 55$.



Our next result explores how the optimal price, $p_t^*(y)$, changes in the fraction of

loyalty consumers and the reimbursement rate.

Proposition III.7. *Given y units of inventory and t periods to go until the end of the season:*

(a) *The amount of price adjustment (either premium or discount) from considering a reward sale becomes greater if more consumers are eligible to purchase with points in period t (the higher $\bar{\beta}_t$).*

(b) *The optimal price $p_t^*(y)$ and the resultant expected revenue increase if R increases. Although more consumers buy with points, no consumer's utility increases.*

As expected, the seller makes a more aggressive price adjustment if there are more *loyalty consumers* (Part a). This is simply because consumer substitution (from cash to points, or *vice versa*) has a greater impact when there are more *loyalty consumers*. Part (b) states that as the reimbursement rate goes up, the seller increases the price so that more consumers buy with points. One interesting observation is that no consumers are better off in spite of the fact that more consumers buy products with points (not paying in cash). The reason is that, as the point requirement q being fixed, the only way the seller can induce more reward sales is to increase price. While the consumer's utility from a point purchase is independent of price, a high price negatively affects the utility of consumers who previously purchased with cash. For instance, a consumer who switched from cash purchase to point purchase because of price increasing, will have a lower utility than before.

Remark 2. In Proposition III.7, we assume that the change of a parameter (β or R) is limited only to period t . If the change is global (the parameter changes for all periods), it affects the marginal value of inventory in each period and, as a result, value function. This makes comparing two different policies difficult as one needs to know the exact difference of the value functions for two separate dynamic programming problems. However, we have conducted an extensive numerical study and found that Proposition III.7 continues to hold globally. In our numerical exper-

iment, we consider three different combinations of reservation price and point-worth distributions ($F(\cdot)$ and $G(\cdot)$ follow uniform over $[0, 100]$, exponential over $[0, 100]$ with mean 60, and truncated normal over $[0, 100]$ with mean 60 and standard deviation 20; where the point requirement given at $q = 1$), three different values for the arrival probability ($\lambda \in \{0.3, 0.6, 0.9\}$), ten different values of *loyalty consumers'* fraction ($\bar{\beta} \in \{0.1, 0.2, \dots, 1.0\}$), and ten different values of reimbursement rate ($R \in \{10, 20, \dots, 100\}$), with inventory levels ranging from 1 to 20 with $T = 20$ periods to go until the end of the season. In all instances, we note that the result of Proposition III.7 continues to hold even when we change the parameters (β and R) globally.

Figure 3 shows that the optimal price is monotone in the inventory level and remaining periods. The following proposition summarize this result and it's implication on the customers' redemption behavior.

Proposition III.8. *The optimal price, $p_t^*(y)$, is decreasing in the inventory level y and increasing in the remaining periods t . Consequently, consumers are more likely to use points when the inventory level is low and/or more periods remain.*

As in the dynamic pricing literature (without loyalty program consideration), the proposition above shows that the seller is willing to move inventory faster by lowering price as the inventory level increases or the remaining time decreases. In other words, the price adjustment (premium or discount) in Proposition 1, while it reduces the seller's mark-up and mark-down, does not change the monotonicity of price in inventory and time. The result is simply a consequence of Lemma 2-(a) that the probability of (cash + reward) sales is decreasing in cash price for a given point requirement. However, this result has an interesting implication on how the reward channel is utilized. *Kim et al.* (2004) show that the seller uses reward sales to reduce excess capacity (e.g., high inventory level and short periods left). Given the constant reimbursement rate, this implies that the seller is more willing to sell with points when

the marginal value of inventory is very low. However, the seller can only incur more consumers to use points by increasing price, which decreases the likelihood of total sales. In this case, the seller needs to decrease price and can not induce more reward sales. In other words, consumers use points more likely when it is less desirable (when the seller can earn high revenue from cash sales) and use points less likely when it is more desirable (when the seller can earn high revenue from reward sales). This implies the limitation of seller's price-only control on reward sales.

3.4 The Black-Out Model

In Section 3, we consider the seller who does not have any discretion on the terms of reward sales. However, we observe from practice that sellers have different degrees of discretion. One of such discretion is the ability to block out the reward sales (i.e., disallowing consumers to use points for purchase). We now consider a seller who decides the price along with whether to allow consumers to use points (*open*) or not (*close*) in each period.

We modify the dynamic programming problem of (3.1) for this case (*black-out seller*). We let $V_t^B(y)$ be the optimal value function of the black-out seller with y units of inventory in period t , which is given by the following optimality equation:

$$\begin{aligned} V_t^B(y) &= \max \left\{ \tilde{V}_t^c(y), \tilde{V}_t^o(y) \right\} \quad \text{for } y > 0, t = 1, 2, \dots, T, \text{ and} \\ V_0^B(y) &= 0 \text{ for } y \geq 0, \quad \text{and} \quad V_t^B(0) = 0 \text{ for } t = 1, 2, \dots, T, \end{aligned} \quad (3.4)$$

where $\tilde{V}_t^c(y)$ and $\tilde{V}_t^o(y)$, respectively, represent the value functions when the seller

closes and opens the reward-sale channel in period t :

$$\tilde{V}_t^c(y) = \max_{p \geq 0} \left\{ \lambda \bar{F}(p) \left(p + V_{t-1}^B(y-1) \right) + \left(1 - \lambda \bar{F}(p) \right) V_{t-1}^B(y) \right\} \quad \text{and} \quad (3.5)$$

$$\begin{aligned} \tilde{V}_t^o(y) = \max_{p \geq 0} \left\{ \lambda \bar{\beta} \left[\bar{F}(p) \bar{G}^{(p/q)} \left(p + V_{t-1}^B(y-1) \right) + \int_0^{p/q} \bar{F}(qx) dG(x) \left(R + V_{t-1}^B(y-1) \right) \right. \right. \\ \left. \left. + \left(1 - \bar{F}(p) \bar{G}^{(p/q)} - \int_0^{p/q} \bar{F}(qx) dG(x) \right) V_{t-1}^B(y) \right] \right. \\ \left. + \lambda \beta \left[\bar{F}(p) \left(p + V_{t-1}^B(y-1) \right) + F(p) V_{t-1}^B(y) \right] + (1 - \lambda) V_{t-1}^B(y) \right\}. \end{aligned} \quad (3.6)$$

Analogous to the *cash-only price*, $p_t^c(y)$, in Section 3, we denote the solution of (3.5) by $\tilde{p}_t^c(y)$ to represent the *cash-only price* that maximizes the revenue if a reward sale is not available (*closed*) in period t . (We note that the solution of (3.3) is different from the solution of (3.5). We assume that the seller always allows reward sales from period $t-1$ onward in (3.3) while, in (3.5), the seller optimally controls the reward sale availability from period $t-1$ onward. From now on, we use the *cash-only price* to denote the solution of (3.5), $\tilde{p}_t^c(y)$.) Also, let $\tilde{p}_t^o(y)$ be the solution of (3.6), the best price when a reward sale is allowed in period t . We denote the optimal price by $p_t^B(y)$, which is either $\tilde{p}_t^c(y)$ or $\tilde{p}_t^o(y)$ depending on the open/close decision.

The first question to ask is when the seller should allow reward sales and when to block. The following proposition and Figure 3 answer this question and specify the effect of reward sales on price depending on the inventory level y and remaining periods t .

Proposition III.9. *Given t remaining periods, there exist two (inventory) thresholds $h_0(t)$ and $h_1(t)$, $h_0(t) \leq h_1(t)$, such that*

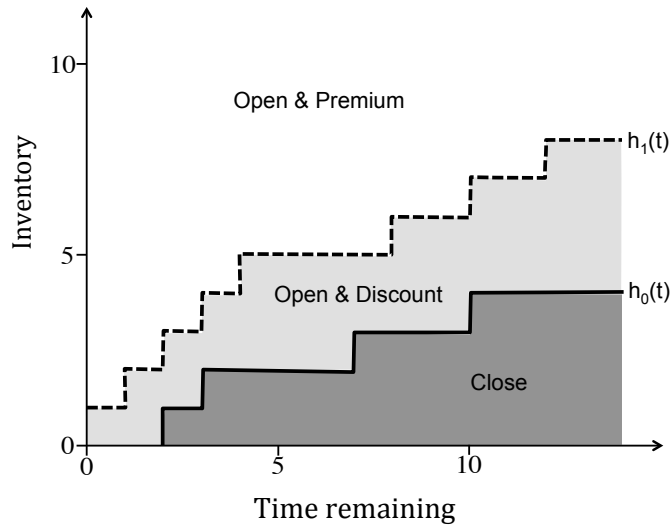
- (a) *if $y < h_0(t)$, it is optimal to close reward sales and set $p_t^B(y) = \tilde{p}_t^c(y) > \tilde{p}_t^o(y)$.*
- (b) *if $h_0(t) \leq y < h_1(t)$, it is optimal to open and offer a discount, i.e., $p_t^B(y) =$*

$$\tilde{p}_t^o(y) < \tilde{p}^c(t).$$

(c) if $h_1(t) \leq y$, it is optimal to open reward sales and add a premium, i.e., $p_t^B(y) = \tilde{p}_t^o(y) > \tilde{p}^c(t)$.

Furthermore, the two thresholds, $h_0(t)$ and $h_1(t)$, are increasing in t .

Figure 3.4: The optimal strategy when the seller has y units of inventory with t periods to go until the end of the season: $\lambda = 0.8$, $\bar{\beta} = 0.5$, $q = 10$, $R = 55$, $F(\cdot) \sim \text{Uniform}[0, 100]$, $G \sim \text{Uniform}[0, 10]$.



Proposition III.9 implies that the seller blocks out reward sales when the inventory level is sufficiently low compared to the remaining time, $y < h_0(t)$. From earlier discussion, the *cash only price* likely exceeds R in this case, making a cash buyer more valuable to the seller. To induce more cash sales, the seller has to offer a discount (Proposition III.5-b) to make sure not too many consumers buy with points. If a discount becomes too deep, the seller instead blocks reward sales and focuses only on cash sales. It should be noted that, however, it can still be optimal to open even when R is smaller than $p_t^o(y)$ (when $h_0(t) \leq y < h_1(t)$). Although the seller gets a lower reimbursement than a cash sale, opening reward sales increases the overall sales (Lemma III.1-b). As long as the discount is not too big, the benefit of boosting sales still outweighs the discounted margin from cash sales. In this case, it is still

optimal to open reward sales with a discounted price. When the inventory level is high relative to the number of periods left, opening reward sales is obviously optimal and the seller adds a premium to induce even more reward sales.

Our next result explores how the optimal strategy changes in the reimbursement rate and the fraction of *loyalty consumers*.

Proposition III.10. *For a given period t , the thresholds, $h_0(t)$ and $h_1(t)$, decrease :*

(a) *if the reimbursement in period t (R_t), increases.*

(b) *if the fraction of loyalty consumers in period t (β_t) increases for any reimbursement rate $R > 0$.*

First note that the fact that two thresholds decrease (in Figure 3) implies that a region of open decision expands (the seller tends to open more likely). Given that, the first result is obvious since reward sales become more desirable as the reimbursement rate increases. On the other hand, the second result is surprising as one may think that the seller may want to avoid the reward sales if the reimbursement rate is low enough and there are more *loyalty consumers*. However, the result states the opposite, that is, the seller opens more likely if there are more *loyalty consumers* even for the low enough reimbursement rate. To see why it is, suppose that it is optimal to open, i.e., $V_t^o(y) > \tilde{V}_t^c(y)$. Notice that the seller's revenue is a weighted sum of the expected revenue from a *cash-only consumer* and a *loyalty consumer*. Since $\tilde{p}_t^c(y)$ maximizes the revenue from a *cash-only consumer*, we know that the expected revenue from a *cash-only consumer* at $\tilde{p}_t^o(y)$ is lower than that from a *loyalty consumer*. That is, when the seller opens reward sales with price adjustment (premium or discount), the seller indeed sacrifices the revenue from a *cash-only consumer* to get more from a *loyalty consumer*. In other words, when $V_t^o(y) > \tilde{V}_t^c(y)$, the revenue from a *loyalty consumer* at $p = \tilde{p}_t^o(y)$ is higher than that from a *cash-only consumer*. In such case, if there are more *loyalty consumers* (the higher $\bar{\beta}$), the overall expected revenue (which is a weighted sum) increases, which makes the open decision more profitable. This

is true even when the reimbursement rate is sufficiently low. The reason behind this result is that, one of the benefits of reward sales is that it increases the likelihood of total sales (Lemma 1-b), which further increases in the fraction of *loyalty consumer*. Thus, regardless of the reimbursement rate, the seller tends to open reward sales more likely if there are more *loyalty consumers*.

Observe from Figure 3 that the optimal policy changes from *close* to *open* as inventory level increases (or remaining time decreases). Hence, the optimal price switches from $\tilde{p}_t^c(y)$ to $\tilde{p}_t^o(y)$. It can be shown that $\tilde{p}_t^c(y)$ and $\tilde{p}_t^o(y)$ are decreasing in y (increasing in t), which immediately implies that the optimal price is indeed decreasing in y within each of the two regimes – *open* and *close*. However, it is not clear whether the monotonicity is preserved for all inventory level because $p_t^o(y)$ can be higher or lower than $p_t^c(y)$ depending on the inventory level (Proposition III.9). Notice that closing is beneficial only when R is sufficiently small so that the seller has to offer a deep discount in order to curb reward sales: $\tilde{p}_t^c(y) > \tilde{p}_t^o(y)$ for $y < h_0(t)$. That is, the optimal price (either of $\tilde{p}_t^c(y)$ or $\tilde{p}_t^o(y)$) decrease in y . A similar result holds for the monotonicity of price in the number of periods left as summarized in the following result.

Proposition III.11. *(Analogous to Proposition III.8) The optimal price, $p_t^B(y)$, is decreasing in the inventory level y and increasing in the number of remaining periods t .*

One interesting question to ask is when the seller gains the most from the ability to block out reward sales. We conduct a numerical study to answer this, and compare the following three different settings – (i) no reward sales in any period, (ii) reward sales in every period (denoted with O), and (iii) discretionally opening/closing reward sales in each period (denoted with B). For the sellers who allow reward sales, we assume that point requirement is given at 10. We consider a 20-period selling season in which at most one consumer arrives in each period with a

Table 3.1: Summary statistics for percentage revenue change from allowing reward sales in every periods (O) and allowing only when it is optimal (B) compared to the bench-mark in which there is no reward sale. The number in the bracket represents the proportion of open decisions for the black-out seller.

$\bar{\beta}$	R	uniform		exponential		normal	
		O	B (open)	O	B (open)	O	B (open)
0.2	10	-13.20	0.00(0.00)	-12.73	0.00(0.00)	-11.66	0.00(0.00)
0.2	20	-9.92	0.00(0.00)	-7.42	0.00(0.00)	-9.48	0.00(0.00)
0.2	30	-6.58	0.00(0.00)	-2.00	2.13(0.65)	-7.19	0.00(0.00)
0.2	40	-3.17	0.96(0.55)	3.50	5.96(0.75)	-4.78	0.00(0.00)
0.2	50	0.30	2.93(0.70)	9.10	10.39(0.85)	-2.25	0.60(0.55)
0.2	60	3.85	5.33(0.80)	14.79	15.33(0.90)	0.41	1.76(0.75)
0.5	10	-27.76	0.00(0.00)	-26.94	0.00(0.00)	-22.21	0.00(0.00)
0.5	20	-20.87	0.00(0.00)	-15.40	0.00(0.00)	-18.41	0.00(0.00)
0.5	30	-13.68	0.00(0.00)	-3.49	5.18(0.60)	-14.17	0.00(0.00)
0.5	40	-6.20	2.35(0.55)	8.97	14.27(0.75)	-9.46	0.00(0.00)
0.5	50	1.55	7.08(0.65)	21.80	24.81(0.80)	-4.23	1.47(0.50)
0.5	60	9.59	12.81(0.75)	34.96	36.45(0.85)	1.51	4.31(0.70)
0.8	10	-38.59	0.00(0.00)	-37.40	0.00(0.00)	-28.91	0.00(0.00)
0.8	20	-29.11	0.00(0.00)	-21.44	0.00(0.00)	-24.27	0.00(0.00)
0.8	30	-18.97	0.00(0.00)	-4.10	8.15(0.60)	-18.89	0.00(0.00)
0.8	40	-8.22	3.72(0.50)	14.17	21.88(0.70)	-12.65	0.00(0.00)
0.8	50	3.11	10.95(0.65)	33.27	37.85(0.80)	-5.42	2.32(0.50)
0.8	60	15.01	19.71(0.70)	53.00	55.47(0.85)	2.82	6.75(0.70)

probability $\lambda = 0.9$. For each setting, we vary the starting inventory levels (ranging from 1 to 20). We also consider three different reservation and point-worth distributions ($F(\cdot), G(\cdot)$) – uniform, exponential, normal. (Specifically, we consider (i) $F(\cdot) \sim U[0, 100], G(\cdot) \sim U[0, 10]$, (ii) $F(\cdot) \sim Exp(\text{mean} = 60), G(\cdot) \sim Exp(\text{mean} = 6)$, $F(\cdot) \sim N(60, 20), G(\cdot) \sim N(6, 2)$. For normal and exponential, we consider the truncated versions.) We also vary fraction of *loyalty consumers* ($\bar{\beta} \in \{0.2, 0.5, 0.8\}$), and reimbursement rate ($R \in \{10, 20, 30, 40, 50, 60\}$). For each scenario, we measure the percentage revenue increase/decrease compared to the seller with no reward sales. The result is summarized in Table 1.

Notice that reward sales do not always increase the seller’s revenue. As expected,

when R is low, allowing reward sales can be quite costly to the seller. This shows once again that reward sales with a low reimbursement rate cannibalizes the seller's cash-sale revenue. In such case, the seller has every intention to not use reward sales. If R is high, allowing reward sales is indeed beneficial as it increases the total sales and enables further price discrimination. Giving the ability to black out reward sales helps the seller only when R is low: the black-out seller takes advantage of her discretion by disallowing reward sales if the reimbursement is low. On the other hand, if R is high, blocking out is hardly used by the seller, which makes the additional benefit of discretion relatively small.

