Chapter 8
Is Assortment Selection a Popularity Contest?
A Study of Assortment, Return Policy, and Pricing Decisions of a Retailer

Aydın Alptekinoğlu, Alex Grasas, and Elif Akçalı

Abstract Should retailers take product returns into account when choosing their assortments? And, when doing so, should they consider assortment selection as a *popularity contest* – by carrying products that they think will be popular among consumers? Or, is there ever a case for carrying *eccentric* products – those that are least likely to be purchased by a typical consumer? In search of answers to these questions, we explore in this chapter the interactions between product assortment, return policy, and pricing decisions of a retailer. We consider a category of horizontally differentiated products delivered in two alternative supply modes: make-to-order (MTO) and make-to-stock (MTS). In the MTO mode, products are supplied after demand materializes, whereas in the MTS mode, the retailer stocks products prior to the selling season. Underlying our demand model, consumer choice behavior follows a nested multinomial logit model, with the first stage involving a product choice, and the second stage involving a keep-or-return decision. 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. We also

Aydın Alptekinoğlu
SMU Cox School of Business, Dallas, Texas, USA, e-mail: aalp@cox.smu.edu

Alex Grasas
Economics and Business, Universitat Pompeu Fabra, Barcelona, SPAIN,
e-mail: alex.grasas@upf.edu

Elif Akçalı
Industrial and Systems Engineering, University of Florida, Gainesville, FL, USA,
e-mail: akcali@ufl.edu

S. Netessine, C.S. Tang (eds.), *Consumer-Driven Demand and Operations Management Models*, International Series in Operations Research & Management Science 131,
DOI 10.1007/978-0-387-98026-3_8,

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

8.1 Introduction

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).

Given their financial importance, 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.

Is assortment selection a popularity contest? When choosing their product assortments for a particular category, say different colors and styles of a golf shirt, should retailers always prefer what they think will be popular among consumers? Or, is there ever a case for carrying eccentric products? In this chapter, we argue that relatively strict return policies (with less than full refunds) can render eccentric products more profitable than popular ones. Our argument is moderated by the retailer's basic operational mode: make-to-order (MTO – the retailer does not keep its own inventory but rather buys and delivers the product after a consumer places an order) versus make-to-stock (MTS – the retailer keeps the product in stock). When refunds are sufficiently low, it is optimal for retailers to carry nothing but most eccentric products in the MTO case, and a mix of most popular and most eccentric products in the MTS case. We also find that more lenient return policies (higher refunds) may sometimes require deeper assortments (larger variety of products). In view of our analytical results and numerical observations, we conclude that retailers should not only carefully consider their return policy in assortment planning, but also take their basic operational mode (MTO versus MTS) into account.

The rest of this chapter begins with an abridged version of a model that we developed and analyzed in a recent working paper (Grasas et al., 2008). We then describe three specific extensions of this model. The main purpose of the chapter is to demonstrate (by numerical experimentation) that the analytical results of our working paper, which we summarize above and in greater detail in Section 8.4.1, are robust to these extensions. 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.

8.2 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 (e.g., Cachon and Kök, 2007, Alptekinoğlu and Corbett, 2008b); impact of product variety on consumer behavior (e.g., Kim et al., 2002, Borle et al., 2005); and interactions between product variety and operational considerations such as inventory and leadtime (e.g., 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). The presence of product returns obviously complicates assortment planning further, yet it has not been addressed in this literature so far. 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 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 (e.g., 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 (e.g., Pasternack, 1985, Emmons and Gilbert, 1998), another stream concentrates on the influence of a retailer's return policy on consumers (e.g., Yalabik et al., 2005, 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.

8.3 Models

We first provide a compact description of our base model. For a more complete discussion of the key features and assumptions, we refer the reader to our working paper (Grasas et al., 2008).

8.3.1 Base Model: Assortment Decision for Exogenous Price and Return Policy

Motivated with the question of whether retailers should consider returns when merchandising (as they compose their product assortments), we explore in our working paper 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. These 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, a retailer specialized in high-end gear and apparel for outdoors, carries all of its products in inventory at its central warehouse in Utah).

8.3.1.1 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 the 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*.

The 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.

We assume exogenous prices. 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 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). 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. (We discuss in Sections 8.4.2 and 8.5 potential implications of endogenizing price.)

The types of returns we consider involve products returned in resalable condition. 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, which focuses on a single horizontally differentiated product line), or dictated by common industry practice. (We discuss in Sections 8.4.3 and 8.5 potential implications of endogenizing refund fraction.) 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] 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. (We discuss the implications of extending our model to multiple periods in Sections 8.4.4 and 8.5.) 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).

[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.

8.3.1.2 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 a two-stage nested multinomial logit (N-MNL) model. The nests are products, and they each contain two post-purchase alternatives: keep and return.

