MODELS FOR ASSORTMENT PLANNING UNDER PRODUCT RETURNS

By

ALEXANDRE GRASAS

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2009

To my wife Mireia

ACKNOWLEDGMENTS

This represents the finale of one of the most enriching experiences in my life, and I would like to thank all people who helped along the way. First, I am indebted to Jaume Ribera, who saw in me potential to pursue an academic career and encouraged me to start this journey. Second, I wish to express my sincere gratitude to both of my advisors, Elif Akçalı and Aydın Alptekinoğlu for their immense help and guidance. *Bana verdiğiniz desteğin kıymetini biliyor, her ikinize de çok teşekkür ediyorum.* I am thankful to the other members of my Supervisory Committee, Professors Joseph Geunes, H. Edwin Romeijn, Asoo J. Vakharia and Janice E. Carrillo for their time and helpful comments. I also would like to acknowledge the financial support provided by the National Science Foundation under Grants No. 0522960 and 0536026.

Gainesville has been my home for more than four years, and I have had the opportunity to meet wonderful people with whom I shared my ups and downs during this adventure. They have been my family *across the pond.* Of all of them, I would like to mention Vanessa, Ignasi, Caro, Guille, Laura, Martín, Jess, Guðbjört, Árni, Gaia, Mehmet and Richard. I also thank all ISE grad students for all gatherings, BBQ's and good moments in Weil Hall. My thankfulness also goes to all my soccer buddies for keeping me in shape and allowing me to win all those Gator T-shirts.

Next, I wish to thank my family for their unconditional love and support (and their visits): my parents Joan and Laura, my brothers Edu and Joan, my beloved sister Laia, my parents-in-law Modest and Chus and my brother-in-law David. They were always understanding despite distance and the time we have been abroad.

Vull agrair a la meva família la seva estima i suport incondicional (i les seves visites): els meus pares Joan i Laura, els meus germans Edu i Joan, la meva estimada Laia, els meus sogres Modest i Chus, i el meu cunyat David. Han estat sempre comprensius malgrat la distància i el temps que hem estat fora.

Finally, my most heartfelt appreciation goes to my dearly loved Mireia, *la meva Mir*. None of this would have been possible without her love, support and encouragement. Here we start another chapter in our lives, and I am immensely happy to share it with her.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

MODELS FOR ASSORTMENT PLANNING UNDER PRODUCT RETURNS

By

Alexandre Grasas

May 2009

Chair: Elif Akçalı
Cochair: Aydın Alptekinoğlu
Major: Industrial and Systems Engineering

In the past decade, internet and flexible manufacturing have revolutionized some of the basic principles of retailing. Two such aspects relate to product assortment and return policies. With the aid of advanced production technology, companies continuously increase their product assortments to reach more customers and satisfy their specific needs better. Higher product variety, however, typically raises operational complexity and costs. In addition, these costs can even be more significant when product returns are considered.

We integrated return policies into a multiproduct model, where assortment, inventory, and/or pricing decisions were made in an integrated manner. Our research agenda focused on an expected-profit-maximizing firm that offers a set of horizontally differentiated products. The firm accepts product returns that are in resalable condition. We characterized the firm's return policy by the money refunded to the customer in case of return. We have a demand model that is based on individual consumer behavior, conceptualized to fit a well established two-stage utility maximization framework (nested multinomial logit model). Consumers decide which product (if any) out of a given assortment to buy in the first stage, and then decide to keep or return the item in the second stage.

Our study shows an interesting interaction between product assortment and return policy. We explored the implications of return policies on product assortment planning. We showed that the structure of the optimal assortment fundamentally changes depending

on the amount refunded and/or operational mode (make-to-order versus make-to-stock). Surprisingly, there are situations where a retailer is better off by offering eccentric products (i.e., those that are least likely to be purchased by a typical consumer). We also explored return policies for customized products. We determined that customizing firms should aim for product returns that are neither a net cost nor a net benefit. This can be achieved by partial refunding when designing their consumer return policies. In addition, we were the first to investigate return policies in a multiperiod environment. Restricting return policies to a single period analysis only, as other authors do, may lead to wrong conclusions. We show, for example, that firms can reduce their price when they consider multiple periods. In this multiperiod setting, we found that a salvage-down-to level inventory policy is optimal.

CHAPTER 1
INTRODUCTION

1.1 Background and Motivation

Retailing is the second-largest industry in the United States by number of businesses and number of employees. Plunkett Research, Ltd., an industry sector analysis and research provider, estimated retail sales to be $4.49 trillion in 2007. In the past decade, internet and flexible manufacturing have revolutionized some of the basic principles of retailing. Two such aspects relate to product assortment and consumer return policies. With the aid of advanced production technology, companies continuously increase their product assortments and offer product customization opportunities. The possibility of reaching more customers and satisfying their specific needs better, have led firms to expand their assortments both in terms of breadth and depth. Higher product variety, however, typically raises operational complexity and costs. In addition, these costs can even be more significant when consumer returns are considered. It is estimated that overall, customer returns represent 7% of sales in the United States, according to the National Retail Federation, and may run as high as 15% for mass merchandisers and up to 35% for catalog and e-commerce retailers (Rogers and Tibben-Lembke 1998, pp. 6–8). The annual value of returned goods in the United States is approximately $100 billion, and companies spend more than $40 billion annually on their reverse logistics processes for handling and disposition of returns (Blanchard 2005, Enright 2003).

Return policies are usually thought of as micro and more operational, whereas product assortment is usually thought of as strategic and more marketing related. Therefore, decisions associated with each are often made separately (see Stock et al. 2006, and Olavson and Fry 2006). In our research, we integrate return policies into a multiproduct model, where assortment, inventory, and/or pricing decisions are made in an integrated manner; something that has never been explored before in the literature. We show that integrating these decisions Our research agenda focuses on an expected-profit-maximizing firm that offers a set of horizontally differentiated products. The firm accepts product returns that are in resalable condition. These returns result from a late realization of misfit with consumer preferences (e.g., a garment not feeling right). We characterize the firm's return policy by the money refunded to the customer in case of return. We have a demand model that is based on individual consumer behavior, conceptualized to fit a well established two-stage utility maximization framework (nested multinomial logit model). Consumers decide which product (if any) out of a given assortment to buy in the first stage, and then decide to keep or return the item in the second stage.

1.2 Impact

As we discussed earlier, the issues that relate to product returns and assortment planning are studied independently in operations management and marketing literatures. In our work, we attempt to merge these two streams of research to study the influence of product returns on assortment planning. Our objective is to develop a basic set of models and analyze these models to investigate the interaction between assortment planning and product return policies. By concentrating on these models, we can generate some helpful insights for retailers. For example, product returns should be considered when designing product assortments as they may change their optimal composition. For customizing firms, establishing efficient return policies and reaching profitable secondary markets improve both customer purchasing experience and firms' benefits.

CHAPTER 2
LITERATURE REVIEW

The topic of this dissertation merges two streams of literature that have been traditionally separate: assortment planning and product returns. In this Chapter we provide a general overview of the two streams of literature. In Chapters 3 and 5 we discuss literature related to each specific problem in more detail.

2.1 Assortment Planning

The impact of product assortment on consumer behavior has been studied in the marketing literature extensively (Broniarczyk et al. 1998, Hoch et al. 1999, Boatwright and Nunes 2001, van Herpen and Pieters 2002, Borle et al. 2005, de Vries-van Ketel 2006, Berger et al. 2007). There is also a newly burgeoning stream of product variety literature in operations management (OM). In a seminal paper, van Ryzin and Mahajan (1999) introduce operational costs (i.e., inventory costs) to the assortment planning problem. Using the multinomial logit (MNL) model for the consumer choice process, they show that the optimal assortment has a very simple structure: it consists of some number of most popular products. As we prove analytically in Chapter 3, consideration of product returns can reverse this intuitive result, also shown in various other contexts (Aydin and Ryan 2000, Hopp and Xu 2005, Maddah and Bish 2007, Li 2007). Later, Cachon et al. (2005) introduce the consumer search to the assortment problem. In their model, the consumers can opt to leave without purchasing, and search for the product elsewhere. The authors find that, when the different products within a category are limited (e.g., digital cameras), the firm expands its assortment to prevent the customer from balking.

Gaur and Honhon (2006) also study the assortment planning problem with inventory costs, but, unlike the other papers cited, they use a locational choice model to characterize the demand process, and point out interesting differences it makes. Under static substitution (that is, costumers do not substitute in the event of a stockout), Gaur and Honhon find that the optimal assortment contains products that are equally spaced out and sufficiently far apart from one another to eliminate substitution. They also provide two heuristics for the dynamic substitution version of the problem. Maddah and Bish (2007) extend the work of van Ryzin and Mahajan (1999) by endogenizing the pricing decision; they derive the structure of the optimal assortment when all products have the same unit cost and different endogenous prices. Li (2007) proposes an assortment and inventory joint optimization problem where the cost parameters for the products are allowed to be different. He determines the optimal structure assuming the store traffic is continuous. The optimal assortment includes some number of products with the highest profit rate (i.e., expected profit from a product if it were to attract 100% of the store traffic). Finally, the reader is referred to Kök et al. (2006) for a complete review of the assortment planning literature.

2.2 Product Returns

Literature on product returns is very extensive. From a marketing and/or an operational perspective, there are many different research lines that deal with product returns. An important bulk of research in operations management (OM) literature focuses on the management of end-of-use / end-of-life returns. Operational, tactical and strategic decisions associated with used product returns have been well-studied in the closed-loop supply chain management literature (for an overview, see Guide and van Wassenhove 2003, and Dekker et al. 2004). Involving unused products, Pasternack (1985), Emmons and Gilbert (1998), and Tsay et al. (1999) analyze return policies offered by manufacturers to retailers (sometimes also framed as *buy-back contracts*) for products that remain unsold at the end of a selling season.

In our research, we focus on consumer product returns, that is, those returns offered by retailers to consumers, where a returned product is usually in resalable condition. In OM, arguing that returns need to be taken into account in inventory management, since they can act as a supplementary source to satisfy demand, the existing research focuses on characterizing the optimal ordering policy of a retailer (Vlachos and Dekker

2003, Mostard et al. 2005, and Mostard and Teunter 2006). These papers assume a given consumer return policy. To the best of our knowledge, only three papers allow endogenous refund decisions in a retailing context. Mukhopadhyay and Setoputro (2005) adopt a deterministic linear demand model that depends on price, refund and modularity, where the last two are decision variables. Yalabik et al. (2005) integrate logistics and marketing decisions into the return system, and Su (2008) endogenizes price and order quantity in addition to refund. In these two papers, demand is driven by consumer valuation of the product.

In marketing literature, product returns research concentrates on the influence of a retailer's return policy on consumers (Wood 2001). For example, Davis et al. (1995), Che (1996), Davis et al. (1998) and Heiman et al. (2002) study the implications of full money-back guarantees on consumers' behavior. Others take refund as decision variable, and show the benefits of stricter return policies (Hess et al. 1996, Mukhopadhyay and Setoputro 2004, Shulman et al. 2007, and Shulman et al. 2008).

CHAPTER 3
ASSORTMENT PLANNING UNDER EXOGENOUS PRICE AND REFUND FRACTION IN A SINGLE PERIOD SETTING

3.1 Introduction

Should retailers take product returns into account when merchandising (choosing their product assortments)? Return policies are usually thought of as micro and more operational, whereas product assortment is usually thought of as strategic and more marketing related. Therefore, decisions associated with each are often made separately (see Stock et al. 2006, and Olavson and Fry 2006). Our theoretical model counters this conventional thinking by showing that optimal assortment decisions fundamentally change in the presence of returns.

If a consumer decides to return a product due to quality problems (i.e., as it is damaged or does not work), then typically the retailer returns this product to the manufacturer and is compensated for it. However, if the consumer decides to return the product due to a late realization of misfit with her preferences (i.e., the product is fine but the consumer simply does not want it anymore, e.g., a garment not feeling right), then the retailer has to handle the return. Our focus in this dissertation is the latter type of returns involving a product in resalable condition.

Financial impact of return policies can be quite large for a retailer. Overall customer returns are estimated to be 6% of sales in the United States, and may run as high as 15% for mass merchandisers and up to 35% for catalog and e-commerce retailers (Rogers and Tibben-Lembke 1998, pp. 6–8). The annual value of returned goods in the United States is approximately $100 billion, and companies spend more than $40 billion annually on their reverse logistics processes for handling and disposition of returns (Blanchard 2005, Enright 2003).

Motivated with the question of whether retailers should consider returns when merchandising, in this chapter we explore the interactions between product assortment decision and return policy of a price-taking retailer under both make-to-order (MTO) and make-to-stock (MTS) environments. In some product categories, or with particular brands, many retailers do not dictate prices, but rather sell their products at MSRP, *manufacturer suggested retail price* (e.g., backcountry.com, a retailer specialized in high-end gear and apparel for outdoors, sells many of its products at MSRP; Crocs Shoes, a manufacturer and online retailer of shoes and other footwear, exercises a very high degree of control over the retail price of its products available in many online and brick-and-mortar retailers). The two basic operational modes, MTO and MTS, allow us to draw a distinction between cases where supply decision is made after and before the demand materializes, respectively. In the MTO case, the retailer procures the product after consumers make their purchase decisions (e.g., many of the sports gears sold online at REI.com are drop-shipped directly from a third-party supplier). Whereas in the MTS case, the reverse happens (e.g., backcountry.com carries all of its products in inventory at its warehouse in Utah).

We have a demand model that is based on individual consumer behavior, conceptualized to fit a well established two-stage utility maximization framework (nested multinomial logit model). Consumers decide which product (if any) out of a given assortment to buy in the first stage, and then decide to keep or return the item in the second stage. On the supply side, the retailer makes an assortment decision by choosing a subset of all potential product offerings that fall within a particular product line of horizontally differentiated items. In the MTS case, the retailer also makes an inventory decision for each product offered. Price is exogenous (i.e., dictated by the manufacturer through MSRP), and products differ only in terms of their attractiveness (defined precisely in §3.3). We call products with high (low) attractiveness *popular* (*eccentric*), because they are more (less) likely to be purchased by a typical consumer.

We exclusively focus on one aspect of return policies: refund amount, which we parameterize by *refund fraction*, the percentage of price refunded in the event of a return. Like price, we assume refund fraction to be exogenous, possibly driven by a category-or store-wide analysis (beyond the scope of ours that focuses on a single horizontally differentiated product line), or dictated by common industry practice. While it is common to offer refunds for the full purchase price in some settings (e.g., backcountry.com allows customers to send products back for a full refund with no questions asked), offering partial refunds and retaining some portion of the price in restocking fees is common in others (e.g., buydig.com, a retailer of consumer electronics, charges a processing fee of 10% of the value of all merchandise returned for a refund)[1] .

We show that the structure of the optimal assortment, which maximizes the retailer's expected profit, critically depends on the refund fraction and whether the products are supplied on an MTO or MTS basis. More specifically, we have two major results:

- For a strict return policy (with a sufficiently low refund fraction), the optimal assortment has a counterintuitive structure. In the MTO case, it is composed of some number of most eccentric products; whereas, in the MTS case, some number of most popular and some number of most eccentric products.
- For a lenient return policy (with a sufficiently high refund fraction), the optimal assortment is composed of some number of most popular products in the MTO case. Although we could not analytically prove that the same structure is optimal for the MTS case as well, our extensive numerical experiments reported in §3.5 confirm this. Including only the most popular products in an assortment agrees with common intuition, previous results in the literature (van Ryzin and Mahajan 1999, Aydin and Ryan 2000, Hopp and Xu 2005, Maddah and Bish 2007, Li 2007, and Cachon and Kök 2007), and some industry practice (Cargille et al. 2005, and Olavson and Fry 2006). As indicated above, we show that the presence of returns can reverse this intuitive result.

The basic rationale for including an eccentric product in the optimal assortment is to benefit from the processing and resale of returned items. This benefit is higher for low refund fractions, and eccentric products have a higher likelihood of being returned. The

[1] Newegg.com charges 15% for all returned items. Best Buy and Target charge 15% for many consumer electronics items. Returning a home theater set to Circuit City open box, even if not used at all, incurs 25% restocking fee. See van Riper and Nolan (2008) for more examples.

case for popular products, on the other hand, is two-fold. If returns are a net loss to the retailer, popular products become desirable because they minimize the likelihood of return. If the retailer is operating in an MTS mode, popular products also have the advantage of lower relative demand variability (measured by coefficient of variation) and therefore reduced operational risks.

In light of our analytical results (presented in more detail in §3.4) and numerical observations (reported in §3.5), we conclude that: retailers should not only carefully consider their return policy when merchandising, they should also take their basic operational mode (MTO versus MTS) into account.

3.2 Literature Review

Product assortment planning or product variety management has attracted considerable interest in the literature from various different angles: strategic/competitive aspects of product variety (Shugan 1989, Bayus and Putsis 1999, Cachon and Kök 2007, and Alptekinoğlu and Corbett 2008b); impact of product variety on consumer behavior (Hoch et al. 1999, Kim et al. 2002, and Borle et al. 2005); and interactions between product variety and operational considerations such as inventory and leadtime (van Ryzin and Mahajan 1999, Smith and Agrawal 2000, Aydin and Ryan 2000 , Cachon et al. 2005, Hopp and Xu 2005, Gaur and Honhon 2006, Li 2007, Maddah and Bish 2007, and Alptekinoğlu and Corbett 2008a). Presence of product returns obviously complicates assortment planning further, yet it has not been addressed in this literature so far. To the best of our knowledge, our work is the first in posing an assortment planning problem that incorporates returns. We demonstrate a specific setting when returns make a fundamental difference for assortment decisions - beyond just complicating them.

Although operational, tactical and strategic decisions associated with used product returns have been well-studied in the closed-loop supply chain management literature (for an overview, see Guide and van Wassenhove 2003, and Dekker et al. 2004), research on resalable product returns has been somewhat limited. Arguing that returns need to be taken into account in inventory management, since they can act as a supplementary source to satisfy demand, the existing research focuses on characterizing the optimal ordering policy of a retailer (Vlachos and Dekker 2003, Mostard et al. 2005, and Mostard and Teunter 2006). Guide et al. (2006) note the value that can be recovered from returns is time sensitive and focus on identifying the preferred reverse supply chain structure for a manufacturer. This entire line of work exclusively treats single product systems. Therefore, by considering assortment planning, we tackle a host of issues that have been ignored by the current literature on operations management of returns.

Another line of research that is closely related to our work pertains to product return policies. While a stream of research focuses on return policies between a manufacturer and a retailer (Pasternack 1985, Padmanabhan and Png 1997, and Emmons and Gilbert 1998), another stream concentrates on the influence of a retailer's return policy on consumers (Wood 2001, Yalabik et al. 2005, and Shulman et al. 2008). Our work is similar to some of the work in the latter stream in that we have an explicit model of consumer choice, and limit attention to a single aspect of return policies: refund amount. The difference is that we explore how return policy interacts with product assortment, an issue none of these papers address.

3.3 Model

We consider an expected-profit-maximizing retailer that takes product assortment decisions in a single period setting. Our main research objective is to explore the interactions between optimal assortment and return policy in make-to-order (MTO) and make-to-stock (MTS) environments. We use these terms in a broader sense than their traditional use in the literature; the firm in our model does not necessarily 'make' what it is selling. MTO refers to retailing environments where the quantity decision is made *after* the realization of demand; the retailer does not stock the item but requests it from the manufacturer (or any other type of supplier) once an order is placed by a consumer (e.g., REI practices drop-shipping). In the case of MTS, the quantity decision is made *before* the realization of demand (e.g., backcountry.com fulfills orders from its central warehouse).

3.3.1 Product Assortment and Return Policy

When choosing its product assortment, we assume that the retailer exclusively considers an exogenous set of potential product designs; this set may represent a supplier's catalog of different variants in a given product line. Let $N = \{1, 2, ..., n\}$ denote set of all products that the retailer can potentially offer, and let S be the subset of products actually offered by the retailer ($S \subseteq N$), termed *assortment*.

A standard assumption in the literature with regard to assortment decisions is that the firm incurs a fixed cost per product included in S (see, for instance, Smith and Agrawal (2000) for a discussion of what this fixed cost may entail in retailing, p. 55). We do not make this assumption for two reasons: parsimony and accent. Analytically speaking, such a fixed cost can be easily incorporated in our model, and it would not notably influence any of our analytical results or managerial insights. Therefore, we choose to drop it for ease of exposition. Secondly, fixed cost itself would be a reason to offer less variety. Since we already have such a reason in the model, inventory risks (detailed below), we wish not to confound the effect of MTO/MTS environment and the associated production/inventory policies on variety and returns. In other words, the current model without a fixed cost for variety gives more prominence to our results, some of which may otherwise be perceived as driven by fixed costs.

Assortment decision (S) considered here is for a narrow category of products, which are horizontally differentiated along a taste attribute such as color or some other component of fashion. All products in N are assumed to have the same unit production cost c, the same retail price p, and the same salvage value v. There is only one difference among the products in question: their *attractiveness* (a's introduced below). Following standard practice, we assume that $v < c < p$. The latter inequality, $c < p$, is necessary for the market to be profitable. The former inequality, $v < c$, says that any amount of leftovers can be sold below-cost in a secondary market for v per unit; if $v \geq c$ were to hold true, the retailer's quantity decision would be riskless and thus uninteresting.

Furthermore, as discussed in the Introduction, we assume exogenous prices. Allowing prices to be decision variables would be clearly useful, but also analytically very difficult (see Maddah and Bish (2007) for an attempt at endogenizing price in an MNL-choice-based assortment problem that also considers inventories but omits product returns). Yet, as pointed out by van Ryzin and Mahajan (1999) in the context of a closely related model, there are "realistic cases in which a retailer's pricing flexibility is quite limited" (p. 1498). We limit our analysis to such a case, as they also do, with the retailer exercising little or no control over prices, e.g., it sells the product line in question at MSRP.

The types of returns we consider involve products returned in resalable condition. Again, as discussed in the Introduction, we assume an exogenous return policy, and focus on one aspect of it: percentage of price refunded by the retailer when a consumer returns a product. Let α denote the *refund fraction* ($0 \leq \alpha \leq 1$), which makes the refund amount per unit return αp. We assume that this single refund fraction applies to all products in S, which is how almost all retailers operate in practice (especially within a given narrow product category, as in our model). The retailer incurs a reverse logistics cost l for each unit of returned products. This figure includes such cost items as sorting, repackaging, and restocking.

Finally, consistent with common practice in retailing, we omit the possibility of product exchange. Many retailers, including backcountry.com (sports gear), Lids.com (baseball caps), Steve Madden (shoes), and buydig.com (consumer electronics), allow returns and ask consumers to place a new order if they want to do an exchange even for another product in the same product line. Excluding exchanges from consideration is not without loss of generality, of course, because those new orders would go to subsequent periods, which we do not model. Allowing exchanges is akin to dynamic substitution, which is known to pose great difficulties in assortment optimization (more about this in the discussion of MTS environment).

3.3.2 Individual Consumer Choice Behavior and Aggregate Demand

Any given consumer's consideration set comprises all the products in S offered by the retailer and the possibility of not purchasing any of those products, termed the *outside option*, which we denote by 0. We conceptualize the consumers' choice among $S \cup \{0\}$ and their subsequent decision to keep or return the purchased product by the nested multinomial logit (N-MNL) model. The nests are products, and they each contain two post-purchase alternatives: keep and return.

In the N-MNL framework, consumer choice can be viewed as a sequential process in which the consumer first chooses a product in S or the outside option with probability P_i^S, where $i \in S \cup \{0\}$. Then, conditional on this first choice, the consumer chooses to keep or return the purchased product (if any) with respective probabilities $P_{keep|i}$ and $P_{return|i}$ for $i \in S$. Hence the joint probability of choosing $i \in S$ and $t \in \{keep, return\}$ is $P_{it}^S = P_i^S P_{t|i}$. We now describe this two-stage choice process in more detail.

***Stage* 2.** Conditional on purchasing product $i \in S$ in the first stage, we model the consumer's post-purchase decision to keep or return the product by utility maximization. Let the *attractiveness* of product i be a_i, which may differ across the products but not across consumers. Suppose the utilities associated with purchasing product i and keeping or returning it are given by: $u_{i,keep} = a_i - p + \epsilon_{i,keep}$, and $u_{i,return} = -(1-\alpha)\,p + \epsilon_{i,return}$, where $\epsilon_{i,keep}$ and $\epsilon_{i,return}$ are independent and identically distributed (*iid*) Gumbel random variables with mean zero and scale $1/\mu_2$ ($\mu_2 > 0$)[2] . Note that the deterministic portion of

[2] The cumulative distribution function (*cdf*) of a Gumbel random variable X with mean zero and scale $1/\mu$ is given by $P(X \leq x) = \exp\left[-\exp\left(-\frac{x}{\mu} - \gamma\right)\right]$, and has a variance of $\mu^2\pi^2/6$, where γ is Euler's constant ($\gamma \approx 0.5772$) and μ is a positive constant. Gumbel distribution is also known as double-exponential distribution.

$u_{i,keep}$ is the attractiveness minus the price; and the deterministic portion of $u_{i,return}$ is the negative of the dollar amount not refunded by the retailer. (If returns involve a fixed cost or disutility for the consumer, we could incorporate a deterministic parameter in $u_{i,return}$ to account for that; none of our findings would change as a result.)

By the principle of utility maximization, the probability that a typical consumer chooses the return option in the second stage is then $\Pr\{u_{i,return} > u_{i,keep}\}$, which yields the following formula [3] :

$$P_{return|i} = \frac{1}{1 + \exp\left(\frac{a_i - \alpha p}{\mu_2}\right)}$$

And, of course, $P_{keep|i} = 1 - P_{return|i}$. Should the consumer choose the outside option in the first stage, there is no further choice to make in the second stage. Note that $P_{return|i}$ is non-zero even if the retailer offers no refund ($\alpha = 0$). This is largely a matter of scaling; the model should be calibrated such that $P_{return|i}$ is negligibly small when $\alpha = 0$, because most consumers would probably not return the product for no refund.

***Stage* 1.** For a consumer who is grappling with the first stage decision of which product to purchase (if any), the *expected utility* of product $i \in S$ (or nest i) is $A_i \equiv E\left[\max\left(u_{i,keep}, u_{i,return}\right)\right]$, which can be derived as (see Anderson et al. (1992) for a generic proof):

$$A_i = \mu_2 \ln\left[\exp\left(\frac{a_i}{\mu_2}\right) + \exp\left(\frac{\alpha p}{\mu_2}\right)\right] - p$$

Furthermore, we assume without loss of generality that the outside option is a nest with zero expected utility, that is, $A_0 = 0$.

We model the consumer's purchase decision also by utility maximization. Suppose the utility of choosing nest $i \in S \cup \{0\}$ is given by: $U_i = A_i + \varepsilon_i$, where ε_i are *iid* Gumbel

[3] We use the fact that the difference of two Gumbel random variables, ϵ_1 and ϵ_2, with scale $1/\mu$ follows a logistic distribution with *cdf* given by: $\Pr\{\epsilon_2 - \epsilon_1 \leq x\} = \left[1 + \exp\left(\frac{-x}{\mu}\right)\right]^{-1}$. See Johnson and Kotz (1970) for a proof.

random variables with mean zero and scale $1/\mu_1$ ($\mu_1 > 0$). (ε_i are also independent from $\epsilon_{j,keep}$ and $\epsilon_{j,return}$ for all $i, j \in N$.) Again by the principle of utility maximization, the probability that nest $i \in S \cup \{0\}$ is chosen in the first stage is $\Pr\left\{U_i = \max_{j \in S\cup\{0\}} U_j\right\}$, which yields the following logit formula[4] :

$$P_i^S = \frac{\exp\left(\frac{A_i}{\mu_1}\right)}{\sum_{j\in S\cup\{0\}} \exp\left(\frac{A_j}{\mu_1}\right)}$$

where P_0^S denotes the probability of choosing the outside option or not buying. Note that while conditional probability of return $P_{return|i}$ only depends on a_i (i.e., it is independent of the rest of the products in S), the unconditional probability of return $P_{return} = \sum_{j\in S} P_{return|j} P_j^S$ does depend on the retailer's assortment S.

In sum, we represent consumers' choice process with a two-stage random utility model. Consumers are *a priori* homogenous, but *ex post* heterogeneous on their tastes, preferences, and outside factors that may shape their pre- and post-purchase decisions. The random terms capture this heterogeneity. In particular, ε_i reflect consumers' diverse preferences for products and return policies, their diverse circumstances in which they need this product, their diverse information states, etc. They also differ in their post-purchase inclinations, as summed up in $\epsilon_{i,keep}$ and $\epsilon_{i,return}$. Heterogeneity at this stage stems from how different consumers deal with keep and return options given a purchase decision in the first stage. For instance, among two consumers who are considering to keep an apparel item, their spouses may give them different feedback. And, among two consumers who are considering to return a pair of hiking shoes, their experience with the product may differ due to their different backgrounds (or lack thereof) in hiking. Larger μ_1 and μ_2 mean higher variance for the random terms and thus higher heterogeneity. For the

[4] This is a standard result that comes from the fact that maximum of Gumbel random variables has a Gumbel distribution (see Anderson et al. (1992) for a proof).

N-MNL model to be technically consistent, we require $\mu_1 \geq \mu_2$ (McFadden 1978), which is plausible in our context. Consumers' pre-purchase heterogeneity is generally higher than their post-purchase heterogeneity, because presumably those who buy the same product will know more about what they want (or do not want) based on first-hand experience with a given product, and will differ less from each other due to this common experience.

We will make a semantic distinction between products with high and low values of attractiveness. The higher the attractiveness of a product, the higher the expected utility of consuming it (i.e., buying and keeping it), and thus higher the probability of purchase. In view of utility maximization behavior described above, every consumer buys what they consider to be the best or most 'attractive' product. So, the magnitude of a_i does not so much reflect the attractiveness of a product in the common sense of the word, but rather determines the likelihood of purchase for product i. We will thus refer to products with high attractiveness values as *popular* products (in the sense that a typical consumer is more likely to buy them); and, those with low attractiveness values as *eccentric* products (in the sense that consumers with rare tastes will buy them).