3.5 Dynamic Adjustment Model

We now extend our model to the case of a seller who changes price and point requirement in each period. In particular, we consider the seller who chooses a point requirement from a set, $Q = \{q_1, q_2, \dots, q_N\}$ where $q_i < q_j$ for any $i < j$. We also allow that the seller earns different reimbursement revenue depending on the amount of points redeemed: $R(q_i) \leq R(q_j)$ for all $i < j$. To capture general property, we did not make a specific assumption about $R(q)$ except that it is non-decreasing. This allows us to capture many cases. For instance, if there is a fixed conversion rate, $R(q)$ is linear in q . If $q_N = \infty$ (or sufficiently large), this is equivalent to blocking out reward sales entirely. If $N = 1$ and $q_1 < \infty$, this becomes our basic model considered in Section 3. Thus, the dynamic adjustment model can represent the previous two models as special cases.

The seller's problem of choosing optimal price and point requirement, $(p_t^Q(y), q_t^Q(y))$, with remaining period t and inventory level y is as follows:

$$V_t^Q(y) = \max_{q \in Q, p \geq 0} J_t^Q(p, q, y) \quad \text{for } y > 0, t > 0, \text{ and } V_t^Q(y) = 0 \text{ if } t = 0 \text{ or } y = 0.$$

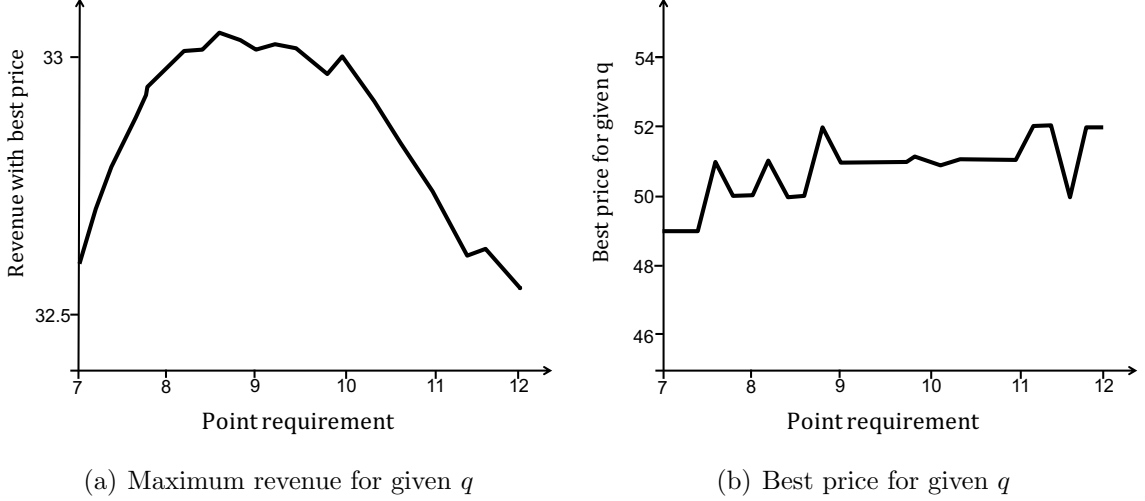
where

$$\begin{aligned}
J_t^Q(p, q, y) = & \lambda \bar{\beta}(q) \left[\bar{F}(p) \bar{G}(p/q) \left(p + V_{t-1}^Q(y-1) \right) \right. \\
& + \int_0^{p/q} \bar{F}(qx) dG(x) \left(R(q) - V_{t-1}^Q(y-1) \right) \\
& \left. + \left(1 - \bar{F}(p) \bar{G}(p/q) - \int_0^{p/q} \bar{F}(qx) dG(x) \right) V_{t-1}^Q(y) \right] \\
& + \lambda \beta(q) \left[\bar{F}(p) \left(p + V_{t-1}^Q(y-1) \right) + F(p) V_{t-1}^Q(y) \right] + (1 - \lambda) V_{t-1}^Q(y). \quad (3.7)
\end{aligned}$$

From Theorem 1, $J_t^Q(p, q, y)$ is unimodal in p for a given q under the conditions $A1-A3$ hold for all q . Thus, for each $q \in \{q_1, q_2, \dots, q_N\}$, there exists a unique price (denoted by $\tilde{p}_t(y, q)$) that maximizes the revenue. However, since q itself is a decision variable, conditions $A1-A3$ are no longer sufficient to characterize the property of an optimal policy, $(p_t^Q(y), q_t^Q(y))$. This is due to the fact that a change in q triggers non-trivial changes in both consumer's and seller's decisions. For instance, suppose that a seller increase a point requirement (q increases). Then, a smaller fraction of consumers will buy with points for two reasons: (i) the proportion of eligible consumers decreases (as $\bar{\beta}(q)$ decreases in q) and (ii) buying with points becomes more expensive (as $G(p/q)$ decreases in q). On the other hand, the seller receives a higher reimbursement rate for a reward sale as $R(q)$ increases in q , changing the optimal price to charge. Because of complex interactions among the three attributes, the induced revenue function does not necessarily have nice analytical properties (see Figure 4).

Let $\tilde{p}_t(q, y)$ be the best price for given q that maximizes the revenue function, $J_t^Q(p, q, y)$. Figure 3.5-(a) and (b) illustrate that not only the revenue is multi-modal but also the best price for each q changes non-monotonically, which makes the seller's problem difficult to analyze. Note that we incorporate multi-dimensional attributes

Figure 3.5: Examples of best revenue ($\max_p J_t^Q(p, q, y)$) and best price ($\arg \max_p J_t^Q(p, q, y)$) for each given point requirement when $F \sim N(60, 30)$, $G \sim N(6, 3)$, $q \in \{7.0, 7.2, 7.4, \dots, 12.0\}$, $\bar{\beta}(q) = \frac{1}{\log q}$, $R(q) = 8q$, $t = 1$, and $y = 1$.



of consumers' behavior – i.e., reservation price $F(\cdot)$, point balance $\beta(\cdot)$, and point worth $G(\cdot)$. With dynamic adjustment, now the reimbursement rate (the revenue from a reward sale) is no longer a constant. The interaction among these three attributes makes the optimal policy complicated and analysis difficult as $J_t^Q(p, q, y)$ is not jointly unimodal in (p, q) . However, since $|Q|$ is finite, one can utilize the fact that the revenue function is unimodal for given q and find the optimal price and point requirement combination, $(p_t^Q(y), q_t^Q(y))$, by comparing the revenue at $p = \tilde{p}_t(q_i, y)$ for each q_i . In addition, we provide a condition under which we can characterize the seller's optimal policy.

Theorem III.12. *The best price for a given q , $\tilde{p}_t(q, y)$, increases in q if the conditions A1–A3, and the following two conditions hold:*

A4. *For a given p , $\frac{1}{q} \cdot \frac{\bar{\beta}(q)g(p/q)}{1 - \bar{\beta}(q)G(p/q)}$ decreases in $q \in Q$. That is, the substitution effect (from cash to points) decreases in q in the hazard rate order.*

A5. *The reimbursement rate is smaller than the cash-only price, i.e., $\max_{q \in Q} R(q) \leq \tilde{p}_t^c(y)$. This condition is automatically satisfied when $R(q_N)$ is smaller than the cash-*

only price for a single-period problem, i.e., $\max_{q \in Q} R(q) \leq \arg \max_p p\bar{F}(p)$.

The condition $A4$ reflects the fact that, as q increases, a point purchase become less and less attractive (compared to paying p in cash). In particular, $A4$ implies that the proportion of marginal consumers who are indifferent between using q points and paying p in cash to consumers who purchases with cash decreases in q . In other words, the proportion of consumers who will buy with points over cash decreases sharply as the point requirement increases.


Condition $A5$ stipulates that, as long as the reimbursement rate is low enough (thus, the margin from a reward sale is considerably smaller than the margin from a cash sale), the best price for given q is increasing in q . Note that low reimbursement rates can be observed quite frequently in a number of examples. For instance, Marriott program reimburses \$201 for one award night at Ritz-Carlton South Beach (a premier Marriott property) while the cash price of the same room is \$599 (Ollila 2012 and Ollila 2013). Note that, when the reimbursement rate is sufficiently low, the seller wants to discourage the reward sale by offering a discount from a *cash-only price* ($p < p^c$). In this region, increasing q has two reinforcing effects. Since $R(q)$ increases, the seller needs to offer a smaller discount. In addition, as q increases, the proportion of consumers who are eligible to buy with points decreases. Thus, the seller needs to worry less about the effect of point redemption, thus can charge a higher price. Both effects contribute to an increase in price as q increases.

Now consider the case when condition $A5$ does not meet. This happens in a region where the reimbursement rate is higher than the optimal price. From earlier results, this is precisely the instance in which the seller needs to offer a premium (compared to a *cash-only price*) in order to induce more reward sales. As q increases, $R(q)$ increases further, which will induce the seller to increase a premium and charge a even higher price. However, the proportion of consumers who are eligible to buy with points decreases as well, and so does the importance of revenues from reward

sales. If this effect is sufficiently strong, then the seller’s price will decrease toward a *cash-only price*. Thus, depending on which effect is strong, the seller’s price may increase or decrease in q . Given these conflicting effects, the best price for each q can change non-monotonically if condition *A5* is violated, as illustrated in figure 4-(b).

In cases where two additional conditions hold, we can further characterize how the optimal price and point requirement change in remaining inventory and time, as given in the following Proposition and Figure 3.6.

Figure 3.6: The optimal price and point requirement depending on the remaining inventory level and period when $\lambda = 0.9$, $Q = \{7, 8, 9, 10, 12\}$, $\bar{\beta}(q) = 0.6 - 0.05q$, $R(q) = 5q$, $F \sim \text{Uniform } [0, 100]$, $G \sim \text{Uniform } [0, 10]$, $t = 9$ (left), and $y = 6$ (right).



Proposition III.13. *If conditions $A1 \sim A5$ hold, the optimal price and point requirement, $p_t^Q(y)$ and $q_t^Q(y)$, are both decreasing in the inventory level, y , and increasing in the remaining periods, t .*

We note that Proposition III.13 generalizes Proposition III.9 and show that both price and point requirement move in the same monotonic direction. For instance, as inventory level y increases (or the number of remaining period t decreases), the seller wants to move inventory fast, which makes the seller to choose the lower price and point requirement. In fact, this result explains the current dynamic pricing patterns in the field, including hotels using a fixed ratio between cash price and point requirement.

One pertaining question to ask is whether a full capability of dynamic adjustment is necessary. Note that a numerical experiment in Section 4 shows that the seller can substantially accrue a higher revenue with an option to block out. We further conduct a numerical study to examine the effect of seller's discretion (in multiple level) on its revenue. To this end, we compare 4 different cases, ranging from no discretion to full discretion (see below for details). For each case, we consider the same settings from Section 4 (20 periods of selling season, $\lambda = 0.9$, starting inventory = $\{1, \dots, 20\}$, (F, G) = uniform, exponential, and normal), but now the point balance and reimbursement rate depend on what point requirement is chosen. Thus, we consider three different point-balance distributions ($\bar{\beta}(q) = 0.6 - q/20$, $1.35 - q/8$, and $2.1 - q/5$) and four different reimbursement-rate function ($R(q) = 40+q$, $30+2q$, $20+3q$, and $10+4q$). For each setting, we consider the following different cases to study the effect of discretion at multiple levels:

- Q_0 : No black-out and the static point requirement $\bar{q} = \arg \min_{q \in \{6,7,8,9,10\}} V_t(y; q)$,
- Q_1 : No black-out and the static point requirement $\bar{q} = \arg \max_{q \in \{6,7,8,9,10\}} V_t(y; q)$,
- Q_2 : Black-out and the static point requirement $\bar{q} = \arg \max_{q \in \{6,7,8,9,10\}} V_t^B(y; q)$,
and
- Q_3 : Black-out and the dynamic point control among 5 point levels, $\{6, 7, 8, 9, 10, \infty\}$.

The first two sellers (Q_0 and Q_1) have no discretion at all, and a static point requirement is used throughout a planning horizon (no black-out policy). But, the difference between Q_0 and Q_1 is on which point-requirement is being used: The point requirement of seller Q_0 (which is a bench-mark seller) is set to be the worst among $\{6, 7, 8, 9, 10\}$ (i.e., minimizes its revenue), while that of seller Q_1 is the best static point requirement that maximizes its revenue. By comparing these two sellers, we measure the effect of choosing an optimal static point requirement on the seller's

revenue (e.g., if the seller cannot dynamically control point requirement, what point requirement should he start with?). The seller Q_2 has an option to block out in each period. The seller Q_3 can dynamically control point requirement and price in each period. For each scenario, we solve dynamic programming problems associated with each seller’s scenario, measure the average (over different starting inventories) percentage revenue improvement compared to the bench-mark seller (Q_0), and summarize the results in Table 2.

Table 3.2: Summary statistics for percentage revenue improvement from the seller’s discretion at multiple levels.

$\bar{\beta}(q)$	$R(q)$	Uniform			Exponential			Normal		
		Q_1	Q_2	Q_3	Q_1	Q_2	Q_3	Q_1	Q_2	Q_3
$0.6 - q/20$ $\bar{\beta}(6) = 0.3$ $\bar{\beta}(10) = 0.1$	$40 + q$	3.20	8.98	9.02	5.05	8.19	8.24	3.65	9.96	9.97
	$30 + 2q$	5.05	8.19	8.24	3.65	9.96	9.97	3.02	10.08	10.21
	$20 + 3q$	3.65	9.96	9.97	3.02	10.08	10.21	4.44	8.52	8.73
	$10 + 4q$	3.02	10.08	10.21	4.44	8.52	8.73	3.94	11.75	11.81
$1.35 - q/8$ $\bar{\beta}(6) = 0.6$ $\bar{\beta}(10) = 0.1$	$40 + q$	5.42	16.55	16.63	7.69	13.86	13.96	6.10	17.58	17.60
	$30 + 2q$	7.69	13.86	13.96	6.10	17.58	17.60	5.34	19.10	19.34
	$20 + 3q$	6.10	17.58	17.60	5.34	19.10	19.34	6.91	14.89	15.33
	$10 + 4q$	5.34	19.10	19.34	6.91	14.89	15.33	6.86	21.17	21.29
$2.1 - q/5$ $\bar{\beta}(6) = 0.9$ $\bar{\beta}(10) = 0.1$	$40 + q$	7.23	23.27	23.37	8.86	17.74	17.98	8.27	24.09	24.11
	$30 + 2q$	8.86	17.74	17.98	8.27	24.09	24.11	7.17	27.05	27.40
	$20 + 3q$	8.27	24.09	24.11	7.17	27.05	27.40	8.16	19.71	20.42
	$10 + 4q$	7.17	27.05	27.40	8.16	19.71	20.42	9.58	29.34	29.52
Avg. improvement		5.37	14.99	15.11	5.72	16.44	16.60	6.15	18.29	18.48

A number of observations can be made from Table 2. First, we observe that optimizing a static point requirement (the seller Q_1) can improve the revenue by more than 5 ~ 6% on average than the seller with the worst point requirement (bench-mark Q_0). Second, as we observe from section 4, the seller’s black-out (open/close) policy continues to show a significant extent of improvements even compared to the seller Q_1 with the optimal fixed point requirement. We find that the additional benefit of dynamic adjustment is surprisingly small. That is, the black-out seller Q_2 achieves almost the same revenues of the seller Q_3 who can further refine point requirement. This implies that, compared to a black-out seller, further segmenting reward sales

with multiple tiers of point requirements only marginally increases the revenue.

Remark 3. Note that the dynamic adjustment model allows the reimbursement rate changes in point requirement. If the reimbursement rate is constant, i.e., $R(q) = R$ for any q , changing point requirement decision only affects the consumers' choice, not the seller's compensation scheme. In this simple extension model, all our results in this section hold without conditions $A4$ and $A5$.

3.6 Conclusion

In this paper we study the consumers' point redemption behavior (redeeming accumulated points to pay for a product) and its consequent effects on the seller's dynamic pricing and inventory rationing decisions. Since the seller earns different amounts of revenue from cash payment and point redemption, the loyalty program and consumer's redemption choice affects the seller's bottom line. For instance, if the reimbursement rate is low (compared to the cash price), the loyalty program hurts the seller's revenue (i.e., cannibalization with a lower priced good). On the other hand, loyalty program can benefit the seller if a consumer who can not afford a cash price purchases with points. We first model the consumer choice behavior, capturing three attributes – reservation price, point balance, and perceived value of a loyalty point. Then, we incorporate this choice model into the seller's problem to see how the seller reacts to the point redemption behavior and its impact on the revenue. We particularly modeled three different cases depending on the level of sellers' discretion over the terms of reward sales: (1) the seller who has no discretion, (2) the seller who can black out reward sales, and (3) the seller who can change both the price and point requirement in each period, and analyze the impact of seller's discretion on the seller's strategy and revenue.

We find that the seller either adds a premium or offers a discount depending

on the gap between the reimbursement rate and the *cash-only price*. For instance, when the reimbursement rate is low relative to the *cash-only price*, the seller offers a discount in order to convert some of reward sales to cash sales. In the opposite case, the seller offers a premium to divert some consumers to reward sales. When a seller can further control the availability of reward sales, we find that it is optimal to block out when the inventory level (relative to the remaining selling season) or the reimbursement rate is very low. Notice that, in both cases, the seller can avoid a deep discount on price by simply closing the reward sales. However, unless the gap between the *cash-only price* and the reimbursement rate is significant, allowing reward sales is generally better for the seller as doing so increases the total sales and revenue. Thus, even when the reimbursement rate is lower than the cash-only sales, allowing reward sales with a discount can be still optimal as long as the benefit from increased sales is significant. In the case where the seller can dynamically change both price (for cash sales) and point requirement (for reward sales), we find that the seller change both the price and point requirement in the same monotone direction (e.g., increase price and point requirement as the inventory level decreases or the remaining time increases). We also find that further segmenting reward sales with multiple tiers of point requirements only marginally increase the revenue from the seller with the option to black-out reward sales.