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. Without loss of generality, we sort products in N in non-increasing order of attractiveness levels, 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.

Suppose the utilities associated with purchasing product i and keeping or returning it are given by $u_{i,\text{keep}} = a_i - p + \varepsilon_{i,\text{keep}}$, and $u_{i,\text{return}} = -(1-\alpha)p + \varepsilon_{i,\text{return}}$, where $\varepsilon_{i,\text{keep}}$ and $\varepsilon_{i,\text{return}}$ 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 $u_{i,\text{keep}}$ is the attractiveness minus the price; and the deterministic portion of $u_{i,\text{return}}$ 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 $u_{i,\text{return}}$ 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 $P_{\text{return}|i} \equiv \Pr\left\{u_{i,\text{return}} > u_{i,\text{keep}}\right\}$, which yields the following formula:[3]

$$P_{\text{return}|i} = \frac{1}{1 + \exp[(a_i - \alpha p)/\mu_2]}$$

And, of course, $P_{\text{keep}|i} = 1 - P_{\text{return}|i}$. Should the consumer choose the outside option in the first stage, there is no further choice to make in the second stage. Note that $P_{\text{return}|i}$ 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 $P_{\text{return}|i}$ 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 $A_i \equiv$

---

[2] The cumulative distribution function (*cdf*) of a Gumbel random variable X with mean zero and scale $1/\mu$ is given by $P(X \leq x) = \exp[-\exp(-x/\mu - \gamma)]$, 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.

[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 $\Pr\{\varepsilon_2 - \varepsilon_1 \leq x\} = [1 + \exp(-x/\mu)]^{-1}$.

$E\left[\max\left(u_{i,\text{keep}}, u_{i,\text{return}}\right)\right]$, which can be derived as

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

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 random variables with mean zero and scale $1/\mu_1$ ($\mu_1 > 0$). (ε_i are also independent of $\varepsilon_{j,\text{keep}}$ and $\varepsilon_{j,\text{return}}$ 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 $P_i^S \equiv \Pr\left\{U_i = \max_{j \in S \cup \{0\}} U_j\right\}$, which yields the following logit formula:[4]

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

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

[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).

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 $\sigma(\lambda P_i^S)^\beta$, 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 $\lambda P_{i,\text{return}}^S$ and standard deviation $\sigma(\lambda P_{i,\text{return}}^S)^\beta$. 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 $P_{i,\text{return}}^S$, respectively. Also, Poisson demands and returns constitute a natural special case of our aggregate demand model (i.e., set $\sigma = 1$ and $\beta = 1/2$, and use normal approximation of Poisson).

8.3.1.3 Supply Process and the Timing of Events

We consider two alternative modes of supply: MTO and MTS. In either case, 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:

$$\Pi_{\mathrm{MTO}}(S) = \sum_{j \in S} E\left[(p-c)D_j - (\alpha p + l - v)R_j\right] \tag{8.1}$$

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. We discuss the implications of this in Sections 8.4.4 and 8.5.

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 Web sites 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]

$$\begin{aligned}\Pi_{\mathrm{MTS}}(S) = \sum_{j \in S} \max_{x_j \geq 0} \Big\{ E\Big[& pD_j - cx_j - e(D_j - x_j)^+ \\ & - (\alpha p + l - v)R_j + v(x_j - D_j)^+ \Big] \Big\}\end{aligned} \tag{8.2}$$

[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 equals 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.

Timing of Events

To sum up, events in our base 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, $\Pi_{\text{MTO}}(S)$ or $\Pi_{\text{MTS}}(S)$. 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.

8.3.2 Extension 1: Assortment and Price Decisions for Exogenous Return Policy

In this extension we drop the assumption of exogenous price. So, everything remains the same as in the base model, except now price is also a decision taken by the retailer, simultaneous with the assortment decision (with return fraction still fixed).

8.3.3 Extension 2: Assortment and Return Policy Decisions for Exogenous Price

In this extension we drop the assumption of exogenous refund fraction. So, everything remains the same as in the base model, except now return fraction is also a decision taken by the retailer, simultaneous with the assortment decision (with price still fixed).

8.3.4 Extension 3: Assortment Decision for Multiple Periods

Finally, we take the base model and assume a multiple-period planning horizon. As in the base model, product assortment is decided at the very beginning – the beginning of the first period. To simplify the inventory management problem, and to amplify the impact of inventory on assortment over multiple periods, we assume that returned items are not salvaged until the end of the last period, i.e., they are always re-stocked and possibly used in succeeding periods.

8.4 Analytical Results and Numerical Observations

8.4.1 Optimal Assortment in the Base Model

In this section we seek to optimize the retailer's assortment decision for a given retail price and return policy. This is generally a difficult task as there are 2^n different possibilities. We provide structural results that significantly reduce the search space for accomplishing this task.