We now specify how individual consumer choice behavior described above translates into aggregate demand for each product in S. Let λ denote the average number of consumers going through this choice process. Assuming that the consumers' product choice is purely governed by the set S, and not influenced at all by the details of the retailer's fulfillment process (e.g., MTO versus MTS, inventory status, etc.), we model the demand for product $i \in S$ by a normal random variable D_i with mean λP_i^S and standard deviation $\sigma \left(\lambda P_i^S\right)^\beta$, where $\sigma > 0$ and $0 \leq \beta < 1$. (This model of aggregate demand, dubbed the Independent Population Model, has been first proposed by van Ryzin and Mahajan (1999), and later used by Maddah and Bish (2007), Li (2007) and others.) Furthermore, we model the returns of product i by a normal random variable R_i with mean $\lambda P_{i,return}^S$ and standard deviation $\sigma \left(\lambda P_{i,return}^S\right)^\beta$. Note that the coefficient of variation (defined as standard deviation divided by mean) for D_i and R_i are decreasing in P_i^S and $P_{i,return}^S$, respectively. Also, Poisson demands and returns constitute a natural special case of our aggregate demand model (i.e., $\sigma = 1$ and $\beta = 1/2$).

3.3.3 Supply Process and the Timing of Events

We consider two alternative modes of supply: MTO and MTS. As mentioned before, we use these terms broadly, in essence, to draw a distinction between the cases in which the supply amount (order quantity) needs to be decided after and before random demand is known, respectively. In either case, though, we assume away capacity limitations: the retailer can order as many units as desired of each item in S.

MTO Environment. Under MTO, ordering takes place after demand is realized. Therefore, demand for a given product never goes unsatisfied, which reduces the risk of the supply decision. In fact, in the case of MTO, the supply decision becomes trivial: the order quantity must be equal to the realized demand, because any inventory in excess of demand would certainly not be sold but rather salvaged for a unit loss of $(c - v)$. Nevertheless, due to the presence of returns, the quantity risk does not completely vanish; some products may be returned, and will need to be salvaged, which may involve a net loss (recall that $v < c$).

The expected profit in this case can be expressed as follows:

$$\Pi_{MTO}(S) = \sum_{j \in S} E\left[(p-c)D_j - (\alpha p + l - v)R_j\right] = \sum_{j \in S} \left[p - c - (\alpha p + l - v)P_{return|\,j}\right] \lambda P_j^S \tag{3–1}$$

The first term within expectation is the revenue, net of procurement costs. The second term is the net cost of handling returns: for each unit of returned product, the retailer refunds αp, pays l for reverse logistics activities, and eventually salvages it for v (e.g., sells it in a secondary market, such as a clearance store). We assume that returned items can only be salvaged (sold at a secondary market for a reduced price). A more general model of handling returns would allow resale of returned products in the store (possibly for full price), requiring a multiple-period planning horizon.

MTS Environment. Under MTS, the retailer takes an ordering decision for each product prior to the selling season, before demands realize. The supply decision under MTS is therefore riskier (than that under MTO): there is a chance that the retailer may over- or under-stock each and every product. Let x_j be the quantity of product j ordered and stocked in advance of the selling season.

In the event of a stock-out, the retailer places an emergency order at a unit cost of e ($v < c < e < p$), and we assume that the consumer is willing to wait for the delivery of her most preferred item and does not substitute for another item that happens to be in-stock. Emergency orders are common in retailing. For instance, Express (apparel) and Famous Footwear both have written promises in their websites that if they happen not to have the right size or color of a particular product in their store, they would find and ship it for free. There is no guarantee of course that every consumer would take up this offer. So, we are clearly making a simplifying assumption, which helps us focus on the interaction between the retailer's assortment decision and the return policy in effect. If consumers were allowed to switch from their most preferred product that is out-of-stock to a different product that is in-stock, the model would be significantly more complicated, and quite likely, analytically intractable. Assortment and inventory management under stock-out based substitution (also called dynamic substitution in the literature) is by itself a difficult problem, even if product returns were ignored (see, for instance, Gaur and Honhon (2006) for a near-optimal heuristic approach).

The expected profit under MTS can be expressed as follows[5] :

$$\Pi_{MTS}(S) = \sum_{j \in S} \max_{x_j \geq 0} \left\{ E\left[pD_j - cx_j - e(D_j - x_j)^+ - (\alpha p + l - v)R_j + v(x_j - D_j)^+ \right] \right\}$$

[5] For any real number y, let $(y)^+$ be equal to y if $y > 0$, and to 0 otherwise.

where x_j is the regular (non-emergency) order quantity for product j. The first term within expectation is the revenue; sales equal demand because, by assumption, the retailer can backlog excess demand and satisfy it with emergency orders. The second term is the cost of regular supply; and the third term is the cost of emergency supply. The fourth term is the cost of having to deal with returned items (consistent with the MTO case, returned items are salvaged). The last term is the salvage revenue from excess inventory, items that have never been sold.

Given that the demand for each product has a normal distribution, it is well-known that the optimal order quantity for each product is given by: $x_j^* = \lambda P_j^S + z^* \sigma (\lambda P_j^S)^\beta$ for all $j \in S$, where $z^* = \Phi^{-1}\left(\frac{e-c}{e-v}\right)$, and $\Phi(\cdot)$ is the cumulative distribution function of a standard normal random variable. Plugging the optimal order quantities back into the above profit expression, we obtain:

$$\Pi_{MTS}(S) = \sum_{j \in S} \left[p - c - (\alpha p + l - v) P_{return|\,j}\right] \lambda P_j^S - (e-v)\sigma\phi(z^*)(\lambda P_j^S)^\beta \qquad (3\text{–}2)$$

where $\phi(\cdot)$ is the probability density function of a standard normal random variable.

As mentioned before, our MTS model assumes backlogging of excess demand through emergency orders, which is a standard assumption in the inventory management literature to gain analytical tractability (Tagaras and Vlachos 2001). Lost sales case, where consumers walk away when faced with a stock-out, is much more difficult. The key difference is that the newsvendor critical fractile (z^*) would then depend on P_j^S, and thus both on α and S. Therefore, this compromise is crucial for us to develop analytical results on how product assortment and return policy interact.

Timing of Events. To sum up, events in our model unfold as follows. With a given return policy - defined by refund fraction α - in effect, the assortment decision (S) is taken at the beginning of the period to maximize expected profit, $\Pi_{MTO}(S)$ or $\Pi_{MTS}(S)$. Then, in the case of MTO, random demands realize and the retailer orders the quantity demanded of each product. In the case of MTS, order quantity decisions (x_j for all $j \in S$) are taken first, and then demands realize. Consumers' random choice behavior in the first stage of the N-MNL model (described above) is what drives the realization of demands. Next, consumers who purchase their product of choice decide to keep or return it (following the behavior described in the second stage of the N-MNL model). Finally, the retailer salvages any returned or excess items at the end of the period.

Determining the optimal refund fraction is an interesting problem in its own right. It inevitably requires simultaneous consideration of multiple product lines, which is beyond the scope of our current analysis. Even for a single product line and a given S, it is analytically intractable in our modeling framework. From a practical standpoint, however, optimizing α is in some sense the easy problem. Because, refund fractions in practice are usually round numbers; therefore, one can always compute the expected profit for $\alpha = 0\%, 1\%, 2\%, \ldots, 100\%$, to find the near-optimal refund fraction for a given assortment (we indeed demonstrate this in §3.5, when developing managerial insights about optimal refund fraction based on a numerical study). What is difficult is to find the optimal assortment as there are 2^n different possibilities. We provide structural results in the next section that significantly reduce the search space for accomplishing this task.

3.4 Structure of the Optimal Assortment

In this section we seek to optimize the retailer's assortment decision for a given return policy. Let S^* be the optimal assortment, and $\omega_i = \exp\left(\frac{A_i}{\mu_1}\right)$ be the *preference* of product i (a term adopted from van Ryzin and Mahajan 1999). By assumption, $A_0 = 0$ and $\omega_0 = 1$.

Without loss of generality, we sort products in N in decreasing order of preference, i.e., $\omega_1 \geq \omega_2 \geq \cdots \geq \omega_n$. Since ω_i is increasing in A_i, and A_i is increasing in a_i, this ordering applies to attractiveness levels as well, i.e., $a_1 \geq a_2 \geq \cdots \geq a_n$. Thus, lower-indexed products are more popular, and higher-indexed products are more eccentric.

3.4.1 The MTO Model with Returns

To lay the groundwork for discovering the structure of the optimal assortment, we first need the following thought experiment. Suppose the current assortment is some proper subset S of N. Now consider adding a product with preference δ to the current assortment. Let $P_{return|\delta}$ be the conditional probability of returning this product (with a slight abuse of notation). How does the new expected profit behave as a function of δ? The answer is summarized below.

Lemma 1. *Let $h_{MTO}(\delta)$ for $\delta \in \left[\min_{j\in N\setminus S}(\omega_j), \max_{j\in N\setminus S}(\omega_j)\right]$ be the expected profit function when a product with preference δ is added to an existing assortment $S \subset N$ under the MTO environment. That is, $h_{MTO}(\delta) = g(\delta)/f(\delta)$, where $f(\delta) = \sum_{j\in S}\omega_j + \omega_0 + \delta$ and*

$$g(\delta) = \sum_{j\in S} \lambda\left[p - c - (\alpha p + l - v)P_{return|j}\right]\omega_j + \lambda\left[p - c - (\alpha p + l - v)P_{return|\delta}\right]\delta$$

If $\alpha \geq (v-l)/p$, then $h_{MTO}(\delta)$ is quasiconcave and non-decreasing. Else, if $\alpha < (v-l)/p$, then $h_{MTO}(\delta)$ is quasiconvex.

The additional product considered in Lemma 1 can be thought of as a hypothetical product with attractiveness level a such that its preference $\omega = \exp(A/\mu_1)$ is equal to δ. When δ coincides with the preference ω_i of one of the products $i \in N \setminus S$ potentially considered for inclusion in the assortment, then $h_{MTO}(\delta)$ represents the resulting profit, i.e., $h_{MTO}(\omega_i) = \Pi_{MTO}(S \cup \{i\})$.

Studying the behavior of $h_{MTO}(\cdot)$ allows us to establish a local optimality result on which product (if any) should be added to an existing assortment S. Imagine a two-step procedure for finding the answer: (1) find the additional product i^* that yields the highest profit $\Pi_{MTO}(S \cup \{i^*\})$, and (2) compare it with $\Pi_{MTO}(S)$ to decide if i^* should be added. Lemma 1 settles the first step. It essentially says that: for a sufficiently lenient return policy with refund fraction $\alpha \geq (v-l)/p$, i^* must be the most popular of the remaining products in $N \setminus S$; whereas, for a strict return policy with $\alpha < (v-l)/p$, i^* must be either the most popular or the most eccentric product in $N \setminus S$.

The second step, dealing with the question of whether adding a particular product improves the profit or not, is settled in the following lemma. Let $M_j = p - c - (\alpha p + l - v) P_{return|j}$ be the expected profit margin per unit sales of product j. Note that, since $P_{return|i}$ is decreasing in a_i, $\alpha \geq (v - l)/p$ implies $M_j \geq M_{j+1}$, and $\alpha < (v - l)/p$ implies $M_j \leq M_{j+1}$ for $j = 1, \ldots, n-1$.

Lemma 2. *Adding a product $i \in N \setminus S$ to an existing assortment $S \subset N$ improves profit under the MTO environment, i.e., $\Pi_{MTO}(S \cup \{i\}) \geq \Pi_{MTO}(S)$, if and only if $M_i \geq \sum_{j \in S} P_j^S M_j$, which is equivalent to $M_i \geq \sum_{j \in S \cup \{i\}} P_j^{S \cup \{i\}} M_j$.*

Lemma 2 says that, for product i to be included in an existing assortment S, its expected profit margin must be greater than or equal to both the expected profit margin of the current set S and the new set $S \cup \{i\}$. Rules of thumb similar in nature to this result have been documented in practice (Cargille et al. 2005, and Olavson and Fry 2006).

These two lemmas, regarding the local optimality of adding another product (if any) to an existing assortment, provide building blocks for proving the structure of the optimal assortment. Define $\mathcal{A}_i = \{1, ..., i\}$ and $\mathcal{Z}_j = \{n - j + 1, ..., n\}$ for all positive integers i and j between 1 and n; and define $\mathcal{A}_0 = \mathcal{Z}_0 = \phi$. In words, $\mathcal{A}_i$ is the i most popular products in N, and $\mathcal{Z}_j$ is the j most eccentric products in N.

Theorem 1. *(a) For a sufficiently lenient return policy with return fraction $\alpha \geq (v - l)/p$, the optimal assortment under the MTO environment is composed of some number of most popular products from N, i.e., $S^* = \mathcal{A}_k$ for some $k \in \{0, 1, ..., n\}$.*

(b) For a sufficiently strict return policy with return fraction $\alpha < (v - l)/p$, the optimal assortment under the MTO environment is composed of some number of most eccentric products from N, i.e., $S^* = \mathcal{Z}_k$ for some $k \in \{0, 1, ..., n\}$.

Presence of returns clearly changes the structure of the optimal assortment. If the refund fraction is sufficiently large, reflecting a lenient return policy, carrying only the most popular products is optimal. This result agrees with common intuition and previous results in the literature (van Ryzin and Mahajan 1999, Aydin and Ryan 2000, Hopp and

Xu 2005, Maddah and Bish 2007, Li 2007, and Cachon and Kök 2007). Since high refund fractions are costly, they induce the retailer to be more selective when deciding on variety, and thus to offer products with less chances of being returned, i.e., the popular products.

However, if the refund fraction is low, reflecting a strict return policy, then it is optimal to carry only the most eccentric products. The intuitive reason is that the retailer makes more money from an item that is sold and returned than an item that is sold and not returned. In the former case, net unit profit is $(p - c - \alpha p - l + v)$; whereas in the latter case, it is $(p - c)$. This is akin to the "service escape" model of Xie and Gerstner (2007), in which a firm profits from service cancellations. Other factors that favor popular products, such as higher probability of purchase, seem to be dominated. Note that, by Lemma 1, it can be best to add to an existing assortment the most popular (remaining) product. Even though this is true for incremental additions to an assortment, Theorem 1b establishes most-eccentric assortments as optimal for strict return policies.

3.4.2 The MTS Model with Returns

The analysis proceeds similarly; as in the MTO case, we first consider the question of which product (if any) should be added to an existing assortment.

Lemma 3. *Let $h_{MTS}(\delta)$ for $\delta \in \left[\min_{j\in N\setminus S}(\omega_j), \max_{j\in N\setminus S}(\omega_j)\right]$ be the expected profit function when a product with preference δ is added to an existing assortment $S \subset N$ under the MTS environment. That is, $h_{MTS}(\delta) = \left[g(\delta) + \tilde{g}(\delta)\right]/f(\delta)$, where $f(\delta) = \sum_{j\in S}\omega_j + \omega_0 + \delta$ and*

$$g(\delta) = \sum_{j\in S} \lambda\left[p - c - (\alpha p + l - v)P_{return|j}\right]\omega_j + \lambda\left[p - c - (\alpha p + l - v)P_{return|\delta}\right]\delta$$

$$\tilde{g}(\delta) = -(e - v)\sigma\lambda^{\beta}\phi(z^*)\left(\sum_{j\in S}\omega_j^{\beta} + \delta^{\beta}\right)(f(\delta))^{1-\beta}$$

If $\alpha < (v - l)/p$, then $h_{MTS}(\delta)$ is quasiconvex.

Based on this fact, we establish the following result regarding the structure of the optimal assortment.

Theorem 2. *For a sufficiently strict return policy with return fraction $\alpha < (v - l)/p$, the optimal assortment under the MTS environment is composed of some number of most popular and some number of most eccentric products from N, i.e., $S^* = \mathcal{A}_{k-j} \cup \mathcal{Z}_j$ for some $j \in \{0, ..., k\}$ and $k \in \{0, 1, ..., n\}$. Furthermore, there exist problem instances where the optimal assortment is composed of: (1) most popular products only ($0 = j < k$), (2) most eccentric products only ($0 \leq j = k$), or (3) some most popular and some most eccentric products ($0 < j < k$).*

This result paves the way to showing that the structure of the optimal assortment is fundamentally different under MTS than under MTO. The type of assortments proved to be optimal in Theorem 1b is a subset of those in Theorem 2; setting $j = k$ leaves only those assortments with some number of most eccentric products. Therefore, one example with $0 = j < k$ or $0 < j < k$ in the optimal solution is enough to affirm that in the MTS case, it is possible to have - unlike the MTO case - an optimal assortment with a strictly positive number of most popular products only, or a strictly positive number of most popular products and a strictly positive number of most eccentric products. Such an example is illustrated in Table 3-1; details of the example are described in §3.4.4.

The key reason behind this counterintuitive result is the operational mode itself. Under the MTS environment, the ordering decision for each and every product in the assortment carries risks of over- and under-stocking. As shown in the second term of equation (3–2), the burden of these risks is proportional to the standard deviation of demand. Normalizing by demand size, coefficient of variation (defined as standard deviation divided by mean) as a measure of relative demand variability is generally a good indicator of how risky a product is - operationally speaking. In our model, products with higher attractiveness enjoy a larger probability of purchase and a smaller coefficient of variation. Hence, for strict return policies, the retailer has two opposing goals: (1) choose eccentric products to benefit from their resale (much like in the MTO case); (2) choose popular products to take advantage of their lower relative demand variability and therefore reduce operational risks. The structure of the optimal assortment reflects both of these goals.

Clearly, our analytical results in the MTS case are limited to the strict return policy case only. Although we are unable to prove this, based on extensive numerical studies (only a subset of which is presented in §3.5) we conjecture that the lenient return policy case requires the optimal assortment to include some number of most popular products, just as in the MTO environment. The intuition given above for Theorem 2 also supports our claim, because for lenient return policies the retailer finds popular products more desirable on both counts. They not only have less relative demand variability, but also a smaller chance of return.

3.4.3 The MTO and MTS Models without Returns

Both our MTO and MTS models include as a special case the possibility of the retailer disallowing returns. By a slight abuse of model definition, we can analyze this case by setting $\alpha = -\infty$, which implies that the consumers will choose the 'keep' option with probability 1 in the second stage of our N-MNL model regardless of which product they choose in the first stage (i.e., they never return products). In fact, the N-MNL model reduces to a standard MNL model. The optimal assortment would then be comprised of some number of most popular products under both MTO and MTS environments. (We omit the proof; same result was obtained by van Ryzin and Mahajan (1999) in an MTS model with lost sales and without returns.)

Therefore, by contrasting this result with Theorems 1 and 2, we conclude that if retailers were to ignore product returns when merchandising, they might easily run the risk of composing sub-optimal assortments. This is especially true if they have relatively strict return policies.

3.4.4 Numerical Example

We conclude our analysis of the optimal assortment with a numerical example that illustrates the different kinds of solutions that arise under MTO/MTS environments with strict/lenient return policies. Table 3-1 displays the optimal assortment out of a given set of 10 potential products (sorted in decreasing order of attractiveness levels) for different values of refund fraction α and for both MTO and MTS models. The optimal assortment in each of these instances is computed by complete enumeration. Note that the threshold refund fraction that separates strict and lenient return policies in this example is $(v - l)/p = 0.8$.

As expected, optimal variety ($|S^*|$) is lower under MTS. The reason is that higher variety costs more under MTS due to operational risks of over- and under-stocking of products in an assortment.

3.5 Insights and Discussion

In this section we investigate the following research questions by mostly computational means.

- If a retailer ignored the presence of product returns when composing its assortment, or it assumed that the best assortment is always composed of most popular products, what would be the magnitude of its profit loss relative to the optimal profit?
- How is the *depth* of the optimal assortment, the cardinality of S^*, influenced by changes in refund fraction? Is it necessarily the case that more lenient return policies imply less variety?
- How does the degree of differentiation among the potential products considered by the firm (set N) influence its profit and depth of assortment? If the firm had any influence over this degree of differentiation, would it prefer higher or lower differentiation?
- What is the effect of post-purchase heterogeneity (μ_2) on the optimal profit for a given refund fraction? Can more heterogeneity be ever beneficial for the firm?
- How does larger aggregate demand variability influence optimal assortment (number and composition of products)?

The base parameter values used throughout this section, unless otherwise noted, are displayed in Table 3-2. The corresponding threshold refund fraction that separates strict return policies ($\alpha < (v - l)/p$) from lenient return policies ($\alpha \geq (v - l)/p$) is

$(v-l)/p = 0.8$; we indicate this threshold with a vertical line in Figures 3-1, 3-2, 3-5, 3-6, and 3-7. All of the observations we make in this section appear to be robust; equivalent experiments with different sets of parameters yield qualitatively similar results.

3.5.1 Profit Loss from Ignoring Product Returns or Assuming the Wrong Structure for Optimal Assortment

It is intuitively very appealing to include only the most popular products in an assortment. In models that do not consider product returns, this was indeed shown to be optimal by van Ryzin and Mahajan (1999) and others (cited earlier). However, as we show in Theorems 1 and 2, including only the most popular products would be sub-optimal for relatively strict return policies. In this subsection we quantify the profit loss the retailer would suffer by doing so. More precisely, we compare the best most-popular assortment with the optimal assortment. We also quantify the profit loss from ignoring product returns when making the product assortment decision.

For values of refund fraction α ranging from 0 to 1 with 0.1 increments, we compute the expected profits from the optimal assortment, from the best possible assortment composed of most popular products only (this is what we call 'assuming the wrong structure'), and from the most profitable assortment with product returns ignored (i.e., with S optimized by setting $\alpha = -\infty$, but the resulting profit calculated - as with the other two scenarios - from $\Pi_{\cdot}(S)$ for a fixed α value). Let $\Pi_{\cdot}^{*}$, $\Pi_{\cdot}^{W}$, and $\Pi_{\cdot}^{I}$ denote these profits, respectively. Figure 3-1 plots the percentage profit loss that results from assuming the wrong structure $\left(\left(\Pi_{\cdot}^{*} - \Pi_{\cdot}^{W}\right)/\Pi_{\cdot}^{*}\right)$, and from ignoring the returns $\left(\left(\Pi_{\cdot}^{*} - \Pi_{\cdot}^{I}\right)/\Pi_{\cdot}^{*}\right)$ for both operational modes.

In this particular example, the loss of profit from assuming the wrong structure can be up to 12% (7%) in the MTS (MTO) environment for strict return policies. The loss is 0% for lenient return policies (which is due to Theorem 1a in the case of MTO). Ignoring product returns can be much more harmful (profit losses of up to 23% are possible), especially under MTS with strict return policies, and under MTO with lenient return policies.

The lesson from a managerial perspective is that a retailer should be generally careful about assuming that most-popular assortments are always the most profitable. For relatively strict return policies, this commonsensical assumption can be quite misleading, and more so for MTS - due to inventory risk - than for MTO. Also, not very surprisingly, ignoring product returns when making assortment decisions can result in a loss of opportunity of making substantially more profits (the question of how substantial can only be addressed with real data, of course).

3.5.2 Does More Lenient Return Policy Mean Less Variety?

Intuitively speaking, more lenient return policies with higher refund fractions must lead to less variety. Because, higher refunds are costly, and they will induce the retailer to be more careful about expanding its assortment and thereby increasing the total volume of returns. Our numerical experiments, detailed below, show that this is indeed true for lenient return policies. In fact, under an MTO environment, we can provide a mild sufficient condition for a more lenient return policy to always result in a reduction in optimal variety.

Proposition 1. *Assume that the firm is operating in an MTO environment and all potential products yield positive expected utility (if kept), i.e., $a_i \geq p$ for all $i \in N$. For all lenient return policies with $\alpha \geq (v-l)/p$, the cardinality of the optimal assortment $|S^*|$ is decreasing in refund fraction α. That is, the more lenient the return policy, the less the optimal variety.*

To further explore the relationship between variety and return policy, we plot in Figure 3-2 the cardinality of the optimal assortment $|S^*|$ as a function of α (ranging from 0 to 1 with 0.01 increments). Clearly, for sufficiently high and sufficiently low α values, higher refund fraction leads to less variety. Yet there is also a range of α values for which the variety is increasing in α; that is, more lenient return policies result in more variety.

This range typically includes relatively higher α values within the strict return policy region.

We observe in Figure 3-2, that starting at $\alpha = 0$, there is first a decrease and then an increase in the number of products in S^* as α approaches to $(v - l)/p$. At the extremes, when α is close to either 0 or $(v - l)/p$, the expected profit margin in the MTO case (M_j) approaches $(p - c)$ for all j, making all products almost *equally* profitable. (Because, when α is close to 0, $P_{return|j} \simeq 0$; and, when α is close to $(v - l)/p$, $\alpha p + l - v \simeq 0$.) As a consequence, we can expect higher variety around these bounds to capture more demand without much cannibalization. For α values in between, the margins become unequal, leading to more cannibalization concerns and thus less variety. For MTS, we observe based on our numerical experiments that a similar pattern holds (except, changes in variety are usually less steep, and $|S^*|$ peaks earlier).

The managerial take away from this experiment is the fact that more lenient return policies may sometimes call for deeper assortments with higher variety. This happens especially when the refund fraction is at neither extreme (0% or 100%), but just below a certain threshold ($(v - l)/p$). Therefore, for product categories with good secondary markets ($v \approx p$) and low reverse logistics costs (l), i.e., when this threshold is close to 1, this observation is likely to be more salient.

3.5.3 Impact of Product Differentiation on Optimal Variety and Profit

In this section, we explore how the degree of differentiation among potential products considered by the retailer (set N) influences both the number of products in the optimal assortment and the optimal expected profit under MTO and MTS environments. Retailers, by virtue of selecting their suppliers, may have some control over this degree of differentiation. One relevant question for them would then be: is higher differentiation always better?

To investigate this effect, we first consider a problem where all 10 potential products in set N have identical attractiveness levels. For two specific α values, 0.5 and 1, we compute the optimal assortment and the corresponding expected profit. We then compare the results with those obtained by differentiating the potential product set. The differentiation is achieved by varying the ω's, and thus a's, for the products in set N. We keep the sum of ω's constant in all cases such that if all products were offered, the probability of buying some product would be the same. This is to avoid the confounding factor of demand expansion or contraction. Table 3-3 shows the preferences (ω's) that correspond to five sets of products considered in this experiment. We maintain the rest of the parameter values as in Table 3-2.

Impact on optimal variety. As the set of potential products becomes more differentiated, the optimal assortment consists of fewer products. The intuition behind this lies in the cannibalization effect. When products are similar, their expected margins are also similar, so it is in the retailer's best interest to offer higher variety to capture more market share. The specific product within the set that the consumer ends up buying has little effect on the total expected profit. On the other hand, when products are differentiated, the retailer tends to offer only the most profitable products, which are among the most popular, most eccentric or a combination of both according to Theorems 1 and 2. In larger assortments the cannibalization has a more negative impact because it lessens the demand for the most profitable products.

Figure 3-3 plots the number of products in the optimal assortment, $|S^*|$, for $\alpha = 0.5$ and $\alpha = 1$ under MTO and MTS environments. Other values of α yield similar curves. The descent (if any) in the MTS case is usually not as steep as in MTO.

Impact on optimal profit. We also observe that more differentiation increases the optimal expected profit in almost all circumstances (see Figure 3-4). Since differentiation increases the differences in margins, the retailer can choose the most profitable products, as discussed above, to increase its overall profit. With regard to the magnitude of this increase, we first note that the impact in MTS is always more significant than it is in MTO. In MTS, the possibility of choosing among more differentiated products, allows the retailer to better control inventory costs because the relative demand variability for more popular products is lower. Second, for strict return policies, the differentiation has almost no effect on the optimal profit for the MTO case. Third, especially for larger values of α, higher differentiation yields higher profit increments. A larger refund fraction increases the probability of return; and, from a more differentiated product set, the retailer is able to offer products that are more popular and with less chances of being returned.

From a managerial point of view, a retailer moving from an MTO to an MTS environment should seek higher product differentiation in its consideration set (N), because it will matter more. That effort is even more worthwhile when the retailer's return policy is more lenient.

3.5.4 Effect of Post-purchase Heterogeneity on Optimal Profit

Firms can usually exert some control on the chances of product return (e.g., by detailing product characteristics better as described in Shulman et al. (2008)). In other words, they may be able to directly influence post-purchase heterogeneity (characterized by μ_2 in our model) by their actions. One question of practical interest is then whether retailers should always prefer reducing it.

In N-MNL, it can be shown that "only the ratio [of μ_1/μ_2] can be identified from the data" (Ben-Akiva and Lerman 1985, p. 287). Therefore, it is common to normalize μ_1 to 1. This normalization gives us a natural range for μ_2; in this experiment we set $\mu_1 = 1$ and vary μ_2 in the $(0, 1)$ interval. More specifically, Figures 3-5 and 3-6 plot the optimal expected profit for given values of refund fraction, and for different levels of post-purchase heterogeneity ($\mu_2 = 0.25$, 0.50, 0.90), under MTO and MTS, respectively.

We learn from Figures 3-5 and 3-6 that lower post-purchase heterogeneity yields higher profit when return policies are lenient. The rationale is quite transparent: since giving higher refunds is more costly for the retailer, it will try to minimize returns by reducing the post-purchase heterogeneity. Nonetheless, for strict return policies, the effect is opposite. Since the retailer can obtain additional profit from returns, as pointed out in

§3.4.1, higher heterogeneity in the keep/return decision increases the probability of return, and therefore it benefits the retailer.

The managerial insight from this experiment is clear: it is possible that higher post-purchase heterogeneity can be beneficial for a retailer. The reasonable presumption that higher heterogeneity about consumers' keep/return decisions will lead to lower profits is wrong for strict return policies.