CHAPTER IV

On (Rescaled) Multi-Attempt Consumer Choice Model and Its Implication on Assortment Optimization

4.1 Introduction

Assortment optimization is one of the most important problems in operations and marketing; it is both mathematically challenging and practically prevalent. Despite a few decades of research on the topic, the pursuit of a new approach that can efficiently solve a general assortment optimization problem that takes into account a wide variety of real-world business constraints is still very vibrant. There are many reasons why assortment optimization is difficult. First, the estimation of customer choice behavior itself is far from trivial — it continues to be one of the most important topics in the academic literature (*Hess and Daly* 2014). Second, even after a customer choice model has been successfully estimated, the resulting model is sometimes difficult to optimize, which limits how a decision maker can operationalize assortment, pricing, and inventory decisions based on the solution of a model. Our work is motivated by these very concerns. Working under the framework of mixed logit model (this assumption is without loss of generality since *McFadden et al.* (2000) show that any random utility model can be approximated to any degree of accuracy by a mixed

logit model), we propose an approximation scheme that improves the approximation quality of the so-called *exogenous* demand model (see below) and show that this approximation can be potentially used to solve a general assortment optimization problem with a wide variety of real-world constraints.

Exogenous demand model is perhaps the most popular choice model used in the operations literature. It assumes that each customer behaves in the following way: when faced with an assortment of products (an offer set), she first looks for her favorite product in the assortment; if this product is not available, she considers a substitute product, and if this substitute product is also not available, she will not purchase anything. Since a customer is only making two attempts when purchasing a product, we also call this 2-attempt model. As explained in *Kök and Fisher (2007)*, the assumption that a customer does not consider further substitutions in her search is not necessarily restrictive, at least in some settings. The strengths of exogenous demand model are obvious: Not only it is intuitively appealing, it also provides a *tractable* estimation and optimization framework. (We are not aware of works that study the theoretical complexity of assortment optimization under exogenous demand model. However, per our experience on running numerical experiments with exogenous demand model, its Mixed Integer Program (MIP) formulation can be solved very efficiently within a few seconds for a reasonable sized problem; see Table 2.) The main weakness of exogenous demand model is that, when many customers are willing to consider more than one substitute, it does not necessarily provide an accurate approximation of the true choice model. Thus, an important research question is how to improve the approximation quality of exogenous demand model without significantly compromising its strengths.

The most relevant advancement to the above question that we are aware of is made recently by *Blanchet et al. (2013)*. (There exist other approximation schemes such as ranking-based model, see *Bertsimas and Mišić (2016)* for literature review;

however, these models are not the focus of our work.) The authors propose an iterative Markov search model where a customer does not stop after the first substitution attempt but continues searching until either she finds a product that she likes within the assortment or she hits the *no-purchase* option. More precisely, they interpret the substitution probability as a transition probability in a Markov chain where both the no-purchase option and the set of products in the assortment act as the absorbing states. The authors show that their proposed model approximates the true choice model well, and they develop a polynomial-time algorithm to solve the corresponding unconstrained assortment optimization problem. Although they do not benchmark the performance of Markov chain model against the exogenous demand model, we show using numerical experiments in Section 2.3 that the former significantly improves the accuracy of the later. Moreover, since Markov chain model requires exactly the same number of parameters as exogenous demand model, it is as tractable as exogenous demand model from the estimation perspective.

The main drawback of Markov chain model is that its corresponding choice probability cannot be explicitly written. This makes it rather difficult for practitioners to use Markov chain model in conjunction with off-the-shelf optimization solvers, especially for the setting of assortment optimization with constraints. As noted in *Bertsimas and Mišić* (2016), firms typically have many business rules that limit the set of possible assortments. To name a few, a firm may have a limited shelf space which dictates that only a finite number of products can be displayed at any time; a firm may require that some products be offered together; and, a firm may also require that only a number of products within a certain category to be offered at any time, etc. There are two typical approaches taken by researchers to solve assortment optimization problem with constraints. The first approach is what *Bertsimas and Mišić* (2016) call the *fix-then-exploit* approach, where the researchers first *fix* a particular choice model and then *exploit* the structure of the resulting assortment problem to

develop either an exact or approximate solution (e.g., *Rusmevichientong et al.* 2010 and *Désir et al.* 2015). The second approach is the so-called *Mixed Integer Optimization* (MIO) approach where an assortment problem is formulated as an MIP (or its variant) and then solved using an off-the-shelf MIP solver; this approach typically requires that the corresponding choice probability can be explicitly written. Note that while the fix-then-exploit approach allows researchers to develop a highly efficient algorithm for a specific model, the MIO approach is highly flexible in the sense that no problem-specific effort to develop a specialized algorithm is required and practitioners can simply *declare* their constraints to the solver. Since the approximate choice probability under Markov chain model cannot be explicitly written, a specific algorithm needs to be developed to solve a constrained assortment optimization problem under Markov chain model. Indeed, this is the approach taken by *Feldman and Topaloglu* (2014) and *Désir et al.* (2015), where the authors focus on specific forms of constraints. In particular, *Feldman and Topaloglu* (2014) develop a linear programming algorithm in the context of the network revenue management problem, and *Désir et al.* (2015) develop constant factor approximations for assortment optimization problem with the cardinality and capacity constraints. In contrast to this, constrained assortment optimization problems under exogenous demand model, at least for the types of constraint discussed above, can be easily formulated as a Mixed Integer Linear Program (MILP) and solved using an off-the-shelf solver. (Per our numerical experiments in Section 3, the resulting MILP can be solved within a few seconds for a reasonable sized problem.)

Our contribution. In this work, we wish to bridge the gap between the classical exogenous demand model and the recently introduced Markov chain model. The central question we ask is whether it is possible to improve the approximation quality of exogenous demand model without sacrificing its tractability and versatility in dealing with real-world constraints. We are particularly interested in a type of approximation

whose corresponding choice probability can be explicitly written as it allows practitioners to simply use off-the-shelf optimization solvers to solve a variety of constrained assortment problems without having to develop a specific algorithm for a specific set of constraints. Thus, our work shares the same spirit as the recent work of *Bertsimas and Mišić* (2016). (Our work differs from theirs in that they use a ranking-based approximation whereas we use a new multi-attempt approximation.) We first study the approximation quality of a natural generalization of exogenous demand model, called *multi-attempt* model. To be precise, assuming that all customers are willing to consider at most $k - 1$ substitutes, how much improvement does this extra flexibility give, in general, as a function of k ? We show that the approximation error of multi-attempt model relative to the true choice probability decreases exponentially in k . This confirms our intuition that capturing higher substitution dynamics leads to a better approximation. Unfortunately, while the theoretical bound of multi-attempt model is encouraging, its empirical performance is somewhat discouraging as it heavily depends on the number of products n . (Per our results in Table 4.1, for $n = 10$, 4-attempt model is better than Markov chain model; for $n = 100$, even 5-attempt model is still a lot worse than Markov chain model. This is not satisfactory because k -attempt model with $k \geq 3$ requires a lot more parameters than Markov chain model.) Upon a closer examination, however, it turns out that multi-attempt model consistently underestimates the true choice probability, which leads to its poor empirical performance. This motivates us to construct a modified multi-attempt model, which we call the *rescaled multi-attempt* model. The idea is to start with the original k -attempt model and then re-scale it with a non-constant factor to make the sum of probability equals one. The proposed re-scaling significantly improves the performance of the original multi-attempt model: If the true choice model is Multinomial Logit (MNL), we show that the rescaled k -attempt model is *exact* for all $k \geq 1$ (this result is reminiscent of the result in *Blanchet et al.* (2013) that Markov chain model is

exact for MNL); if, on the other hand, the true choice model is not MNL, our numerical experiments show that the approximation quality of the rescaled 2-attempt model is very close to Markov chain model and the approximation quality of the rescaled 3-attempt consistently dominates the Markov chain model (see Table 4.1).

Both the rescaled 2-attempt and the Markov chain models share exactly the same number of parameters; and yet, the corresponding choice probability under the rescaled 2-attempt model can be explicitly written. This allows us to more easily formulate an assortment optimization problem with constraints. In Section 3, we show that the resulting constrained assortment optimization problems (with typical constraints discussed before) under the rescaled 2-attempt model can be written as a Mixed Integer Linear Fractional Program (MILFP). Although MILFP in general is difficult to solve, many important problems in engineering and science can be formulated as MILFPs; these have motivated intensive researches in the scientific community to develop efficient methods (both exact and approximate) for solving large-scale MILFPs (e.g., *Tawarmalani and Sahinidis* 2002 and *Yue et al.* 2013). On another note, the MILFP formulation of assortment optimization under the rescaled 2-attempt model can be equivalently transformed into a 0-1 quadratic programming. Again, although 0-1 quadratic programming is in general difficult to solve (i.e., from theoretical complexity perspective), we do have a 50-year deep of literature on the topic of approximation algorithm for 0-1 quadratic programming (e.g., *Kochenberger et al.* 2014). Thus, we are not lacking of sophisticated algorithms that can be used to solve the resulting assortment problem under our proposed approach. Indeed, this is another advantage of having an explicit expression of approximate choice probability as it allows us to borrow tools from existing literature in addition to using off-the-shelf solvers. For the purpose of numerical illustrations, in this work, we will only focus on one approach, the so-called *Dinkelbach algorithm*. We discuss this in more detail in Section 3.

4.2 Choice Approximation Models

In this section, we describe both the *multi-attempt* and the *rescaled multi-attempt* models. In addition, we also provide results from numerical experiments to compare the approximation accuracy of these models with Markov chain model. We denote the universe of n products by the set $\mathcal{N} = \{1, \dots, n\}$ and the no-purchase alternative as product 0. Since *McFadden et al.* (2000) show that any random utility choice model can be approximated by a mixture of Multinomial Logits (MNLs) at any degree of accuracy, we will assume that the underlying true model is a mixture of M MNL models. Let θ_m , $m = 1, \dots, M$, denote the probability that a random customer belongs to segment m (by construction, we must have $\theta_1 + \dots + \theta_M = 1$) and let the MNL parameters for segment m be denoted by $u_{im} \geq 0$ for $i \in \mathcal{N}_0 = \mathcal{N} \cup \{0\}$ and $m = 1, \dots, M$. Then, for any offer set $S \subset \mathcal{N}$, the true choice probability of product $i \in S_0 := S \cup \{0\}$ is given by

$$\pi(i, S) = \sum_{m=1}^M \theta_m \frac{u_{im}}{\sum_{j \in S_0} u_{jm}}.$$

4.2.1 Multi-Attempt Model

Per our discussions in Section 1, under the k -attempt model, each customer considers up to $k - 1$ substitutes, beyond her favorite product, before she decides to not purchase anything. To illustrate, suppose that $\mathcal{N} = \{1, 2, 3, 4\}$ and $S = \{1, 2\}$. Under 2-attempt model, a customer will purchase product 1 if either (1) it is her favorite product among all four products and it is preferred to the no-purchase alternative, or (2) she likes either product 3 or 4 best but unfortunately neither of these is included in S and her next favorite product is 1. Let U_i denote the utility of product i and let \bar{S} denote the complement of S . Mathematically, we can write the probability that a customer will purchase product 1 as follows: $\hat{\pi}_2(1, S) = P(U_1 > \max\{U_0, U_2, U_3, U_4\}) + P(U_3 > U_1 > \max\{U_0, U_2, U_4\}) + P(U_4 > U_1 > \max\{U_0, U_2, U_3\}) := \lambda_1 + \lambda_{31} + \lambda_{41}$.

Note that this choice probability is the same as the choice probability under the classic exogenous demand model. Similarly, under 3-attempt model, a customer will purchase product 1 if either (1) it is her favorite product among all four products and the no-purchase alternative, or (2) it is her second favorite product after either product 3 or 4, or (3) it is her third favorite product after both products 3 and 4. We can write the probability that a customer will purchase product 1 as follows:

$$\begin{aligned}
\hat{\pi}_3(1, S) &= P(U_1 > \max\{U_0, U_2, U_3, U_4\}) \\
&\quad + P(U_3 > U_1 > \max\{U_0, U_2, U_4\}) + P(U_4 > U_1 > \max\{U_0, U_2, U_3\}) \\
&\quad + P(\min\{U_3, U_4\} > U_1 > \max\{U_0, U_2\}) \\
&:= \lambda_1 + \lambda_{31} + \lambda_{41} + \lambda_{\{3,4\}1}.
\end{aligned}$$

More generally, given a set of products \mathcal{N} and an offer set S , the probability that a customer will purchase product $i \in S_0$ under k -attempt model is given by

$$\hat{\pi}_k(i, S) = \lambda_i + \sum_{j_1 \in \bar{S}} \lambda_{j_1 i} + \sum_{\{j_1, j_2\} \subseteq \bar{S}} \lambda_{\{j_1, j_2\} i} + \cdots + \sum_{\{j_1, j_2, \dots, j_{k-1}\} \subseteq \bar{S}} \lambda_{\{j_1, j_2, \dots, j_{k-1}\} i},$$

where $\lambda_{\{j_1, j_2, \dots, j_{k-1}\} i}$ is the probability that a customer values product $j \in \{j_1, j_2, \dots, j_{k-1}\}$ better than i and product $j' \in \mathcal{N} - \{j_1, j_2, \dots, j_{k-1}\} \cup \{0\}$ worse than i . That is,

$$\begin{aligned}
\lambda_{\{j_1, j_2, \dots, j_{k-1}\} i} &= P\left(\min\{U_{j_1}, \dots, U_{j_{k-1}}\} > U_i \right. \\
&\quad \left. \geq \max\{U_l : l \in \mathcal{N} \setminus \{j_1, j_2, \dots, j_{k-1}\} \cup \{0\}\}\right).
\end{aligned}$$

Since a customer makes a purchase as soon as her next favorite product is in S , she only needs to consider at most $|\bar{S}|$ substitutes (beyond her most favorite product) before making a purchase. This means that, under multi-attempt model, we must have: $\hat{\pi}_{|\bar{S}|+1}(i, S) = \pi(i, S)$ for all $i \in S_0$. Moreover, by construction, we also have

$\hat{\pi}_k(i, S) < \pi(i, S)$ for all $k < |\bar{S}| + 1$ and $i \in S_0$.

Error bound for multi-attempt model. We now derive an error bound for k -attempt model. Let $u_{max}(\bar{S})$ be the maximum probability that the most favorite product of a random customer from any segment $m = 1, \dots, M$ is included in a compliment of offer set S , $\bar{S} := \mathcal{N} \setminus S \cup \{0\}$. That is, $u_{max}(\bar{S}) = \max_m \sum_{i \in \bar{S}} u_{im}$. The following theorem tells us that the relative error of multi-attempt model decreases exponentially with the number of attempts k .

Theorem IV.1. *For any $S \subset \mathcal{N}$ and $i \in S_0$, we have:*

$$(1 - u_{max}(\bar{S})^k) \cdot \pi(i, S) \leq \hat{\pi}_k(i, S) \leq \pi(i, S). \quad (4.1)$$

Note that multi-attempt model best approximates the true choice model when $u_{max}(\bar{S})$ is small. Intuitively, this is likely to happen when S is large. As mentioned in Section 1, although the theoretical bound of multi-attempt model is encouraging, we will show that its empirical performance is not satisfactory: see numerical results in Table 4.1. This motivates us to construct a modified multi-attempt model, called the rescaled multi-attempt model which we discuss next.

4.2.2 The Rescaled Multi-Attempt Model

Under the rescaled k -attempt model, we approximate $\pi(i, S)$ with $\hat{\pi}_k^R(i, S)$ defined below:

$$\hat{\pi}_k^R(i, S) = \frac{\hat{\pi}_k(i, S)}{\sum_{j \in S_0} \hat{\pi}_k(j, S)}.$$

Two comments are in order. First, since the rescaled k -attempt model uses the k -attempt model as its primitive, they share the same set of parameters. In particular, all three models – the 2-attempt model, the rescaled 2-attempt model, and the Markov

chain model – share exactly the same set of parameters. Second, the rescaled 1-attempt model is identical to MNL approximation. Thus, if the underlying true choice model is MNL (i.e., there is only 1 segment of customer), the rescaled 1-attempt model is exact.

Analogous to $u_{max}(\bar{S})$, we define $u_{min}(\bar{S})$, the minimum probability that the most favorite product of a random customer from any segment $m = 1, \dots, M$ is included in \bar{S} . That is, $u_{min}(\bar{S}) = \min_m \sum_{i \in \bar{S}} u_{im}$. The following result is an immediate corollary of Theorem 1.

Corollary IV.2. *For any $S \subset \mathcal{N}$ and $i \in S_0$, we have:*

$$\frac{1 - u_{max}(\bar{S})^k}{1 - u_{min}(\bar{S})} \cdot \pi(i, S) \leq \hat{\pi}_k^R(i, S) \leq \frac{1}{1 - u_{max}(\bar{S})} \cdot \pi(i, S). \quad (4.2)$$

While multi-attempt model consistently underestimates the true choice probability, the rescaled multi-attempt model may sometimes overestimate the true probability. Note that the lower bound in Corollary IV.2 is larger than the lower bound in Theorem IV.1. This suggests that the rescaled multi-attempt model improves the underestimation error while admitting the overestimation error. The important question is whether this is a good compromise overall. Our numerical results in Table 4.1 show that the rescaled multi-attempt model significantly improves the empirical accuracy of multi-attempt model. Theoretically, we are also able to show the exactness of the rescaled multi-attempt model when the true choice probability is MNL. This result is reminiscent of the result in *Blanchet et al. (2013)* that Markov chain model is exact in the case of MNL.