8.4.1.1 MTO Model with Returns

To lay the groundwork for discovering the structure of the optimal assortment, we first conduct a thought experiment. Suppose the current assortment is some proper subset S of N. Consider adding a product with a certain attractiveness to the current assortment. How does the new expected profit behave as a function of the attractiveness of the "new" product? Does adding this particular product to the assortment improve the profit? These two questions are resolved in two lemmas reported in our working paper, and they provide building blocks for proving the structure of the optimal assortment.

Theorem 1 (Grasas et al., 2008).
(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.
(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.

The 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, previous results in the literature (e.g., 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 (e.g., Cargille et al., 2005, and Olavson and Fry, 2006). 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. We note that 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.

8.4.1.2 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. Based on this finding (reported in our working paper), we establish the following result regarding the structure of the optimal assortment.

Theorem 2 (Grasas et al., 2008).
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. There exist problem instances where the optimal assortment is composed of: (1) most popular products only, (2) most eccentric products only, or (3) some most popular and some most eccentric products.

This result paves the way to showing that the structure of the optimal assortment is fundamentally different under MTS than under MTO. 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 8.1; details of the example are described in Section 8.4.1.4.

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 usual with newsvendor costs, 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); and (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.

Table 8.1 Optimal assortment S^*, composed of products that correspond to shaded cells, for the problem instance in Table 8.2 with threshold refund fraction, $(v-l)/p = 0.825$.

		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	■	■	■	■	■	■	■	■	■		

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 our working paper), 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.

8.4.1.3 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 suboptimal assortments. This is especially true if they have relatively strict return policies.

8.4.1.4 A Numerical Example

We conclude our analysis of the base model with a numerical example that illustrates the different kinds of solutions that arise under MTO/MTS environments with strict/lenient return policies. Table 8.1 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 $(v-l)/p = 0.825$. As expected, optimal variety (number of products in the optimal assortment) is lower under MTS.

Our working paper (Grasas et al., 2008) contains an extensive numerical study section that explores the following research questions. We provide a brief summary of our most interesting findings here, and refer the reader to the paper for a full exposition.

- 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, number of products offered, influenced by changes in refund fraction? Is it necessarily the case that more lenient return policies imply less variety? We find that the answer is no. 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 ($(v-l)/p$).
- How does the degree of differentiation among the potential products considered by the retailer (spread of a-values for products in N) influence its profit and depth of assortment? If the retailer had any influence over this degree of differentiation, would it prefer higher or lower differentiation? 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.
- 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 retailer? Somewhat surprisingly, yes. The reasonable presumption that higher heterogeneity about consumers' keep/return decisions will lead to lower profits is wrong for strict return policies.

- How does the optimal refund fraction depend on the structure of the assortment? We find 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.

In the rest of this section, we report results from our numerical study of the three extensions of the base model. Unless we state otherwise, in all experiments we use a set of base parameter values displayed in Table 8.2. Also, we report only the MTO case (the MTS case does not reveal any notably different insight).

Table 8.2 Base parameter values.

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

8.4.2 Optimal Assortment and Price in Extension 1

In this subsection, 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.

8.4.2.1 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 (see Figure 8.1). 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. 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

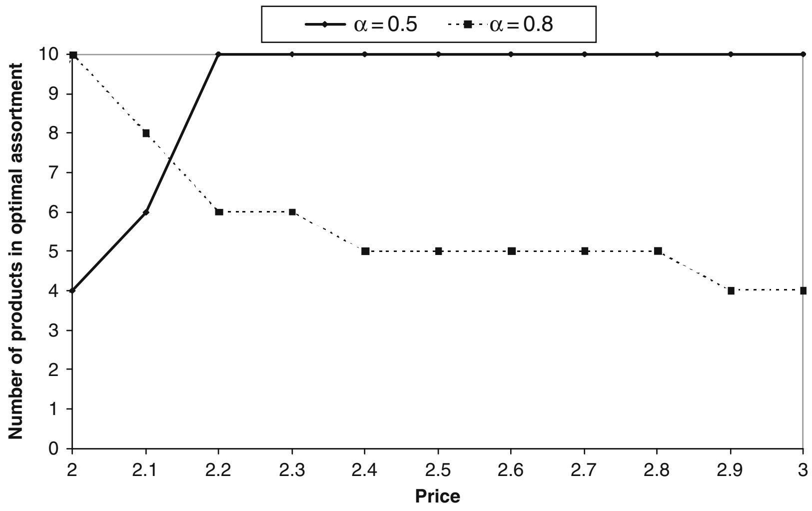

Fig. 8.1 Variety versus price: Number of products in the optimal assortment ($|S^*|$) as price (p) varies for different values of refund fraction (α).