3.5.5 Impact of Demand Variability on Optimal Assortment

How would a retailer modify its assortment both in number and composition of products, under different levels of aggregate demand variability? Recall that our aggregate demand model, the Independent Population Model (van Ryzin and Mahajan 1999), is parameterized by σ and β, where σ is an indicator of demand variability. Obviously, in a MTO context, demand variability does not change the optimal assortment because all products are ordered after demand is realized. In the MTS case, however, higher variability will have an impact since there exists inventory risk for overstocking and understocking. Using the base case example in Table 3-2, we compute the optimal assortment for different values of σ and varying α from 0 to 1 in 0.01 increments. As seen in Figure 3-7, higher variability leads to narrower assortments.

It is quite intuitive that in a highly variable market where inventory related costs are more relevant, the firm reduces its assortment. If we take a closer look to the composition of the optimal assortment when demand variability is high, i.e., $\sigma = 5$ (see Table 3-4), we observe that the firm is more inclined to offer just the most popular product because it has lower relative demand variability, and therefore less inventory risk.

3.6 Concluding Remarks

We believe that this chapter highlights an interesting interaction between product assortment and return policy, which is critically moderated by operational mode (MTO versus MTS). For a lenient return policy (with a sufficiently high refund fraction), the optimal assortment is composed of some number of most popular products in both MTO and MTS environments. For a strict return policy (with a sufficiently low refund fraction), however, the optimal assortment has distinct and counterintuitive structures in MTO and MTS environments. In particular, the optimal assortment is composed of some number of most eccentric products under MTO, whereas it can be composed of some number of most popular and most eccentric products under MTS. These results underscore the need to consider consumer returns and operational mode in making product assortment decisions.

Our analytical and numerical results amply illustrate that assortment and refund fraction can exhibit interactions that are not easily predictable. Therefore, endogenizing the return policy decision analytically would be a worthwhile extension of our work. An equally important direction would be to endogenize the pricing decision.

Table 3-1. Optimal assortment S^*, composed of products that correspond to shaded cells, for the problem instance in Table 3-2 with threshold refund fraction, $(v-l)/p = 0.8$

		α										
	i	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
MTO	1	■								■	■	■
	2	■								■	■	■
	3	■	■							■	■	■
	4	■	■							■	■	
	5	■	■	■					■	■		
	6	■	■	■	■				■	■		
	7	■	■	■	■	■	■	■	■	■		
	8	■	■	■	■	■	■	■	■	■		
	9	■	■	■	■	■	■	■	■	■		
	10	■	■	■	■	■	■	■	■	■		
MTS	1	■	■							■	■	■
	2	■								■	■	■
	3									■		
	4											
	5											
	6			■					■			
	7		■	■	■	■	■	■	■			
	8	■	■	■	■	■	■	■	■			
	9	■	■	■	■	■	■	■	■			
	10	■	■	■	■	■	■	■	■			

Table 3-2. Base parameter values for the numerical study in Chapter 3

Parameter	Value	Product, i	a_i
λ	100	1	4.00
p	2	2	3.72
e	1.9	3	3.44
c	1.8	4	3.17
v	1.7	5	2.89
l	0.1	6	2.61
μ_1	1	7	2.33
μ_2	0.5	8	2.06
σ	1	9	1.78
β	0.5	10	1.50

Table 3-3. The preferences (ω values) for five sets of products with different degrees of differentiation: *I* (identical), *VS* (very similar), *S* (similar), *D* (different), and *VD* (very different).

	I	VS	S	D	VD
ω_1	3	3.30	3.60	4.20	4.90
ω_2	3	3.23	3.47	3.93	4.48
ω_3	3	3.17	3.33	3.67	4.06
ω_4	3	3.10	3.20	3.40	3.63
ω_5	3	3.03	3.07	3.13	3.21
ω_6	3	2.97	2.93	2.87	2.79
ω_7	3	2.90	2.80	2.60	2.37
ω_8	3	2.83	2.67	2.33	1.94
ω_9	3	2.77	2.53	2.07	1.52
ω_{10}	3	2.70	2.40	1.80	1.10
$\sum_{i \in N} \omega_i$	30	30	30	30	30

Table 3-4. Optimal assortment S^*, composed of products that correspond to shaded cells, for the problem instance in Table 3-2 with demand variability $\sigma = 5$

		α										
	i	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
MTS	1	■	■	■					■	■	■	■
	2											
	3											
	4											
	5											
	6											
	7											
	8				■	■	■	■				
	9				■	■	■	■				
	10			■	■	■	■	■				

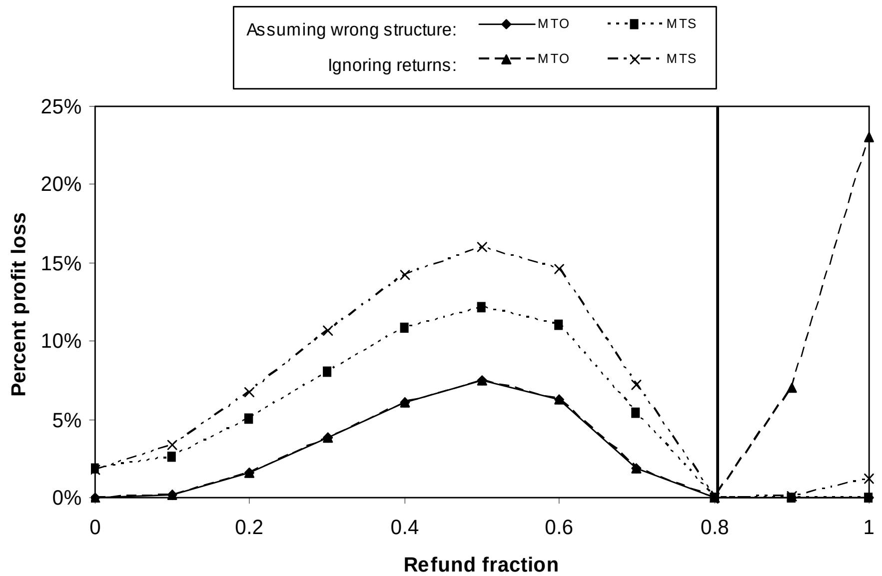

Figure 3-1. Profit loss from assuming the wrong structure (for optimal assortment) and from ignoring returns

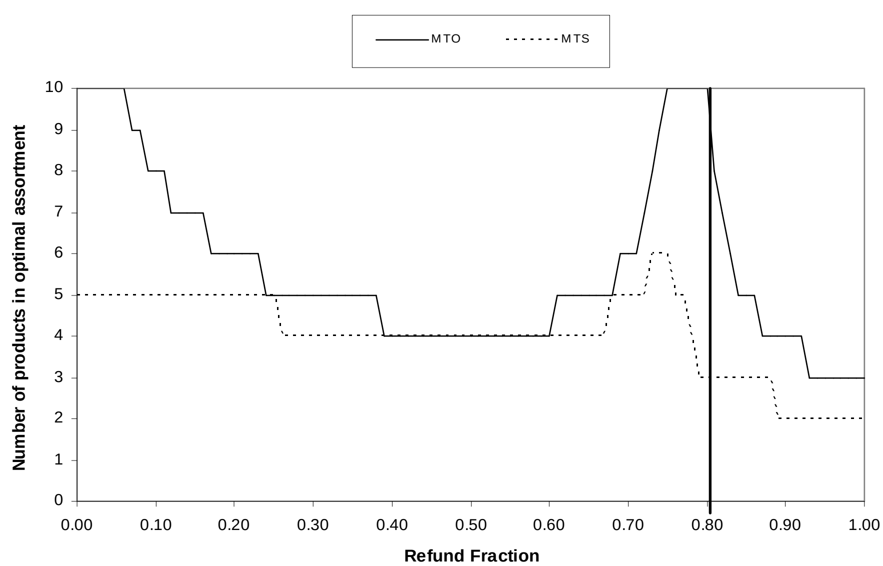

Figure 3-2. Variety versus return policy: Number of products in the optimal assortment ($|S^*|$) as refund fraction (α) varies

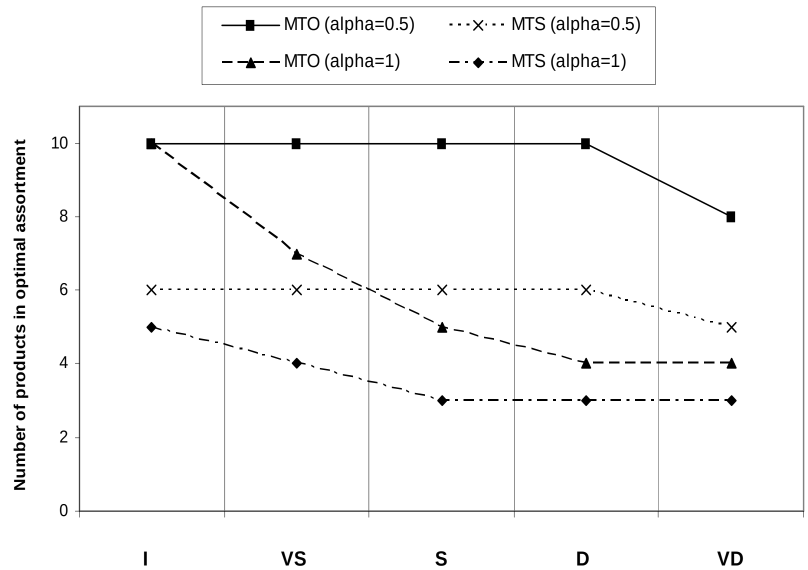

Figure 3-3. Number of products in the optimal assortment ($|S^*|$) at five degrees of product differentiation (data given in Table 3-3)

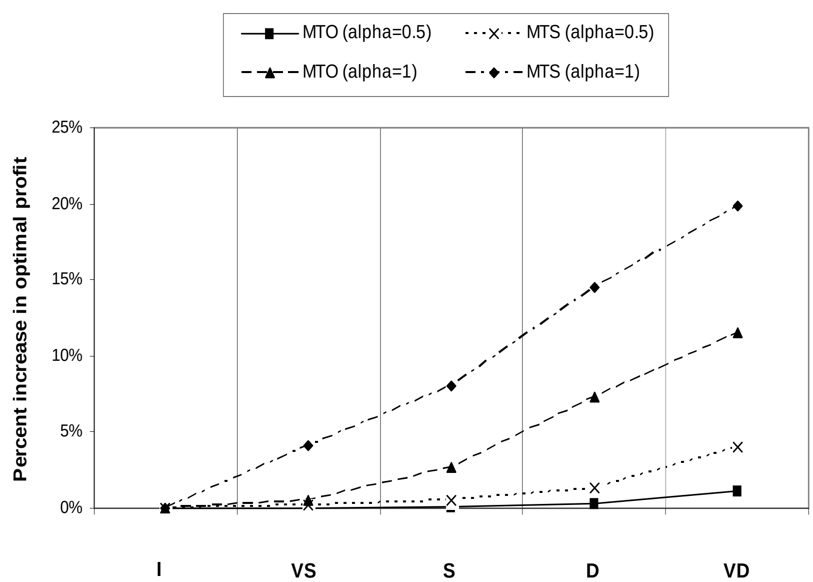

Figure 3-4. Percent increase in optimal expected profit (with respect to scenario I) at five degrees of product differentiation (data given in Table 3-3)

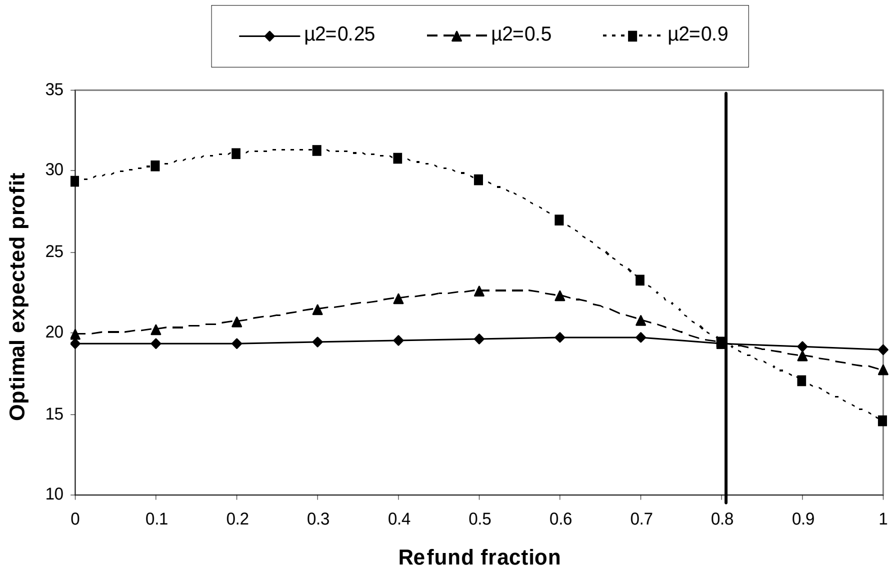

Figure 3-5. Optimal expected profit versus refund fraction (α) for different levels of post-purchase heterogeneity (μ_2) under MTO environment

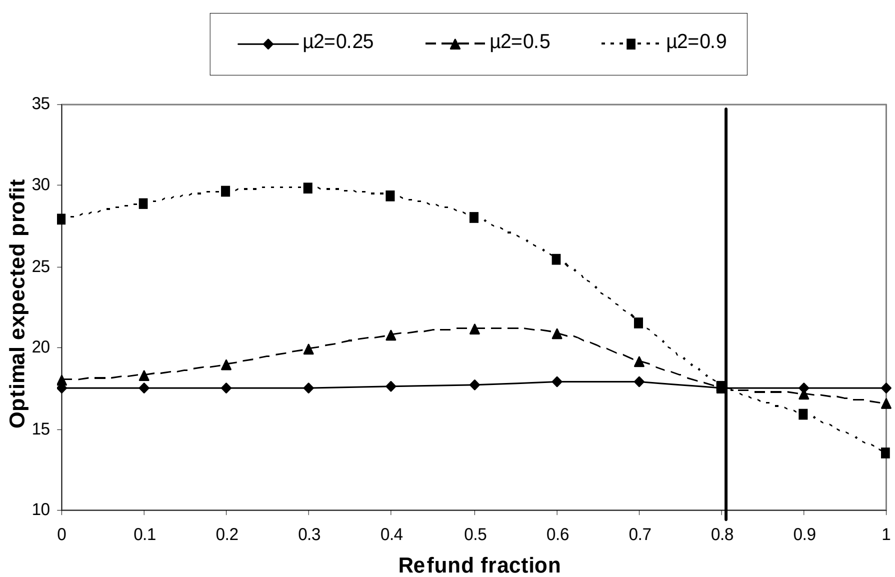

Figure 3-6. Optimal expected profit versus refund fraction (α) for different levels of post-purchase heterogeneity (μ_2) under MTS environment

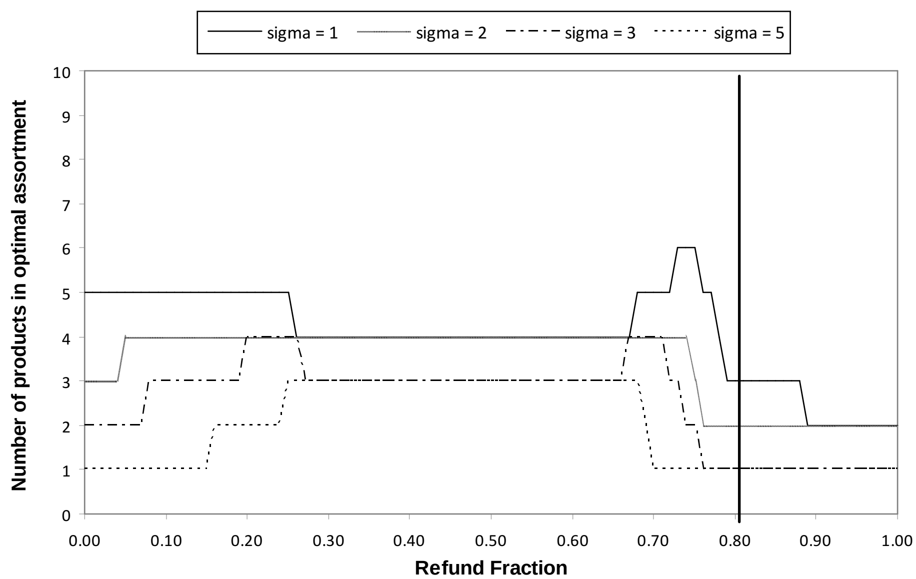

Figure 3-7. Number of products in the optimal assortment ($|S^*|$) with different levels of aggregate demand variability (σ)

CHAPTER 4
ASSORTMENT PLANNING WITH ENDOGENOUS PRICE, REFUND FRACTION, AND IN A MULTIPLE-PERIOD SETTING

4.1 Introduction

In Chapter 3, we were able to determine the structure of the optimal assortment for a given price and refund fraction in a single period setting. The main purpose of this chapter is to demonstrate (by numerical experimentation) that the analytical results of the previous chapter are robust to the following extensions: (1) endogenous price, (2) endogenous refund fraction, and (3) multiple periods. That is, interesting aspects of our results regarding when a retailer should carry eccentric products survive these extensions, which - we have good reasons to believe - are analytically intractable.

Unless we state otherwise, in all experiments we use a set of base parameter values displayed in Table 4-1.

4.2 Endogenous Price

In this section, allowing price to be a decision variable, we explore how pricing decisions interact with the optimal assortment and the return policy in effect. As in the base model, return fraction is considered exogenous.

4.2.1 Variety versus Price

Does higher price lead to more or less variety? The answer depends on the refund fraction. For two values of refund fraction, $\alpha = 0.5$ and $\alpha = 0.8$, we compute the optimal assortment while varying price from 2 to 3 in MTO and MTS (see Figures 4-1 and 4-2). For $\alpha = 0.5$, and all prices within the range considered (from 2 to 3), we are in the strict return policy region, i.e., $\alpha < (v - l)/p$. As price increases, the unit cost of returns $(\alpha p + l - v)$ approaches to 0, and that makes all products more similar in terms of their profitability. The retailer then opts to offer full assortment to capture more demand. For $\alpha = 0.8$, the effect is opposite and more interesting, especially in the MTO case. Increasing price increases the probability of return. Since we are in the lenient return policy region $(\alpha \geq (v - l)/p)$ for all price points except $p = 2$, and the unit cost of returns

$(\alpha p + l - v)$ increases in price, returns become increasingly more costly. The retailer then reduces its assortment by offering less number of most-popular products, which effectively reduces the likelihood of return. It is interesting that, in a monopoly setting, lower variety can coincide with higher prices. This is not uncommon in competitive environments (Alptekinoğlu and Corbett 2008b), but in monopoly environments price and variety are usually positively related (in fact, we do not know of a counterexample to this rule, beside the one caused by product returns in this work). For MTS, the number of products in the optimal assortment is almost constant.

4.2.2 Behavior of Expected Profit with Respect to Price

We now study the behavior of the expected profit with respect to price. Among other things, we want to understand if the expected profit is generally unimodal, which would make numerical optimization of price relatively easy.

For different values of refund fraction, Figures 4-3 and 4-4 plot the expected profit as price varies from 2 to 6 in MTO and MTS cases, respectively. For every data point shown in the charts, the assortment is optimized. We observe that the expected profit is unimodal for these problem instances. In fact, we have not seen any problem instance to the contrary. Note also from the graphs that the optimal price increases as refund fraction decreases. We examine this in more detail in the next subsection.

4.2.3 Optimal Price with Respect to Refund Fraction

Figures 4-5 and 4-6, for MTO and MTS, respectively, show how optimal price changes as refund fraction (α) varies between 0 and 1 by increments of 0.1. A dashed line separates the strict return policy region ($\alpha p < v - l$) from the lenient return policy region ($\alpha p \geq v - l$).

Again, the optimal assortment is computed for every data point. The optimization over S takes advantage of structural results presented in Chapter 3 for the base model, whereas the optimization over p is done numerically by line search. For all problem instances that we have seen, we observe that the expected profit is generally unimodal in p, which makes the line search easy.

The optimal price increases very slightly for strict return policies, and then suddenly drops for lenient return policies as refund fraction approaches to 1. This is because the retailer tries to reduce the probability of return by lowering the price. With a lenient policy, the retailer would rather charge less and obtain a final sale than salvage a product for a lower revenue. It is surprising that optimal price would drop for increasingly more lenient return policies (higher α). Even from the perspective of absolute refund amount, the consumer enjoys a more favorable return policy as α increases, because αp^* also keeps increasing, albeit at a diminishing rate. The optimal price for the MTS case is always higher than that of the MTO case, although the difference is minimal as it can be observed in the graphs.

4.2.4 Optimal Price for Most-popular and Most-eccentric Assortments

In this subsection we explore how optimal price changes for two assortment structures proved to be optimal under some of the various settings we analyze in §3.4.1 and §3.4.2. We consider the assortments that include i most popular (eccentric) products, $\mathcal{A}_i$ ($\mathcal{Z}_i$) for $i = 1, \ldots, 10$, and compare how the optimal price changes as we keep adding the next most popular (eccentric) product to the assortment. Intuitively, as we add more products, the assortment improves in the sense that a customer is more likely to buy a product. We dub this the *broader assortment effect.* Does it always lead to higher prices in consequence? The answer is not always.

Starting from an assortment with the most eccentric product, as we keep adding products the firm is able to charge a higher price (see Figures 4-7 and 4-8) for two reasons. First, added products have higher expected utility, and the assortment becomes better in general. Second, cannibalization of demand to these more popular products reduces probability of return, also allowing the firm to increase the price. This explains why 'Most-eccentric assortment' curves are increasing. For higher values of α (i.e., $\alpha = 1$ in the graphs), the effect is opposite when we start from an assortment with the most popular product. As we add products with lower expected utility, some demand is cannibalized to these products with higher probability of return. As a result, the firm reduces price to control returns (note the decreasing 'Most-popular assortment ($\alpha = 1$)' curves). However, when the return policy is stricter (i.e., $\alpha = 0.5$ in the graphs), even in the 'Most-popular assortment' case, the optimal price increases with more products in the assortment. In this case, the broader assortment effect overcomes the *return effect*. Since α is low, both $(\alpha p + l - v)$ and $P_{return|i}$ are low compared to $(p - c)$. The firm charges more purely because of broader assortment even if it increases the probability of return (return effect).

4.3 Endogenous Refund Fraction

The refund fraction, α, plays a key role in our model; it not only affects the probability of return $P_{return|\,i}$, but also the expected utility A_i (and therefore probability of purchase), and ultimately the expected profit margin per unit sales. Moreover, these interactions depend on the attractiveness levels; a variation in α may cause opposite changes on purchase probabilities of popular versus eccentric products in the assortment. When the retailer operates in an MTS environment, the situation is further complicated by inventory risk, also affected by α. All of these factors make optimizing α a very challenging problem. In this subsection, we investigate how endogenizing refund fraction influences our assortment problem.

4.3.1 Behavior of Expected Profit with Respect to Refund Fraction

Figures 4-9 and 4-10 plot the expected profit for several α values from 0 to 1 (with 0.05 increments) at three different price points. For every data point we optimize the assortment, therefore different data points may correspond to different product assortments. For both MTO and MTS cases, at $p = 2$ the optimal refund fraction is $\alpha^* = 0.55$; at $p = 2.25$, $\alpha^* = 0.5$; and at $p = 2.5$, $\alpha^* = 0.45$. So, for the three price points considered in this experiment, the optimal refund fraction is lower for higher prices.

This result, which we further explore in the next subsection, complements the price versus refund analysis in §4.2.3.

4.3.2 Optimal Refund with Respect to Price

In this subsection, we study how optimal refund fraction is affected by changes in price. We vary the price from 2 to 6, and compute the optimal refund fraction and optimal assortment. As before, the optimization over S takes advantage of structural results presented in Chapter 3 for the base model, whereas the optimization over α is done numerically by line search. For all problem instances that we have seen, we observe that the expected profit is generally unimodal in α, which makes the line search easy.

Does higher price imply higher refund fraction? The answer is not necessarily. As seen in Figures 4-11 and 4-12, the optimal refund fraction represented by square dots, first decreases and then increases in price. The intuition is the following. Starting from a low price, an increase in price raises the probability of return, forcing the retailer to reduce α to discourage returns. For low values of p, the expected profit margin per unit sales, $p - c - (\alpha p + l - v)P_{return|i}$, is more sensitive to returns since $(p - c)$ is small relative to the cost of returns term. As we keep increasing price, $(p - c)$ increases and returns become less relevant for the profit margin. Since the retailer is extracting enough profit from $(p - c)$, it can afford increasing α to make its value proposition more attractive. Note that a dashed line separates the strict return policy region ($\alpha p < v - l$) from the lenient return policy region ($\alpha p \geq v - l$) in the graph. Also note that refund amount, $\alpha^* p$, does consistently increase in price; thus at higher prices the retailer is effectively charging more for a more generous return policy. Last, the optimal refund for the MTO case is always higher than that of the MTS case, although the difference is minimal as it can be observed in the graphs.

4.3.3 Optimal Refund Fraction for Most-popular and Most-eccentric Assortments

As we did with price, in this subsection we explore how optimal refund fraction, α^*, changes for two assortment structures: the assortments that include i most popular (eccentric) products, $\mathcal{A}_i$ ($\mathcal{Z}_i$) for $i = 1, \ldots, 10$. For each assortment, the optimal α is computed through a line search over the interval $[0, 1]$ and the results are displayed in Figure 4-13. Note that when the retailer keeps adding the next most popular product to its assortment, the optimal refund fraction decreases (see 'Most-popular assortment' curves in the graph). As products with lower expected utility are added (i.e., higher likelihood of being returned), it is desirable for the retailer to make its refund policy stricter to dissuade consumers from returning.

In the case of eccentric products (see 'Most-eccentric assortment' curves in the graph), however, there is first a decrease and then an increase in the evolution of α^*. The decrease is due to the *demand expansion effect*, and the increase due to a *popularity effect*. When the retailer only offers the most eccentric product (product n), it can expect a relatively low demand. Adding a second (or a third) product increases the total demand considerably, allowing the retailer to decrease the refund fraction as it feels less pressure to offer better refunds to attract more demand. At some point, when additional products yield only small gains in demand (i.e., demand expansion effect becomes much less relevant), α^* starts going up, because more and more consumers switch to relatively more popular products within S, which are less likely to be returned. Due to this popularity effect, in turn, the retailer is able to afford higher refund fraction values that make the assortment more desirable overall. Finally, note that Figure 4-13 exhibits very similar curves for the MTO and MTS models. That is, for a given assortment, the effect of operational mode on α^* appears to be almost negligible.

The main insight from this experiment is that sometimes higher variety requires a higher refund fraction. This essentially complements our observation earlier that moving toward more lenient return policies and deeper assortments simultaneously can be optimal.

4.4 Multiple-Period Problem

In this subsection, we extend the problem to a multiple period setting. We assume that all returns are kept as inventory to satisfy future demand. Only returns from the last period (and any remaining inventory) are salvaged at the very end. Using the same base parameters shown in Table 4-1, we compute the optimal assortment for different values of α as we did in Chapter 3. We use an inventory cost of 0.05 per period for the returns kept in stock. In order to compute the expected profit for multiple periods, we use Monte Carlo simulation methods. The procedure is as follows: for every product in the assortment we generate random demand and return strings of size T, the length of the planning horizon. With known demands and returns, we easily compute the actual profit. We then estimate the expected profit by averaging the profits at a sufficiently large sample of realizations, $1,000$ in our case (Robert and Casella 1999, p. 208). By the Law of Large Numbers, this estimation converges with probability 1 to the expected profit as the sample size goes to infinity. For every possible assortment (i.e., $2^{10} - 1 = 1023$), we compute the approximate expected profit, and choose the one that yields the maximum.

We observe that the assortments that yield maximum expected profit have the same structures found to be optimal in the single period setting (see Chapter 3). Tables 4-2 and 4-3 show the optimal assortment for MTO and MTS cases for the multiple period problem with $T = 3$ and $T = 10$, respectively.

An interesting question that arises in a multiple-period context is whether the retailer includes more products as the length of the planning horizon T increases. Tables 4-2 and 4-3 suggest that, the longer planning horizon (and multiple re-selling opportunities it brings) changes neither the structure of the assortment nor the composition in any significant fashion.