Lemma IV.3. *Suppose that the underlying true choice model is an MNL. For any $k > 0$, $S \subset \mathcal{N}$ and $i \in S_0$, the rescaled k -attempt model is exact, i.e., $\hat{\pi}_k^R(i, S) = \pi(i, S)$ for any $k > 0$.*

4.2.3 Numerical Experiments

We conduct numerical experiments with respect to a mixture of M MNLs to compare the performance of the multi-attempt, the rescaled multi-attempt, and the Markov chain (MC) models. Let n denote the number of products and M denote the number of customer segments in the MMNL model. For a fixed combination of number of products ($n = 10, 20, 50, 100$) and number of segments ($M = 3, 5, 10, 20$), we generate 100 instances. The probability distribution over different MNL segments, $\theta_1, \dots, \theta_M$, are first generated using i.i.d samples of the uniform distribution in $[0, 1]$ and then normalized such that $\theta_1 + \dots + \theta_M = 1$. For each segment $m = 1, \dots, M$, the MNL parameters of segment m , u_{0m}, \dots, u_{nm} are randomly sampled from the uniform distribution in $[0, 1]$. For each instance, we generate a random offer set of size between $n/3$ and $2n/3$, and compute the choice probabilities under the three models. We report both the average and maximum relative errors defined as: $\text{avg.Error} = \frac{1}{400} \sum_{a=1}^{400} \text{Error}(S_a)$ and $\text{max.Error} = \max_{1 \leq a \leq 400} \text{Error}(S_a)$, where $\text{Error}(S) = 100\% \cdot \max_{i \in S} \frac{|\hat{\pi}(i, S) - \pi(i, S)|}{\pi(i, S)}$. The results can be seen in Table 4.1.

Table 4.1: Comparison of approximation accuracy of various models

		MC	k -attempt					rescaled k -attempt				
			$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$
n=10	max.Error	11.69	72.06	48.89	30.42	16.58	7.20	21.95	15.44	9.68	5.10	1.95
	avg.Error	2.45	48.36	21.47	8.58	2.99	0.82	5.64	3.35	1.70	0.71	0.22
n=20	max.Error	8.40	68.55	45.76	29.59	18.43	10.96	17.46	11.26	6.64	4.14	2.77
	avg.Error	2.13	50.82	25.02	11.91	5.47	2.40	4.56	2.91	1.72	0.95	0.49
n=50	max.Error	8.44	63.41	39.76	24.64	15.07	9.10	13.99	10.44	7.41	5.02	3.24
	avg.Error	1.76	50.68	25.58	12.84	6.41	3.18	3.66	2.40	1.51	0.91	0.53
n=100	max.Error	3.17	61.17	37.16	22.41	13.42	7.97	6.03	4.23	2.86	1.87	1.18
	avg.Error	1.35	51.10	26.18	13.44	6.90	3.55	2.30	1.51	0.96	0.59	0.35

A number of observations can be made from Table 4.1. First, although the accuracy of multi-attempt model improves as k increases, its rate of improvement is not satisfactory. For example, when $n = 100$, the average relative error of 5-attempt model is 3.55%; in contrast, the average relative error of Markov chain model is only about 1.35%. Considering the fact that 5-attempt model requires much more param-

eters than Markov chain model, this level of performance is not acceptable. Second, the rescaled 2-attempt model significantly improves the accuracy of 2-attempt model and its relative error is very close to the relative error of Markov chain model. Moreover, the rescaled 3-attempt consistently performs better than Markov chain model. This highlights the benefit of re-scaling.

4.3 Assortment Optimization

We now discuss how to use the rescaled multi-attempt model in assortment optimization. Since the approximation quality of the rescaled 2-attempt model is very close to Markov chain model, in this work, we will only focus our discussions on the rescaled 2-attempt model. (Our approach for the rescaled 2-attempt model is also generalizable to re-scaled k -attempt model.) We show that assortment optimization under the rescaled 2-attempt model is not much harder than assortment optimization under exogenous (2-attempt) demand model. In particular, it can be formulated as a Mixed Integer Fractional Linear Program (MILFP) and can be solved using the so-called Dinkelbach algorithm.

4.3.1 Optimization Formulation

Let r_i denote the revenue of product i and $x_i \in \{0, 1\}$ be a binary decision variable for product i . We first consider unconstrained assortment optimization problem under exogenous demand model. This can be written as a Mixed Integer Linear Program

(MILP) formulation below:

$$\begin{aligned}
J_2 &= \max_{\vec{x} \in \{0,1\}^n} \sum_{i=1}^n r_i \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot x_i \cdot (1 - x_j) \right] \\
&= \max_{\substack{\vec{x} \in \{0,1\}^n \\ \vec{y} \in [0,1]^{n(n-1)}}} \sum_{i=1}^n r_i \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot y_{ji} \right] \\
&\quad \text{s.t. } y_{ji} \leq x_i, \quad y_{ji} \leq 1 - x_j, \quad y_{ji} \geq x_i - x_j \quad \forall i \neq j
\end{aligned}$$

We next consider unconstrained assortment optimization under the rescaled 2-attempt model:

$$\begin{aligned}
J_2^R &= \max_{\vec{x} \in \{0,1\}^n} \sum_{i=1}^n \frac{r_i \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot x_i \cdot (1 - x_j) \right]}{\lambda_0 + \sum_{j \neq 0} \lambda_{j0} \cdot (1 - x_j) + \sum_{i=1}^n \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot x_i \cdot (1 - x_j) \right]} \\
&= \max_{\substack{\vec{x} \in \{0,1\}^n \\ \vec{y} \in [0,1]^{n(n-1)}}} \sum_{i=1}^n \frac{r_i \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot y_{ji} \right]}{\lambda_0 + \sum_{j \neq 0} \lambda_{j0} \cdot (1 - x_j) + \sum_{i=1}^n \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot y_{ji} \right]} \\
&\quad \text{s.t. } y_{ji} \leq x_i, \quad y_{ji} \leq 1 - x_j, \quad y_{ji} \geq x_i - x_j \quad \forall i \neq j
\end{aligned}$$

Note that J_2^R is an MILFP. As discussed in Section 1, although MILFP is in general difficult to solve, it appears in many applications in engineering and science (*Tawarmalani and Sahinidis 2002*). Consequently, there is a deep and ever-growing literature on different algorithmic approaches to solve MILFP, either exactly or approximately. When it comes to large-scale MILFP, one popular approach is based on Dinkelbach algorithm, first developed in *Dinkelbach (1967)*. In the context of our assortment problem above, Dinkelbach algorithm works as follows. First, we define $N(\vec{x}, \vec{y}) = \sum_{i=1}^n r_i \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot y_{ji} \right]$ and $D(\vec{x}, \vec{y}) = \lambda_0 + \sum_{j \neq 0} \lambda_{j0} \cdot (1 - x_j) + \sum_{i=1}^n \left[\lambda_i \cdot x_i + \sum_{j \neq i} \lambda_{ji} \cdot y_{ji} \right]$, and let $F(q) = \max\{N(\vec{x}, \vec{y}) - q \cdot D(\vec{x}, \vec{y}) : (\vec{x}, \vec{y}) \in A^*\}$, where A^* is the set of feasible (\vec{x}, \vec{y}) . Now, we proceed in three steps:

Step 1. Choose an arbitrary feasible (\vec{x}^1, \vec{y}^1) , set $q_2 = \frac{N(\vec{x}^1, \vec{y}^1)}{D(\vec{x}^1, \vec{y}^1)}$, and let $t = 2$

Step 2. Compute $F(q_t)$ and denote its optimal solution as (\bar{x}^t, \bar{y}^t) .

Step 3. If $F(q_t) \leq \epsilon$ (optimality tolerance), stop and output (\bar{x}^t, \bar{y}^t) ;

Otherwise, let $q_{t+1} = \frac{N(\bar{x}^t, \bar{y}^t)}{D(\bar{x}^t, \bar{y}^t)}$, set $t = t + 1$, and go back to Step 2.

Note that computing $F(q_t)$ in Step 2 requires solving an MILP with similar size as J^2 . So, the running time of Dinkelbach algorithm approximately equals the running time for solving J^2 multiplies the number of iterations for $F(q_k)$ to be sufficiently close to 0. It has been shown that $F(q_k) \rightarrow 0$ at a super-linear rate (You et al. 2009); in fact, when all the variables are binary, in the worst case scenario, Dinkelbach algorithm only requires about $\log(\text{number of variables})$ iterations. This highlights the practicality of Dinkelbach algorithm for solving MILFP, especially when the corresponding inner optimization can be quickly solved.

Dealing with constraints. Our optimization model can further accommodate a variety of constraints on the assortment. For example, the following types of constraint from *Bertsimas and Mišić* (2016) can be easily included: (1) At most U products can be chosen from a subset of size B (maximum subset, also called as cardinality constraints); (2) the number of offered products from a subset of size B cannot be greater than that from the other subset of size B (precedence type 1); (3) a specific product must be offered to include any product from a subset of size $B - 1$ (precedence type 2). Any of these constraints can be formulated as a linear constraint, and adding linear constraints still results in an MILFP (under the rescaled 2-attempt model). Thus, we can still use Dinkelbach algorithm.

4.3.2 Numerical Experiments

To compare the performance of the multi-attempt choice model in assortment optimization, we conduct numerical experiments using the same random instances of the mixture of M MNLs as in Section 2.3. In addition, we also generate a random

number between 0 and 1 for the revenue of each product(i.e., r_i for product i). We then compute the optimal assortment under the Markov chain, 2-attempt, and rescaled 2-attempt models, and calculate the expected revenue of each solution under the true choice model. Table 4.2 summarizes the average relative gap in expected revenue from the true optimal revenue, including the average running time, for each model. We note that all the computational experiments are carried out on a Mac with Intel Core i5 @ 2.7 GHz and 16-GB RAM. All models and solution procedures are coded in Matlab 2011 and the MILP problems in the proposed algorithm are solved using CPLEX 12 with optimality tolerance of 10^{-5} .

Table 4.2: Average relative gap in expected revenue for Markovian model and multi-attempt models with its computing time in second.

		Markov Chain		2-attempt		rescaled 2-attempt	
		gap(%)	time(s)	gap(%)	time (s)	gap(%)	time (s)
n=10	$M = 3$	0.0243	0.0011	9.5785	0.0062	0.0723	0.0193
	$M = 5$	0.1077	0.0009	9.5440	0.0061	0.1246	0.0196
	$M = 10$	0.0601	0.0009	9.4490	0.0064	0.0998	0.0201
n=20	$M = 6$	0.0558	0.0028	14.6167	0.0260	0.0751	0.0859
	$M = 10$	0.0217	0.0027	14.5592	0.0252	0.0217	0.0834
	$M = 20$	0.0478	0.0029	14.6049	0.0250	0.0492	0.0859
n=50	$M = 10$	0.0380	0.0134	21.2468	0.2997	0.0498	1.2327
	$M = 20$	0.0380	0.0134	21.2468	0.2997	0.0498	1.2327
	$M = 50$	0.0069	0.0169	20.7945	0.2882	0.0080	1.2406
n=100	$M = 10$	0.0218	0.0407	24.5002	3.3212	0.0234	14.8783
	$M = 20$	0.0179	0.0439	23.8145	3.2690	0.0186	14.6535
	$M = 50$	0.0044	0.0555	24.4001	3.2699	0.0045	14.5374

Observe that re-scaling significantly improves the performance of 2-attempt model. Moreover, the difference between the relative gap of Markov chain model and rescaled 2-attempt model is negligible. As expected, assortment optimization under Markov chain model can be solved extremely quickly. Although the running time of assortment optimization under the rescaled 2-attempt model is not as short as the running time under Markov chain model, it is nevertheless still quite tractable. Note that the running time under 2-attempt model is only about 3 seconds for $n = 100$. In the

case of the rescaled 2-attempt, we use about 5 iterations in the Dinkelbach algorithm, which explains the approximate running time of 15 seconds for $n = 100$. The number of iterations in Dinkelbach algorithm is dictated by the optimality tolerance ϵ (see Step 3). Practically, by adjusting the desired optimality level, one can further reduce the running time under the rescaled 2-attempt model.

Table 4.3: Average relative gap in expected revenue for constrained (non-scaled and rescaled) 2-attempt models with its computing time in second.

		2-attempt		rescaled 2-attempt	
		gap(%)	time(s)	gap(%)	time(s)
n = 10, M = 5	No constraints	9.5541	0.0068	0.1218	0.0191
	Max.subset, C = 2, B = 5, U = 3	7.8698	0.0077	0.1250	0.0203
	Prec.type 1, C = 1, B = 5	9.6111	0.0076	0.1392	0.0210
	Prec.type 2, C = 2, B = 5	7.1255	0.0073	0.0846	0.0218
n = 20, M = 10	No constraints	14.8914	0.0268	0.0440	0.0839
	Max.subset, C = 4, B = 5, U = 3	12.2535	0.0387	0.0420	0.0944
	Prec.type 1, C = 3, B = 5	13.7742	0.0373	0.0534	0.0959
	Prec.type 2, C = 4, B = 5	11.2007	0.0334	0.0668	0.1110
n = 50, M = 20	No constraints	21.3102	0.2994	0.0229	1.2814
	Max.subset, C = 5, B = 10, U = 5	14.6982	0.4824	0.0209	1.3596
	Prec.type 1, C = 4, B = 10	20.3282	0.3591	0.0154	1.3101
	Prec.type 2, C = 5, B = 10	17.1013	0.3500	0.0148	1.5295
n = 100, M = 20	No constraints	24.3847	3.3763	0.0111	15.2983
	Max.subset, C = 10, B = 10, U = 5	18.2043	5.5751	0.0119	16.9237
	Prec.type 1, C = 9, B = 10	24.7584	4.4081	0.0125	16.8627
	Prec.type 2, C = 10, B = 10	20.0725	3.8482	0.0143	18.3797

Constrained problem. To see the effect of constraints in optimization performance, we solve the optimization instances that we used in Table 2 with a combination of constraints that we discussed in Section 3.1. For each constraint set, we create C constraints by randomly partitioning a set of n products into (mutually exclusive) subsets of size B . The average relative gap and computing time for each constraint are summarized in Table 3. We confirm that the constrained problems also can be solved quickly (less than 20 seconds when $n = 100$) and its performance is very close to the true optimal performance (less than 0.2% of relative average gap).

4.4 Conclusion

In this work, we provide a new approach to approximate a general mixed-logit-based choice model. We show that the classic exogenous demand model can be significantly improved by re-scaling. The resulting approximation is exact for MNL and has an empirical performance that is very close to the performance of the recently developed Markov chain model. Moreover, since the proposed approximation model has an explicit mathematical expression, it can be immediately used in an assortment optimization with a variety real-world constraints. Our numerical experiments show that our model is quite tractable for a reasonable sized problem.

CHAPTER V

Conclusions

To summarize the dissertation, we study consumer behavior and its effects on the firm's operations management strategies. We consider various aspects of consumer behavior – strategic joining behavior, loyalty point redemption, and substitution behavior – in the contexts of the capacity investment, pricing adjustment, inventory rationing, and product line selection problems. Incorporating such behavior into the traditional operations management problem, the dissertation extends insights about demand or market reactions to the firm's decisions, and highlights directions on how to properly fix the firm's strategies accordingly.

In Chapter II, we study the customers' strategic joining behavior in a service system. In particular, we consider three scenarios for consumers – selfish, collective, and social – for their decisions to join the service system after observing the queue length. Depending on how consumers join, the firm chooses the service rate, which also depends on their types – profit-maximizing and welfare-maximizing. We find that there can be a “*Benefit of Anarchy*” in contrast to the price of anarchy literature. That is, the utility of selfish (non-strategic) customers can be higher than that of collective customers who are coordinated to maximize their average utility. In particular, this happens when the profit-maximizing firm endogenously chooses the lower service rate to collective customers knowing that collective customers join less

to resolve the unnecessary traffic cost. Furthermore, we also find that the self-interest behavior can generate a greater welfare of the system. Despite that we consider a very simple queueing context, this work confirms the interesting effects of strategic (or selfish) behavior of consumers and the firm on the system performance through their interactions in the equilibrium.

In Chapter III, we study the consumers' redemption behavior of loyalty points and its impacts on the seller's pricing and inventory rationing strategies. In widely used loyalty programs, consumers actively manage and redeem their accumulated loyalty points to acquire a product or service. While consumers strategically compare when to use loyalty points instead of paying the posted price, the seller accrues different amount of revenues depending on these consumers' payment decisions. Taking account for its effect on the seller's revenue, we characterize how the seller should adjust the posted price and manage the terms of reward sales. In particular, we find that the seller should either add a premium or offer a discount critically depending on whether the seller receives a higher or lower revenue from reward sales compared to when there are no reward sales. We also find that the seller blocks out the availability of point redemption (or increases the point requirement) when the seller accrues the significantly lower revenue from reward sales.

In Chapter IV, we study the consumers' substitution behavior when their preferred product is not available, and its implications on the seller's assortment optimization problem. Given the combinatorial properties of the problem itself, the estimation and optimization regarding these substitutions are mathematically challenging but also very important. Motivated by the classic exogenous demand model and the recently developed Markov chain model, we develop a new approximation to the random utility-based customer choice model called *rescaled multi-attempt* model, in which a customer may consider several substitutes before finally deciding to not purchase anything. The rescaled multi-attempt model provides an explicit form of the

approximate choice probability. From a practical perspective, this allows the decision maker to use off-the-shelf solver to solve a general assortment optimization problem with a variety of real-world constraints.

APPENDICES

APPENDIX A

Proofs and Technical Details for Chapter II: Effect of Strategic Behavior In a Service System

Proof of Proposition II.2. (Comparison of three thresholds)

(i) Suppose that $n^I(\mu) < n^C(\mu)$ for some μ . Then, any collective customer who joins the queue of length greater than $n^I(\mu)$ receives negative utility from service. Moreover, such a customer can only decrease the utility of other customers. We therefore have a contradiction to the optimality of $n^C(\mu)$ for customer utility.

That $n^C(\mu) \leq n^S(\mu)$ is trivially true because the definitions of $n^C(\mu)$ and $n^S(\mu)$ share the same function, $f(\mu, n)$, on the LHS of the condition which is increasing in n , and the RHS for the collective threshold is less than for the social threshold ($b \leq b_0$).

(ii) In this proof, we relax the integer constraint on the thresholds, and use the unrounded thresholds, $\tilde{n}^I(\mu)$ and $\tilde{n}^S(\mu)$. We want to show that there exists an interval (μ_L, μ_H) such that the unrounded individual threshold is larger than the unrounded social threshold if and only if the service rate falls in this interval, i.e., $\tilde{n}^I(\mu) \geq \tilde{n}^S(\mu)$ if and only if $\mu \in (\mu_L, \mu_H)$. This is equivalent to the following:

$$\frac{\tilde{n}^I \mu^{\tilde{n}^I(\mu)+1} - (\tilde{n}^I(\mu) + 1) \mu^{\tilde{n}^I(\mu)} + 1}{\mu^{\tilde{n}^I(\mu)} (\mu - 1)^2} \geq b_0 \quad \text{if and only if } \mu \in (\mu_L, \mu_H). \quad (\text{A.1})$$

After plugging in $\tilde{n}^I(\mu) = b\mu$, (A.1) for $\mu \neq 1$ is equivalent to

$$g(\mu) := (b_0 - b)(\mu - 1) \left(\mu - \frac{b_0}{b_0 - b} \right) + 1 - \mu^{-b\mu} \leq 0 \quad \text{if and only if } \mu \in (\mu_L, \mu_H). \quad (\text{A.2})$$

The first derivative of $g(\mu)$ with respect to μ is

$$g'(\mu) = (b_0 - b) \left(2\mu - \frac{2b_0 - b}{b_0 - b} \right) + b\mu^{-b\mu}(1 + \log \mu). \quad (\text{A.3})$$

Note that $g(0) > 0$, $\lim_{\mu \rightarrow \infty} g(\mu) > 0$, $g(1) = 0$, and $g'(1) = 0$. Thus, (A.2) will hold if we can show that $g'(x)$ has at most 3 zeros including one at $\mu = 1$. We have that $g'(\mu) = 0$ is equivalent to

$$(b_0 - b) \left(2\mu - \frac{2b_0 - b}{b_0 - b} \right) = -b\mu^{-b\mu}(1 + \log \mu). \quad (\text{A.4})$$

A number of observations about (A.4):

- (i) The LHS is an increasing linear function in μ because $b_0 \geq b$.
- (ii) The RHS is positive only if $\mu < \frac{1}{e}$ and negative otherwise.
- (iii) The first derivative of the RHS with respect to μ is

$$b\mu^{-1-b\mu}(b\mu(1 + \log \mu)^2 - 1).$$

Thus, there exists a unique μ_a such that $b\mu_a(1 + \log \mu_a)^2 = 1$, and the RHS decreases in μ if $\mu < \mu_a$, and increases to 0 if $\mu > \mu_a$.

- (iv) The second derivative of the RHS with respect to μ is

$$b\mu^{-2-b\mu} (1 + 3b\mu(1 + \log \mu) - b^2\mu^2(1 + \log \mu)^3).$$

There exists μ_b such that the RHS is convex if $\frac{1}{e} < \mu < \mu_b$ and concave if $\mu > \mu_b$.

(v) We have $\mu_b > \mu_a$ because the function is convex around the minimum at μ_a .

Then, we have that the RHS is convex decreasing if $\frac{1}{e} < \mu < \mu_a$, convex increasing if $\mu_a < \mu < \mu_b$, and concave increasing if $\mu > \mu_b$. Depending on the values of μ_a , μ_b , and 1, we have the following three cases:

(vi-a) $\mu_a \leq \mu_b \leq 1$: If the slope of the LHS is greater than that of the RHS (i.e., $2(b_0 - b) > b\mu^{-2-b\mu}(1 + 3b\mu(1 + \log \mu) - b^2\mu^2(1 + \log \mu)^3)$), we can have at most three solutions to (A.4) where all the non-trivial solutions are smaller than 1, i.e., $\mu_1 \leq \mu_2 \leq \mu_3 = 1$. If the slope of the LHS is smaller, we also have three solutions ($\mu_1 \leq \mu_2 = 1 \leq \mu_3$).

(vi-b) If $\mu_a \leq 1 \leq \mu_b$, we have the three solutions to (A.4): $\mu_1 \leq \mu_2 = 1 \leq \mu_3$.

(vi-c) If $1 \leq \mu_a \leq \mu_b$, we can have at most three solutions to (A.4): $\mu_1 = 1 \leq \mu_2 \leq \mu_3$.

We have shown that $g'(\mu)$ has at most three zeros, which implies that $g(\mu)$ has at




Figure A.1: Illustration of the LHS and RHS of (A.4) where $\mu = 1$ is the largest solution (left: $b = 5, b_0 = 16$), the middle solution (middle: $b = 3, b_0 = 4$), or the smallest solution (right: $b = 1, b_0 = 1.25$).

most three stationary points. If there are two different intervals of μ within which $g(\mu) \leq 0$, this is only possible if $g(\mu_1) < 0$, $g(\mu_2) > 0$ and $g(\mu_3) < 0$ (i.e., $g(\mu)$ has a shape of W). Given that one of the three stationary points is $\mu = 1$ and $g(1) = g'(1) = 0$, it follows that one of the $g(\mu_1)$, $g(\mu_2)$, and $g(\mu_3)$ equals zero. Thus, there exists a unique interval (μ_L, μ_H) such that (A.2) holds for $\mu \neq 1$, and $\tilde{n}^S(\mu) \geq \tilde{n}^I(\mu)$ if and only if $\mu \neq 1$ is within that interval. This further implies that, after applying the integer constraint on the threshold, there exists an interval within which $n^S(\mu) \geq n^I(\mu)$, otherwise $n^S(\mu) \leq n^I(\mu)$. ■

Proof of Corollary II.3. (Comparison of Π , \mathcal{U} , \mathcal{W} by customer type)

(i) The effective arrival increases in n , thus, the profit function increases in n .

(ii) For given μ , the customer utility function is discretely unimodal in n . Given that $n^C(\mu)$ is a unique optimal threshold, it follows that customer utility increases in n if $n < n^C(\mu)$, and decreases if $n > n^C(\mu)$.

(iii) and (iv): As with (ii), the welfare function is discretely unimodal in n for given μ while $n^S(\mu)$ achieves the optimum. ■

Proof of Lemma II.4. (Profit function properties)

(i) The profit function for given n is $\Pi(\mu, n) = p\lambda_e(\mu, n) - c(\mu)$ where the effective arrival rate (λ_e) is given by

$$\lambda_e(\mu, n) = \begin{cases} 1 - \frac{\mu-1}{\mu^{n+1}-1}, & \mu \neq 1 \\ \frac{n}{n+1}, & \mu = 1. \end{cases} \quad (\text{A.5})$$

For any $\mu \neq 1$, the effective arrival rate is increasing in μ for given $n > 0$:

$$\begin{aligned} \frac{\partial \lambda_e(\mu, n)}{\partial \mu} &= \frac{n\mu^{n+1} - (n+1)\mu^n + 1}{(\mu^{n+1} - 1)^2} \\ &= \frac{(\mu - 1) \left\{ n\mu^n - \frac{\mu^n - 1}{\mu - 1} \right\}}{(\mu^{n+1} - 1)^2} \\ &= \frac{(\mu - 1) \{ (\mu^n - 1) + (\mu^n - \mu) + \dots + (\mu^n - \mu^{n-1}) \}}{(\mu^{n+1} - 1)^2} \geq 0. \end{aligned} \quad (\text{A.6})$$

And the effective arrival rate is concave in μ for given n :

$$\begin{aligned} \frac{\partial^2 \lambda_e(\mu, n)}{\partial \mu^2} &= - \frac{(n+1)\mu^{n-1} \left(n(\mu-1)(\mu^{n+1}+1) - 2\mu(\mu^n-1) \right)}{(\mu^{n+1}-1)^3} \\ &= - \frac{(n+1)\mu^{n-1}(\mu-1) \left(n(\mu^{n+1}+1) - 2\mu \frac{\mu^n-1}{\mu-1} \right)}{(\mu^{n+1}-1)^3} \leq 0. \end{aligned} \quad (\text{A.7})$$

The last inequality holds because $n(\mu^{n+1} + 1) - 2\mu \frac{\mu^n - 1}{\mu - 1} > 0$ (see the Proof of Lemma 1 for details). Thus, the effective arrival rate is concave in μ . Given the convex operating cost, we complete the proof.

(ii) The effective arrival rate is increasing in n for given μ :

$$\lambda_e(\mu, n+1) - \lambda_e(\mu, n) = -\frac{\mu-1}{\mu^{n+2}-1} + \frac{\mu-1}{\mu^{n+1}-1} > 0.$$

This implies that the profit function jumps upward when the threshold increases: $\lim_{\epsilon \rightarrow 0^+} \lambda_e(\mu - \epsilon, n-1) = \lambda_e(\mu, n-1) < \lambda_e(\mu, n)$ for $\mu < \infty$. ■

Proof of Lemma II.5. (Upper-envelope function)

(i) The relaxed effective arrival rate is given by

$$\lambda_e(\mu, \tilde{n}(\mu)) = \begin{cases} 1 - \frac{\mu-1}{\mu^{\tilde{n}(\mu)+1}-1}, & \mu \neq 1 \\ \frac{\tilde{n}(\mu)}{\tilde{n}(\mu)+1}, & \mu = 1. \end{cases}$$

To show the continuity of the relaxed profit function, we only need to show that $\lambda_e(\mu, \tilde{n}(\mu))$ is continuous at $\mu = 1$. Using L'Hôpital's rule, we have

$$\lim_{\mu \rightarrow 1} \left(1 - \frac{\mu-1}{\mu^{\tilde{n}(\mu)+1}-1} \right) = \lim_{\mu \rightarrow 1} \left(1 - \frac{1}{(\tilde{n}(\mu)+1)\mu^{\tilde{n}(\mu)} + \tilde{n}'(\mu)\mu^{\tilde{n}(\mu)} \log \mu} \right) = \frac{\tilde{n}(1)}{\tilde{n}(1)+1}.$$

Hence, the relaxed profit function is continuous.

To show that the relaxed profit function is an upper envelope, note that $\lambda_e(\mu, \tilde{n}(\mu)) \geq \lambda_e(\mu, n(\mu))$ holds because $\lambda_e(\mu, n)$ is increasing in n and $\tilde{n}(\mu) \geq n(\mu)$ for any μ . And the equality holds only for every $\underline{\mu}(n)$ because $\tilde{n}(\mu) = n(\mu)$ only for $\mu = \underline{\mu}(n)$.

(ii) Consider a profit maximizer is μ^* with the corresponding threshold n^* . That is, $\mu^* \in [\underline{\mu}(n^*), \underline{\mu}(n^*+1))$. Suppose that there exists no stationary point of the relaxed profit function within $(\underline{\mu}(n^*-1), \underline{\mu}(n^*+1))$. This implies that the relaxed profit

function is monotone within $[\underline{\mu}(n^* - 1), \underline{\mu}(n^* + 1)]$. If the relaxed profit function is decreasing, $\Pi(\underline{\mu}(n^* - 1), n^* - 1)$ is greater than $\Pi(\mu^*, n^*)$, which is a contradiction. Likewise, if the relaxed profit function is increasing, $\Pi(\underline{\mu}(n^* + 1), n^* + 1)$ is greater than $\Pi(\mu^*, n^*)$, which is a contradiction. ■

Proof of Lemma II.6. (Upper-envelope function with individual customers)

We first show the relaxed profit function with individual customers is concave for any μ if $\tilde{n} \geq 2$. Then, we will show that (2.12) is a sufficient condition for the equilibrium threshold to be greater than or equal to 2.

The effective arrival rate with the individual customers' relaxed threshold (i.e., $\tilde{n}^I(\mu) = b\mu$) is

$$\lambda_e(\mu, \tilde{n}^I(\mu)) = 1 - \frac{\mu - 1}{\mu^{b\mu+1} - 1} \quad \text{for } \mu > 1/b,$$

where $b = \frac{R-p}{h}$. The first and second derivatives of $\lambda_e(\mu, \tilde{n}(\mu))$ with respect to μ are given by

$$\frac{d\lambda_e(\mu, \tilde{n}^I(\mu))}{d\mu} = \frac{b\mu^{b\mu+1}(\mu - 1)(1 + \log \mu) - (\mu^{b\mu} - 1)}{(\mu^{b\mu+1} - 1)^2} \geq 0,$$

and

$$\begin{aligned} \frac{d^2\lambda_e(\mu, \tilde{n}^I(\mu))}{d\mu^2} = & -\frac{\mu^{b\mu}}{(\mu^{b\mu+1} - 1)^3} \left[-2(\mu^{b\mu} - 1) \right. \\ & - b\left((\mu^{b\mu+2} - 1) + (\mu^{b\mu+1} - 1) - 5(\mu - 1)\right) \\ & + b^2\mu(\mu - 1)(\mu^{b\mu+1} + 1) + \left(-2b(\mu^{b\mu+1} - 2\mu + 1) \right. \\ & \left. \left. + 2b^2\mu(\mu - 1)(\mu^{b\mu+1} + 1)\right) \log \mu + b^2\mu(\mu - 1)(\mu^{b\mu+1} + 1)(\log \mu)^2 \right]. \end{aligned}$$

We want to show that the second derivative is non-positive for any b and $\mu > 1/b$.

After replacing b by \tilde{n}/μ , it is equivalent to show the following for any $\tilde{n} > 1$ and μ :

$$\begin{aligned} F(\mu, \tilde{n}) &= -\tilde{n}(1 + \mu^{\tilde{n}+1})(1 + \log \mu)^2 + 2 \left(\frac{\mu^{\tilde{n}+1} - 1}{\mu - 1} - 2 \right) (1 + \log \mu) \\ &\quad + (\mu^{\tilde{n}+1} - 1) + \frac{2}{\tilde{n}} \frac{\mu^{\tilde{n}+1} - \mu}{\mu - 1} \leq 0 \end{aligned}$$

for any $\mu > 0$ and $\tilde{n} > 1$.

$F(\mu, \tilde{n})$ can be shown to be decreasing in \tilde{n} as follows,

$$\begin{aligned} \frac{\partial F(\mu, \tilde{n})}{\partial \tilde{n}} &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[F(\mu, \tilde{n} + \epsilon) - F(\mu, \tilde{n}) \right] \\ &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[(-\tilde{n} + \epsilon)\mu^{\tilde{n}+1+\epsilon} + \tilde{n}\mu^{\tilde{n}+1} - \epsilon \right] (1 + \log \mu)^2 + 2\mu^{\tilde{n}+1} \frac{\mu^\epsilon - 1}{\mu - 1} (1 + \log \mu) \\ &\quad + \mu^{\tilde{n}+1}(\mu^\epsilon - 1) + \frac{2\mu}{(\mu - 1)} \left(\frac{\mu^{\tilde{n}+\epsilon} - 1}{\tilde{n} + \epsilon} - \frac{\mu^{\tilde{n}} - 1}{\tilde{n}} \right) \Big] \\ &\leq \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[(-\tilde{n} + \epsilon)\mu^{\tilde{n}+1+\epsilon} + \tilde{n}\mu^{\tilde{n}+1} - \epsilon \right] (1 + \log \mu)^2 + 2\mu^{\tilde{n}+1} \frac{\mu^\epsilon - 1}{\mu - 1} (1 + \log \mu) \\ &\quad + \mu^{\tilde{n}+1}(\mu^\epsilon - 1) + \mu^{\tilde{n}+1}(\mu^\epsilon - 1) \Big] \tag{A.8} \end{aligned}$$

$$\begin{aligned} &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[(-\tilde{n} + \epsilon)\mu^{\tilde{n}+1+\epsilon} + \tilde{n}\mu^{\tilde{n}+1} - \epsilon \right] (1 + \log \mu)^2 + 2\mu^{\tilde{n}+1} \frac{\mu^\epsilon - 1}{\mu - 1} (\mu + \log \mu) \Big] \\ &= \lim_{\epsilon \rightarrow 0} \frac{\mu^{\tilde{n}+1}}{\epsilon} \left[(-\tilde{n} + \epsilon)\mu^\epsilon + \tilde{n} - \epsilon\mu^{-\tilde{n}-1} \right] (1 + \log \mu)^2 + 2\frac{\mu + \log \mu}{\mu - 1} (\mu^\epsilon - 1) \Big] \\ &\leq \lim_{\epsilon \rightarrow 0} \frac{\mu^{\tilde{n}+1}}{\epsilon} \left[(-\tilde{n} + \epsilon)\mu^\epsilon + \tilde{n} - \epsilon\mu^{-\tilde{n}-1} \right] (1 + \log \mu)^2 + \frac{4\mu - 2}{\mu - 1} (\mu^\epsilon - 1) \Big] \tag{A.9} \end{aligned}$$

$$\leq \lim_{\epsilon \rightarrow 0} \frac{\mu^{\tilde{n}+1}}{\epsilon} \left[-(1 + \epsilon)\mu^\epsilon + 1 - \epsilon\mu^{-2} \right] (1 + \log \mu)^2 + \frac{4\mu - 2}{\mu - 1} (\mu^\epsilon - 1) \Big] \tag{A.10}$$

$$\begin{aligned} &= \lim_{\epsilon \rightarrow 0} \mu^{\tilde{n}+1} \left[-(\mu^\epsilon + (1 + \epsilon)\mu^\epsilon \log \mu + \mu^{-2}) (1 + \log \mu)^2 + \frac{4\mu - 2}{\mu - 1} \mu^\epsilon \log \mu \right] \\ &\tag{A.11} \end{aligned}$$

$$= \mu^{\tilde{n}+1} \left\{ -(1 + \log \mu + \mu^{-2}) (1 + \log \mu)^2 + \frac{4\mu - 2}{\mu - 1} \log \mu \right\} \leq 0.$$

where the upper bounds in (A.8), (A.9), and (A.10) are derived from the followings:

(i) The upper bound, (A.8), is derived from the fact that $\frac{\partial \mu^{\tilde{n}}}{\partial \tilde{n}}$ is greater than $\frac{\partial \frac{2}{\tilde{n}} \frac{\mu^{\tilde{n}} - 1}{\mu - 1}}{\partial \tilde{n}}$:

$$\mu^{\tilde{n}} \log \mu \geq -\frac{2(\mu^{\tilde{n}} - 1)}{\tilde{n}^2(\mu - 1)} + \frac{2\tilde{n}\mu^{\tilde{n}} \log \mu}{\tilde{n}^2(\mu - 1)},$$

which is equivalent to

$$\log \mu \begin{cases} > \frac{2(\mu^{-\tilde{n}}-1)}{\tilde{n}^2(\mu-1)-2\tilde{n}} & \text{if } \mu < 1 \text{ or } \mu > 1 + \frac{2}{\tilde{n}}, \\ \leq \frac{2(\mu^{-\tilde{n}}-1)}{\tilde{n}^2(\mu-1)-2\tilde{n}} & \text{if } 1 < \mu < 1 + \frac{2}{\tilde{n}}. \end{cases} \quad (\text{A.12})$$

If $\mu > 1 + \frac{2}{\tilde{n}}$, (A.12) is trivially true since the LHS is positive while the RHS is negative.