Table 4-1. Base parameter values for the numerical study in Chapter 4

Parameter	Value	Product, i	a_i
λ	100	1	4.00
p	2	2	3.72
e	1.9	3	3.44
c	1.8	4	3.17
v	1.7	5	2.89
l	0.05	6	2.61
μ_1	1	7	2.33
μ_2	0.5	8	2.06
σ	1	9	1.78
β	0.5	10	1.50

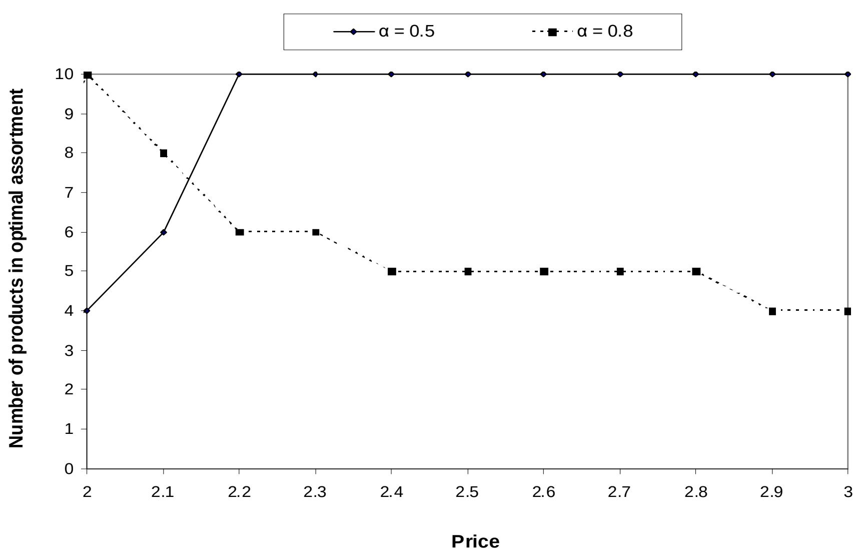

Figure 4-1. Variety versus price (MTO): Number of products in the optimal assortment ($|S^*|$) as price (p) varies for different values of refund fraction (α)

Table 4-2. Optimal assortment S^*, composed of products that correspond to shaded cells, for a multiple period problem with 3 periods

	i	α = 0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
MTO	1									■	■	■
	2	■	■							■	■	■
	3	■	■							■	■	■
	4	■	■							■	■	
	5	■	■	■						■		
	6	■	■	■	■	■			■	■		
	7	■	■	■	■	■	■	■	■	■		
	8	■	■	■	■	■	■	■	■	■		
	9	■	■	■	■	■	■	■	■	■		
	10	■	■	■	■	■	■	■	■	■		
MTS	1									■	■	■
	2										■	■
	3	■									■	
	4	■	■									
	5	■	■									
	6	■	■	■								
	7	■	■	■	■	■	■	■	■	■		
	8	■	■	■	■	■	■	■	■	■		
	9	■	■	■	■	■	■	■	■	■		
	10	■	■	■	■	■	■	■	■	■		

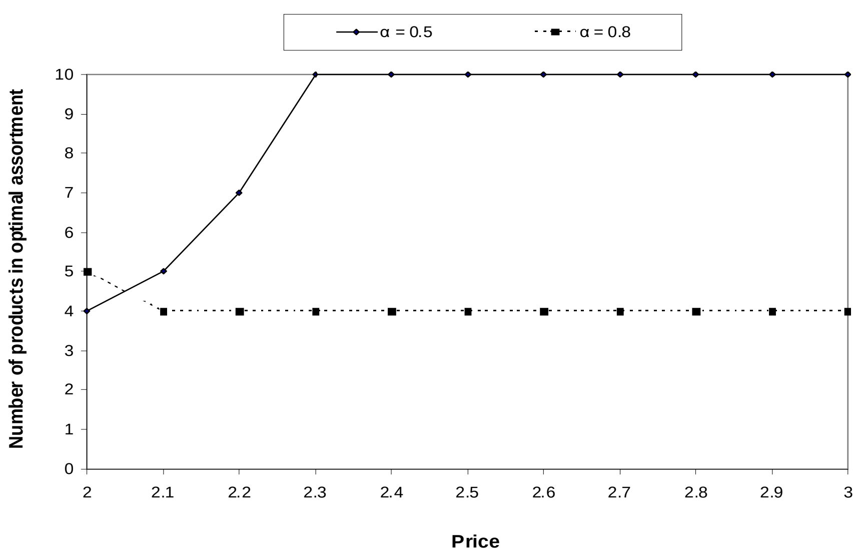

Figure 4-2. Variety versus price (MTS): Number of products in the optimal assortment ($|S^*|$) as price (p) varies for different values of refund fraction (α)

Table 4-3. Optimal assortment S^*, composed of products that correspond to shaded cells, for a multiple period problem with 10 periods

		α										
	i	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
MTO	1									■	■	■
	2	■								■	■	■
	3	■	■							■	■	■
	4	■	■							■	■	
	5	■	■	■						■		
	6	■	■	■	■				■	■		
	7	■	■	■	■	■	■	■	■	■		
	8	■	■	■	■	■	■	■	■	■		
	9	■	■	■	■	■	■	■	■	■		
	10	■	■	■	■	■	■	■	■	■		
MTS	1									■	■	■
	2										■	■
	3										■	
	4	■										
	5	■	■									
	6	■	■	■	■					■		
	7	■	■	■	■	■	■	■	■	■		
	8	■	■	■	■	■	■	■	■	■		
	9	■	■	■	■	■	■	■	■	■		
	10	■	■	■	■	■	■	■	■	■		

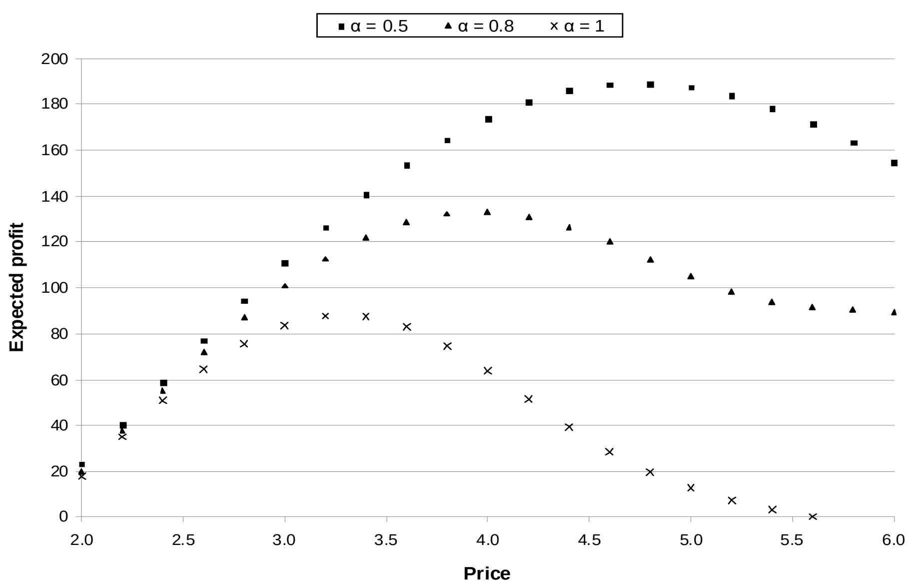

Figure 4-3. Profit versus price (MTO): Expected profit as price (p) varies for different refund fractions ($alpha$) under optimal assortment (S^*)

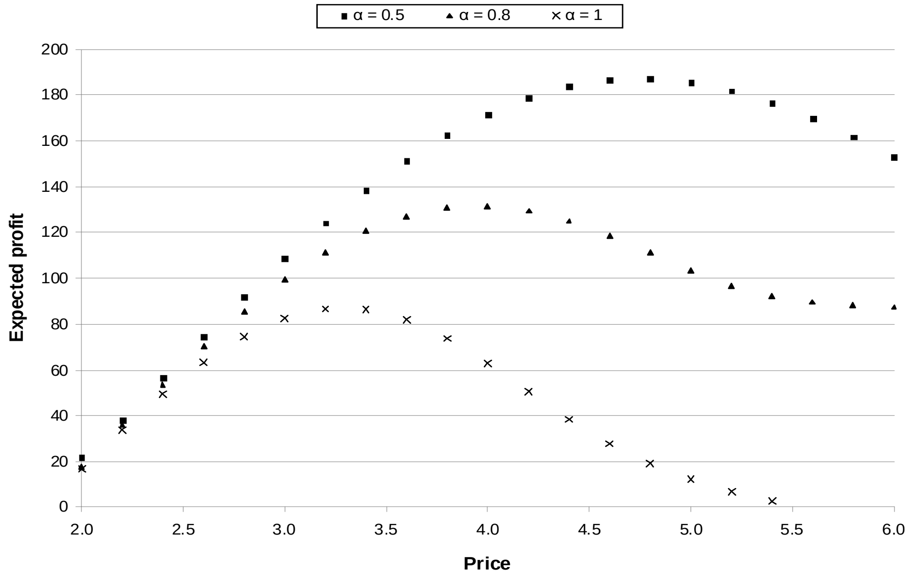

Figure 4-4. Profit versus price (MTS): Expected profit as price (p) varies for different refund fractions ($alpha$) under optimal assortment (S^*)

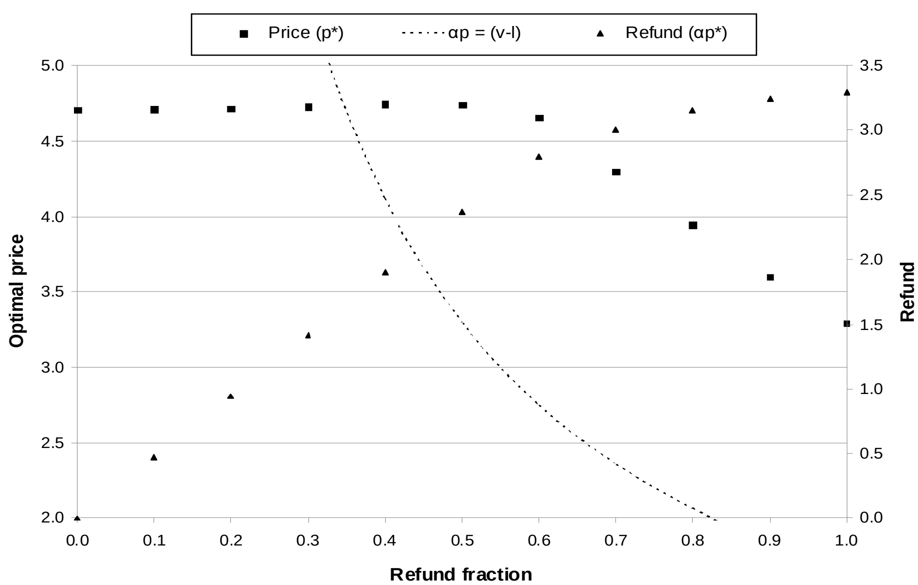

Figure 4-5. Price versus refund (MTO): Optimal price (p^*) and refund (αp^*) for different values of refund fraction (α) under optimal assortment (S^*)

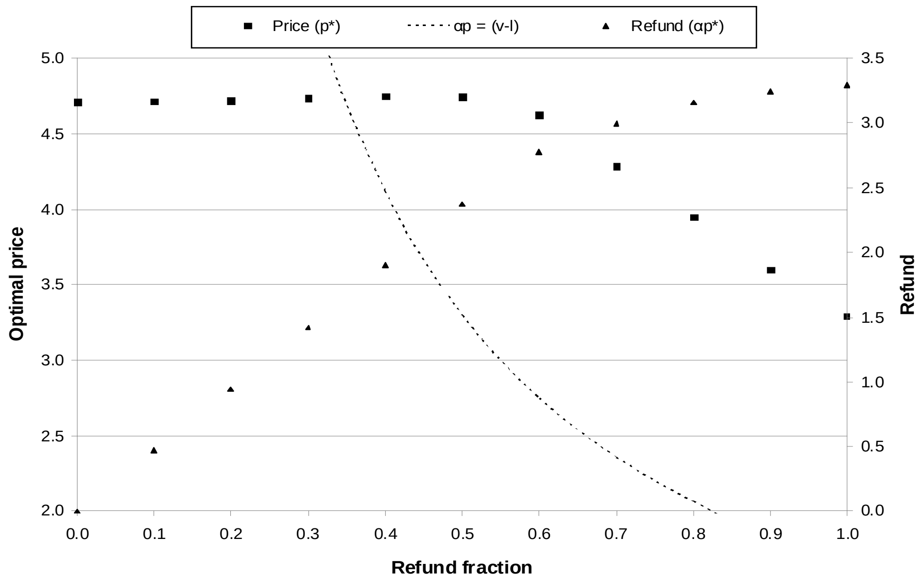

Figure 4-6. Price versus refund (MTS): Optimal price (p^*) and refund (αp^*) for different values of refund fraction (α) under optimal assortment (S^*)

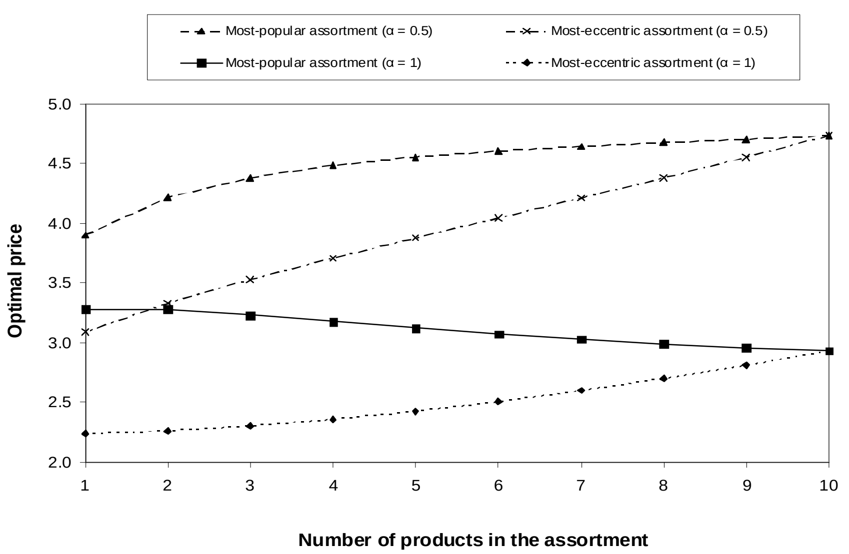

Figure 4-7. Optimal price (p^*) with different assortment structures for MTO case

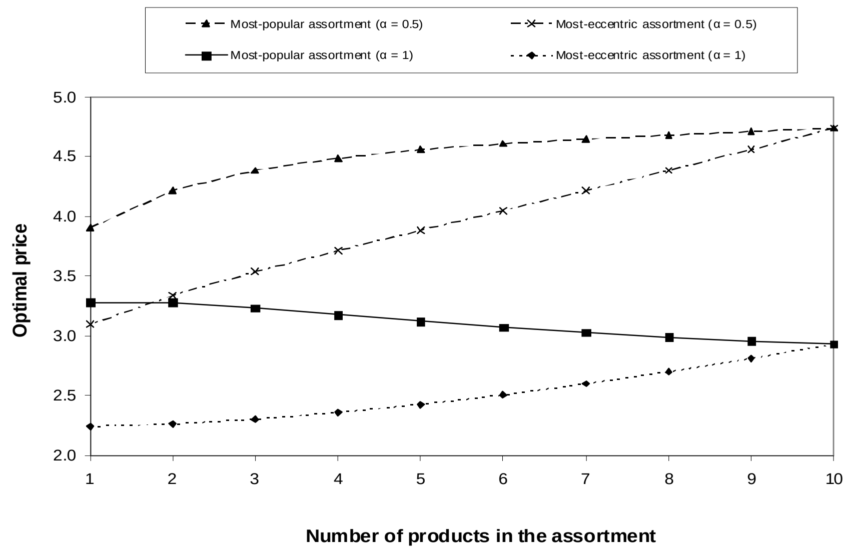

Figure 4-8. Optimal price (p^*) with different assortment structures for MTS case

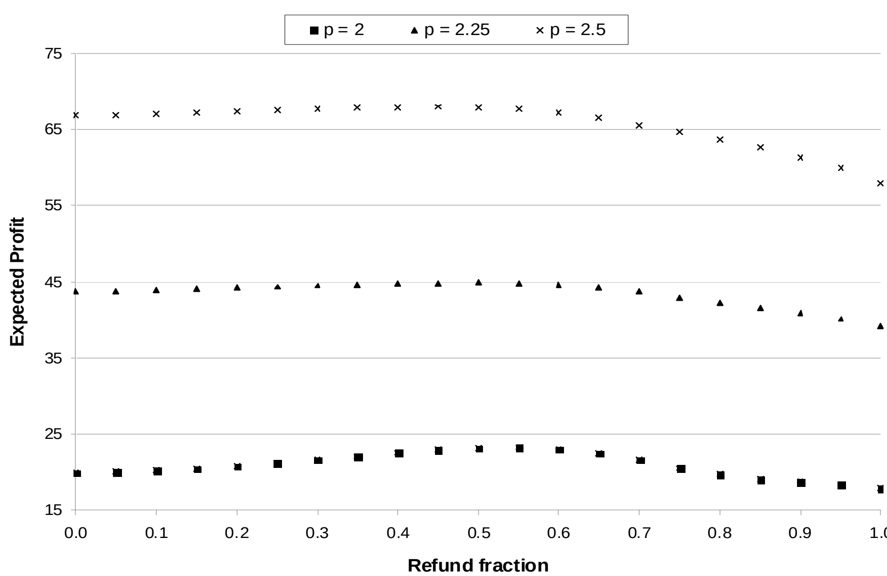

Figure 4-9. Profit versus refund fraction (MTO): Expected profit as refund fraction (α) varies for different prices (p) under optimal assortment (S^*)

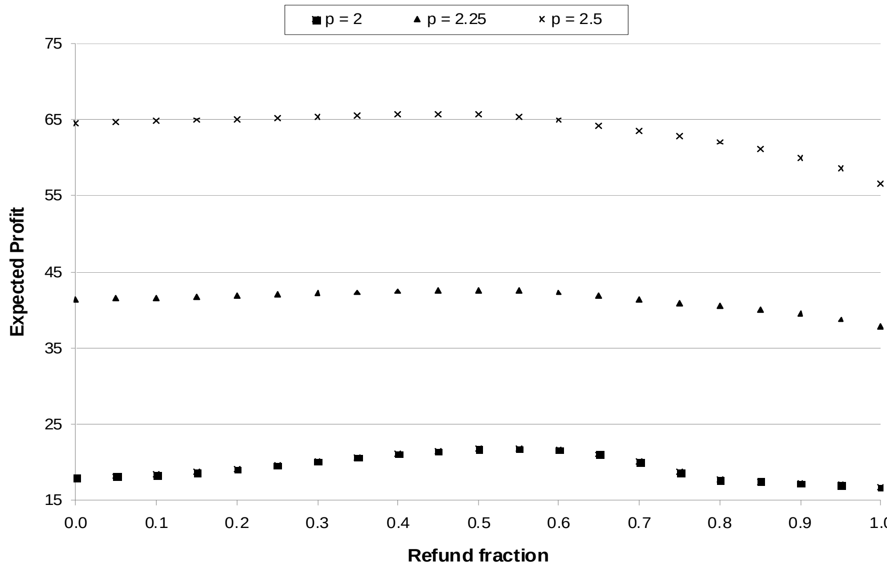

Figure 4-10. Profit versus refund fraction (MTS): Expected profit as refund fraction (α) varies for different prices (p) under optimal assortment (S^*)

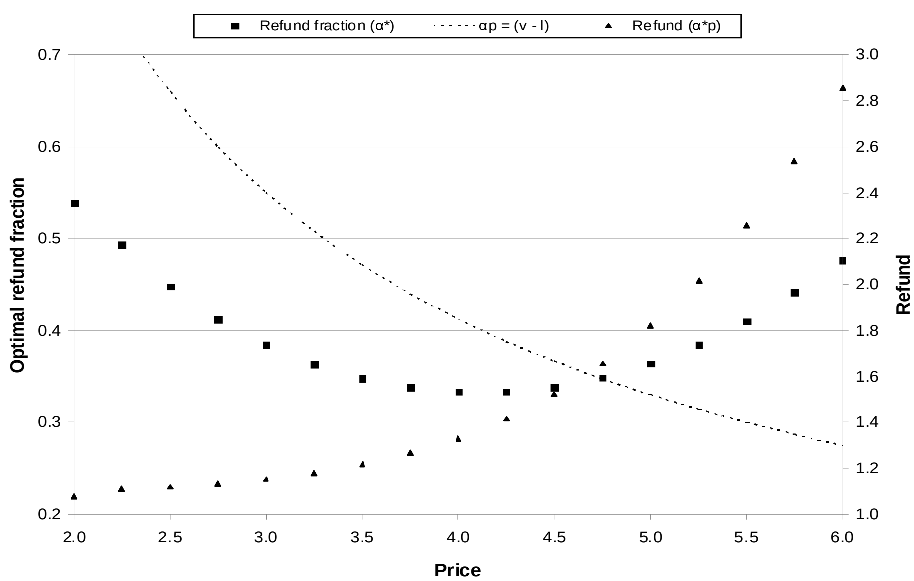

Figure 4-11. Refund versus price (MTO): Optimal refund fraction (α^*) and refund ($\alpha^* p$) for different prices (p) under optimal assortment (S^*)

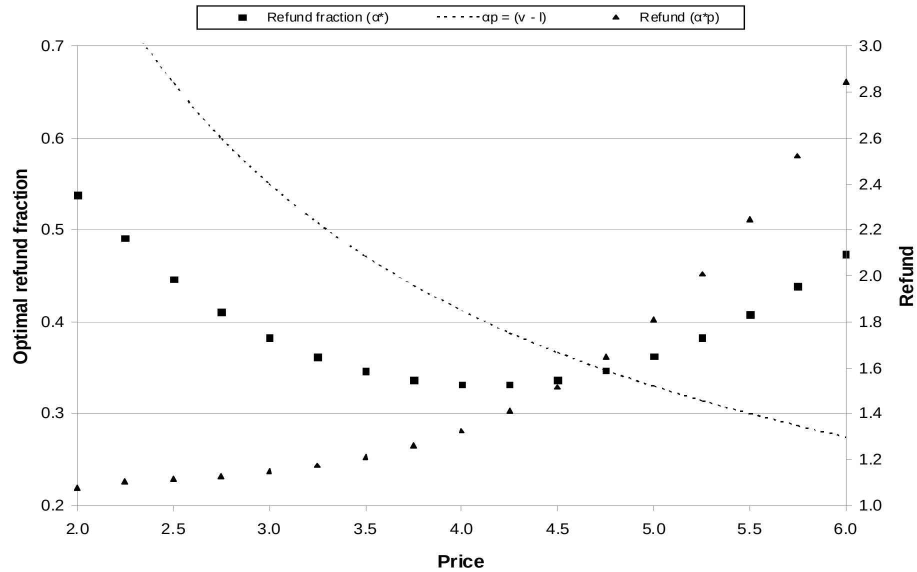

Figure 4-12. Refund versus price (MTS): Optimal refund fraction (α^*) and refund ($\alpha^* p$) for different prices (p) under optimal assortment (S^*)

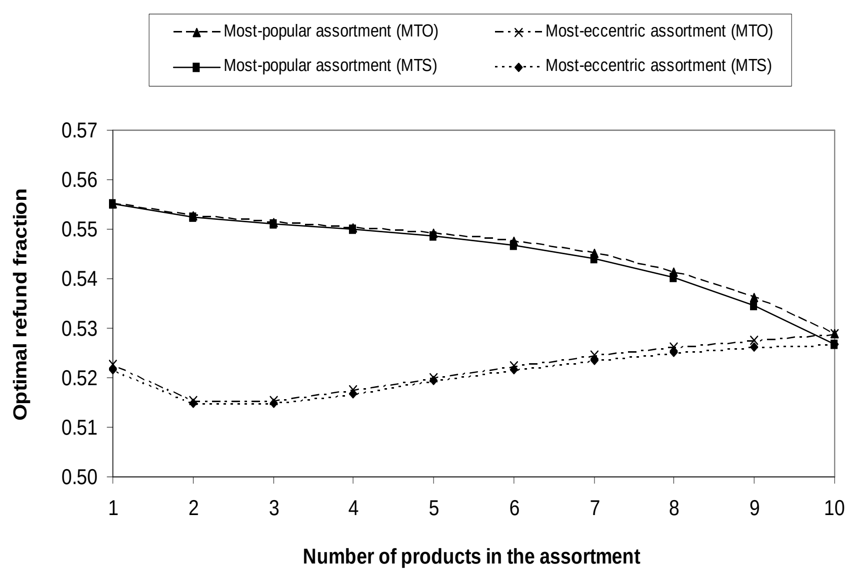

Figure 4-13. Optimal refund fraction (α^*) with different assortment structures

CHAPTER 5
OPTIMAL PRICE AND REFUND FOR A GIVEN ASSORTMENT IN A SINGLE PERIOD SETTING

5.1 Introduction

In the past decade, technological developments along with the expansion of the internet use worldwide have allowed firms to tighten links with their customers (Johnson 2002), establishing direct channels to identify consumer preferences. Today, manufacturers are able to elicit first-hand information about individual consumers' tastes without costly intermediary sales agents. An increasing number of companies from a myriad of industries are using online channels to offer customized goods to their consumers (Berman (2002) and Alptekinoğlu and Corbett (2008b) cite some specific examples). According to *The State of Retailing Online 2007* report[1] , 49% of the 170 retailers surveyed offered customized products, and a third considered product customization very effective. For example, Lands' End, a direct merchant of traditionally styled clothing, offers individually tailored garments ranging from pants to shirts and sweaters. Costumers just need to choose from a set of styles and fabrics, and answer a few questions about fit preferences and body type. Then, Lands' End delivers the customized product to their door. Also, best selling candle brand in the United States (U.S.), Yankee Candle, includes a collection of customized jar candles in their online catalog. Customers can choose their favorite fragrance and label design along with its accompanying set of adornments. Another example is furniture customizer Robb & Stucky, from which costumers can order pieces of furniture with several fabric and frame options.

Selling customized products makes inventory management in the forward supply chain easier, because production occurs after demand realizes (i.e., *make-to-order*). Customizing companies enjoy the advantages of reduced demand uncertainty and product

[1] Shop.org study conducted by Forrester Research, Inc.

postponement (Gupta and Benjaafar 2004). However, when consumer product returns, a habit that averages 7% of online sales in the U.S.[2] , are considered, these benefits may fade away. For instance, returns run two to three times the cost of outbound shipments (Stock et al. 2006). In addition, customization aggravates further the already costly return handling operations, since it leads to countless product variants that may, in the extreme, fit single individuals only, making a future resale impossible. Despite the seemingly troublesome combination of customization and returns, some customizers, in pursuit of market share or driven by competition, are willing to accept returns as well. After all, allowing returns can enhance the shopping experience and boost customer satisfaction. As a matter of fact, nine out of ten direct shoppers (i.e., online and catalog shoppers) indicated a convenient return policy as determinant to shop online, according to a 2007 survey commissioned by returns management provider, Newgistics, Inc. (Campanelli 2007).

Consumer return policies vary across industries and retailers, both with regard to money refunded and conditions for return (e.g., grace period, opened/unopened package). We also observe more significant differences in return policies between *brick-and-mortar* and online retailers. While lenient return policies (e.g., 100% money-back guarantees) are common for the former, the same cannot be said for the latter. Online retailers in general, and customizers in particular, have stricter return policies in the form of nonrefundable charges for shipping and handling costs or *restocking fees* (Hess et al. 1996). Hess et al. (1996) find that roughly 90% of the catalog retailers surveyed had nonrefundable shipping charges. For example, Lands' End and Yankee Candle allow returns for a full refund of the selling price minus shipping costs. Customization master Dell does not refund shipping costs either, and may charge up to 15% of the purchase price in restocking fees. Yet another example, Robb & Stucky imposes a *pick up* fee and a nonrefundable shipping and

[2] Source: *The State of Retailing Online 2007*

processing fee in case of return. Some customizing firms do not even accept returns. For example, clothing manufacturer Ralph Lauren customizes its polos allowing the customer to choose size, color of both polo and embroidered logo, and pony type (their brand icon). They do not accept returns for their Create Your Own collection.

The presence of product returns, then, raises interesting questions for customizing firms. Should these companies accept product returns? If so, what is the optimal return policy? How would their prices compare under different return policies? Prices and return policies are determinant factors in an online purchase of customized products. To this end, we consider a single firm, a customizing manufacturer who has online retailing presence on the internet, and offers a type of product that can be customized into a wide (finite) variety of end products, which we call product variants. We are particularly interested in customized products with variants that may each be of interest to multiple customers, i.e., customization is accomplished over a finite set of options as in the examples mentioned above. In our work we do not consider customization via personalized features such as monogramming, embroidering or engraving. Suitable customization examples for our study would involve choice of color (apparel, electronics, furniture), style and design (hats, apparel), frame and fabric (furniture), and scent (candles), to name a few.

The customizing firm accepts product returns, which we assume are in as-good-as-new condition. We characterize the firm's return policy by the money refunded to the customer in case of return (as in Hess et al. 1996, Yalabik et al. 2005, and Shulman et al. 2008). The purchase of a customized product is clearly a two-stage decision process, where the customer chooses a product variant first and, upon receipt, she decides to keep it or return it (Wood 2001). As in Chapter 3, we model purchase and subsequent keep/return decisions with a nested multinomial logit model. In this model, consumers with heterogeneous tastes seek to maximize their expected utility by either choosing a feasible product variant or an outside option (i.e., not buying from the firm). Therefore, although we study the operations of a single firm, competition is also present (but passive) in our model through this outside option.

In this chapter, we address the problem of product returns for customized products in a single period setting. We study the multiperiod setting in Chapter 6. The former is appropriate in contexts where there is a single selling season (e.g., Christmas, merchandise for special events). Moreover, it will lay the groundwork for the study of the latter, suitable for other products for which multiple selling periods exist (e.g., classic clothing). In a single period context, all returns can be salvaged at the end of the season (e.g., sold in a secondary market). However, when there are multiple periods, the disposition decision is not as straightforward. Returns can either be salvaged, or be stocked for future periods. Therefore, the presence of returns gives rise to a make-to-order (MTO) environment where inventory policy may play a critical role. In Chapter 6, we also tackle this inventory problem, and show the optimality of a salvage-down-to policy. We characterize a closed form expression for the threshold inventory level beyond which all returns should be salvaged. We show the benefits of employing such inventory policy for customizing firms.

5.2 Literature Review

Motivated by an increasing trend towards tightened return policies in practice, consumer return policies have gained considerable attention in research literature lately. Research related to our work characterizes return policies by the amount refunded to the customer. Leaving generous 100% money-back guarantees (Davis et al. 1995) behind, a number of papers in both marketing and operations management explore stricter return policies offered by retailers. From a purely marketing perspective, several authors find that partial refunds are optimal under a vast number of different scenarios. Hess et al. (1996), for example, show that nonrefundable charges are optimal in direct marketing, and support their findings with empirical data. Mukhopadhyay and Setoputro (2004) and Shulman et al. (2007) optimize price and refund in a single firm and a duopoly scenarios, respectively. Matthews and Persico (2005) introduce product information acquisition by costumers, and investigate how it affects optimal price and refund. Shulman et al. (2008) extend this last work to two products where an exchange is also possible.