If $\mu < 1 + \frac{2}{\tilde{n}}$, we first show that the RHS of (A.12) is decreasing in \tilde{n} if $\mu < 1$ and increasing in \tilde{n} if $\mu > 1$. The derivative of the RHS with respect to \tilde{n} is

$$\frac{4(\mu^{\tilde{n}}-1)(\tilde{n}(\mu-1)-1)+2\tilde{n}(2-\tilde{n}(\mu-1))\log \mu}{\tilde{n}^2\mu^{\tilde{n}}(\tilde{n}(\mu+1)^2-2)^2}. \quad (\text{A.13})$$

The derivative of the numerator in (A.13) with respect to μ is non-negative :

$$\begin{aligned} & \frac{2\tilde{n}}{\mu} \{(\mu-1)(2(\tilde{n}+1)\mu^{\tilde{n}}-\tilde{n}-2)-\tilde{n}\mu \log \mu\} \\ & \geq \frac{2\tilde{n}}{\mu} \{(\mu-1)(2(\tilde{n}+1)\mu^{\tilde{n}}-\tilde{n}-2)-\tilde{n}\mu(\mu-1)\} \\ & = \frac{2\tilde{n}(\mu-1)}{\mu} \{2(\tilde{n}+1)(\mu^{\tilde{n}}-1)-\tilde{n}(\mu-1)\} \\ & \geq 0. \end{aligned}$$

Given that the numerator of (A.13) is increasing in μ and it equals 0 when $\mu = 1$, it follows that the RHS of (A.12) is decreasing in \tilde{n} if $\mu < 1$ and increasing in \tilde{n} if $\mu > 1$. Thus, it suffices to show (A.12) is true when $\tilde{n} = 1$, and can be shown as follows,

$$\begin{aligned} \log \mu > \frac{-2(\mu-1)}{\mu(\mu-3)} &\geq \frac{2(\mu^{-\tilde{n}}-1)}{\tilde{n}^2(\mu-1)-2\tilde{n}} && \text{if } \mu < 1, \\ \log \mu \leq \frac{-2(\mu-1)}{\mu(\mu-3)} &\leq \frac{2(\mu^{-\tilde{n}}-1)}{\tilde{n}^2(\mu-1)-2\tilde{n}} && \text{if } 1 < \mu < 1 + \frac{2}{\tilde{n}}. \end{aligned}$$

(ii) The upper bound, (A.9), is derived from $\log \mu \leq \mu - 1$.

(iii) The upper bound, (A.10), holds because the value in the square bracket of (A.9) can be shown to be decreasing in n . Let $l(\mu, \tilde{n}) := -(\tilde{n} + \epsilon)\mu^\epsilon + \tilde{n} - \epsilon\mu^{-\tilde{n}-1}$ and $l_{\tilde{n}}(\mu, \tilde{n}) = 1 - \mu^\epsilon + \epsilon\mu^{-\tilde{n}-1} \log \mu$, the derivative of $l(\mu, \tilde{n})$ with respect to \tilde{n} . For small enough $\epsilon > 0$,

both $l(\mu, \tilde{n})$ and $l_{\tilde{n}}(\mu, \tilde{n})$ are negative. Thus, $\tilde{n} = 1$ gives the upper bound.

Since $F(\mu, \tilde{n})$ is decreasing in \tilde{n} , the upper bound of $F(\mu, \tilde{n})$ can be derived by plugging in a small n . If $\mu > 1$, we can show that $F(\mu, \tilde{n})$ is negative for any n :

$$\begin{aligned}
F(\mu, \tilde{n}) &\leq F(\mu, 1) = -(1 + \mu^2)(1 + \log \mu)^2 + 2(\mu - 1)(1 + \log \mu) + \mu^2 + 2\mu - 1 \\
&= -(1 + \mu^2)(\log \mu)^2 - 2(\mu^2 - \mu + 2) \log \mu + 4(\mu - 1) \\
&= (2 + \log \mu) \left[2(\mu - 1) - (\mu^2 + 1) \log \mu \right] \\
&\leq (2 + \log \mu) \left[2(\mu - 1) - (\mu^2 + 1)(\mu - 1)/\mu \right] \\
&= (2 + \log \mu) \left(-(\mu - 1)^3/\mu \right) \leq 0.
\end{aligned}$$

If $\mu < 1$, we need $\tilde{n} \geq 2$ for the concavity of the relaxed profit function:

$$\begin{aligned}
F(\mu, \tilde{n}) &\leq F(\mu, 2) \\
&= -2(\mu^3 + 1)(1 + \log \mu)^2 + 2(\mu^2 + \mu - 1)(1 + \log \mu) + (\mu^3 + \mu^2 + \mu - 1) \leq 0.
\end{aligned}$$

Thus, for any $\mu > 0$, the sufficient condition for the concave relaxed profit function is that the equilibrium threshold, n^{IP} , is greater than or equal to 2. A sufficient condition for $n^{IP} \geq 2$ is that any profit in the threshold interval for $n = 1$ is less than the starting point of the threshold interval for $n = 2$, given by

$$\max_{\mu \in [1/b, 2/b)} \Pi(\mu, 1) < p \left(1 - \frac{2/b - 1}{(2/b)^2 - 1} \right) - c(1/b) \leq p \left(1 - \frac{2/b - 1}{(2/b)^3 - 1} \right) - c(2/b).$$

With some algebra, this is equivalent to (2.12). Then, $n^{IP} \geq 2$ and the envelope method specifies the unique interval that contains the profit-maximizer with individual customers.

■

Proof of Lemma II.7. (Differentiability and concavity of the welfare function)

(1) Trivially true from its components' differentiability.

(2) Given n fixed, the effective arrival rate is concave from (A.7). Thus, it is sufficient to show the average number of customers, $L(\mu, n)$, is convex in μ for $\mu \geq 1$. The second-order derivative with respect to μ is

$$\frac{\partial^2 L(\mu, n)}{\partial \mu^2} = \frac{2}{(\mu - 1)^3} - \frac{(n + 1)^2(n + 2)\mu^{2n} + n(n + 1)^2\mu^{n-1}}{(\mu^{n+1} - 1)^3}. \quad (\text{A.14})$$

We will first show that (A.14) is increasing in n for any $\mu > 1$. That is, for any $n \geq 1$ and $\mu > 1$,

$$\begin{aligned} \frac{\partial^3 L(\mu, n)}{\partial \mu^2 \partial n} = \frac{(n + 1)\mu^{n-1}}{(\mu^{n+1} - 1)^4} & \left[\left((n + 1)(n + 2)\mu^{2n+2} + 4(n + 1)^2\mu^{n+1} + n(n + 1) \right) \log \mu \right. \\ & \left. - \left((3n + 5)\mu^{2n+2} - 4\mu^{n+1} - (3n + 1) \right) \right] \geq 0. \end{aligned} \quad (\text{A.15})$$

This is equivalent to

$$\log \mu - \frac{(3n + 5)\mu^{2n+2} - 4\mu^{n+1} - (3n + 1)}{(n + 1)(n + 2)\mu^{2n+2} + 4(n + 1)^2\mu^{n+1} + n(n + 1)} \geq 0.$$

Letting $x = \mu^{n+1}$, this is equivalent to

$$k(x, n) := \log x - \frac{(3n + 5)x^2 - 4x - (3n + 1)}{(n + 2)x^2 + 4(n + 1)x + n} \geq 0.$$

The derivative of $k(x, n)$ with respect to x is positive for any $x > 1$, i.e.,

$$\frac{\partial k(x, n)}{\partial x} = \frac{(x - 1)^3 (n^2(x - 1) + 4x(n + 1))}{x((n + 2)x^2 + 4(n + 1)x + n)^2} > 0,$$

and $k(1, n) = 0$. Thus, (A.15) is true, and (A.14) is increasing in n for any $\mu > 1$. Then, from

$$\frac{\partial^2 L(\mu, n)}{\partial \mu^2} \geq \frac{\partial^2 L(\mu, 1)}{\partial \mu^2} = \frac{2}{(\mu - 1)^3} - \frac{12\mu^2 + 4}{(\mu^2 - 1)^3} = \frac{2}{(\mu + 1)^3} > 0, \quad (\text{A.16})$$

we complete the proof. ■

Proof of Lemma II.8. (Discontinuities in the welfare function for individual customers)

Recall that the welfare is discretely unimodal in n for given μ . For $\mu \notin (\mu_L(p), \mu_H(p))$, in which the individual threshold is smaller than or equal to the social threshold, the threshold increase makes the throughput closer to the socially optimal level, thus, it increases the welfare. The argument for $\mu \notin (\mu_L(p), \mu_H(p))$ is similar. ■

Proof of Lemma II.9. (Discontinuities in the welfare function for collective customers)

We only need to show that the customers utility is continuous and increasing in μ . First, if the service rate changes, the collective threshold changes if and only if it improves the resultant customer utility. Given that the customer utility is continuous in μ for any given n , it follows that the utility of collective customers is also continuous in μ . Second, it is trivial that the faster the service is, the higher average utility collective customers receive. Thus the discontinuity of the welfare function when the threshold changes only comes from the profit part, which increases in n . ■

Proof of Lemma II.10. (Continuities in the welfare function for social customers)

Social customers optimize the social welfare. In other words, if the service rate changes, their threshold changes if and only if it improves the resultant welfare. Suppose that the threshold changes from $n - 1$ to n and the resultant welfare jumps upward at $\mu = x$, then, for $\mu = x - \epsilon$ for small enough $\epsilon > 0$, the social welfare can be larger if they take the threshold n instead of $n - 1$. The same explanation also holds when the social welfare jumps downward. ■

Proof of Proposition II.11. (Collective customers and selfish or social firm)

The social firm will produce higher social welfare and lower profit, and hence higher utility, than the profit-maximizing firm. The utility can be shown to be increasing in μ within any threshold interval, and because collective customers change their threshold if and only if it

improves their utility, this implies that overall the collective customers' utility is increasing in μ . Thus, with collective customers, the social firm sets a higher service rate than the profit-maximizing firm's, i.e., $\mu^{CP} \leq \mu^{CS}$. ■

Proof of Proposition II.13. (Social customers and selfish or social firm)

Note that the social threshold does not change in p . That is, the threshold interval remains the same for any change of price. Suppose that μ^{SP} is a profit-maximizer for some fixed p . If p increases, revenue increases. Then, it is trivial that the profit at μ^{SP} still dominates the profit with any $\mu < \mu^{SP}$ because the throughput is increasing in μ . ■

Proof of Proposition II.14. (Price coordination with a social firm)

Since the social firm chooses the service rate to maximize the welfare, we only need to show that there exist a set of prices such that selfish customers choose the social threshold $n^S(\mu^{SS})$ when μ^{SS} is given.

(i) By definition, $n^C(\mu^{SS}) = n^S(\mu^{SS})$ if $p = 0$. From Proposition II.2 (iii), we know $n^C(\mu^{SS})$ is a decreasing step function of price. Thus, the smallest price such that $n^C(\mu^{SS}) \leq n^S(\mu^{SS})$ is \bar{p}^{CS} . (In case of multiple equilibria for the SS case, i.e., multiple pairs of n^{SS} and μ^{SS} exist, the price upper bound, \bar{p}^{CS} , is simply the maximum of the bounds. We further show that welfare is decreasing in price. For any given μ , the collective threshold $n^C(\mu)$ moves away from the social threshold $n^S(\mu)$ as price p increases. This implies that the welfare using the collective threshold and fixed service rate, $\mathcal{W}(\mu, n^C(\mu))$, is decreasing in price. Since the social firm chooses the service rate that maximizes the welfare, the equilibrium welfare is also decreasing in price with collective customers.

(ii) Suppose that μ^{SS} is given. From Proposition II.2 (iii), we know $n^I(\mu^{SS})$ is a decreasing step function in price. We only need to show that $n^I(\mu^{SS}) = 0 \leq n^S(\mu^{SS})$ when p is high and $n^I(\mu^{SS}) \geq n^S(\mu^{SS})$ for $p = 0$. For $p = R$, the individual customer will not join no matter how fast service is because $U(i) = R - p - h \frac{i+1}{\mu} < 0$ for any $i \geq 0$. When $p = 0$, $n^S(\mu) = n^C(\mu)$, and from Proposition II.2, we know $n^I(\mu^{SS}) \geq n^S(\mu^{SS})$ if $p = 0$. ■

Proof of Proposition II.15. (No price regulation for the selfish firm)

First, we show that the welfare-maximizing service rate is a stationary point. From (2.13), the socially optimal welfare, $\mathcal{W}(\mu^{SS}, n^S(\mu^{SS}))$ can be defined as $\max_{n \in \mathbb{N}} \mathcal{W}(\mu^S(n), n)$ where $\mu^S(n)$ is the welfare maximizer assuming the threshold is fixed at n . Because the welfare function of μ for given n is differentiable, the welfare-maximizing function is a stationary point (i.e., it is not on the boundary of any threshold interval).

Suppose that the profit and welfare maximizers are the same for some p , i.e., $\mu^{SP} = \mu^{SS}$. Given the fact that the social threshold does not change in price, this implies that μ^{SP} is not on the boundary of any threshold interval, and it is a stationary point.

From Lemma II.7 (i), the difference between the welfare and the profit functions (i.e., the customer utility function) is strictly increasing in μ within the threshold interval. This contradicts the assumption that both μ^{SS} and μ^{SP} are stationary points. ■

APPENDIX B

Proofs and Technical Details for Chapter III: Dynamic Pricing and Loyalty Programs

Proof of Lemma III.1

$$P(\text{cash purchase}|\text{cash-only}) - P(\text{cash purchase}|\text{loyalty}) = \bar{F}(p)G(p/q) > 0 \text{ and}$$

$$P(\text{total purchase}|\text{loyalty}) - P(\text{cash purchase}|\text{cash-only}) = \int_0^{p/q} \bar{F}(qx)dG(x) > 0.$$

Proof of Lemma III.2

(a) It can be shown by the derivatives of probabilities from Lemma III.1 with respect to p :

$$\begin{aligned} \frac{\partial P(\text{cash purchase}|p)}{\partial p} &= -f(p)\bar{G}(p/q) - \bar{F}(p)g(p/q) < 0 \\ \frac{\partial P(\text{point purchase}|p)}{\partial p} &= \bar{F}(p)g(p/q) > 0 \\ \frac{\partial P(\text{total purchase}|p)}{\partial p} &= -f(p)\bar{G}(p/q) < 0. \end{aligned}$$

(b) Suppose that the point requirement is increased from q to $q + \epsilon$ for any $\epsilon > 0$. Under the higher point requirement, the proportion of loyalty consumer is smaller, i.e., $\bar{\beta}(q) \geq \bar{\beta}(q + \epsilon)$. $P(\text{cash purchase})$ increases and $P(\text{point purchase})$ decreases, as there are more cash-only consumers (who buy with cash more likely than loyalty consumers) and

less loyalty consumers (who can buy with points). Also, for loyalty consumers, increasing q makes $P(\text{point purchase}|\text{loyalty})$ smaller. ■

Proof of Theorem III.3

First, we only need to consider any p greater than or equal to $\Delta_{t-1}(y)$. If $\Delta_{t-1}(y) > p$, it also implies that $\Delta_{t-1}(y) > R$. Together, it means that the seller makes the negative contributions to the expected revenue by making any sales (both the cash sales and the reward sales). That means, $V_t(y) < V_{t-1}(y)$, which contradicts to the property of the value function. Note that the first order condition is given by

$$\frac{\partial J_t(p, y)}{\partial p} = \lambda \left(1 - \bar{\beta}G(p/q)\right) \bar{F}(p) \left[1 - \frac{f(p)}{\bar{F}(p)} \left(p - \Delta_{t-1}(y)\right) - \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \left(p - R\right) \right] = 0.$$

and the second order condition (the second-order derivative is negative for any stationary point) is given by

$$\begin{aligned} \frac{\partial^2 J_t(p, y)}{\partial p^2} \Big|_{p=p^*} &= \lambda \left(1 - \bar{\beta}G(p/q)\right) \bar{F}(p) \left[- \left(\frac{f'(p)}{\bar{F}(p)} + \frac{f(p)^2}{\bar{F}(p)^2} \right) \left(p - \Delta_{t-1}(y)\right) - \frac{f(p)}{\bar{F}(p)} \right. \\ &\quad \left. - \left(\frac{\bar{\beta}g'(p/q)/q^2}{1 - \bar{\beta}G(p/q)} + \frac{\bar{\beta}^2 g(p/q)^2/q^2}{(1 - \bar{\beta}G(p/q))^2} \right) \left(p - R\right) - \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \right] \\ &= \lambda \left(1 - \bar{\beta}G(p/q)\right) \bar{F}(p) \left[- \left(\frac{f'(p)}{\bar{F}(p)} + \frac{f(p)^2}{\bar{F}(p)^2} \right) \left(p - \Delta_{t-1}(y)\right) - \frac{f(p)}{\bar{F}(p)} \right. \\ &\quad \left. - \left(\frac{g'(p/q)/q}{g(p/q)} + \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \right) \left(1 - \frac{f(p)}{\bar{F}(p)} \left(p - \Delta_{t-1}(y)\right)\right) - \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \right] \\ &< 0, \end{aligned}$$

which is equivalent to

$$\begin{aligned} (p - \Delta) \frac{f(p)}{\bar{F}(p)} \left\{ \left(\frac{g'(p/q)/q}{g(p/q)} + \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \right) - \left(\frac{f'(p)}{\bar{F}(p)} + \frac{f(p)}{\bar{F}(p)} \right) \right\} \\ < \frac{f(p)}{\bar{F}(p)} + \frac{g'(p/q)/q}{g(p/q)} + 2 \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)}. \end{aligned}$$

Note that the assumption $A3$ implies that $\left(\frac{g'(p/q)/q}{g(p/q)} + \frac{\bar{\beta}g(p/q)/q}{1 - \bar{\beta}G(p/q)} \right) - \left(\frac{f'(p)}{\bar{F}(p)} + \frac{f(p)}{\bar{F}(p)} \right)$ is non-positive, which further implies that the LHS is non-positive because $p > \Delta_{t-1}(y)$. On the

other hand, the RHS can be shown to be positive under the assumption *A1* and *A2*. Particularly, the assumption *A2* implies that the hazard rate of $F(\cdot)$ is no less than that of $G(\cdot)$. Then, after replacing $\frac{f(p)}{F(p)}$ by $\frac{g(p/q)/q}{G(p/q)}$, the RHS is shown to be positive by the IFR condition of $G(\cdot)$. ■

Proof of Corollary III.4

The proof is immediate from Theorem III.3 and the fact that a marginal value of inventory is increasing in t and decreasing in y . ■

Proof of Proposition III.5

The optimal price, p^* , satisfies the following first-order condition:

$$1 - \frac{f(p)}{F(p)}(p - \Delta) - \frac{\bar{\beta} \cdot g(p/q)/q}{1 - \bar{\beta}G(p/q)}(p - R) = 0, \quad (\text{B.1})$$

and the cash-only price, p^c , satisfies the following first-order condition:

$$1 - \frac{f(p^c)}{F(p^c)}(p^c - \Delta) = 0. \quad (\text{B.2})$$

Suppose that $p^c < R$ and the seller chooses the cash-only price p^c , then, the optimal first-order condition (B.1) will be positive. On the other hand, any $p > R$ makes (B.1) to be negative. From the unimodal property, this implies that $p^* \in (p^c, R)$.

The opposite case when $p^c > R$ can be proved by the same logic. ■

Proof of Corollary III.6

The proof is immediate from Proposition III.5 with the fact that the marginal value of inventory, $\Delta_t(y)$, is increasing in t and decreasing in y . ■

Proof of Proposition III.7

(a) From Proposition III.5, we have two cases to consider, in which the seller either adds a premium or offers a discount considering the reward sales. In the case of adding a premium

($p^c < p^* < R$), note that $1 - \frac{f(p^*)}{F(p^*)}(p^* - \Delta) < 0$ and $-\frac{\bar{\beta}g(p^*/q)/q}{1-\bar{\beta}G(p^*/q)}(p^* - R) > 0$. If $\bar{\beta}$ increases, the first order of the revenue when $p = p^*$ becomes positive, which means that the optimal price should be also higher as $\bar{\beta}$ increases.

In the opposite case when the seller offers a discount ($R < p^* < p^c$), we have $1 - \frac{f(p^*)}{F(p^*)}(p^* - \Delta) > 0$ and $-\frac{\bar{\beta}g(p^*/q)/q}{1-\bar{\beta}G(p^*/q)}(p^* - R) < 0$. Then, increasing $\bar{\beta}$ makes the first order to be negative. Therefore, the amount of discount increases in $\bar{\beta}_t$.

(b) Suppose that the optimal price is \bar{p} when the reimbursement rate and the marginal value of inventory are R_t and Δ , respectively, and consider that R_t increases by $\epsilon > 0$. Then, the first order of the expected revenue at $p = \bar{p}$ becomes positive: $1 - \frac{f(\bar{p})}{F(\bar{p})}(\bar{p} - \Delta) - \frac{\bar{\beta}g(\bar{p}/q)/q}{1-\bar{\beta}G(\bar{p}/q)}(\bar{p} - R_t - \epsilon) = \frac{\bar{\beta}g(\bar{p}/q)/q}{1-\bar{\beta}G(\bar{p}/q)}\epsilon > 0$. From Lemma III.3, this implies that the optimal price should be higher as R_t increases. It is also trivial to see that the expected revenue increases in R_t .

From a consumer's perspective, their utilities are the maximum of three options. Given that the price increases, the utility from a cash purchase decreases while the utilities from a point purchase or no purchase remain the same. This implies the the consumer's utility remains the same or gets worse as a consequence to the increase of reimbursement rate. In particular, the utility of a consumer who previously purchase with cash decreases. ■

Proof of Proposition III.8

We need to show that p^* increases in marginal value of inventory since the marginal value of inventory ($\Delta_t(y)$) increases in t and decreases in y . Suppose that p^* satisfies the first-order condition, given by (B.1), when the marginal value of inventory is Δ . If Δ increases by $\epsilon > 0$, then, the first order at p^* becomes positive : $1 - \frac{f(p^*)}{F(p^*)}(p^* - \Delta - \epsilon) - \frac{\bar{\beta}g(p^*/q)/q}{1-\bar{\beta}G(p^*/q)}(p^* - R) = \frac{f(p^*)}{F(p^*)}\epsilon > 0$. Thus, the optimal price should also increase if the marginal value of inventory increases. ■

Proof of Proposition III.9

From XXX, we know that the marginal value of inventory increases in t and decreases in y . Thus we only need to show that there exist three regions separated by the two threshold

values of Δ . Suppose that Δ is low enough such that the cash-only price is lower than the reimbursement rate (i.e., $p^c(\Delta) < R$), then, it is trivially true that allowing reward sales gives a higher expected revenue by $J^o(p^*) \geq J^o(p^c) > J^c(p^c)$. On the other hand, if Δ is high enough so that $\Delta > R$, the seller is always better off to block reward sales because any point-sale transaction will be a loss. In the middle of these two cases where $\Delta < R < p^* < p^c$,

$$\frac{dJ^o(p^*(\Delta), \Delta)}{d\Delta} = \frac{\partial J}{\partial p^*} \frac{\partial p^*}{\partial \Delta} + \frac{\partial J}{\partial \Delta} = 0 - \left(\bar{F}(p^*) + \bar{\beta} \int_0^{p^*/q} \bar{F}(xq) dG(x) \right) \quad (\text{B.3})$$

while

$$\frac{dJ^c(p^c(\Delta), \Delta)}{d\Delta} = \frac{\partial J}{\partial p^c} \frac{\partial p^c}{\partial \Delta} + \frac{\partial J}{\partial \Delta} = 0 - \bar{F}(p^c) \quad (\text{B.4})$$

Since $p^c > p^*$, it is obvious that the value-to-go function decreases faster in the marginal value of inventory when it is open compared to when closed, i.e., $J_\Delta^o < J_\Delta^c$. This implies that if the value of inventory increases (as the longer periods or the less inventories left), the value function with open decision becomes less favorable than the value function with black-out decision. ■

Proof of Proposition III.10

(a) From Proposition III.7 (b), we know that the expected revenue with reward sales increases in the reimbursement rate while the cash-only revenue remains the same. Thus, it is obvious that for some value of Δ , the optimal decision changes from close to open.

(b) Suppose that it is optimal to open when $\bar{\beta}$ is given, that is, $J^o(p^*) > J^c(p^c)$. Note that the expected revenue with allowing reward sales is a weighted sum of revenue from loyalty consumers, $\bar{F}(p^*)\bar{G}(p^*/q)(p^* - \Delta) + \int_0^{p^*/q} \bar{F}(qx) dG(x)(R - \Delta)$, and cash-only consumers, $\bar{F}(p^*)(p^* - \Delta)$. Let us write the expected revenue with black-out by $\bar{\beta}\bar{F}(p^c)(p^c - \Delta) + \beta\bar{F}(p^c)(p^c - \Delta)$ and compare part by part. Since p^c maximizes the revenue from cash-only

consumers,

$$\beta \bar{F}(p^c)(p^c - \Delta) > \beta \bar{F}(p^*)(p^* - \Delta).$$

With the assumption that the open decision is optimal ($J^o(p^*) > J^c(p^c)$), this implies that the expected revenue from loyalty consumers is greater when allowing reward sales:

$$\bar{\beta} \bar{F}(p^c)(p^c - \Delta) < \bar{\beta} \left[\bar{F}(p^*) \bar{G}(p^*)(p^* - \Delta) + \int_0^{p^*/q} \bar{F}(xq) dG(x)(R - \Delta) \right].$$

Suppose that $\bar{\beta}$ increases. It is obvious that the expected revenue with allowing reward sales increases as $\bar{\beta}$ even under the same price p^* while the expected revenue with black-out decision remains the same. The further adjustment of the optimal price for changing $\bar{\beta}$ will make even higher revenue from allowing reward sales, and this completes the proof. ■

Proof of Proposition III.11

We know that both $p_t^c(y)$ and $p_t^*(y)$ are non-decreasing in t and non-increasing in y (Gallego and Van Ryzin (1994) for p^c and Proposition III.8 for p^*). From Proposition III.9, we also know that the optimal price changes from $p_t^B(y) = p_t^c(y)$ to $p_t^B(y) = p_t^*(y)$ as t decreases and y increases. Suppose that the transition is made between t and $t - 1$, that is, $p_t^B(y) = \tilde{p}_t^c(y)$ and $\tilde{p}_t^o(y + 1) = p_t^B(y + 1)$. From Proposition III.9 (a), the cash-only price is higher than the point-selling price when it is optimal to close and we know that $p_t^o(y)$ decreases as t decreases, thus, we have $\tilde{p}_t^c(y) > \tilde{p}_t^o(y) > \tilde{p}_{t-1}^o(y)$. We can show the same result when the transition is made between y and $y - 1$. ■

Proof of Theorem III.12

Note that $\tilde{p}_t(q, y)$ satisfies the following first-order condition:

$$1 - \frac{f(p)}{\bar{F}(p)}(p - \Delta) - \frac{1}{q} \cdot \frac{\bar{\beta}(q) \cdot g(p/q)}{1 - \bar{\beta}(q) \cdot G(p/q)}(p - R(q)) = 0. \quad (\text{B.5})$$

By A4 and A5, we know that the third term, $-\frac{1}{q} \frac{\bar{\beta}(q) \cdot g(p/q)}{1 - \bar{\beta}(q)G(p/q)}(p - R(q))$, increases in q for a given p . This implies that the best price $\tilde{p}_t(q, y)$ should also increase to satisfy the first-order condition as q increases. ■

Proof of Proposition III.13

First note that any (t, y) can be summarized by the corresponding marginal value of inventory. For simplicity of notation, we will use Δ for replacing t and y , i.e., we use $J(p, q, \Delta)$ for $J_t^Q(p, q, y)$ and $\tilde{p}_t(q, y)$ for $\tilde{p}(q, \Delta)$. We also assume that Δ is a continuous variable for this proof despite the fact that it takes the discrete value because t and y are integral. Now we want to show that, for any pair of point requirements (q_l and q_h with $q_l < q_h$), there exists a unique $\bar{\Delta}_{l,h}$ such that

$$\begin{aligned} J(\tilde{p}(q_l, \Delta), q_l, \Delta) &> J(\tilde{p}(q_h, \Delta), q_h, \Delta) \text{ if } \Delta < \bar{\Delta}_{l,h} \text{ and} \\ J(\tilde{p}(q_l, \Delta), q_l, \Delta) &< J(\tilde{p}(q_h, \Delta), q_h, \Delta) \text{ if } \Delta > \bar{\Delta}_{l,h}. \end{aligned}$$

There are three different cases to consider depending on the value of Δ .

(i) Suppose that Δ is high enough such that $R(q_l) < R(q_h) < \Delta$. In this case, any reward sale (regardless of whether $q = q_l$ or q_h) negatively affects the seller's revenue. Since the probability of a reward sale decreases in q , the higher point requirement is better, i.e., $J(\tilde{p}(q_h, y), q_h, \Delta) < J(\tilde{p}(q_l, y), q_l, \Delta)$.

(ii) Suppose that Δ has some intermediate value such that $R(q_l) < \Delta < R(q_h)$. In this case, only a reward sale at $q = q_l$ negatively affects the seller's revenue. Thus, the value function with a higher point requirement is better, i.e., $J(\tilde{p}(q_h, y), q_h, \Delta) < J(\tilde{p}(q_l, y), q_l, \Delta)$.

(iii) Suppose that Δ is low enough such that $\Delta < R(q_1) < R(q_2)$. The derivative of value-to-go function with respect to Δ is given by the negative of the probability of a purchase:

$$\frac{dJ^Q(\tilde{p}(q, \Delta), q, \Delta)}{d\Delta} = - \left\{ \bar{F}(\tilde{p}(q, \Delta)) - \bar{\beta}(q) \int_0^{\tilde{p}(q, \Delta)} G(x/q) dF(x) \right\} < 0.$$

From $\bar{\beta}(q_l) > \bar{\beta}(q_h)$ and $\tilde{p}(q_l, \Delta) < \tilde{p}(q_h, \Delta)$, we know that the probability of total sales is higher when the point requirement is low. That is, the value-to-go function decreases faster

in Δ with a lower point requirement. With the continuity of the value-to-go function in Δ , this implies that there exists at most one threshold, $\bar{\Delta}_{q_l, q_h}$, such that q_h is better than q_l if $\Delta > \bar{\Delta}_{q_l, q_h}$, and vice versa. Along with the fact that the marginal value of inventory is increasing in t and decreasing in y , this completes the proof. ■

APPENDIX C

Proofs and Technical Details for Chapter IV: On (Rescaled) Multi-Attempt Consumer Choice Model and Its Implication on Assortment Optimization

Proof of Theorem IV.1

Per our discussions above, $\hat{\pi}_k(i, S) = \pi(i, S)$ for $k \geq |\bar{S}|+1$. So, we only need to consider the case $k \leq |\bar{S}|$. We first consider the case where the true choice model is MNL with parameters $\{u_0, u_1, \dots, u_n\}$, $\sum_{i=0}^n u_i = 1$. Note that, for any preference sequence $j_1, j_2, \dots, j_l, i \in \mathcal{N}$, we have:

$$\begin{aligned} P(U_{j_1} > U_{j_2} > \dots > U_{j_l} > U_i \geq \max\{U_m : m \in \mathcal{N} \setminus \{j_1, j_2, \dots, j_l\} \cup \{0\}\}) \\ = \left(\frac{u_{j_1}}{1 - u_{j_1}} \right) \left(\frac{u_{j_2}}{1 - u_{j_1} - u_{j_2}} \right) \dots \left(\frac{u_{j_l}}{1 - u_{j_1} - \dots - u_{j_l}} \right) \cdot u_i. \end{aligned}$$

The above probability is an immediate consequence of the assumption of i.i.d noises with Gumbel distribution in the construction of MNL model and not difficult to prove (we

omit the details). Given the above formula, we can bound $\hat{\pi}_k(i, S)$ as follows:

$$\begin{aligned}
\hat{\pi}_k(i, S) &= \lambda_i + \sum_{j_1 \in \bar{S}} \lambda_{j_1 i} + \sum_{\{j_1, j_2\} \subseteq \bar{S}} \lambda_{\{j_1, j_2\} i} + \cdots + \sum_{\{j_1, j_2, \dots, j_{k-1}\} \subseteq \bar{S}} \lambda_{\{j_1, j_2, \dots, j_{k-1}\} i} \\
&= \sum_{l=0}^{k-1} \sum_{j_1, \dots, j_l \in \bar{S}} \left(\frac{u_{j_1}}{1 - u_{j_1}} \right) \left(\frac{u_{j_2}}{1 - u_{j_1} - u_{j_2}} \right) \cdots \left(\frac{u_{j_l}}{1 - u_{j_1} - \cdots - u_{j_l}} \right) \cdot u_i \\
&\geq \sum_{l=0}^{k-1} \sum_{j_1, \dots, j_l \in \bar{S}} \left(\frac{u_{j_1}}{1 - u_{j_1}} \right) \left(\frac{u_{j_2}}{1 - u_{j_2}} \right) \cdots \left(\frac{u_{j_l}}{1 - u_{j_l}} \right) \cdot u_i \\
&= \sum_{l=0}^{k-1} l! \cdot \left[\sum_{\substack{\{j_1, \dots, j_l\} \subseteq \bar{S} \\ j_1 < j_2 < \cdots < j_l}} \left(\frac{u_{j_1}}{1 - u_{j_1}} \right) \left(\frac{u_{j_2}}{1 - u_{j_2}} \right) \cdots \left(\frac{u_{j_l}}{1 - u_{j_l}} \right) \cdot u_i \right] \\
&= \sum_{l=0}^{k-1} l! \cdot \left[\sum_{\substack{\{j_1, \dots, j_l\} \subseteq \bar{S} \\ j_1 < j_2 < \cdots < j_l}} (u_{j_1} + u_{j_1}^2 + \cdots) (u_{j_2} + u_{j_2}^2 + \cdots) \cdots (u_{j_l} + u_{j_l}^2 + \cdots) \cdot u_i \right] \\
&= u_i \cdot \left[0! + 1! \sum_{j_1 \in \bar{S}} (u_{j_1} + u_{j_1}^2 + \cdots) + 2! \sum_{\substack{\{j_1, j_2\} \subseteq \bar{S} \\ j_1 < j_2}} (u_{j_1} + u_{j_1}^2 + \cdots) (u_{j_2} + u_{j_2}^2 + \cdots) \right. \\
&\quad + 3! \sum_{\substack{\{j_1, j_2, j_3\} \subseteq \bar{S} \\ j_1 < j_2 < j_3}} (u_{j_1} + u_{j_1}^2 + \cdots) (u_{j_2} + u_{j_2}^2 + \cdots) (u_{j_3} + u_{j_3}^2 + \cdots) + \cdots \\
&\quad \left. + (k-1)! \sum_{\substack{\{j_1, \dots, j_{k-1}\} \subseteq \bar{S} \\ j_1 < j_2 < \cdots < j_{k-1}}} (u_{j_1} + u_{j_1}^2 + \cdots) \cdots (u_{j_{k-1}} + u_{j_{k-1}}^2 + \cdots) \right] \\
&\geq u_i \cdot \left[1 + \sum_{j_1 \in \bar{S}} u_{j_1} + \left(\sum_{j_1 \in \bar{S}} u_{j_1}^2 + 2! \sum_{\substack{\{j_1, j_2\} \subseteq \bar{S} \\ j_1 < j_2}} u_{j_1} u_{j_2} \right) \right. \\
&\quad + \left(\sum_{j_1 \in \bar{S}} u_{j_1}^3 + 2! \sum_{\substack{\{j_1, j_2\} \subseteq \bar{S} \\ j_1 < j_2}} (u_{j_1}^2 u_{j_2} + u_{j_1} u_{j_2}^2) + 3! \sum_{\substack{\{j_1, j_2, j_3\} \subseteq \bar{S} \\ j_1 < j_2 < j_3}} u_{j_1} u_{j_2} u_{j_3} \right) + \cdots \\
&\quad + \left(\sum_{j \in \bar{S}} u_{j_1}^{k-1} + 2! \sum_{\{j_1, j_2\} \subseteq \bar{S}} \prod_{\substack{t_1 + t_2 = k-1 \\ t_a \in \mathbb{N}}} u_{j_1}^{t_1} u_{j_2}^{t_2} + \cdots \right. \\
&\quad \left. + (k-1)! \sum_{\substack{\{j_1, \dots, j_{k-1}\} \subseteq \bar{S} \\ j_1 < \cdots < j_{k-1}}} u_{j_1} \cdots u_{j_{k-1}} \right) \left. \right] \\
&= u_i \cdot \left[1 + \left(\sum_{j \in \bar{S}} u_j \right) + \left(\sum_{j \in \bar{S}} u_j \right)^2 + \left(\sum_{j \in \bar{S}} u_j \right)^3 + \cdots + \left(\sum_{j \in \bar{S}} u_j \right)^{k-1} \right],
\end{aligned}$$

where the fourth equality follows from identity $\frac{x}{1-x} = \sum_{n=1}^{\infty} x^n$ for all $x \in [0, 1]$ and the last inequality follows by collecting polynomial terms with the same degree.