In the operations management literature, business-to-business return policies have been studied (Emmons and Gilbert 1998). But current literature on consumer return policies is somewhat scarce. To the best of our knowledge, only three papers allow endogenous refund decisions in a retailing context. Mukhopadhyay and Setoputro (2005) adopt a deterministic linear demand model that depends on price, refund and modularity, where the last two are decision variables. Yalabik et al. (2005) use a demand model with two consumer types that are either a match or a mismatch. They integrate logistics and marketing decisions into the return system, and conclude that optimal refunds are not unique. On the other hand, Su (2008) considers all customers ex post heterogeneous. He endogenizes price and order quantity in addition to refund, and finds that partial refunds are optimal. In our model, as in Yalabik et al. (2005) and Su (2008), demand is driven by consumer valuation of the product, and we consider all customers ex post heterogeneous as in Su (2008). In contrast to these three papers, our demand model entails product choice and an outside option. We are the first to study consumer return policies in a multiple period setting with multiple products.

Our work also draws from three other literature streams: customization, assortment planning, and pricing and inventory coordination. Product customization has been gathering growing interest in the past decade. Strategic (Pine 1993, Murthi and Sarkar 2003), inventory- and production-related (see Swaminathan and Tayur (2003) for a review), and competitive issues involving price and product mix (Dewan et al. 2003, Syam and Kumar 2006, Alptekinoğlu and Corbett 2008b) have been addressed to date. For a review of assortment planning literature, the reader is referred to Kök et al. (2006). In inventory theory, Chan et al. (2004) and Yano and Gilbert (2005) provide comprehensive reviews of models that jointly optimize pricing and inventory decisions. Inventory models with disposal or salvaging option are also related to our work. In the former, discarding inventory is costly (Simpson 1978, Inderfurth 1997), whereas it is profitable (Lovejoy 1990, Petruzzi and Monahan 2003) in the latter.

Our work merges different research topics that have been traditionally separate in the past. We contribute to the returns literature with a model that includes an inventory strategy stemming from our multiperiod approach, and a stochastic demand model with multiple (customized) products.

5.3 Model Description

5.3.1 Firm

We consider a firm, also called retailer or customizer interchangeably hereafter, that offers a product that can be configured into many variants by several horizontal differentiation attributes. The finite number of different product variants n constitutes the retailer's assortment. Although our model is valid for any $n \geq 1$, we consider n to be quite large for a customizing firm. Each product variant, indexed by $i = 1, \ldots, n$, is horizontally differentiated, and for simplicity, we assume that all of them have the same *attractiveness* a. For instance, in the apparel or furniture industries, a different color or pattern would determine horizontally differentiated product variants, whereas materials of different quality would not, because everyone would agree to pay less for a lower quality item. The parameter a, identical for all customers, represents the deterministic component of the value or utility that a customer obtains from a product. Other authors also call it *nominal utility* (Anderson and de Palma 1992, van Ryzin and Mahajan 1999). Under this setting, it is reasonable to assume that all product variants have also identical sales price p, unit production cost c, and salvage value v. This is consistent with institutional practice; for example, Lands' End charges the same price for customized pants regardless of the options chosen (Syam and Kumar 2006).

The retailer accepts product returns which we assume to be in perfectly resalable condition. The retailer refunds b, with $b \leq p$, and incurs a reverse logistics cost l (e.g., sorting, testing). Thus, the difference $(p - b)$ would be the retailer's nonrefundable charge (e.g., shipping cost or restocking fee). A consumer returning a product receives the refund b, and *experiences* a "hassle" cost k (see Davis et al. (1998) for a discussion of what this cost may entail).

In this Chapter we consider a single period problem although we set up the model for a multiple period context, analyzed in Chapter 6. We assume that returns arrive at the end of every period. This is a reasonable assumption when returns are accumulated and processed to verify their condition at the end of the period. The customizer salvages all returns at the end of the period collecting a per-unit revenue of βv, where β is a per-period discount factor ($0 \leq \beta < 1$). As in Hess et al. (1996), we assume that returns are economically efficient, i.e., $\beta v > l$. In a single period context, the customizer salvages all returns. In a multiple period context (Chapter 6), however, the customizer may choose to keep a portion of the returns in stock, and satisfy some of the future demand from this stock. Although products are made to order, the presence of returns forces the firm to adopt an inventory strategy. Specifically, the firm employs a salvage-down-to inventory control policy (which is shown to be optimal in §6.2). That is, at the beginning of every period, if the amount of inventory for a particular product exceeds a certain threshold, salvage-down-to level, then the firm salvages the returns in excess of this threshold, collecting a per-unit salvage revenue v. The firm incurs an inventory holding cost h for each product unit remaining in inventory. Note that setting this threshold to zero would mean the firm salvages all returns, and keeps no finished goods inventory.

We assume that the initial inventory is zero for all product variants, and the firm operates for a long time so that an infinite number of periods, indexed by t, is acceptable. Demand, returns and all parameters are considered stationary in time. Revenues and costs are discounted from period to period by β, as we mentioned earlier. This parameter captures the time value of money. We let x_{it} and y_{it} be the units of product variant i in stock at the beginning of period t before and after the salvage decision, respectively.

The units of product variant i salvaged in period t are then given by $x_{it} - y_{it} \geq 0$, with $x_{i1} = y_{i1} = 0$.

From a supply standpoint, the firm customizes every product after demand is realized at a unit cost c (with $c > \beta v$ for the problem to be well posed). Within this make-to-order scenario then, the ordering decision becomes trivial since there is no need to plan orders ahead in time. We assume that lead times are not an issue in the model, and customers are willing to wait for their customized product. We assume that products can be customized and delivered to the customer within the same period. Considering positive deterministic lead times larger than a period is unlikely to change anything structurally as long as we assume that production of an order starts right after demand is realized.

The customizer, who seeks to maximize its expected profit, needs to make three decisions: (1) *selling price*, p, (2) *refund*, b, and (3) *inventory policy* characterized by the number of products salvaged prior to every selling period, $(x_{it} - y_{it})$. For the single period problem the inventory decision is trivial and all product returns will be salvaged at the beginning of next period.

5.3.2 Demand and Return Processes

In a typical period t, customers arrive according to a Poisson process with rate λ. Upon their arrival, they observe the set of possible product variants, and choose the one that maximizes their expected utility. They can also choose not to buy at all, the so-called outside option denoted with subscript 0. As in Chapter 3, the product choice process, as well as the subsequent return behavior are modeled using the nested multinomial logit (N-MNL) model (Ben-Akiva 1973). That is, in the first stage of the N-MNL, the consumer chooses the product variant that maximizes her utility (if any) with probability q, same for all variants because their a was assumed to be the same. With probability q_0, she chooses the outside option and therefore, does not purchase a customized product from the firm. Obviously, $q_0 + nq = 1$. In the second stage then, given that a product variant has been purchased, the costumer decides whether to keep it or return it with probabilities $q_{keep|\,i}$

and $q_{return|\,i}$, respectively (with $q_{keep|\,i} + q_{return|\,i} = 1$). For notational convenience, we use q_r hereafter for the conditional probability of return $q_{return|\,i}$ since products are identical in terms of all problem parameters.

Keeping or returning a product. Given that a product variant has been chosen in the first stage, the costumer faces the post-purchase decision of keeping or returning the product. She associates a utility with a deterministic and a random component for each case: $u_{keep} = a - p + \epsilon_{keep}$, and $u_{return} = b - p - k + \epsilon_{return}$, where ϵ_{keep} and ϵ_{return} are independent and identically distributed (*iid*) Gumbel random variables with mean zero and scale $1/\mu_2$ ($\mu_2 > 0$). The customer then chooses the option with maximum utility, $\max(u_{keep}, u_{return})$. This leads to the following probability of return after a product is purchased in the first stage:

$$q_r = \frac{1}{1 + \exp\left(\frac{a+k-b}{\mu_2}\right)}$$

Choosing a product (or the outside option). The first stage is also constructed based on the principle of utility maximization. Hence, each customer when selecting the features of her customized product, associates a utility U_i with each product variant $i = 1, \ldots, n$, and a no-purchase utility U_0. The deterministic component of these utilities is given by the *pre-purchase expected utility* $\tilde{u}$ and $\tilde{u}_0$, respectively. $\tilde{u}_0$ is an exogenous parameter that accounts for alternative options to the purchase such as substitutable products or products by competitors. In other words, competition is also present in our model through this parameter although it is passive in the sense that does not respond to the firm's decisions. The pre-purchase expected utility for each product, given by $\tilde{u} \equiv E\left[\max(u_{keep}, u_{return})\right]$, as follows

$$\tilde{u}\,(p, b) = \mu_2 \ln\left[\exp\left(\frac{a}{\mu_2}\right) + \exp\left(\frac{b-k}{\mu_2}\right)\right] - p$$

Whenever it is necessary, we will write $\tilde{u}\,(p, b)$ to make the dependence explicit. Note that a more lenient return policy (i.e., higher refund b), makes the product more *utile* (i.e., higher pre-purchase expected utility $\tilde{u}$). Obviously, a lower price also increases the expected utility. The customer therefore chooses the option that maximizes her utility, i.e., $\max_{i\in\{1,\ldots,n\}}(\tilde{u}+\varepsilon_i, \tilde{u}_0+\varepsilon_0)$, with respective probabilities

$$q = \frac{\exp\left(\frac{\tilde{u}}{\mu_1}\right)}{n\exp\left(\frac{\tilde{u}}{\mu_1}\right)+\exp\left(\frac{\tilde{u}_0}{\mu_1}\right)} \qquad q_0 = \frac{\exp\left(\frac{\tilde{u}_0}{\mu_1}\right)}{n\exp\left(\frac{\tilde{u}}{\mu_1}\right)+\exp\left(\frac{\tilde{u}_0}{\mu_1}\right)}$$

Since a, and thus $\tilde{u}$, are equal for all product variants, they all have the same probability q of being purchased.

Aggregated Demand and Returns. We assume that in period t customers arrive (or enter the retailer's website with the purpose of buying) following a Poisson process with rate λ. Each arrival will correspond to a sale with probability nq and a no-purchase with probability q_0. The two events, purchase and no-purchase, are independent Poisson processes (Ross 2003, p. 296) with rates λnq and λq_0, respectively. By the same argument, every purchase will result in a return with probability q_r, and no return (or final sale) with probability $(1-q_r)$. Therefore, purchases returned and kept are also independent Poisson processes, with rates λnqq_r and $\lambda nq(1-q_r)$, respectively. We denote D_t and R_t the total demand and total returns in period t, respectively. Poisson processes are used extensively in stochastic inventory literature with product returns (de Brito and Dekker 2003).

5.4 Analysis

To study the pricing and refund policy for a single period problem we drop the subscript t. As mentioned earlier, this one period setting is suitable for products that become obsolete (or have their value reduced considerably) once the selling season is over. Fashion products or merchandise for events would fall into this category, for example. We assume that the beginning inventory is zero, and all returns arrive at the end of the season, precluding the possibility of a resale. Although this is a simplifying assumption, it is still plausible in our MTO context where products can take up to weeks to be customized and shipped to the costumer. For instance, Lands' End customized products may take up to four weeks to be delivered. Consider, for example, a single selling period for a Christmas season which starts after Thanksgiving holiday in the U.S. If products can take several weeks to be customized, by the time customers receive their product plus the return grace period, the returned good will arrive too late to fulfill another order. In addition, our refund and price results for the single period case will shed some light on the analysis for the more involved multiperiod problem. Therefore, the retailer must salvage all returns at the end of the season, incurring a net loss (recall that $\beta v < c$). A product's profit in the single period problem can be expressed as follows:

$$\pi(p, b) = (p - c)D - (b + l - \beta v)R \tag{5–1}$$

where D represents the (random) demand for each product variant, and R represents its (random) returns. The retailer, for each unit demanded, obtains a net margin $(p - c)$, and for each unit returned, incurs a cost $(b + l)$ minus the discounted salvage value βv. The retailer aims to maximize the expected value of (5–1) for all product variants which, after plugging $E[D] = \lambda q$ and $E[R] = \lambda q q_r$, results in

$$\Pi(p, b) = nE[\pi(p, b)] = \lambda n q \left[p - c - (b + l - \beta v)q_r\right] \tag{5–2}$$

It will be convenient for exposition of our analysis to denote M the expected profit margin per unit sales of a product variant, that is, $M = p - c - (b + l - \beta v)q_r$. Then, $\Pi = \lambda n q M$. We are then interested in finding the price p and refund b that maximize the expected profit in (5–2). The results are shown in the following theorem:

Theorem 3. *The optimal price p^* and optimal refund b^* for the single period problem (5–2) are given by*

$$p^* = \left[n \exp\left(\frac{\tilde{u}^*(p^*) - \tilde{u}_0}{\mu_1}\right) + 1\right]\mu_1 + c$$

$$b^* = \beta v - l$$

where $\tilde{u}^*(p^*) = \mu_2 \ln\left[\exp\left(\frac{a}{\mu_2}\right) + \exp\left(\frac{\beta v - l - k}{\mu_2}\right)\right] - p^*$.

Theorem 3 provides an implicit expression for the optimal price that can be easily computed using a simple search. The optimal refund is ensured to be strictly between zero and p^*, since $\beta v > l$ and $\beta v - l < p^*$. We will refer to it as partial refund case hereafter. In the case we ruled out previously where $\beta v \leq l$, clearly the firm would not accept returns because that would suppose an extra cost. Illogically, in the event of a return, it would be more economical for the firm to let the costumer keep the product even after refunding her. This optimal refund has a very neat intuition: the retailer should set its refund policy such that returns are neither a net cost nor a net benefit. This finding agrees with Matthews and Persico (2005) and Su (2008), although the latter does not consider any logistics cost to process the return. Setting such refund "implements an efficient ex post allocation of the good", as Su mentions, and protects the firm from opportunistic returns (e.g., consumers that buy a product, use it and return it). There is a noteworthy observation from Theorem 3 that we synthesize in the following corollary:

Corollary 1. *The optimal price p^* increases with the number of product variants n offered.*

Corollary 1 confirms what many authors have found in the past regarding broader product lines (Kekre and Srinivasan 1990, Hopp and Xu 2005). As the company increases its product assortment, the customer is more likely to find her desired product, and consequently, the firm can increase its price. In product customization, it is generally accepted that customers are willing to pay more for products that better match their preferences (Kotha 1995, Berman 2002, Syam and Kumar 2006). For example, Lands' End costumers can purchase a standard pair of chino pants for $29.50 or a customized pair for $80.

5.4.1 Special Case with No Returns Allowed

We now turn our attention to the case where the retailer does not accept returns (we use subscript NR throughout). We should note that not allowing returns makes a product less attractive for the costumer since $\tilde{u}_{NR}$ is strictly smaller than $\tilde{u}$ for a given price. The expected profit in the no returns case is given by

$$\Pi_{NR}\left(p_{NR}\right) = \lambda nq(p_{NR} - c) \tag{5–3}$$

Maximizing the expected profit above yields the optimal price given in the following theorem:

Theorem 4. *The optimal price p^*_{NR} for the single period problem with no returns allowed (5–3) is uniquely given by the implicit expression*

$$p^*_{NR} = \left[n \exp\left(\frac{\tilde{u}^*_{NR}\left(p^*_{NR}\right) - \tilde{u}_0}{\mu_1}\right) + 1\right] \mu_1 + c$$

where $\tilde{u}^*_{NR}\left(p^*_{NR}\right) = a - p^*_{NR}$.

The expression for the optimal price in Theorem 4 is akin to the one for partial refunds in Theorem 3, the only difference is in the pre-purchase expected utility $\tilde{u}_{NR}$. In fact, by comparing such utilities for both cases, we are able to establish the following proposition:

Proposition 2. *In the single period problem, the optimal price for the partial refund case is strictly greater than the optimal price for the no returns case. That is, $p^* > p^*_{NR}$.*

Accepting returns increases the expected utility of a product, as we mentioned earlier. In addition, returns can be seen as a service measure since they reduce customers' risk in the purchase decision. In consequence, retailers are able to charge higher prices to compensate for the service offered (Davis et al. 1995, Tsay and Agrawal 2000). Proposition 2 is also consistent with the result by Mukhopadhyay and Setoputro (2004) where demand and returns are characterized by a linear model.

Next, we compare the optimal expected profit between the returns and no returns cases in the following proposition.

Proposition 3. *In the single period problem, the optimal profit for partial refunds is strictly greater than the optimal profit for no returns. That is, $\Pi^* > \Pi^*_{NR}$.*

Proposition 3 implies that customizing firms are always better off by accepting returns (under the conditions of our model). We have to be cautious with this finding though. We have seen in Theorem 3 that the optimal refund is tightly related to the salvage value and the logistics cost of the return. As the degree of customization increases, the salvage value of a potential return might decrease considerably (e.g., the extreme case of monogrammed items with almost zero salvage value). Also, for some firms, product returns can be very costly, incurring a large l. Combined, low salvage value and high logistics cost could give rise to a very low optimal refund. For such cases, although the firm would benefit more from a (low) partial refund, it may go for the no refund alternative instead.

The main managerial insight for such firms would be to search for profitable secondary sales channels (e.g., outlet retail stores) to increase v, and develop efficient reverse logistics processes to decrease l. Effective returns both alleviate the consumers' purchase risk, and increase sales allowing the retailer to generate extra profit as seen in Proposition 3.

5.4.2 Special Case with Full Refund

We conclude the analysis of the single period model studying the special case where the retailer, in the event of a return, fully refunds the selling price and other costs charged to the customer in the selling transaction, that is, $b = p$. This 100% money-back guarantee policy is still standard in some industries, and especially for brick-and-mortar retailers (Davis et al. 1995). Nevertheless, in online retailing, and for customized products specifically, such generous policies are unusual. In fact, it can be shown as a consequence of Theorem 3, that full refund is never optimal (as in Mukhopadhyay and Setoputro 2004, and Su 2008). In this section, given that refund $b = p$, we are interested in finding the optimal price. First, we derive a very intuitive upper bound for the optimal price. We use the subscript FR (full refund) throughout this section. For full refund, the expected profit margin per unit sales of product is $M_{FR} = p_{FR} - c - (p_{FR} + l - \beta v)q_r$. We can then write the retailer's expected profit as $\Pi_{FR}(p_{FR}) = \lambda nqM_{FR}$. We assume that there exists some p_{FR} such that $M_{FR} > 0$, otherwise it would not be profitable for the firm to operate. The study of M_{FR} yields the following lemma:

Lemma 4. *The upper bound $\bar{p}_{FR}$ for the optimal price in the single period problem with full refunds is the maximizer of M_{FR}, and is given by*

$$(\bar{p}_{FR} + l - \beta v)q_r = \mu_2$$

Moreover, $\bar{p}_{FR}$ exists and is unique as long as $M_{FR} > 0$ for all price points considered.

The upper bound $\bar{p}_{FR}$ has a simple interpretation: the expected profit Π_{FR} is the product of the expected demand, λnq, and the expected margin, M_{FR}. Since probability of purchase q decreases in price, the optimal p_{FR} will never be larger than $\bar{p}_{FR}$, the maximizer of M_{FR}. We are now just one step away from finding an expression for the optimal price:

Theorem 5. *The optimal price p^*_{FR} in the single period problem with full refund is uniquely given by the implicit expression*

$$1 - \frac{q_r}{\mu_2}(p^*_{FR} + l - \beta v) = \frac{q_0 M_{FR}}{\mu_1}$$

Again, we provide an implicit expression to compute the optimal price for the full refund case. A simple search on p_{FR} will lead us to the optimal price p^*_{FR}. The interpretation of the expression provided by Theorem 5 is as follows: a raise in the price increases the probability of no purchasing, which drives the expected profit down (reflected on the right side term above). In the absence of this loss, we would set the price such that the expected margin M_{FR} was maximized (i.e., setting the left side term to 0). Considering both effects, the optimality expression trades off the decrease in total sales versus the decrease in expected profit margin per unit sales.

Note that from Theorem 5, either when $(p^*_{FR} + l - \beta v)$ or q_r are close to zero, the optimality condition tends toward that of the general case (see Theorem 3's proof).

From a managerial point of view, this is equivalent to a customizing firm that is able to efficiently process returns and extract a high value from them ($p^*_{FR} + l - \beta v \simeq 0$), or simply, a customizing firm from an industry where product returns rarely occur ($q_r \simeq 0$). Under either circumstance, the firm might consequently adopt a full refund policy. Although it is not the purpose of our study, this result may justify partly why 100% money-back guarantees are used in other contexts such as brick-and-mortar retailers who can easily place returns back on the shelves for a resale.

5.5 Conclusion

In this chapter, we study the optimal pricing and consumer return policy of a multiproduct customizing firm. The firm offers a type of product that can be customized over a set of different features, establishing a finite but possibly large assortment of end product variants. One of our main conclusions, from a managerial perspective, is that retailers should aim for costless returns when designing their consumer return policies. This could be achieved by efficient reverse logistics processes and profitable secondary markets. We also show that offering partial refunds is optimal. The firm is also able to raise its selling price while offering partial refunds. This can be seen as a service premium since returns reduce customers' risk in the purchase decision.

CHAPTER 6
OPTIMAL PRICE, REFUND AND INVENTORY POLICY FOR A GIVEN ASSORTMENT IN A MULTIPLE PERIOD SETTING

6.1 Introduction

In Chapter 5, we were interested in finding the optimal price and refund for a customizing firm in a single period setting. The inventory decision was trivial in that case since all returns were salvaged at the end of the period. However, the presence of inventory due to returns complicates the characterization of the optimal price and refund considerably in a multiperiod setting (as we will demonstrate later). The firm must first choose a price and refund at the beginning of the selling horizon (static decision), and then decide the inventory status dynamically. We analyze the inventory problem first; our infinite horizon stationary approach yields a structural result on the optimal inventory policy, and an easy-to-compute solution. Later, we analyze the optimal price and refund. We use the same notation as in Chapter 5.

6.2 Optimal Inventory Policy

Given a price p and a refund b, the study of the inventory policy for an assortment of n product variants can be carried out for individual product variants since they are only interrelated through demand, and there is no storage constraint of any kind. Therefore, we drop subscript i for ease of exposition. The profit obtained by a product variant in the assortment with initial inventory x_t in period t is given by

$$\pi_t(x_t) = \; v(x_t - y_t) - hy_t + pD_t - c[D_t - y_t]^+ - (b + l)R_t \tag{6–1}$$

where $[a]^+ = \max\{0, a\}$. Recall that the inventory analysis consists of deciding how many units to salvage every period, that is, $x_t - y_t$. The sequence of events in each period, which corresponds to the order of terms in (6–1), is as follows. (i) The salvage decision is made and a revenue of $v(x_t - y_t)$ is obtained. (ii) The inventory holding cost hy_t is incurred for the items stocked. (iii) Demand D_t is realized and its corresponding revenue pD_t is collected. (iv) Demand is then satisfied with on-hand inventory (if any), and with additional production (if necessary) at a production cost of $c[D_t - y_t]^+$. (v) Products are received by customers and after examination, some (R_t) will be eventually returned in perfect condition, and then, the firm makes refunds and incurs reverse logistics costs, totaling $(b+l)R_t$.

The present value of the expected profit of all product variants over infinite periods is then:

$$\Pi_\infty = n \sum_{t=1}^{\infty} \beta^{t-1} E[\pi_t(x_t)] \tag{6–2}$$

where the successive period's initial inventory levels are related by $x_{t+1} = [y_t - D_t]^+ + R_t$ and y_t are the decision variables. Our infinite horizon inventory problem can be formulated as

$$\begin{aligned} \max \quad & \Pi_\infty = \max \left\{ n \sum_{t=1}^{\infty} \beta^{t-1} E\left[v(x_t - y_t) - hy_t + pD_t - c[D_t - y_t]^+ - (b+l)R_t \right] \right\} \\ \text{subject to} \quad & 0 \le y_t \le x_t \qquad t = 1, 2, \ldots \\ & x_{t+1} = [y_t - D_t]^+ + R_t \qquad t = 1, 2, \ldots \end{aligned} \tag{6–3}$$

Note that this formulation is quite complicated since it deals with an expectation of infinitely many random variables. As in Heyman and Sobel (1984, p. 65), we first transform (6–3) into a much easier problem.

Proposition 4. *The inventory problem in* (6–3) *is equivalent to*

$$\begin{aligned} \max \quad & n \sum_{t=1}^{\infty} \beta^{t-1} E\left[pD_t - (b + l - \beta v)R_t \right] - n \sum_{t=1}^{\infty} \beta^{t-1} E[G(y_t)] \\ \textit{subject to} \quad & 0 \le y_t \le x_t \qquad t = 1, 2, \ldots \end{aligned} \tag{6–4}$$

where $G(y_t) = (v+h)y_t + cE[D_1 - y_t]^+ - \beta v E[y_t - D_1]^+$.

Note that maximizing the objective function in (6–4) is equivalent to minimizing its last term. The key of this result, as in Heyman and Sobel (1984), is being able to collect all terms involving y_t, evaluate their expectations, and obtain expressions, namely $G(\cdot)$, that depend only on the y_t's and not on the x_t's. As shown in Lemma 5 below, such minimization can be easily computed each period (as a single period problem) suppressing temporal dependence. This type of policy is called *myopic policy*, and it turns out to be optimal in our case (see Heyman and Sobel (1984) for a discussion of other optimal myopic inventory policies).

Lemma 5. *The function $G(y)$ is convex and has a minimum which is given by the smallest integer s^* such that*

$$F(s^*) \geq \frac{c - v - h}{c - \beta v} \tag{6–5}$$

where $F(\cdot)$ is the cumulative distribution function of the demand for a product variant.

We assume that $c - v - h > 0$ for the problem to be of interest. Otherwise, there would be no incentive to retain inventory, and all returns would be immediately salvaged. Under this assumption, the critical fractile provided in Lemma 5 is always between 0 and 1. Note that this critical fractile has a familiar newsvendor interpretation. The numerator represents the cost of having one unit less in stock to satisfy demand (i.e., underage cost), that is, the unit production cost incurred (c) minus what the firm saves on holding (h), minus the revenue obtained by salvaging (v), which is negative because it is a revenue. The cost of having one extra unit in stock (i.e., overage cost) is the holding cost (h) plus the difference between the current and the subsequent (i.e., discounted) salvage values ($v - \beta v$). Adding the cost of underage and overage provides the expression in the denominator. Lemma 5 sets the basis for the optimal inventory policy detailed below.

Theorem 6. *The salvage-down-to inventory policy is optimal for the infinite horizon stationary problem formulated in (6–4). It has the form*

$$y_t^* = \begin{cases} x_t & \text{if } x_t \leq s^* \\ s^* & \text{if } x_t > s^* \end{cases}$$

where the integer s^* is denoted the optimal salvage-down-to level.

Theorem 6 states that if the inventory at the beginning of period t is below the *salvage-down-to level* s^*, we salvage nothing (i.e., $x_t - y_t^* = 0$), and conversely, if the inventory is above s^*, we salvage $(x_t - s^*)$ units to bring it down to s^*. This result is very general and independent of the demand distribution. Therefore, it could be applied to numerous situations where there exists a flow of returned products. It can be shown that such salvage-down-to inventory policy is also optimal for a finite horizon problem. In that case, optimal salvage-down-to levels are decreasing over time, and a closed form expression can be obtained only for the last two periods.

6.3 Optimal Price and Refund

Having resolved the inventory problem, we now analyze the optimal price and refund for the multiperiod problem. Under the optimal inventory policy found in §6.2, the firm decides price and refund at the beginning of the selling horizon in order to maximize the present value of the expected profit, which is given by

$$\Pi_\infty(p,b) = n\sum_{t=1}^{\infty}\beta^{t-1}E\left[pD_t - (b+l-\beta v)R_t\right] - n\sum_{t=1}^{\infty}\beta^{t-1}E\left[(v+h)y_t^* + cE[D_1 - y_t^*]^+ - \beta v E[y_t^* - D_1]^+)\right] \quad (6\text{–}6)$$

Unfortunately, the values of y_t^* for every period are unknown a priori, which complicates the analysis extremely. The truncated expectations in the second summation term in (6–6) impede us from finding an analytical solution for the optimal price. To give a sense of the analytical challenge we face here, Karlin and Carr (1962) study the joint optimization of inventory and price for a single product with no returns. Their case is easier in the sense that inventory position can be set to the optimal order-up-to level for every period. In our case, however, if initial inventory is below the optimal salvage-down-to level, we cannot increase it. Hence, the inventory position is a random variable as well. Despite the intractability of the expected profit under the optimal inventory policy, we are able to establish lower and upper bounds for Π_∞, that will allow us to derive the optimal refund

(b^*) and a lower and upper bound for the optimal price (p^*). Clearly, if we employ a policy of salvaging all returns before the beginning of each period, that is, we set $y_t = 0$ for all t, the resulting expected profit is a lower bound for the optimal expected profit. As for the upper bound, if we could relax the constraint $y_t \leq x_t$ from the inventory problem in (6–4), we would set the beginning inventory in each period to $y_t = s^*$ in order to maximize profit. To make this possible, we assume that the firm would be obtaining additional units, i.e., $[s^* - x_t]^+$, whenever necessary at no cost. Obviously this is not feasible, but, for computational purposes, this *ideal* scenario provides an upper bound for the optimal expected profit.

Lemma 6. *For a given price p and refund b, the expected profit under the optimal inventory policy $\Pi_\infty(p,b)$ is bounded below by $\underline{\Pi}_\infty(p,b)$ and above by $\overline{\Pi}_\infty(p,b)$, that is, $\underline{\Pi}_\infty \leq \Pi_\infty \leq \overline{\Pi}_\infty$, where*

$$\underline{\Pi}_\infty(p,b) = \frac{\lambda n q}{1-\beta}\left[p - c - (b + l - \beta v)q_r\right] \tag{6–7}$$

$$\overline{\Pi}_\infty(p,b) = \frac{\lambda n q}{1-\beta}\left[p - (b + l - \beta v)q_r\right] - \frac{n}{1-\beta}\left[(v+h)s^* + cE[D_1 - s^*]^+ - \beta v E[s^* - D_1]^+\right] \tag{6–8}$$

Lemma 6 provides two natural bounds for the expected profit that will be helpful in delimiting the optimal price. Note that the lower bound is nothing else than infinite discounted single period problems, i.e., $\underline{\Pi}_\infty = \Pi/(1-\beta)$. In addition, if we assume λ to be large enough, the Poisson demand with mean λq can be approximated by a Normal distribution with mean λq and standard deviation $\sqrt{\lambda q}$. Then, it is well-known that $s^* = \lambda q + z^* \sqrt{\lambda q}$, where $z^* = \Phi^{-1}\left(\frac{c-v-h}{c-\beta v}\right)$, and $\Phi(\cdot)$ is the cumulative distribution function of a standard normal random variable. Plugging the expression for s^* in (6–8) we obtain

$$\overline{\Pi}_\infty(p,b) = \frac{\lambda n q}{1-\beta}\left[p - (v+h) - (b + l - \beta v)q_r - \frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}\right] \tag{6–9}$$

where $\phi(\cdot)$ is the probability density function of a standard normal random variable.

The difference between the upper and lower bound expressions yields an upper bound on the additional expected profit the customizer can obtain from carrying inventory from period to period. We summarize this result in the following proposition:

Proposition 5. *The additional expected profit that can be obtained from allowing inventory carry-over from period to period (as opposed to salvaging all returned items immediately every period) is bounded above by*[1]

$$\frac{\lambda nq}{1-\beta}\left[c-(v+h)-\frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}\right]$$

As discussed above, finding an exact expression for the optimal price is not possible. Nonetheless, by analyzing the upper and lower bounds for the expected profit under the optimal inventory policy, we are able to bracket its value. The following Lemma is an intermediate step to find the lower and upper bounds for the optimal price.

Lemma 7. *(a) The price $\underline{p}^*$ and refund $\underline{b}^*$ that maximizes the lower bound of the expected profit $\underline{\Pi}_\infty$ are given, respectively, by*

$$\underline{p}^* = \left[n\exp\left(\frac{\tilde{u}\left(\underline{p}^*\right)-\tilde{u}_0}{\mu_1}\right)+1\right]\mu_1+c$$

$$\underline{b}^* = \beta v - l$$

(b) The price $\overline{p}^*$ and refund $\overline{b}^*$ that maximizes the upper bound of the expected profit $\overline{\Pi}_\infty$ are given, respectively, by

$$\overline{p}^* = \left[n\exp\left(\frac{\tilde{u}\left(\overline{p}^*\right)-\tilde{u}_0}{\mu_1}\right)+1\right]\mu_1+v+h+\frac{(c-\beta v)\phi(z^*)}{2\sqrt{\lambda}}\sqrt{n+\exp\left(\frac{\tilde{u}_0-\tilde{u}\left(\overline{p}^*\right)}{\mu_1}\right)}$$

$$\overline{b}^* = \beta v - l$$

(c) In addition, $\underline{p}^* \geq \overline{p}^*$.

[1] The upper bound is non-negative for reasonable values of λq.

Note that by Lemma 6, we can write the expected profit under the optimal inventory policy as a convex combination of the lower and upper bound, i.e., $\Pi_\infty = \theta\underline{\Pi}_\infty + (1-\theta)\overline{\Pi}_\infty$, with $0 \leq \theta \leq 1$. Being able to do so and after having found the optimal price for these bounds (Lemma 7), the analysis of the expected profit leads us to the Theorem below:

Theorem 7. *(a) The optimal refund b^* for the infinite horizon multiperiod problem is given by*

$$b^* = \beta v - l$$

(b) The optimal price p^* for the infinite horizon multiperiod problem satisfies

$$\overline{p}^* \leq p^* \leq \underline{p}^*$$

The results of this Section leave us with some important managerial implications for customizing firms. It is quite intuitive that a firm may potentially increase its profit by keeping some of the returns in inventory to satisfy demand in subsequent periods rather than salvaging them. The additional profit from allowing inventory carry-over, bounded from above in Proposition 5, is particularly relevant when keeping inventory is relatively cheap, or the customizer does not benefit from profitable secondary markets. What is more surprising is that customers may also benefit from this multiperiod scenario. Going from a single period setting to a multiple period setting complicates the inventory management of the firm. Having to deal with random inventory loads from period to period may increase the firm's overall operational costs. We show, however, that an effective salvage-down-to policy may allow the firm to pass some of the savings to its customers by lowering price as shown in Theorem 7b, where optimal p^* is shown to be lower than $\underline{p}^*$, the price from the salvage-everything case.

6.4 Expected Returns Heuristic for the Optimal Price

In Chapter 5, we optimized price and refund for the single period case. For the multiperiod case, we could only optimize the refund and provide lower and upper bounds for the optimal price. Unfortunately, we could not find an analytical solution for the optimal price. Even for a given price and refund, the expected profit value is not directly computable. We propose a simple heuristic to approximate the optimal price for the multiperiod problem. The heuristic is based on the approximation of the initial inventory level by the expected returns value.

From Theorem 7, we know that the optimal price will be of the form:

$$p^* = \left[n \exp\left(\frac{\tilde{u}^*(p^*) - \tilde{u}_0}{\mu_1}\right) + 1\right]\mu_1 + \theta^* c + (1-\theta^*)\left[v + h + \frac{(c-\beta v)\phi(z^*)}{2\sqrt{\lambda}}\sqrt{n + \exp\left(\frac{\tilde{u}_0 - \tilde{u}(p^*)}{\mu_1}\right)}\right] \quad (6\text{–}10)$$

The only missing piece which cannot be computed analytically is θ^*. Numerically, we can resort to Monte Carlo simulation methods to generate strings of random demands and random returns for every given price as we did in §4.4. With known demands and returns, we can easily compute the actual profit then. Averaging the profits of a sufficiently large sample of realizations, we can estimate the expected profit (Robert and Casella 1999, p. 208). By the Law of Large Numbers, this estimation would converge with probability 1 to the expected profit Π_∞ as the sample size goes to infinity. Therefore, for every θ, we can estimate the expected profit and choose θ^* that maximizes it. Although we may have to generate several samples of demand and returns for many different prices, this methodology would be effective especially when the price range (i.e., difference of price bounds) is relatively low.

To avoid the possibly large computational burden of having to search for θ^*, we propose a heuristic to estimate it. Let θ^H be that estimation, and p^H its corresponding price. The lower and upper bound expressions for the optimal expected profit are obtained by assuming that the inventory level in every period is 0 and s^*, respectively. In reality, the inventory level will not necessarily be at either of these bounds, and will take a value in between. Since initial inventory is directly related with returns from the previous period, our heuristic uses the expected returns, $E[R_t]$, as approximation of the initial inventory level y_t^*. Clearly, if expected returns are above the salvage level s^*, we can expect to salvage some returns in every period and begin with an inventory after salvaging $y_t = s^*$. For such a case, the upper bound $\overline{\Pi}_\infty$ will approximate fairly well the optimal expected profit, and therefore, θ^H will be close to 0, establishing $p^* \simeq \overline{p}^*$. At the other extreme, when expected returns are close to 0, the lower bound $\underline{\Pi}_\infty$ will be a good approximation of the expected profit, setting $\theta^* \simeq 1$ and $p^* \simeq \underline{p}^*$. This is also true when $s^* \simeq 0$, of course. For the rest of the cases, when $E[R_t]$ falls in between 0 and s^*, our heuristic (in Step 4) will measure the relative differences between $E[R_t]$ and 0 and s^*, respectively, providing a θ^H value proportionally. The steps of our *Expected Returns Heuristic* are outlined as follows:

1. Initialize: $\theta^0 = 0.5$. Set $j = 0$.
2. Compute price p^j with θ^j using equation (6–10).
3. Determine salvage level, s^j, and expected returns, $E[R]^j$, for p^j.
4. Compute $\theta^{j+1} = 1 - \min\left\{\frac{E[R]^j}{s^j}, 1\right\}$.
5. If $\theta^j \neq \theta^{j+1}$, go to Step 2 with $j \to j+1$. Otherwise, go to Step 6.
6. Set $\theta^H = \theta^j$ and $p^H = p^j$.

In the next subsection we evaluate the performance of our heuristic for several problem instances.

6.4.1 Heuristic Performance

Our Expected Returns Heuristic is very intuitive and quite simple to compute, eliminating the computational burden of simulation methods. We run several experiments varying the parameters to study the performance of our heuristic. In order to generate different scenarios, we take a base case example (see Table 6-1), and modify one of the parameters while keeping the rest constant. For every set of parameters, we then compute the near optimal solution using Monte Carlo simulation, and the heuristic solution. Using the fact that the optimal price is between the bounds provided by Theorem 7, a near optimal solution can be obtained by evaluating the expected profit for different values of θ. In particular, we discretize θ into 101 points corresponding to values between 0 and 1 with 0.01 increments. We find the expected profit in each case by averaging the profits over a sufficiently large sample of realizations, 1,000 in our case. Since we are considering an infinite horizon problem, we allow the number of periods simulated to be sufficiently large so that the discounted profit contribution from the last period is negligible (usually 200 periods is enough). Figure 6-1 illustrates the expected profit for prices between the two bounds provided by Theorem 7 using the parameters in Table 6-1. The expected profit seems to be concave in price, so obtaining a near optimal price is relatively easy using simulation methods.

We therefore compare the near optimal price with the price provided by our heuristic. In all our experiments, the relative difference between the near optimal price and the heuristic price, computed as $(|p^* - p^H|)/(p^*)$ was less than 2%. Table 6-2 displays the results from a typical subset of our experiments. The first and second columns specify the parameter modified and its corresponding value (the rest of the parameters are given in Table 6-1).

As observed in the last column of the table, all prices provided by our heuristic are very close to the near optimal solution. We conclude that our heuristic works very well.

6.4.2 Multi-single Period versus Multiple Period

We already showed in Section 6.3 that the company can obtain additional profit by keeping some of the returns in inventory according to a salvage-down-to inventory policy to satisfy future demand. But, how much extra profit can be obtained? Next, we compare the expected profit obtained by salvaging all returns and keeping no inventory to the expected profit of the multiple period problem with an optimal salvage-down-to policy. Table 6-3 shows optimal prices, $\underline{p}^*$ and p^*, and optimal expected profits, $\underline{\Pi}^*_\infty$ and Π^*_∞, for the two cases, respectively. In the last column, we provide the percentage difference in profit between the two approaches, computed as $(\Pi^*_\infty - \underline{\Pi}^*_\infty)/(\Pi^*_\infty)$, for the same experiments shown above.

The percentage difference can be significant for some extreme cases (in excess of 40%). Obviously, for problem instances with higher probability of return, the benefits of a multiperiod approach are more relevant since the firm can gain from the resale of a returned product. Note, for example, the problem instance with $a = 1$ where there is a 44.79% difference between the two cases. The product's low attractiveness leads to a high probability of return whose negative consequences are relieved by the possibility of a resale in a multiple period problem. On the contrary, the example with $\mu_2 = 0.1$ has almost no difference between the multi-single period and multiple period problems because the low post-purchase heterogeneity μ_2 makes the probability of return negligible.

6.5 Approximate Solution

The Expected-Returns Heuristic is based on the approximation of the initial inventory by the expected returns from the previous period. Its excellent performance suggests that it can be a good approximation of the optimal expected profit for the multiple period problem formulated in §6.3. In this Section we provide an approximate analytical solution to the optimal price in the multiperiod problem, and show numerically its usual proximity with to near optimal solution obtained by Monte Carlo simulation. The approximation of the expected profit, denoted $\widetilde{\Pi}_\infty$, is obtained as follows. We take the present value of expected profit given by equation (6–6), and substitute $E[R_t]$ for y^*_t, $(E[D_1] - E[R_t])$ for $E[D_1 - y^*_t]^+$, and 0 for $E[y^*_t - D_1]^+$. In other words, we suppose that the initial inventory is always equal to the expected returns from previous period.After plugging the optimal refund $b^* = \beta v - l$ and computing the expected values, the resulting equation is given by:

$$\widetilde{\Pi}_\infty(\tilde{p}, b^*) = \frac{\lambda n q}{1-\beta}\left[\tilde{p} - c - (c - v - h)q_r\right]$$

where $\tilde{p}$ is the new price. The approximate expected profit $\widetilde{\Pi}_\infty$ can be analyzed as we did previously, leading to the following Proposition:

Proposition 6. *The price $\tilde{p}^*$ that maximizes the approximate expected profit $\widetilde{\Pi}_\infty$ is uniquely given by the following implicit expression*

$$\tilde{p}^* = \left[n \exp\left(\frac{\tilde{u}\left(\tilde{p}^*\right) - \tilde{u}_0}{\mu_1}\right) + 1\right]\mu_1 + c - (c - v - h)q_r$$

The optimization of the approximate expected profit provides a similar implicit expression to compute the price like the ones we found previously. As we did with the performance of our heuristic, we now compare in Table 6-4 the pricing results for both the approximate and exact objective functions for the same set of experiments. Last column displays the relative difference between the prices, computed as $(|p^* - \tilde{p}^*|)/(p^*)$. The differences in all cases are rather small, so the approximation of the expected profit provides near-optimal price for the multiple period problem.

6.6 Conclusion

In this Chapter, we extended our analysis in Chapter 5 to a multi-period setting, where we investigate optimal price, refund and inventory decisions for a given product assortment. We analyze the inventory problem first using an infinite horizon stationary approach. We show that a salvage-down-to level inventory policy is optimal. That is, instead of salvaging all returns (e.g., by selling them in a secondary market), as in the single period setting, the firm decides to keep some of them up to a salvage level to satisfy future demand. Second, we address the price and refund optimization.

As in the single period setting, we also find that the firm should set its refund such that returns are costless for the firm. Although price optimization is analytically intractable, we can establish upper and lower bounds on the optimal price, and we provide a heuristic and an approximate analysis that perform well. Our analysis provides an interesting result. One can intuitively expect that a customizing firm being able to extract some additional value from a returned product by selling it in the next period would refund a higher fraction of the selling price, which can be computed as refund divided by price. This is indeed true in our case. But surprisingly, the higher fraction is not due to a higher refund amount, which stays constant, but due to a lower selling price. The managerial lesson we can extract from this is the following: offering higher refunds might be an incentive for customers to return more. Instead, the company prefers to reward the customer by lowering its price, thereby encouraging her not to return the product.

Table 6-1. Base parameter values for studying the Expected Returns Heuristic

Parameter	Value
a	2.7
c	3
v	2
h	0.2
l	0.2
k	0
λ	5,000
n	10
μ_1	1
μ_2	0.7
β	0.95
$\tilde{u}_0$	2

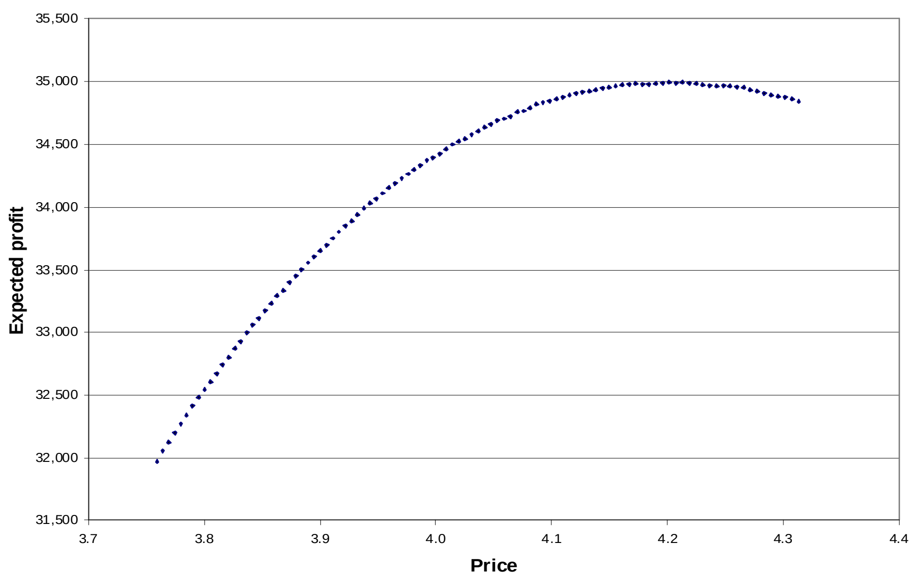

Figure 6-1. Expected profit for different prices using Monte Carlo simulation methods

Table 6-2. Heuristic performance. Base case in boldface

Parameter	Value	b^*	$\underline{p}^*$	$\overline{p}^*$	θ^*	p^*	θ^H	p^H	Difference (%)
	1.0	0.750	4.286	3.131	0.94	4.204	0.948	4.214	0.24
	1.5	1.225	4.295	3.425	0.91	4.207	0.900	4.197	0.23
v	**2.0**	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	2.3	1.985	4.327	3.979	0.76	4.241	0.742	4.234	0.15
	2.7	2.365	4.358	4.292	0.74	4.341	0.591	4.331	0.23
	1.0	1.700	4.143	3.497	0.29	3.678	0.312	3.692	0.38
	1.5	1.700	4.167	3.536	0.41	3.786	0.462	3.819	0.86
a	2.0	1.700	4.208	3.603	0.63	3.975	0.627	3.974	0.04
	2.5	1.700	4.275	3.706	0.73	4.113	0.771	4.138	0.59
	2.7	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	3.0	1.700	4.376	3.852	0.89	4.314	0.871	4.303	0.24
	2.5	1.700	3.950	3.756	0.84	3.919	0.807	3.912	0.17
	3.0	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
c	3.5	1.700	4.711	3.761	0.80	4.503	0.822	4.526	0.51
	4.0	1.700	5.137	3.762	0.82	4.862	0.828	4.874	0.24
	4.5	1.700	5.588	3.763	0.80	5.179	0.834	5.246	1.31
	0.1	1.700	4.313	3.693	0.80	4.180	0.821	4.194	0.33
	0.2	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
h	0.3	1.700	4.313	3.826	0.83	4.226	0.812	4.217	0.21
	0.4	1.700	4.313	3.894	0.80	4.226	0.809	4.230	0.09
	0.5	1.700	4.313	3.963	0.85	4.259	0.805	4.243	0.38
	0	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	0.1	1.700	4.307	3.753	0.83	4.207	0.837	4.211	0.09
k	0.2	1.700	4.303	3.747	0.87	4.226	0.855	4.217	0.21
	0.3	1.700	4.300	3.742	0.88	4.228	0.872	4.224	0.11
	0.4	1.700	4.299	3.738	0.88	4.227	0.887	4.231	0.10
	0.5	1.700	4.294	3.734	0.89	4.228	0.900	4.235	0.15
	0.1	1.800	4.318	3.766	0.81	4.207	0.794	4.198	0.22
l	**0.2**	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	0.3	1.600	4.309	3.753	0.82	4.203	0.837	4.212	0.23
	0.4	1.500	4.303	3.747	0.85	4.214	0.855	4.217	0.07
	0.5	1.400	4.300	3.742	0.87	4.222	0.872	4.224	0.03
	5	1.700	4.177	3.548	0.79	4.039	0.815	4.055	0.41
n	**10**	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	25	1.700	4.591	4.145	0.83	4.509	0.819	4.504	0.12
	50	1.700	4.883	4.511	0.87	4.831	0.822	4.812	0.39
	500	1.700	4.313	3.779	0.78	4.190	0.835	4.220	0.73
	1,000	1.700	4.313	3.770	0.78	4.187	0.827	4.214	0.64
λ	**5,000**	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	10,000	1.700	4.313	3.757	0.83	4.213	0.814	4.203	0.23
	50,000	1.700	4.313	3.753	0.81	4.200	0.810	4.200	0.00

Table 6-2. (continued from previous page)

Parameter	Value	b^*	$\underline{p}^*$	$\overline{p}^*$	θ^*	p^*	θ^H	p^H	Difference (%)
	0.10	0.000	4.282	3.728	0.92	4.235	0.978	4.269	0.80
	0.50	0.800	4.289	3.733	0.87	4.213	0.936	4.251	0.92
β	0.90	1.600	4.309	3.755	0.82	4.203	0.834	4.211	0.20
	0.95	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	0.99	1.780	4.317	3.763	0.80	4.199	0.802	4.201	0.03
	-5	1.700	8.622	8.489	0.80	8.594	0.812	8.596	0.02
	-2	1.700	6.314	6.094	0.78	6.262	0.813	6.270	0.12
u_0	0	1.700	5.078	4.733	0.85	5.022	0.814	5.008	0.27
	2	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	5	1.700	4.021	3.284	0.85	3.911	0.836	3.901	0.26
	0.1	1.700	4.279	3.708	0.99	4.273	1.000	4.279	0.14
	0.3	1.700	4.281	3.712	1.00	4.281	0.967	4.261	0.46
μ_2	0.5	1.700	4.293	3.729	0.88	4.221	0.887	4.225	0.10
	0.7	**1.700**	**4.313**	**3.759**	**0.81**	**4.201**	**0.816**	**4.205**	**0.09**
	0.9	1.700	4.339	3.797	0.74	4.190	0.764	4.204	0.33

Table 6-3. Comparison between Multi-single Period and Multiple Period Problems. Base case in boldface

Parameter	Value	$\underline{p}^*$	$\underline{\Pi}^*_\infty$	p^*	Π^*_∞	Difference (%)
	1.0	4.286	28,818	4.204	31,192	7.61
	1.5	4.295	29,675	4.207	32,951	9.94
v	**2.0**	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	2.3	4.327	32,893	4.241	36,200	9.14
	2.7	4.358	36,123	4.341	37,132	2.72
	1.0	4.143	12,872	3.678	23,316	44.79
	1.5	4.167	15,722	3.786	24,332	35.39
a	2.0	4.208	20,599	3.975	26,893	23.40
	2.5	4.275	27,526	4.113	32,043	14.10
	2.7	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	3.0	4.376	37,801	4.314	40,747	7.23
	2.5	3.950	45,030	3.919	46,791	3.76
	3.0	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
c	3.5	4.711	20,650	4.503	25,571	19.24
	4.0	5.137	13,213	4.862	18,283	27.73
	4.5	5.588	7,908	5.179	12,840	38.41
	0.1	4.313	31,113	4.180	35,475	12.29
	0.2	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
h	0.3	4.313	31,242	4.226	34,526	9.51
	0.4	4.313	31,241	4.226	34,057	8.27
	0.5	4.313	31,300	4.259	33,593	6.83

Table 6-3. (continued from previous page)

Parameter	Value	$\underline{p}^*$	$\underline{\Pi}^*_\infty$	p^*	Π^*_∞	Difference (%)
k	**0**	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	0.1	4.307	30,850	4.207	34,168	9.71
	0.2	4.303	30,521	4.226	33,404	8.63
	0.3	4.300	30,196	4.228	32,704	7.67
	0.4	4.299	29,833	4.227	32,037	6.88
	0.5	4.294	29,655	4.228	31,579	6.09
l	0.1	4.318	31,671	4.207	35,999	12.02
	0.2	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	0.3	4.309	30,763	4.203	34,184	10.01
	0.4	4.303	30,497	4.214	33,401	8.70
	0.5	4.300	30,186	4.222	32,707	7.71
n	5	4.177	17,808	4.039	20,320	12.36
	10	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	25	4.591	59,241	4.509	65,004	8.87
	50	4.883	88,555	4.831	95,686	7.45
λ	500	4.313	3,114	4.190	3,507	11.19
	1,000	4.313	6,227	4.187	7,005	11.10
	5,000	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	10,000	4.313	62,419	4.213	69,980	10.80
	50,000	4.313	311,753	4.200	350,013	10.93
β	0.10	4.282	1,567	4.235	1,579	0.79
	0.50	4.289	2,883	4.213	2,950	2.27
	0.90	4.309	15,382	4.203	16,976	9.39
	0.95	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	0.99	4.317	157,745	4.199	179,808	12.27
u_0	-5	8.622	462,180	8.594	474,378	2.57
	-2	6.314	231,307	6.262	241,735	4.31
	0	5.078	107,710	5.022	115,607	6.83
	2	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	5	4.021	2,077	3.911	2,426	14.37
μ_2	0.1	4.279	27,902	4.273	27,919	0.06
	0.3	4.281	28,133	4.281	28,723	2.05
	0.5	4.293	29,254	4.221	31,437	6.95
	0.7	**4.313**	**31,180**	**4.201**	**35,002**	**10.92**
	0.9	4.339	33,637	4.190	38,961	13.67

Table 6-4. Comparison between prices. Base case in boldface

Parameter	Value	p^*	$\tilde{p}^*$	Difference (%)
v	1.0	4.204	4.206	0.04
	1.5	4.207	4.188	0.46
	2.0	**4.201**	**4.197**	**0.10**
	2.3	4.241	4.229	0.27
	2.7	4.341	4.331	0.22
a	1.0	3.678	3.653	0.66
	1.5	3.786	3.790	0.09
	2.0	3.975	3.955	0.52
	2.5	4.113	4.127	0.33
	2.7	**4.201**	**4.197**	**0.10**
	3.0	4.314	4.298	0.36
c	2.5	3.919	3.911	0.20
	3.0	**4.201**	**4.197**	**0.10**
	3.5	4.503	4.507	0.08
	4.0	4.862	4.838	0.50
	4.5	5.179	5.186	0.15
h	0.1	4.180	4.183	0.06
	0.2	**4.201**	**4.197**	**0.10**
	0.3	4.226	4.212	0.34
	0.4	4.226	4.226	0.00
	0.5	4.259	4.240	0.43
k	**0**	**4.201**	**4.197**	**0.10**
	0.1	4.207	4.203	0.08
	0.2	4.226	4.211	0.36
	0.3	4.228	4.218	0.25
	0.4	4.227	4.226	0.02
	0.5	4.228	4.230	0.04
l	0.1	4.207	4.189	0.42
	0.2	**4.201**	**4.197**	**0.10**
	0.3	4.203	4.205	0.06
	0.4	4.214	4.211	0.08
	0.5	4.222	4.218	0.11
n	5	4.039	4.047	0.21
	10	**4.201**	**4.197**	**0.10**
	25	4.509	4.496	0.30
	50	4.831	4.803	0.57
λ	500	4.190	4.197	0.18
	1,000	4.187	4.197	0.24
	5,000	**4.201**	**4.197**	**0.10**
	10,000	4.213	4.197	0.36
	50,000	4.200	4.197	0.06

Table 6-4. (continued from previous page)

Parameter	Value	p^*	$\tilde{p}^*$	Difference (%)
β	0.10	4.235	4.269	0.80
	0.50	4.213	4.251	0.90
	0.90	4.203	4.205	0.05
	0.95	**4.201**	**4.197**	**0.10**
	0.99	4.199	4.191	0.21
u_0	-5	8.594	8.595	0.01
	-2	6.262	6.268	0.09
	0	5.022	5.005	0.34
	2	**4.201**	**4.197**	**0.10**
	5	3.911	3.870	1.07
μ_2	0.1	4.273	4.279	0.14
	0.3	4.281	4.260	0.50
	0.5	4.221	4.220	0.02
	0.7	**4.201**	**4.197**	**0.10**
	0.9	4.190	4.194	0.10

CHAPTER 7
CONCLUSION

7.1 Summary

This dissertation concentrates on assortment planning problems with the objective of developing a set of analytical decision making models for the effective management of retail operations when product returns are taken into account. Return policies are usually thought of as micro and more operational, whereas product assortment is usually thought of as strategic and more marketing related. Therefore, decisions associated with each are often made separately. In our research, we integrate return policies into a multiproduct model, where assortment, inventory, return policy and/or pricing decisions are made in an integrated manner.

In Chapter 3, motivated by whether retailers should consider product returns when they compose their product assortments, we analyze the retailer's assortment decision under two basic operational modes, make-to-order (MTO) and make-to-stock (MTS). We show that the structure of the optimal assortment strongly depends on both the return policy, which we parameterize by refund fraction (percentage of price refunded upon return), and the supply mode (MTO vs. MTS). For relatively *strict* return policies with a sufficiently low refund fraction, it is optimal for the retailer to offer most eccentric products in the MTO mode, and a mix of most popular and most eccentric products in the MTS mode. For relatively *lenient* return policies, on the other hand, conventional thinking applies: the retailer selects most popular products. Therefore, when merchandising as part of their product strategy, retailers should not only carefully consider their return policy, but also take their basic operational mode (MTO versus MTS) into account.

In Chapter 4, we study three extensions of our previous base model to incorporate: (1) endogenous price, (2) endogenous refund fraction, and (3) multiple periods. We demonstrate numerically that interesting aspects of our results, especially regarding strict return policies, prevail under all of these extensions.

In Chapters 5 and 6, we focus on the characterization of the optimal pricing, consumer return, and inventory policies of a customizing firm under a given product assortment. We analyze the problem in single- and multiple-period settings, and we conclude that retailers should aim for costless returns when designing their consumer return policies. In a single period context, we show that offering partial refunds is optimal. The firm is also able to increase its selling price while offering partial refunds. This can be seen as a service premium since returns reduce customers' risk in the purchase decision. In a multiple period setting, we prove that a salvage-down-to inventory policy is optimal, which enables the firm to increase its expected profit compared to the one obtained in multiple single periods. We also develop practically implementable heuristics to determine the optimal price and refund in the analytically intractable multiple-period setting.

7.2 Future Research

In Chapters 5 and 6 we concentrated on customized products that are made to order after demand is realized. A natural extension is to consider products that have to be ordered or produced in advance. Make-to-stock products not only complicate the inventory policy, but also the return policy analysis since returns would now depend on sales rather than on demand.

In the context of online versus brick-and-mortar retailers, there exists a myriad of potential problems related to these areas that are worth exploring. The popularity of online shopping with a proliferation of online retailers over these years has redefined retailing industry. An ever increasing number of consumers consider the online channel as their first shopping option. The coexistence of traditional retailing and *e-tailing* brings up very intriguing research questions. Among those, questions related to product assortment and return policies are very interesting. These two components together with price are the main drivers of retailer choice among consumers. Two possible research directions would be the following. First, from a single-firm perspective, what are the optimal prices, assortments and return policies for a retailer that owns both channels? How do its decisions compare against each other for online and brick-and-mortar stores? Second, we could investigate a competitive game under this 'online vs traditional' setting. What if each of the channels belongs to different firms? How does the online retailer compete with the traditional retailer in terms of price, assortment and return policy? From a consumer choice perspective, this problem is very rich. Not only should we consider three aspects like price, assortment, and return policy in the demand model (as we do in this thesis), but we should also account for the consumer's propensity to online versus traditional shopping.