Now, if the true choice probability is a mixture of M MNL models with parameters $\{u_{im}\}$ and $\{\theta_m\}$ for all $i \in S_0$ and $m \in \{1, \dots, M\}$, applying the result above, we can bound $\hat{\pi}_k(i, S)$ as follows:

$$\begin{aligned}
\hat{\pi}_k(i, S) &\geq \sum_{m=1}^M \theta_m u_{im} \cdot \left[1 + \sum_{j \in \bar{S}} u_{jm} + \dots + \left(\sum_{j \in \bar{S}} u_{jm} \right)^{k-1} \right] \\
&= \sum_{m=1}^M \theta_m u_{im} \cdot \frac{1 - \left(\sum_{j \in \bar{S}} u_{jm} \right)^k}{1 - \left(\sum_{j \in \bar{S}} u_{jm} \right)} \\
&\geq \left(1 - \max_m u_m(\bar{S})^k \right) \cdot \sum_{m=1}^M \theta_m \cdot \frac{u_{im}}{1 - \left(\sum_{j \in \bar{S}} u_{jm} \right)} \\
&= \left(1 - u_{max}(\bar{S})^k \right) \cdot \pi(i, S).
\end{aligned}$$

This completes the proof. \blacksquare

Proof of Theorem IV.2

Let $\hat{\pi}_k^m(i, S)$ denote the choice probability under k -attempt model by a customer that belongs to segment m . We can write:

$$\hat{\pi}_k^R(i, S) = \frac{\sum_m \theta_m \hat{\pi}_k^m(i, S)}{\sum_m \sum_{j \in S_0} \theta_m \hat{\pi}_k^m(j, S)}.$$

Given the lower bound of $\hat{\pi}_k(i, S)$ in Theorem 1, we can bound:

$$\hat{\pi}_k^R(i, S) \geq \frac{(1 - u_{max}(\bar{S})^k) \cdot \pi(i, S)}{\sum_m \sum_{j \in S_0} \theta_m \hat{\pi}_k^m(j, S)} \geq \frac{(1 - u_{max}(\bar{S})^k) \cdot \pi(i, S)}{\max_m \sum_{j \in S_0} \hat{\pi}_k^m(j, S)} = \frac{1 - u_{max}(\bar{S})^k}{1 - u_{min}(\bar{S})} \cdot \pi(i, S).$$

Similarly, we also have:

$$\hat{\pi}_k^R(i, S) \leq \frac{\pi(i, S)}{\sum_m \sum_{j \in S_0} \theta_m \hat{\pi}_k^m(j, S)} \leq \frac{\pi(i, S)}{\min_m \sum_{j \in S_0} \hat{\pi}_k^m(j, S)} = \frac{1}{1 - u_{max}(\bar{S})} \cdot \pi(i, S).$$

This completes the proof. \blacksquare

Proof of Theorem IV.3

Let $\alpha_l(\bar{S}) = \sum_{j_1, \dots, j_l \in \bar{S}} \frac{u_{j_1}}{1-u_{j_1}} \frac{u_{j_2}}{1-u_{j_1}-u_{j_2}} \dots \frac{u_{j_l}}{1-u_{j_1}-\dots-u_{j_l}}$. Per our note in the proof of Theorem 1, $\alpha_l(\bar{S}) \cdot u_i$ is the probability that a customer values product $j \in \{j_1, j_2, \dots, j_l\}$ better than i and product $j' \in \mathcal{N} \setminus \{j_1, j_2, \dots, j_l\} \cup \{0\}$ worse than i . Since a customer only purchases product i if her other favorite products (which rank higher than i) are not in the offer set, by definition of random utility model, we must have: $\pi(i, S) = \sum_{l=0}^{|\bar{S}|} \alpha_l(\bar{S}) u_i$. As for k -attempt model, since customers only consider up to $k - 1$ substitutes, we can write: $\pi_k(i, S) = \sum_{l=0}^{k-1} \alpha_l(\bar{S}) u_i$. Putting all things together,

$$\pi(i, S) - \hat{\pi}_k^R(i, S) = \sum_{l=0}^{|\bar{S}|} \alpha_l(\bar{S}) u_i - \frac{\sum_{l=0}^{k-1} \alpha_l(\bar{S}) u_i}{\sum_{l=0}^{k-1} \alpha_l(\bar{S}) u(S_0)} = \frac{u_i}{u(S_0)} - \frac{u_i}{u(S_0)} = 0.$$

This completes the proof. ■

BIBLIOGRAPHY

BIBLIOGRAPHY

- Ahn, H.-s., M. Gümüs, and P. Kaminsky (2007), Pricing and manufacturing decisions when demand is a function of prices in multiple periods, *Operations Research*, 55(6), 1039–1057.
- Aviv, Y., and A. Pazgal (2008), Optimal pricing of seasonal products in the presence of forward-looking consumers, *Manufacturing & Service Operations Management*, 10(3), 339–359.
- Basumallick, D., B. Ozdaryal, and C. Madamba-Brown (2013), Perceived value of a mile, *Journal of Revenue & Pricing Management*, 12(1), 8–15.
- Bertsimas, D., and V. V. Mišić (2016), Data-driven assortment optimization, *submitted*.
- Bijmolt, T. H., M. Dorotic, and P. C. Verhoef (2010), Loyalty programs: generalizations on their adoption, effectiveness and design, *Foundation and Trends in Marketing*, 5(4), 197–258.
- Bitran, G., and R. Caldentey (2003), An overview of pricing models for revenue management, *Manufacturing & Service Operations Management*, 5(3), 203–229.
- Bitran, G. R., and S. V. Mondschein (1997), Periodic pricing of seasonal products in retailing, *Management Science*, 43(1), 64–79.
- Blanchet, J. H., G. Gallego, and V. Goyal (2013), A markov chain approximation to choice modeling., in *EC*, pp. 103–104.
- BoardingArea (2015), Why hhonors has variable award pricing within each category, Web-page.
- Bolton, R. N., P. K. Kannan, and M. D. Bramlett (2000), Implications of loyalty program membership and service experiences for customer retention and value, *Journal of the academy of marketing science*, 28(1), 95–108.
- Borgs, C., J. T. Chayes, S. Doroudi, M. Harchol-Balter, and K. Xu (2014), The optimal admission threshold in observable queues with state dependent pricing, *Probability in the Engineering and Informational Sciences*, 28, 101–119.
- Breugelmans, E., T. H. Bijmolt, J. Zhang, L. J. Basso, M. Dorotic, P. Kopalle, A. Minnema, W. J. Mijnlief, and N. V. Wunderlich (2014), Advancing research on loyalty programs: a future research agenda, *Marketing Letters*, pp. 1–13.
- Carlsson, F., and Å. Löfgren (2006), Airline choice, switching costs and frequent flyer programmes, *Applied Economics*, 38(13), 1469–1475.
- Chapple, S., L. Moerman, and K. Rudkin (2010), Ifric 13: accounting for “customer loyalty programmes”, *Accounting Research Journal*, 23(2), 124–145.
- Chen, H., and M. Frank (2004), Monopoly pricing when customers queue, *IIE Transactions*, 36(6), 569–581.

- Chun, S. Y., D. A. Iancu, and N. Trichakis (2015), Points for peanuts or peanuts for points? setting the optimal value of loyalty currency.
- Colloquy (2011), Loyalty census.
- Colloquy (2015), Loyalty census.
- De Vany, A. (1976), Uncertainty, waiting time, and capacity utilization: A stochastic theory of product quality, *The Journal of Political Economy*, pp. 523–541.
- Désir, A., V. Goyal, D. Segev, and C. Ye (2015), Capacity constrained assortment optimization under the markov chain based choice model, *Operations Research, Forthcoming*.
- Dewan, S., and H. Mendelson (1990), User delay costs and internal pricing for a service facility, *Management Science*, 36(12), 1502–1517.
- Dinkelbach, W. (1967), On nonlinear fractional programming, *Management Science*, 13(7), 492–498.
- Dorotic, M., T. H. Bijmolt, and P. C. Verhoef (2012), Loyalty programmes: Current knowledge and research directions*, *International Journal of Management Reviews*, 14(3), 217–237.
- Dorotic, M., P. C. Verhoef, D. Fok, and T. H. Bijmolt (2014), Reward redemption effects in a loyalty program when customers choose how much and when to redeem, *International Journal of Research in Marketing*, 31(4), 339–355.
- Dreze, X., and J. C. Nunes (2004), Using combined-currency prices to lower consumers’ perceived cost, *Journal of Marketing Research*, 41(1), 59–72.
- Economist (2005a), Frequent-flyer miles: In terminal decline?, magazine.
- Economist (2005b), Frequent-flyer miles: Funny money, magazine.
- Elmaghraby, W., and P. Keskinocak (2003), Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions, *Management Science*, 49(10), 1287–1309.
- Feldman, J. B., and H. Topaloglu (2014), Revenue management under the markov chain choice model.
- Gallego, G., and G. Van Ryzin (1994), Optimal dynamic pricing of inventories with stochastic demand over finite horizons, *Management science*, 40(8), 999–1020.
- Gilboa-Freedman, G., R. Hassin, and Y. Kerner (2014), The price of anarchy in the Markovian single server queue, *Automatic Control, IEEE Transactions on*, 59(2), 455–459.
- Gordon, N., and K. Hlavinka (2011), The 2011 forecast of us consumer loyalty program points value.
- Grassmann, W. K. (1979), The economic service rate, *The Journal of the Operational Research Society*, 30(2), pp. 149–155.
- Hallberg, G. (2004), Is your loyalty programme really building loyalty? why increasing emotional attachment, not just repeat buying, is key to maximising programme success, *Journal of Targeting, Measurement and Analysis for Marketing*, 12(3), 231–241.
- Hassin, R., and M. Haviv (2003), *To queue or not to queue: Equilibrium behavior in queueing system*, vol. 59, Springer Science; Business Media.
- Haviv, M., and T. Roughgarden (2007), The price of anarchy in an exponential multi-server, *Operations Research Letters*, 35(4), 421–426.
- Hess, S., and A. Daly (2014), *Handbook of choice modelling*, Edward Elgar Publishing.

- HiltonWorldwide (2015a), Hilton honor fact sheet.
- HiltonWorldwide (2015b), A world of opportunities, White paper.
- IdeaWorks (2015), Airline ancillary revenue projected to be 59 billion usd worldwide in 2015, White paper.
- Kim, B.-D., M. Shi, and K. Srinivasan (2001), Reward programs and tacit collusion, *Marketing Science*, 20(2), 99–120.
- Kim, B.-D., M. Shi, and K. Srinivasan (2004), Managing capacity through reward programs, *Management Science*, 50(4), 503–520.
- Kivetz, R. (2003), The effects of effort and intrinsic motivation on risky choice, *Marketing Science*, 22(4), 477–502.
- Kivetz, R., and I. Simonson (2002a), Earning the right to indulge: Effort as a determinant of customer preferences toward frequency program rewards, *Journal of Marketing Research*, 39(2), 155–170.
- Kivetz, R., and I. Simonson (2002b), Self-control for the righteous: Toward a theory of precommitment to indulgence, *Journal of Consumer Research*, 29(2), 199–217.
- Knudsen, N. C. (1972), Individual and social optimization in a multiserver queue with a general cost-benefit structure, *Econometrica*, 40(3), pp. 515–528.
- Kochenberger, G., J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang (2014), The unconstrained binary quadratic programming problem: a survey, *Journal of Combinatorial Optimization*, 28(1), 58–81.
- Kök, A. G., and M. L. Fisher (2007), Demand estimation and assortment optimization under substitution: Methodology and application, *Operations Research*, 55(6), 1001–1021.
- Kopalle, P. K., Y. Sun, S. A. Neslin, B. Sun, and V. Swaminathan (2012), The joint sales impact of frequency reward and customer tier components of loyalty programs, *Marketing Science*, 31(2), 216–235.
- Kuo, C.-W., H.-S. Ahn, and G. Aydin (2011), Dynamic pricing of limited inventories when customers negotiate, *Operations research*, 59(4), 882–897.
- Lariviere, M. A. (2006), A note on probability distributions with increasing generalized failure rates, *Operations Research*, 54(3), 602–604.
- Lewis, M. (2004), The influence of loyalty programs and short-term promotions on customer retention, *Journal of Marketing Research*, 41(3), 281–292.
- Lippman, S. A., and J. Stidham, Shaler (1977), Individual versus social optimization in exponential congestion systems, *Operations Research*, 25(2), pp. 233–247.
- Liston-Heyes, C. (2002), Pie in the sky? real versus perceived values of air miles, *Journal of Consumer Policy*, 25(1), 1–27.
- Liu, Y. (2007), The long-term impact of loyalty programs on consumer purchase behavior and loyalty, *Journal of Marketing*, 71(4), 19–35.
- Maglaras, C., and J. Meissner (2006), Dynamic pricing strategies for multiproduct revenue management problems, *Manufacturing & Service Operations Management*, 8(2), 136–148.
- Marriott (2013), Annual report.
- Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995), *Microeconomic Theory*, vol. 1, Oxford university press New York.

- McFadden, D., K. Train, et al. (2000), Mixed mnl models for discrete response, *Journal of applied Econometrics*, 15(5), 447–470.
- Mendelson, H. (1985), Pricing computer services: Queueing effects, *Commun. ACM*, 28(3), 312–321.
- Naor, P. (1969), The regulation of queue size by levying tolls, *Econometrica*, 37(1), pp. 15–24.
- NerdWallet (2015), Nerdwallet consumer credit card report, *Tech. rep.*, NerdWallet.
- Ollila, J. (2012), Unintended feature of hilton’s new website: Award reimbursement rates shown, Webpage.
- Ollila, J. (2013), Marriott’s award night reimbursements to hotels, Webpage.
- PointsGuy (2015), Redeeming chase ultimate rewards points for maximum value, Webpage.
- Roughgarden, T. (2005), *Selfish routing and the price of anarchy*, vol. 174, MIT press Cambridge.
- Rusmevichientong, P., Z.-J. M. Shen, and D. B. Shmoys (2010), Dynamic assortment optimization with a multinomial logit choice model and capacity constraint, *Operations research*, 58(6), 1666–1680.
- Sharp, B., and A. Sharp (1997), Loyalty programs and their impact on repeat-purchase loyalty patterns, *International journal of Research in Marketing*, 14(5), 473–486.
- Shin, J., and K. Sudhir (2010), A customer management dilemma: When is it profitable to reward one’s own customers?, *Marketing Science*, 29(4), 671–689.
- Stidham, J., Shaler (1992), Pricing and capacity decisions for a service facility: Stability and multiple local optima, *Management Science*, 38(8), pp. 1121–1139.
- Stourm, V., E. T. Bradlow, and P. S. Fader (2015), Stockpiling points in linear loyalty programs, *Journal of Marketing Research*, 52(2), 253–267.
- Tawarmalani, M., and N. V. Sahinidis (2002), *Convexification and global optimization in continuous and mixed-integer nonlinear programming: theory, algorithms, software, and applications*, vol. 65, Springer Science & Business Media.
- Uncles, M. D., G. R. Dowling, and K. Hammond (2003), Customer loyalty and customer loyalty programs, *Journal of consumer marketing*, 20(4), 294–316.
- Yechiali, U. (1971), On optimal balking rules and toll charges in the GI/M/1 queuing process, *Operations Research*, 19(2), pp. 349–370.
- Yechiali, U. (1972), Customers’ optimal joining rules for the GI/M/s queue, *Management Science*, 18(7), pp. 434–443.
- You, F., P. M. Castro, and I. E. Grossmann (2009), Dinkelbach’s algorithm as an efficient method to solve a class of minlp models for large-scale cyclic scheduling problems, *Computers & Chemical Engineering*, 33(11), 1879–1889.
- Yue, D., G. Guillén-Gosálbez, and F. You (2013), Global optimization of large-scale mixed-integer linear fractional programming problems: A reformulation-linearization method and process scheduling applications, *AIChE Journal*, 59(11), 4255–4272.
- Zhang, D., and W. L. Cooper (2005), Revenue management for parallel flights with customer-choice behavior, *Operations Research*, 53(3), 415–431.
- Ziya, S., H. Ayhan, and R. D. Foley (2004), Relationships among three assumptions in revenue management, *Operations Research*, 52(5), 804–809.