Furthermore, in online retailing, the virtual nature of stores gives retailers additional possibilities as far as their inventory management is concerned. For example, drop-shipping, where a supplier stocks goods and delivers them directly to costumers, is a very common practice in the industry. Research on the inventory management of such supply chains has been conducted recently with the objective of determining the appropriate channel structure under different market circumstances. However, presence of product returns in this inventory setting has not been addressed in the literature. Consideration of product returns adds interesting reverse flows to the supply chain that managers cannot ignore. Online retailers have several options to handle product returns: keep them to satisfy further demand, send them back to the supplier (or drop-shipper), salvage them in secondary markets (e.g., outlet stores), or outsource their management completely to specialized companies. How does the retailer's inventory strategy change when product returns are taken into account? What are the best channel structures under different circumstances? Besides adding another dimension (and more generality) to supply chains with drop-shipping operations, we believe product returns could reveal interesting implications, which makes this problem very appealing to investigate.

APPENDIX: PROOFS

[*Throughout the appendix we will use the shorthand notation* $P_{r|i}$ *for* $P_{return|i}$.]

Proof of Lemma 1

We first show that $h_{MTO}(\delta)$ is quasiconcave when $\alpha \geq \frac{v-l}{p}$, and quasiconvex when $\alpha < \frac{v-l}{p}$. We use the following result from Mangasarian (1969, p. 148):

The function $h(\cdot) = \frac{g(\cdot)}{f(\cdot)}$ *is quasiconcave (quasiconvex) on a set* $\Gamma \subset \mathbb{R}^n$ *if* $g(\cdot)$ *is concave (convex) on* Γ, $f(\cdot) > 0$ *on* Γ, *and* $f(\cdot)$ *is linear on* $\mathbb{R}^n$.

We show concavity (convexity) of $g(\delta)$ by examining its second derivative:

$$g'(\delta) = \lambda(p-c) - \lambda(\alpha p + l - v)\left[P_{r|\delta} + \delta P'_{r|\delta}\right]$$

$$g''(\delta) = -\lambda(\alpha p + l - v)\left[2P'_{r|\delta} + \delta P''_{r|\delta}\right] \tag{A–1}$$

where $P'_{r|\delta} = -\frac{\mu_1}{\mu_2} P_{r|\delta}\delta^{-1}$ and $P''_{r|\delta} = \frac{\mu_1}{\mu_2}\left(\frac{\mu_1}{\mu_2}+1\right)P_{r|\delta}\delta^{-2}$. The term in brackets in (A–1) simplifies to

$$2P'_{r|\delta} + \delta P''_{r|\delta} = \frac{\mu_1}{\mu_2}\left(\frac{\mu_1}{\mu_2} - 1\right)P_{r|\delta}\delta^{-1}$$

which is positive since $\mu_1 \geq \mu_2$. We can therefore state that $g(\delta)$ is concave when $\alpha \geq \frac{v-l}{p}$, and convex when $\alpha < \frac{v-l}{p}$. Since $f(\delta)$ is strictly positive and linear in δ, it follows from Mangasarian's result that $h_{MTO}(\delta)$ is quasiconcave when $\alpha \geq \frac{v-l}{p}$, and quasiconvex when $\alpha < \frac{v-l}{p}$.

Next, we show that $h_{MTO}(\delta)$ is non-decreasing when $\alpha \geq \frac{v-l}{p}$. To do so, we evaluate the first derivative of $h_{MTO}(\delta)$:

$$h'_{MTO}(\delta) = -\lambda P_0 \sum_{j\in S}\left[p - c - (\alpha p + l - v)P_{r|j}\right]P_j + \lambda P_0 \sum_{j\in S}\left[p - c - (\alpha p + l - v)P_{r|\delta}\right]P_j$$
$$+ \lambda P_0^2\left[p - c - (\alpha p + l - v)P_{r|\delta}\right] - \lambda(\alpha p + l - v)P_\delta P'_{r|\delta}$$

It is easy to see that $p - c - (\alpha p + l - v)P_{r|\delta}$ is increasing in δ since $P_{r|\delta}$ is decreasing in δ. Therefore, there must exist some $\widehat{\delta}$ such that

$$p - c - (\alpha p + l - v)P_{r|\widehat{\delta}} \geq p - c - (\alpha p + l - v)P_{r|j}, \quad \forall j \in S \tag{A–2}$$

We now examine the sign of $h'_{MTO}(\delta)$. Since all products with $p-c-(\alpha p+l-v)P_{r|\delta} < 0$ are unprofitable and can be discarded from the analysis, the third term in $h'_{MTO}(\delta)$ is always positive; the fourth term is also positive since $P'_{r|\delta} \leq 0$. Now, for all $\delta \geq \widehat{\delta}$, inequality (A–2) implies that the first term is always less than the second term, and their sum is therefore positive. Hence, we can conclude that in the interval $[\widehat{\delta}, \infty)$, $h'_{MTO}(\delta)$ is positive, or $h_{MTO}(\delta)$ is increasing when $\alpha \geq \frac{v-l}{p}$.

On the other hand, for $\delta < \widehat{\delta}$, we can show by contradiction that the function $h_{MTO}(\delta)$ is non-decreasing. Let δ_L and δ_H be two values of δ such that $\delta_L < \widehat{\delta} < \delta_H$, and assume that $h_{MTO}(\delta)$ is decreasing for $\delta < \widehat{\delta}$. Then,

$$h_{MTO}(\delta_L) > h_{MTO}(\widehat{\delta}) \tag{A–3}$$

$$h_{MTO}(\widehat{\delta}) < h_{MTO}(\delta_H) \tag{A–4}$$

Furthermore, we have previously shown that $h_{MTO}(\delta)$ is quasiconcave when $\alpha \geq \frac{v-l}{p}$. By the definition of quasiconcavity

$$h_{MTO}(\lambda\delta_L + (1-\lambda)\delta_H) \geq \min\{h_{MTO}(\delta_L), h_{MTO}(\delta_H)\}, \qquad \lambda \in [0,1]$$

Since $\delta_L < \widehat{\delta} < \delta_H$, there exists $\widehat{\lambda} \in [0,1]$ such that $\widehat{\lambda}\delta_L + (1-\widehat{\lambda})\delta_H = \widehat{\delta}$. Using the above inequality for $\lambda = \widehat{\lambda}$, we obtain

$$h_{MTO}(\widehat{\delta}) \geq \min\{h_{MTO}(\delta_L), h_{MTO}(\delta_H)\}$$

which contradicts either inequality (A–3) or (A–4). Hence, when $\alpha \geq \frac{v-l}{p}$, we conclude that $h_{MTO}(\delta)$ is non-decreasing for $\delta < \widehat{\delta}$ as well.

Proof of Lemma 2

Product i should be added to the current assortment *iff* $\Pi_{MTO}(S \cup \{i\}) \geq \Pi_{MTO}(S)$, or

$$\left(P_0^S - P_0^{S\cup\{i\}}\right) M_i \geq \sum_{j\in S} \left(P_j^S - P_j^{S\cup\{i\}}\right)(M_j - M_i)$$

This inequality states that the profit gain made by the additional market share captured by adding product i should be larger than the potential profit loss due to cannibalization. Using the MNL purchase probabilities, we rewrite it as follows:

$$\left[\frac{\omega_0}{\omega_0 + \sum_{k \in S} \omega_k} - \frac{\omega_0}{\omega_0 + \sum_{k \in S \cup \{i\}} \omega_k}\right] M_i \geq$$
$$\sum_{j \in S} \left[\frac{\omega_j}{\omega_0 + \sum_{k \in S} \omega_k} - \frac{\omega_j}{\omega_0 + \sum_{k \in S \cup \{i\}} \omega_k}\right] (M_j - M_i)$$

By substituting $\omega_0 = 1$ and simplifying as follows, we obtain the result.

$$\frac{\left[1 + \sum_{k \in S \cup \{i\}} \omega_k\right] - \left[1 + \sum_{k \in S} \omega_k\right]}{\left[1 + \sum_{k \in S} \omega_k\right]\left[1 + \sum_{k \in S \cup \{i\}} \omega_k\right]} M_i \geq$$
$$\sum_{j \in S} \left[\frac{\omega_j \left[1 + \sum_{k \in S \cup \{i\}} \omega_k\right] - \omega_j \left[1 + \sum_{k \in S} \omega_k\right]}{\left[1 + \sum_{k \in S} \omega_k\right]\left[1 + \sum_{k \in S \cup \{i\}} \omega_k\right]}\right] (M_j - M_i)$$

$$\left[1 + \sum_{j \in S} \omega_j\right] M_i \geq \sum_{j \in S} \omega_j M_j \tag{A–5}$$

$$M_i \geq \sum_{j \in S} P_j^S M_j = \frac{\Pi_{MTO}(S)}{\lambda}$$

Dividing both sides of (A–5) by $\left[1 + \sum_{k \in S \cup \{i\}} \omega_k\right]$, we also obtain:

$$\left[1 - P_i^{S \cup \{i\}}\right] M_i \geq \sum_{j \in S} P_j^{S \cup \{i\}} M_j$$
$$M_i \geq \sum_{j \in S \cup \{i\}} P_j^{S \cup \{i\}} M_j = \frac{\Pi_{MTO}(S \cup \{i\})}{\lambda}$$

Proof of Theorem 1

Proof of Part (a). The proof is by construction. Assume that the optimal assortment has cardinality k, $k \in \{0, 1, ..., n\}$. The theorem holds trivially for $k = 0$ and $k = n$. Take any subset S of N with cardinality $k \in \{1, ..., n-1\}$. Let $n_a = \max\{i \mid \mathcal{A}_i \subseteq S \ , \ i \in \{0, 1, ..., k\}\}$ be the number of most popular products of N that belong to S. Clearly n_a cannot be strictly larger than k. If $n_a < k$, then there must exist some product $i \in S$ such that $i > n_a + 1$. Since $h_{MTO}(\delta)$ is non-decreasing due to Lemma 1, product i can be replaced with product $n_a + 1$ without decreasing the profit. Proceeding recursively with such profit-improving replacements, $n_a = k$ will in the end be satisfied, which implies that $S^* = \mathcal{A}_k$.

Proof of Part (b). This requires a three-step proof. We first show in Step 1 that the structure of the optimal assortment can only conform to one of three possible configurations. In Steps 2 and 3, we then discard two of these configurations to establish the structure of the optimal assortment. Suppose that the optimal assortment has cardinality $k \in \{0, 1, ..., n\}$. The theorem holds trivially for $k = 0$ and $k = n$. Hence, we only consider cases with $k \in \{1, ..., n-1\}$.

Step 1. We first prove that the optimal assortment is composed of some number of most popular and some number of most eccentric products from N, i.e., $S^* = \mathcal{A}_{k-j} \cup \mathcal{Z}_j$ for some $j \in \{0, 1, ..., k\}$ and $k \in \{1, ..., n-1\}$.

The proof is by construction. The claim holds trivially for $k = n - 1$. Take any subset S of N with cardinality $k \in \{1, ..., n-2\}$. Let $n_a = \max\{i \mid \mathcal{A}_i \subseteq S\ ,\ \ i \in \{0, 1, ..., k\}\}$ be the number of most popular products of N that belong to S; and $n_z = \max\{j \mid \mathcal{Z}_j \subseteq S\ ,\ \ j \in \{0, 1, ..., k\}\}$ be the number of most eccentric products of N that belong to S. Clearly $n_a + n_z$ cannot be strictly larger than k. If $n_a + n_z < k$, then there must exist some product $i \in S$ such that $n_a + 1 < i < n - n_z$. Since $h_{MTS}(\delta)$ is quasiconvex due to Lemma 1, product i can be replaced with product $n_a + 1$ or with product $n - n_z$ without decreasing the profit. Proceeding recursively with such profit-improving replacements, $n_a + n_z = k$ will in the end be satisfied, which implies that $S^* = \mathcal{A}_{k-j} \cup \mathcal{Z}_j$ for some $j \in \{0, 1, ..., k\}$. Therefore, the optimal assortment can have one of three structures: (1) it can be composed of some number of most popular products ($0 = j \leq k$), or (2) some number of most eccentric products ($0 < j = k$), or (3) it can contain some most popular and some most eccentric products ($0 < j < k$). Step 2 eliminates the first possibility, and Step 3 eliminates the third possibility.

Step 2. We now claim that the optimal assortment cannot be composed of some number of most popular products.

Assume that the optimal assortment is $S^* = \mathcal{A}_k$ for some $k \in \{1, \ldots, n-1\}$, and set $\lambda = 1$ without loss of generality. We will show that adding the next most popular product is always profit-improving, i.e., $\Pi_{MTO}(\mathcal{A}_{k+1}) > \Pi_{MTO}(\mathcal{A}_k)$, which is a contradiction. Using equation (3–1) and the profit margin notation, the new profit after adding product $k+1$ is as follows:

$$\begin{aligned}\Pi_{MTO}(\mathcal{A}_{k+1}) &= \sum_{j \in \mathcal{A}_{k+1}} P_j^{\mathcal{A}_{k+1}} M_j = \sum_{j \in \mathcal{A}_k} P_j^{\mathcal{A}_{k+1}} M_j + P_{k+1}^{\mathcal{A}_{k+1}} M_{k+1} \\ &= \sum_{j \in \mathcal{A}_k} P_j^{\mathcal{A}_{k+1}} M_j + \left[\sum_{j \in \mathcal{A}_k \cup \{0\}} P_j^{\mathcal{A}_k} - P_j^{\mathcal{A}_{k+1}} \right] M_{k+1}\end{aligned}$$

where the last equality follows from the identity $P_{k+1}^{\mathcal{A}_{k+1}} = \sum_{j \in \mathcal{A}_k \cup \{0\}} P_j^{\mathcal{A}_k} - P_j^{\mathcal{A}_{k+1}}$. Since $\alpha < (v-l)/p$ implies $M_j \leq M_{k+1}$ for all $j \in \mathcal{A}_k$, we obtain:

$$\begin{aligned}\Pi_{MTO}(\mathcal{A}_{k+1}) &\geq \sum_{j \in \mathcal{A}_k} P_j^{\mathcal{A}_{k+1}} M_j + \sum_{j \in \mathcal{A}_k} \left[P_j^{\mathcal{A}_k} - P_j^{\mathcal{A}_{k+1}} \right] M_j + \left[P_0^{\mathcal{A}_k} - P_0^{\mathcal{A}_{k+1}} \right] M_{k+1} \\ &= \Pi_{MTO}(\mathcal{A}_k) + \left[P_0^{\mathcal{A}_k} - P_0^{\mathcal{A}_{k+1}} \right] M_{k+1}\end{aligned}$$

As the assortment grows, the probability of consumers choosing the outside option drops, i.e., $P_0^{\mathcal{A}_{k+1}} < P_0^{\mathcal{A}_k}$. Furthermore, M_{k+1} must be positive; if it was negative, $M_1, \ldots, M_k$ would all have to be negative (because $\alpha < (v-l)/p$ implies $M_j \leq M_{j+1}$), which contradicts our premise that $S^* = \mathcal{A}_k$. Using these two facts, we conclude that $\Pi_{MTO}(\mathcal{A}_{k+1}) > \Pi_{MTO}(\mathcal{A}_k)$.

Step 3. We further claim that the optimal assortment cannot be composed of some most popular products and some most eccentric products.

The proof is due to Lemma 2. Assume that the optimal assortment is $S^* = \mathcal{A}_i \cup \mathcal{Z}_j$ with $i > 0$, $j > 0$, and $i + j < n$. Pick a product $k \notin S^*$ $(i < k < j)$, for which it must be true by Lemma 2a that:

$$M_k < \sum_{j \in S^*} P_j^{S^*} M_j$$

Otherwise, S^* would not be optimal. Because $\alpha < (v-l)/p$ implies $M_l \leq M_k$ for all $l \in \{1, \ldots, k\}$, every product in $\mathcal{A}_i$ must satisfy the following:

$$M_l \leq M_k < \sum_{j \in S} P_j^{S^*} M_j, \qquad \forall l \in \{1, \ldots, i\}$$

Therefore, due to Lemma 2, the firm would improve its profit by removing one of the products in $\mathcal{A}_i$, say product i, from S^*. This constitutes a contradiction; unless all products in $\mathcal{A}_i$ are removed, S^* cannot be optimal.

Proof of Lemma 3

The proof is analogous to the first part of the proof of Lemma 1. We use the same result from Mangasarian (1969) to show the quasiconvexity of $h_{MTS}(\delta)$. We already know from the proof of Lemma 1 that $g(\delta)$ is convex when $\alpha < \frac{v-l}{p}$. We need to show that $\tilde{g}(\delta)$ is also convex when $\alpha < \frac{v-l}{p}$. Taking the first and second derivatives, we verify below that this is indeed the case for any α.

$$\tilde{g}'(\delta) = -(e-v)\sigma\lambda^\beta\phi(z^*)\left[\beta\delta^{\beta-1}f(\delta)^{1-\beta} + \left(\sum_{j \in S}\omega_j^\beta + \delta^\beta\right)(1-\beta)f(\delta)^{-\beta}\right]$$

$$\tilde{g}''(\delta) = (e-v)\frac{\beta(1-\beta)\sigma\lambda^\beta\phi(z^*)}{f(\delta)^\beta}\left[\frac{\delta^{\beta-2}\left(\sum_{j \in S}\omega_j + \omega_0\right)^2 + \sum_{j \in S}\omega_j^\beta}{f(\delta)}\right] > 0$$

Since both $g(\delta)$ and $\tilde{g}(\delta)$ are convex when $\alpha < \frac{v-l}{p}$, their sum is also convex; and, due to Mangasarian's result, $h_{MTS}(\delta)$ is quasiconvex when $\alpha < \frac{v-l}{p}$.

Proof of Theorem 2

Since the proof is analogous to the Step 1 of the proof of Theorem 1b, we only give a sketch of the argument. Due to the quasiconvexity of $h_{MTS}(\delta)$ when $\alpha < \frac{v-l}{p}$, a fact provided by Lemma 3, any assortment that does not conform to the structure stated in the theorem can be improved, and therefore cannot be optimal. The result is tight, because one can construct examples where the optimal assortment has: (1) some number of most popular products only; (2) some number of most eccentric products only; or (3) some most popular and some most eccentric products. See Table 1 for such examples. Note also that Steps 2 and 3 of the proof of Theorem 1b are no longer valid for the MTS case. The main reason is that the expected profit margin per unit sales (defined below) now depends on the assortment S (whereas, in the MTO case, M_j is independent of S):

$$\widetilde{M_j}(S) = p - c - (\alpha p + l - v)P_{r|j} - (e - v)\sigma\phi(z^*)(\lambda P_j^S)^{\beta-1}$$

Proof of Proposition 1

We know from Theorem 1a that, for lenient return policies with $\alpha \geq \frac{v-l}{p}$, the optimal assortment is composed of some number of most popular products, i.e., $S^* = \mathcal{A}_k$ for some $k \in \{0, 1, \ldots, n\}$. Furthermore, according to Lemma 2, the firm is better off by adding product $(i+1)$ to an existing assortment $S \subset N$ that includes the i most popular products, i.e. $S = \mathcal{A}_i$, if and only if the following inequality holds:

$$\left[1 + \sum_{j \in S} \omega_j\right] M_{i+1} \geq \sum_{j \in S} \omega_j M_j \tag{A–6}$$

It is easy to verify that, as α increases, the expected profit margin M_j decreases and preference ω_j increases for all j. Therefore, if M_{i+1} decreases at least as fast as M_j for all $j \in S$, the left-hand-side will be relatively smaller compared to the right-hand-side, as we increase α. This is equivalent to say that in the region where $\alpha \geq \frac{v-l}{p}$, a more lenient return policy (higher α) leads to lower variety since the incentive to add a product to any existing assortment will be lower, i.e., inequality (A–6) will be less likely to be satisfied. For this to be true, it is sufficient that $\left|\frac{\partial M_{i+1}}{\partial \alpha}\right| \geq \left|\frac{\partial M_j}{\partial \alpha}\right|$ for all $j \in S$. The derivative of M_j with respect to α is given by

$$\frac{\partial M_j}{\partial \alpha} = -\frac{\partial P_{r|\,j}}{\partial \alpha}(\alpha p + l - v) - pP_{r|\,j} = \; -pP_{r|\,j}\left[1 + \frac{1}{\mu_2}(\alpha p + l - v)(1 - P_{r|\,j})\right]$$

where the last equality is obtained by substituting $\frac{\partial P_{r|\,j}}{\partial \alpha} = \frac{p}{\mu_2} P_{r|\,j}(1 - P_{r|\,j})$.

Thus, the inequality $\left|\frac{\partial M_{i+1}}{\partial \alpha}\right| \geq \left|\frac{\partial M_j}{\partial \alpha}\right|$ is equivalent to

$$P_{r|\,i+1} + P_{r|\,i+1}\frac{1}{\mu_2}(\alpha p + l - v)(1 - P_{r|\,i+1}) \geq P_{r|\,j} + P_{r|\,j}\frac{1}{\mu_2}(\alpha p + l - v)(1 - P_{r|\,j}).$$

Since $P_{r|\,i+1} \geq P_{r|\,j}$ for all $j \in \mathcal{A}_i$, the condition above is always satisfied if $P_{r|\,i+1} \leq 0.5$, or equivalently $a_{i+1} \geq \alpha p$. Hence, the condition that ensures all products to have a positive expected utility, i.e., $a_j \geq p$ for all $j \in N$, is sufficient for the above inequality to hold.

Proof of Theorem 3

We first find the expression for the optimal refund and then, for the optimal price. The first derivatives of $\Pi\,(p, b)$ with respect to b and p are, respectively,

$$\begin{aligned} \frac{\partial \Pi\,(p,b)}{\partial b} &= \lambda n q q_r \left[\frac{q_0 M}{\mu_1} - 1 - \frac{(1-q_r)}{\mu_2}(b + l - \beta v)\right] \\ \frac{\partial \Pi\,(p,b)}{\partial p} &= -\lambda n q \left[\frac{q_0 M}{\mu_1} - 1\right] \end{aligned}$$

The first order conditions must be satisfied at optimality, that is, the terms in brackets must be zero. Note that there exist extreme cases (when $q = 0$ or $q_r = 0$) where the first derivatives are also 0. Such cases occur when setting p too high or b too low (negative). They are discarded from the analysis because they are not of interest. Therefore we have

$$\frac{q_0 M}{\mu_1} - 1 - \frac{(1-q_r)}{\mu_2}(b + l - \beta v) = 0 \tag{A–7}$$

$$\frac{q_0 M}{\mu_1} - 1 = 0 \tag{A–8}$$

To satisfy both conditions, clearly, $\frac{(1-q_r)}{\mu_2}(b^* + l - \beta v) = 0$. Since $q_r < 1$, we must have $b^* = \beta v - l$. We now show that b^* is indeed a maximum. The second derivative of Π with respect to b is

$$\begin{aligned} \frac{\partial^2 \Pi\,(p,b)}{\partial b^2} = \lambda n q q_r &\left[\frac{q_0 q_r}{\mu_1} + \frac{(1-q_r)}{\mu_2}\right]\left[\frac{q_0 M}{\mu_1} - 1 - \frac{(1-q_r)}{\mu_2}(b + l - \beta v)\right] \\ &- \frac{\lambda n q q_0 q_r^2}{\mu_1}\left[\frac{nqM}{\mu_1} + 1 + \frac{(1-q_r)}{\mu_2}(b + l - \beta v)\right] \\ &+ \frac{\lambda n q q_r (1-q_r)}{\mu_2}\left[\frac{q_r}{\mu_2}(b + l - \beta v) - 1\right] \end{aligned}$$

Using the optimal b^* and the first order condition in (A–7), the second derivative at b^* is negative

$$\frac{\partial^2 \Pi\left(p, b^*\right)}{\partial b^2} = -\frac{\lambda n q q_0 q_r^2}{\mu_1}\left[\frac{nqM}{\mu_1}+1\right] - \frac{\lambda n q q_r(1-q_r)}{\mu_2} < 0$$

Therefore, b^* is a maximum. Now, we proceed similarly to find the optimal price. The expected profit margin plugging the optimal refund is given by $M = p - c$. Since (A–8) must be satisfied at optimality, $q_0(p^* - c) = \mu_1$. Substituting the probability of not buying, we obtain the implicit expression for the optimal price, $p^* = \left[n \exp\left(\frac{\tilde{u}^*(p^*)-\tilde{u}_0}{\mu_1}\right)+1\right]\mu_1 + c$, where $\tilde{u}^*\left(p^*\right) = \mu_2 \ln\left[\exp\left(\frac{a}{\mu_2}\right)+\exp\left(\frac{\beta v - l - k}{\mu_2}\right)\right] - p^*$. Note that since $\tilde{u}^*$ decreases in p^*, the solution to the equality is unique. Finally, we check the second derivative of Π with respect to p to show that p^* is a maximum:

$$\frac{\partial^2 \Pi\left(p, b\right)}{\partial p^2} = \frac{\lambda n q q_0}{\mu_1}\left[\frac{q_0 M}{\mu_1}-1\right] - \frac{\lambda n q q_0}{\mu_1}\left[\frac{nqM}{\mu_1}+1\right]$$

which is negative since $\frac{q_0 M}{\mu_1} = 1$.

Proof of Theorem 4

Before the proof, we need to comment on a technicality in our consumer choice model. Since Gumbel random variables can take any real value, the no returns case is achieved, by a slight abuse of the model definition, by setting $b = -\infty$. In such a case, the pre-purchase expected utility is $\tilde{u}_{NR} = a - p_{NR}$, and the customer keeps the product with probability 1, which is equivalent to not allowing returns at all. The first and second derivatives of Π_{NR} with respect to price are respectively:

$$\begin{aligned}
\frac{\partial \Pi_{NR}\left(p_{NR}\right)}{\partial p_{NR}} &= \lambda n q \left[1 - \frac{q_0}{\mu_1}(p_{NR} - c)\right] \\
\frac{\partial^2 \Pi_{NR}\left(p_{NR}\right)}{\partial p_{NR}^2} &= -\frac{\lambda n q q_0}{\mu_1}\left[1 - \frac{q_0}{\mu_1}(p_{NR}-c) + \frac{nq}{\mu_1}(p_{NR}-c)+1\right]
\end{aligned}$$

Setting the first derivative to 0 provides the expression for the optimal price, that is, $1 = \frac{q_0}{\mu_1}(p_{NR}^* - c)$. If we substitute the value of q_0 and rearrange terms, we obtain

$p^*_{NR} = \left[n \exp\left(\frac{\tilde{u}_{NR}(p^*_{NR}) - \tilde{u}_0}{\mu_1} \right) + 1 \right] \mu_1 + c$. Plugging this value into the second derivative yields

$$\frac{\partial^2 \Pi_{NR}\left(p^*_{NR}\right)}{\partial p^2_{NR}} = -\frac{\lambda n q q_0}{\mu_1} \left[\frac{nq}{\mu_1 q_0} + 1 \right] < 0$$

Then, p^*_{NR} is a maximum. Note that $\tilde{u}_{NR}$ decreases in p_{NR}, therefore p^*_{NR} is unique.

Proof of Proposition 2

The proof is straightforward and follows directly from the comparison between pre-purchase expected utilities in both cases. It is easy to see that for a given price p and refund b,

$$\tilde{u}\left(p, b\right) = \mu_2 \ln \left[\exp\left(\frac{a}{\mu_2}\right) + \exp\left(\frac{b-k}{\mu_2}\right) \right] - p > \mu_2 \ln \left[\exp\left(\frac{a}{\mu_2}\right) \right] - p = a - p = \tilde{u}_{NR}\left(p\right)$$

since $\exp\left(\frac{b-k}{\mu_2}\right) > 0$. The result follows from the optimal price expressions in Theorems 3 and 4.

Proof of Proposition 3

At optimality, we have $\frac{q_0 M}{\mu_1} = 1$ for both the partial refund case and the no returns case. Let's use the superscript PR and NR, respectively, to distinguish between probabilities in both cases. We must have $q_0^{PR}(p^* - c) = q_0^{NR}(p^*_{NR} - c) = \mu_1$. By Proposition 2, we know that $(p^* - c) > (p^*_{NR} - c)$. Then $q_0^{PR} < q_0^{NR}$, which is equivalent to $nq^{PR} > nq^{NR}$. Since expected profit margin M and probability of purchase nq are higher for the partial refund case, the result follows directly.

Proof of Lemma 4

We examine the first derivative of the expected profit:

$$\frac{\partial \Pi_{FR}\left(p_{FR}\right)}{\partial p_{FR}} = \lambda n q (1 - q_r) \left[1 - \frac{q_r}{\mu_2}(p_{FR} + l - \beta v) - \frac{q_0 M_{FR}}{\mu_1} \right] \tag{A–9}$$

The term outside the brackets, $\lambda n q(1 - q_r)$, is always nonnegative and can be left aside for the analysis. Denote LT and RT the left and the right terms, respectively, in brackets in

(A–9):

$$LT = 1 - \frac{q_r}{\mu_2}(p_{FR} + l - \beta v)$$
$$RT = -\frac{q_0 M_{FR}}{\mu_1}$$

By restricting to only those values of p_{FR} such that the profit is positive, we have $M_{FR} > 0$, and RT always negative in the region of interest. Now observe that LT is positive when $(p_{FR} + l - \beta v)q_r < \mu_2$, negative when $(p_{FR} + l - \beta v)q_r > \mu_2$, and decreasing

$$\frac{\partial LT}{\partial p_{FR}} = -\frac{q_r(1-q_r)}{\mu_2}(p_{FR} + l - \beta v) - \frac{q_r}{\mu_2} = -\frac{q_r}{\mu_2}\left[1 + \frac{(1-q_r)}{\mu_2}(p_{FR} + l - \beta v)\right]$$

Therefore, for values of p_{FR} such that $(p_{FR} + l - \beta v)q_r > \mu_2$, the expected profit Π_{FR} in the region where $M_{FR} > 0$, is always decreasing, i.e., $LT + RT < 0$. Hence, the upper bound on price is given by $(\bar{p}_{FR} + l - \beta v)q_r = \mu_2$.

We next show that $\bar{p}_{FR}$ maximizes M_{FR} and is unique. The derivative of M_{FR} is

$$\frac{\partial M_{FR}}{\partial p_{FR}} = (1-q_r)\left[1 - \frac{q_r}{\mu_2}(p_{FR} + l - \beta v)\right]$$

which is 0 at $\bar{p}_{FR}$. The second derivative at $\bar{p}_{FR}$ is negative:

$$\frac{\partial^2 M_{FR}}{\partial p_{FR}^2} = -\frac{q_r(1-q_r)}{\mu_2}\left[2 + \frac{1-2q_r}{\mu_2}(\bar{p}_{FR} + l - \beta v)\right] = -\frac{(1-q_r)}{\mu_2} < 0$$

To complete the proof, uniqueness can easily be checked since $(p_{FR} + l - \beta v)q_r$ strictly increases in p_{FR} for positive values of $(p_{FR} + l - \beta v)$.

Proof of Theorem 5

The optimal price must satisfy the first order condition. From Lemma 4, setting the first derivative of the expected profit (A–9) to 0, we obtain:

$$1 - \frac{q_r}{\mu_2}(p^*_{FR} + l - \beta v) = \frac{q_0 M_{FR}}{\mu_1} \tag{A–10}$$

We denote LT and RT the left and the right side terms in (A–10), respectively. The second derivative of Π_{FR} evaluated at the optimal price p^*_{FR} is

$$\frac{\partial^2 \Pi_{FR}(p^*_{FR})}{\partial p^2_{FR}} = -\lambda nq(1-q_r)\left\{\frac{q_r}{\mu_2^2}(p^*_{FR}+l-\beta v) + \frac{q_0 M_{FR}}{\mu_1}\left[\frac{q_r}{\mu_2} + \frac{(1-q_r)}{\mu_1}\right]\right\}$$

which is negative. Therefore, p^*_{FR} is indeed a maximum. To prove uniqueness, we show that LT decreases while RT increases in p_{FR} for the region of interest. Their respective derivatives are:

$$\begin{aligned}
\frac{\partial LT}{\partial p_{FR}} &= -\frac{q_r}{\mu_2}\left[1 + \frac{(1-q_r)}{\mu_2}(p_{FR}+l-\beta v)\right] < 0 \\
\frac{\partial RT}{\partial p_{FR}} &= \frac{q_0(1-q_r)}{\mu_1}\left[1 - \frac{q_r}{\mu_2}(p_{FR}+l-\beta v)\right] + \frac{nqq_0(1-q_r)M_{FR}}{\mu_1^2} > 0
\end{aligned}$$

where the last inequality holds because $(p_{FR}+l-\beta v)q_r < \mu_2$ (by Lemma 4). This completes the proof.

Proof of Proposition 4

The proof follows the steps of the optimal ordering policy in Heyman and Sobel (1984). The present value of the expected profit provided in (6–2) can be written as

$$\begin{aligned}
\Pi_\infty = {} & n\left[v(x_1-y_1) - hy_1 + pD_1 - c[D_1-y_1]^+ - (b+l)R_1\right] \\
& + n\sum_{t=2}^{\infty}\beta^{t-1}\left\{v(x_t-y_t) - hy_t + pD_t - c[D_t-y_t]^+ - (b+l)R_t\right\}
\end{aligned}$$

Substituting $([y_{t-1}-D_{t-1}]^+ + R_{t-1})$ for x_t if $t>1$, leads to

$$\begin{aligned}
\Pi_\infty = {} & n\left[v(x_1-y_1) - hy_1 + pD_1 - c[D_1-y_1]^+ - (b+l)R_1\right] \\
& + n\sum_{t=2}^{\infty}\beta^{t-1}\left\{-vy_t - hy_t + pD_t - c[D_t-y_t]^+ - (b+l)R_t\right\} \\
& + n\sum_{t=2}^{\infty}\beta^{t-1}\left\{v[y_{t-1}-D_{t-1}]^+ + vR_{t-1}\right\}
\end{aligned} \qquad \text{(A–11)}$$

where the last term is

$$n\sum_{t=2}^{\infty}\beta^{t-1}\left\{v[y_{t-1}-D_{t-1}]^{+}+vR_{t-1}\right\}=n\sum_{t=1}^{\infty}\beta^{t-1}\left\{\beta v[y_t-D_t]^{+}+\beta vR_t\right\}$$

Using this last expression and rearranging terms in (A–11) we obtain

$$\begin{aligned}\Pi_\infty &= nvx_1+n\sum_{t=1}^{\infty}\beta^{t-1}\left\{-(v+h)y_t+pD_t-c[D_t-y_t]^{+}-(b+l-\beta v)R_t+\beta v[y_t-D_t]^{+}\right\}\\ &= n\sum_{t=1}^{\infty}\beta^{t-1}\left\{pD_t-(b+l-\beta v)R_t\right\}-n\sum_{t=1}^{\infty}\beta^{t-1}\left\{(v+h)y_t+c[D_t-y_t]^{+}-\beta v[y_t-D_t]^{+}\right\}\end{aligned}$$

where $x_1 = 0$ by assumption. Let $\Omega(y_t, D_t) = (v+h)y_t + c[D_t - y_t]^+ - \beta v[y_t - D_t]^+$, and define $G(y) = (v+h)y + cE[D_1 - y]^+ - \beta vE[y - D_1]^+$. Since $D_1, D_2, \ldots, D_{t-1}, D_t$ are independent and identically distributed, and y_t and D_t are stochastically independent, it can be shown that $E\left[\Omega(y_t, D_t)\right] = E\left[G(y_t)\right]$. The inventory problem in (6–3) can then be reformulated as

$$\begin{aligned}&\text{maximize} \qquad n\sum_{t=1}^{\infty}\beta^{t-1}E\left[pD_t-(b+l-\beta v)R_t\right]-n\sum_{t=1}^{\infty}\beta^{t-1}E[G(y_t)]\\ &\text{subject to} \qquad 0\le y_t\le x_t \qquad t=1,2,\ldots\end{aligned}$$

Proof of Lemma 5

To show convexity, we just need to check that the second forward difference is non-negative. To find the minimizer, we check the discrete first order condition. Recall the expression for $G(y)$:

$$\begin{aligned}G(y) =&(v+h)y+cE[D_1-y]^{+}-\beta vE[y-D_1]^{+}\\ =&(v+h)y+c\sum_{d_1=y}^{\infty}(d_1-y)P(D_1=d_1)-\beta v\sum_{d_1=0}^{y-1}(y-d_1)P(D_1=d_1)\end{aligned}$$

Let $\Delta G(y) = G(y+1) - G(y)$ and $\Delta^2 G(y) = \Delta G(y+1) - \Delta G(y)$, be the first and the second forward differences, respectively. Then

$$\begin{aligned}\Delta G(y) = (v+h)(y+1) + c\sum_{d_1=y+1}^{\infty}(d_1-y-1)P(D_1=d_1) - \beta v\sum_{d_1=0}^{y}(y+1-d_1)P(D_1=d_1)\\ - (v+h)y - c\sum_{d_1=y}^{\infty}(d_1-y)P(D_1=d_1) + \beta v\sum_{d_1=0}^{y-1}(y-d_1)P(D_1=d_1)\\ = (v+h) - c\,(1-F(y)) - \beta vF(y)\end{aligned}$$

Similarly, $\Delta G(y+1) = (v+h) - c\,(1-F(y+1)) - \beta vF(y+1)$. The second difference is given by $\Delta^2 G(y) = (c-\beta v)f(d) \geq 0$.

The minimizer s^* is going to be the smallest integer such that the first forward difference is non-negative, that is, $F(s^*) \geq \frac{c-v-h}{c-\beta v}$.

Proof of Theorem 6

The proof follows directly from Lemma 5 and the constraint $y_t \leq x_t$.

Proof of Lemma 6

The first part of the Lemma is straightforward to prove: for the lower bound we substitute $y_t^* = 0$ in (6–6), and for the upper bound we substitute $y_t^* = s^*$ in (6–6). Since s^* minimizes $G(y_t)$ (Lemma 5), the result follows. The second part of the Lemma is also trivial.

Proof of Proposition 5

The proof is straightforward and follows from the subtraction $\overline{\Pi}_\infty - \underline{\Pi}_\infty$.

Proof of Lemma 7

Proof of Part (a). The proof is exactly as Theorem 3's proof.

Proof of Part (b). The proof is analogous to Theorem 3's proof, and we just sketch it for completeness. Let $\overline{M} = p - (v+h) - (b+l-\beta v)q_r - \frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}$. The first derivatives of $\overline{\Pi}_\infty$ with respect to b and p are, respectively,

$$\frac{\partial\overline{\Pi}_\infty(p,b)}{\partial b} = \frac{\lambda nqq_r}{1-\beta}\left[\frac{q_0\overline{M}}{\mu_1} - 1 - \frac{(1-q_r)}{\mu_2}(b+l-\beta v) + \frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$
$$\frac{\partial\overline{\Pi}_\infty(p,b)}{\partial p} = -\frac{\lambda nq}{1-\beta}\left[\frac{q_0\overline{M}}{\mu_1} - 1 + \frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$

To satisfy both first order conditions, $\frac{(1-q_r)}{\mu_2}(\overline{b}^* + l - \beta v) = 0$, then $\overline{b}^* = \beta v - l$. The second derivative of $\overline{\Pi}_\infty$ with respect to b at optimality is

$$\begin{aligned}\frac{\partial^2\overline{\Pi}_\infty\left(p,\overline{b}^*\right)}{\partial b^2} &= -\frac{\lambda nqq_0q_r^2}{(1-\beta)\mu_1}\left[\frac{nq\overline{M}}{\mu_1} + 1 - \frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right] \\ &\quad - \frac{\lambda nqq_r(1-q_r)}{(1-\beta)\mu_2} - \frac{\lambda nqq_0q_r^2(c-\beta v)\phi(z^*)}{2(1-\beta)\mu_1^2\sqrt{\lambda q}}\left[nq + \frac{q_0}{2}\right]\end{aligned}$$

Using the optimality condition $\frac{q_0\overline{M}}{\mu_1} - 1 = -\frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}$, clearly the second derivative is negative, and $\overline{b}^*$ is a maximum.

Similarly for the optimal price, the first order condition determines that $q_0\left[\overline{M} + \frac{(c-\beta v)\phi(z^*)}{2\sqrt{\lambda q}}\right] = \mu_1$. At optimality, $\overline{M} = \overline{p}^* - (v+h) - \frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}$. It follows that $\overline{p}^* = \left[n\exp\left(\frac{\tilde{u}(\overline{p}^*)-\tilde{u}_0}{\mu_1}\right)+1\right]\mu_1 + v + h + \frac{(c-\beta v)\phi(z^*)}{2\sqrt{\lambda}}\sqrt{n+\exp\left(\frac{\tilde{u}_0-\tilde{u}(\overline{p}^*)}{\mu_1}\right)}$. Finally, we check the second derivative of $\overline{\Pi}_\infty$ at optimality:

$$\begin{aligned}\frac{\partial^2\overline{\Pi}_\infty(\overline{p}^*,b)}{\partial p^2} &= -\frac{\lambda nqq_0}{(1-\beta)\mu_1}\left[\frac{nq\overline{M}}{\mu_1} + 1 - \frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right] \\ &\quad - \frac{\lambda nqq_0(c-\beta v)\phi(z^*)}{2(1-\beta)\mu_1^2\sqrt{\lambda q}}\left[nq + \frac{q_0}{2}\right]\end{aligned}$$

which is also negative since $-\frac{(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}} = \frac{q_0\overline{M}}{\mu_1} - 1$ by first order condition. Thus, $\overline{p}^*$ is a maximum.

Proof of Part (c). The proof follows directly from Proposition 5.

Proof of Theorem 7

Proof of Part (a). The proof follows also the steps of Theorem 3's proof. We first write the expected profit under optimal inventory policy as a convex combination of the lower and upper bound expressions:

$$\Pi_\infty(p,b) = \theta\underline{\Pi}_\infty + (1-\theta)\overline{\Pi}_\infty$$
$$= \frac{\lambda nq}{1-\beta}\left\{p-(b+l-\beta v)q_r - \theta c - (1-\theta)\left[v+h+\frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}\right]\right\}$$

Let $M = p-(b+l-\beta v)q_r - \theta c - (1-\theta)\left[v+h+\frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}\right]$. The first derivatives of Π_∞ with respect to b and p are, respectively,

$$\frac{\partial \Pi_\infty(p,b)}{\partial b} = \frac{\lambda nqq_r}{1-\beta}\left[\frac{q_0M}{\mu_1} - 1 - \frac{(1-q_r)}{\mu_2}(b+l-\beta v) + \frac{(1-\theta)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$
$$\frac{\partial \Pi_\infty(p,b)}{\partial p} = -\frac{\lambda nq}{1-\beta}\left[\frac{q_0M}{\mu_1} - 1 + \frac{(1-\theta)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$

To satisfy both first order conditions, $\frac{(1-q_r)}{\mu_2}(b^*+l-\beta v) = 0$, then $b^* = \beta v - l$. Let θ^* be the θ value at the optimal price and refund, that is, $\Pi^*_\infty(p^*,b^*) = \theta^*\underline{\Pi}_\infty + (1-\theta^*)\overline{\Pi}_\infty$. The second derivative of Π_∞ with respect to b at optimality is

$$\frac{\partial^2 \Pi_\infty(p,b^*)}{\partial b^2} = -\frac{\lambda nqq_0q_r^2}{(1-\beta)\mu_1}\left[\frac{nqM}{\mu_1} + 1 - \frac{(1-\theta^*)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$
$$-\frac{\lambda nqq_r(1-q_r)}{(1-\beta)\mu_2} - \frac{\lambda nqq_0q_r^2(1-\theta^*)(c-\beta v)\phi(z^*)}{2(1-\beta)\mu_1^2\sqrt{\lambda q}}\left[nq+\frac{q_0}{2}\right]$$

Using the optimality condition $\frac{q_0M}{\mu_1} - 1 = -\frac{(1-\theta^*)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}$, clearly the second derivative is negative, and b^* is a maximum.

Proof of Part (b). Similarly for the optimal price, the first order condition determines that $q_0\left[M + \frac{(1-\theta^*)(c-\beta v)\phi(z^*)}{2\sqrt{\lambda q}}\right] = \mu_1$. At optimality, $M = p^* - \theta^* c - (1-\theta^*)\left[v+h+\frac{(c-\beta v)\phi(z^*)}{\sqrt{\lambda q}}\right]$. It follows that $p^* = \left[n\exp\left(\frac{\tilde{u}^*(p^*)-\tilde{u}_0}{\mu_1}\right)+1\right]\mu_1 + \theta^* c + (1-\theta^*)\left[v+h+\frac{(c-\beta v)\phi(z^*)}{2\sqrt{\lambda}}\sqrt{n+\exp\left(\frac{\tilde{u}_0-\tilde{u}(p^*)}{\mu_1}\right)}\right]$. Finally, we check the second derivative of Π_∞ at optimality:

$$\frac{\partial^2 \Pi_\infty}{\partial p^2} = -\frac{\lambda nqq_0}{(1-\beta)\mu_1}\left[\frac{nqM}{\mu_1} + 1 - \frac{(1-\theta^*)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}}\right]$$
$$-\frac{\lambda nqq_0(1-\theta^*)(c-\beta v)\phi(z^*)}{2(1-\beta)\mu_1^2\sqrt{\lambda q}}\left[nq+\frac{q_0}{2}\right]$$

which is also negative since $-\frac{(1-\theta^*)(c-\beta v)\phi(z^*)q_0}{2\mu_1\sqrt{\lambda q}} = \frac{q_0 M}{\mu_1} - 1$ by first order condition. Thus, p^* is a maximum, and since $0 \le \theta^* \le 1$, $\overline{p}^* \le p^* \le \underline{p}^*$, and the proof is complete.

Proof of Proposition 6

The first and second derivatives of $\widetilde{\Pi}_\infty$ with respect to price are

$$\begin{aligned}
\frac{\partial \widetilde{\Pi}_\infty(\tilde{p}, b^*)}{\partial \tilde{p}} &= -\frac{\lambda n q}{1-\beta}\left[\frac{q_0}{\mu_1}\left[\tilde{p} - c + (c - v - h)q_r\right] - 1\right] \\
\frac{\partial^2 \widetilde{\Pi}_\infty(\tilde{p}, b^*)}{\partial \tilde{p}^2} &= \frac{\lambda n q q_0}{(1-\beta)\mu_1}\left[\frac{q_0}{\mu_1}\left[\tilde{p} - c + (c - v - h)q_r\right] - 1\right] \\
&\quad - \frac{\lambda n q q_0}{(1-\beta)\mu_1}\left[\frac{nq}{\mu_1}\left[\tilde{p} - c + (c - v - h)q_r\right] + 1\right]
\end{aligned}$$

If we set the first derivative to 0, we obtain, after rearranging terms, the following implicit expression for price

$$\tilde{p}^* = \left[n \exp\left(\frac{\tilde{u}(\tilde{p}^*) - \tilde{u}_0}{\mu_1}\right) + 1\right]\mu_1 + c - (c - v - h)q_r$$

We can check easily that the second derivative at $\tilde{p}^*$ is negative, therefore $\tilde{p}^*$ is optimal.

REFERENCES

Alptekinoğlu, A., C. J. Corbett. 2008a. Leadtime - variety tradeoff in product differentiation. Working paper, SMU, Dallas, TX.

Alptekinoğlu, A., C. J. Corbett. 2008b. Mass customization vs. mass production: Variety and price competition. *Manufacturing Service Oper. Management* **10**(2) 204–217.

Anderson, S. P., A. de Palma. 1992. Multiproduct firms: A nested logit approach. *J. Indust. Econom.* **40**(3) 261–276.

Anderson, S. P., A. de Palma, J. F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.

Aydin, G., J. K. Ryan. 2000. Product line selection and pricing under the multinomial logit choice model. Working paper, Purdue University, West Lafayette, IN.

Bayus, B. L., W. P. Putsis. 1999. Product proliferation: An empirical analysis of product line determinants and market outcomes. *Marketing Sci.* **18**(2) 137–153.

Ben-Akiva, M. 1973. Structure of passenger travel demand models. Ph.D. thesis, Department of Civil Engineering, MIT, Cambridge, MA.

Ben-Akiva, M., S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.

Berger, J., M. Draganska, I. Simonson. 2007. The influence of product variety on brand perception and choice. *Marketing Sci.* **26**(4) 460–472.

Berman, B. 2002. Should your firm adopt a mass customization strategy? *Business Horizons* **45**(4) 51–60.

Blanchard, D. 2005. Moving forward in reverse. *Logistics Today* **46**(7) 7–8.

Boatwright, P., J. C. Nunes. 2001. Reducing assortment: An attribute-based approach. *J. Marketing* **65**(3) 50–63.

Borle, S., P. Boatwright, J. B. Kadane, J. C. Nunes, G. Shmueli. 2005. The effect of product assortment changes on customer retention. *Marketing Sci.* **24**(4) 616–622.

Broniarczyk, S. M., W. D. Hoyer, L. McAlister. 1998. Consumers' perceptions of the assortment offered in a grocery category: The impact of item reduction. *J. Marketing Res.* **35**(2) 166–176.

Cachon, G. P., A. G. Kök. 2007. Category management and coordination in retail assortment planning in the presence of basket shopping consumers. *Management Sci.* **53**(6) 934–951.

Cachon, G. P., C. Terwiesch, Y. Xu. 2005. Retail assortment planning in the presence of consumer search. *Manufacturing Service Oper. Management* **7**(4) 330–346.

Campanelli, M. 2007. The gift of hassle-free returns. *Entrepreneur* (Dec 18, 2007). URL http://www.entrepreneur.com.

Cargille, B., C. Fry, A. Raphel. 2005. Managing product line complexity. *OR/MS Today* **32**(3) 34–41.

Chan, L. M. A., Z. J. Shen, D. Simchi-Levi, J. L. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. D. Simchi-Levi, D. Wu, Z. J. Shen, eds., *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*. Kluwer Academic Publishers.

Che, Y-K. 1996. Customer return policies for experience goods. *J. Indust. Econom.* **44**(1) 17–24.

Davis, S., E. Gerstner, M. Hagerty. 1995. Money back guarantees in retailing: Matching products to consumer tastes. *J. Retailing* **71**(1) 7–22.

Davis, S., M. Hagerty, E. Gerstner. 1998. Return policies and the optimal level of hassle. *J. Econom. Bus.* **50**(5) 445–460.

de Brito, M. P., R. Dekker. 2003. Modelling product returns in inventory control - exploring the validity of general assumptions. *Internat. J. Production Econom.* **81-82** 225–241.

de Vries-van Ketel, E. 2006. How assortment variety affects assortment attractiveness: A consumer perspective. Ph.D. thesis, RSM Erasmus University, Erasmus Research Institute of Management. http://hdl.handle.net/1765/7193.

Dekker, R., M. Fleischmann, K. Inderfurth, L. N. van Wassenhove. 2004. *Reverse Logistics: Quantitative Models for Closed-Loop Supply Chains*. Springer-Verlag, New York.

Dewan, R., B. Jing, A. Seidmann. 2003. Product customization and price competition on the internet. *Management Sci.* **49**(8) 1055–1070.

Emmons, H., S. M. Gilbert. 1998. The role of returns policies in pricing and inventory decisions for catalogue goods. *Management Sci.* **44**(2) 276–283.

Enright, T. 2003. Post-holiday logistics. *trafficWORLD* (Jan 03) 20.

Gaur, V., D. Honhon. 2006. Product variety and inventory decisions under a locational consumer choice model. *Management Sci.* **52**(10) 1528–1546.

Guide, V. D. R., Jr., G. C. Souza, L. N. van Wassenhove, J. D. Blackburn. 2006. Time value of commercial product returns. *Management Sci.* **52**(8) 1200–1214.

Guide, V. D. R., Jr., L. N. van Wassenhove. 2003. *Business Aspects of Closed-Loop Supply Chains*. Carnegie Mellon Univ Press.

Gupta, D., S. Benjaafar. 2004. Make-to-order, make-to-stock, or delay product differentiation? A common framework for modeling and analysis. *IIE Trans.* **36**(6) 529–546.

Heiman, A., B. McWilliams, J. Zhao, D. Zilberman. 2002. Valuation and management of money-back guarantee options. *J. Retailing* **78**(3) 193–205.

Hess, J. D., W. Chu, E. Gerstner. 1996. Controlling product returns in direct marketing. *Marketing Lett.* **7**(4) 307–317.

Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research, vol. 2*. McGraw-Hill, New York.

Hoch, S. J., E. T. Bradlow, B. Wansink. 1999. The variety of an assortment. *Marketing Sci.* **18**(4) 527–546.

Hopp, W. J., X. Xu. 2005. Product line selection and pricing with modularity in design. *Manufacturing Service Oper. Management* **7**(3) 172–187.

Inderfurth, K. 1997. Simple optimal replenishment and disposal policies for a product recovery system with leadtimes. *OR Spektrum* **19**(2) 111–122.

Johnson, L. K. 2002. New views on digital CRM. *MIT Sloan Management Rev.* **44**(1) 10.

Johnson, N. L., S. Kotz. 1970. *Continuous Univariate Distributions*. John Wiley & Sons.

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

Kekre, S., K. Srinivasan. 1990. Broader product line: A necessity to achieve success? *Management Sci.* **36**(10) 1216–1231.

Kim, J., G. M. Allenby, P. E. Rossi. 2002. Modeling consumer demand for variety. *Marketing Sci.* **21**(3) 229–250.

Kök, A. G., M. L. Fisher, R. Vaidyanathan. 2006. *Assortment planning: Review of literature and industry practice*. Kluwer Academic Publishers.

Kotha, S. 1995. Mass customization: Implementing the emerging paradigm for competitive advantage. *Strategic Management Journal* **16** 21–42.

Li, Z. 2007. A single-period assortment optimization model. *Production Oper. Management* **16**(3) 369–380.

Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* **36**(6) 724–738.

Maddah, B., E. K. Bish. 2007. Joint pricing, assortment, and inventory decisions for a retailer's product line. *Naval Res. Logistics* **54**(3) 315–330.

Mangasarian, O. L. 1969. *Nonlinear Programming*. McGraw-Hill, New York.

Matthews, S. A., N. G. Persico. 2005. Information acquisition and the excess refund puzzle. Working paper, University of Pennsylvania, Philadelphia, PA.

McFadden, D. 1978. *Modelling the Choice of Residential Location*. North-Holland Publishing Company, Amsterdam.

Mostard, J., R. de Koster, R. Teunter. 2005. The distribution-free newsboy problem with resalable returns. *Internat. J. Production Econom.* **97**(3) 329–342.

Mostard, J., R. Teunter. 2006. The newsboy problem with resalable returns: a single period model and case study. *Eur. J. Oper. Res.* **169**(1) 81–96.

Mukhopadhyay, S. K., R. Setoputro. 2004. Reverse logistics in e-business: optimal price and return policy. *Internat. J. Physical Distribution & Logistics Management* **34**(1) 70–88.

Mukhopadhyay, S. K., R. Setoputro. 2005. Optimal return policy and modular design for build-to-order products. *Journal of Operations Management* **23**(5) 496–506.

Murthi, B. P. S., S. Sarkar. 2003. The role of the management sciences in research on personalization. *Management Sci.* **49**(10) 1344–1362.

Olavson, T., C. Fry. 2006. Understanding the dynamics of value-driven variety management. *MIT Sloan Management Rev.* **48**(1) 63–69.

Padmanabhan, V., I. P. L. Png. 1997. Manufacturer's returns policies and retail competition. *Marketing Sci.* **16**(1) 81–94.

Pasternack, B. A. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Sci.* **4**(2) 166–176.

Petruzzi, N. C., G. E. Monahan. 2003. Managing fashion goods inventories: dynamic recourse for retailers with outlet stores. *IIE Trans.* **35**(11) 1033–1047.

Pine, B. J. 1993. *Mass Customization: The New Frontier in Business Competition*. Harvard Business School Press, Boston, MA.

Robert, C. P., G. Casella. 1999. *Monte Carlo Statistical Methods*. Springer-Verlag, New York.

Rogers, D. S., R. S. Tibben-Lembke. 1998. Going backwards: Reverse logistics trends and practices. *University of Nevada report, Center for Logistics Management. Reverse Logistics Executive Council.* URL http://www.rlec.org/reverse.pdf.

Ross, S. M. 2003. *Introduction to Probability Models*. 8th ed. Academic Press, San Diego, CA.

Shugan, S. M. 1989. Product assortment in a triopoly. *Management Sci.* **35**(3) 304–320.

Shulman, J. D., A. T. Coughlan, R. C. Savaskan. 2007. Managing consumer returns in a competitive environment. Working paper, University of Washington, Seattle, WA.

Shulman, J. D., A. T. Coughlan, R. C. Savaskan. 2008. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. Working paper, University of Washington, Seattle, WA.

Simpson, Vincent P. 1978. Optimum solution structure for a repairable inventory problem. *Oper. Res.* **26**(2) 270–281.

Smith, S., N. Agrawal. 2000. Management of multi-item retail inventory systems with demand substitution. *Oper. Res.* **48**(1) 50–64.

Stock, J., T. Speh, H. Shear. 2006. Managing product returns for competitive advantage. *MIT Sloan Management Rev.* **48**(1) 57–62.

Su, X. 2008. Consumer returns policies and supply chain performance. Working paper, Haas School of Business, University of California, Berkeley, CA.

Swaminathan, J. M., S. R. Tayur. 2003. Models for supply chains in e-business. *Management Sci.* **49**(10) 1387–1406.

Syam, N. B., N. Kumar. 2006. On customized goods, standard goods, and competition. *Marketing Sci.* **25**(5) 525Ű–537.

Tagaras, G., D. Vlachos. 2001. A periodic review inventory system with emergency replenishments. *Management Sci.* **47**(3) 415–429.

Tsay, A. A., N. Agrawal. 2000. Channel dynamics under price and service competition. *Manufacturing Service Oper. Management* **2**(4) 372–391.

Tsay, A. A., S. Nahmias, N. Agrawal. 1999. *Modeling supply chain contracts: A review*. Kluwer Academic Publishers.

van Herpen, E., R. Pieters. 2002. The variety of an assortment: An extension to the attribute-based approach. *Marketing Sci.* **21**(3) 331–341.

van Riper, T., K. Nolan. 2008. The toughest holiday returns. *Forbes* (Jan 14, 2008). URL http://www.forbes.com.

van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45**(11) 1496–1509.

Vlachos, D., R. Dekker. 2003. Return handling options and order quantities for single period products. *Eur. J. Oper. Res.* **151**(1) 38–52.

Wood, S. L. 2001. Remote purchase environments: The influence of return policy leniency on two-state decision processes. *J. Marketing Res.* **38**(2) 157–169.

Xie, J., E. Gerstner. 2007. Service escape: Profiting from customer cancellations. *Marketing Sci.* **26**(1) 18–30.

Yalabik, B., N. C. Petruzzi, D. Chhajed. 2005. An integrated product returns model with logistics and marketing coordination. *Eur. J. Oper. Res.* **161**(1) 162–182.

Yano, C., S. Gilbert. 2005. Coordinated pricing and production/procurement decisions: A review. A. K. Chakravarty, J. Eliashberg, eds., *Managing Business Interfaces: Marketing and Engineering Issues in the Supply Chain and Internet Domains*. Springer.

BIOGRAPHICAL SKETCH

Alex Grasas was born in Barcelona, Spain. He received his B.S. from the Industrial Engineering department at the Technical University of Catalonia (UPC) in Spain (2001). He worked for three years as research assistant in the Operations Management department at the IESE Business School in Spain. After that, he moved to Gainesville, FL, to pursue his doctorate degree in industrial and systems engineering at the University of Florida. His research interests include retail operations, supply chain management and operations/marketing interface. He is currently working on the interactions between product assortment and return policy of a retailer. He expects to continue his career in academia.