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 (e.g., 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, besides the one caused by product returns in this work).

8.4.2.2 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, Figure 8.2 plots the expected profit as price varies from 2 to 6. For every data point shown in the chart, 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 graph that the optimal price increases as refund fraction decreases. We examine this in more detail in the next subsection.

8.4.2.3 Optimal Price with Respect to Refund Fraction

Figure 8.3 shows 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 ($\alpha p \geq v - l$).

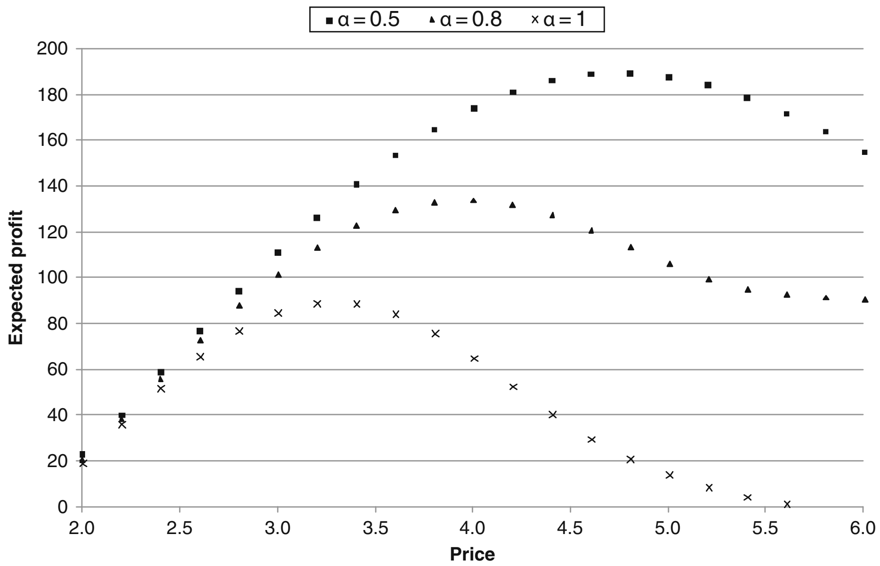

Fig. 8.2 Profit versus price: Expected profit as price (p) varies for different refund fractions (α) under optimal assortment (S^*).

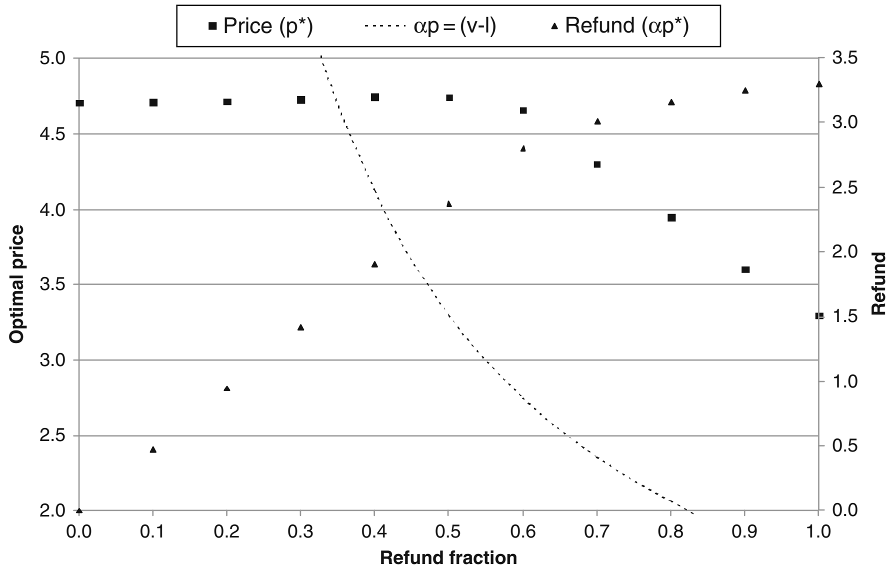

Fig. 8.3 Price versus refund: Optimal price (p^*) and refund (αp^*) for different values of refund fraction (α) under optimal assortment (S^*).

Again, the optimal assortment is computed for every data point. The optimization over S takes advantage of structural results presented earlier 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.

8.4.3 Optimal Assortment and Refund Policy in Extension 2

In this subsection, we investigate how endogenizing refund fraction (α) influences our assortment problem. As in the base model, price is considered exogenous.

8.4.3.1 Behavior of Expected Profit with Respect to Refund Fraction

Figure 8.4 plots 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. 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 Section 8.4.2.3.

8.4.3.2 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. The optimization over S takes advantage of structural results presented earlier 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 Figure 8.5, the optimal refund fraction represented by square dots, first

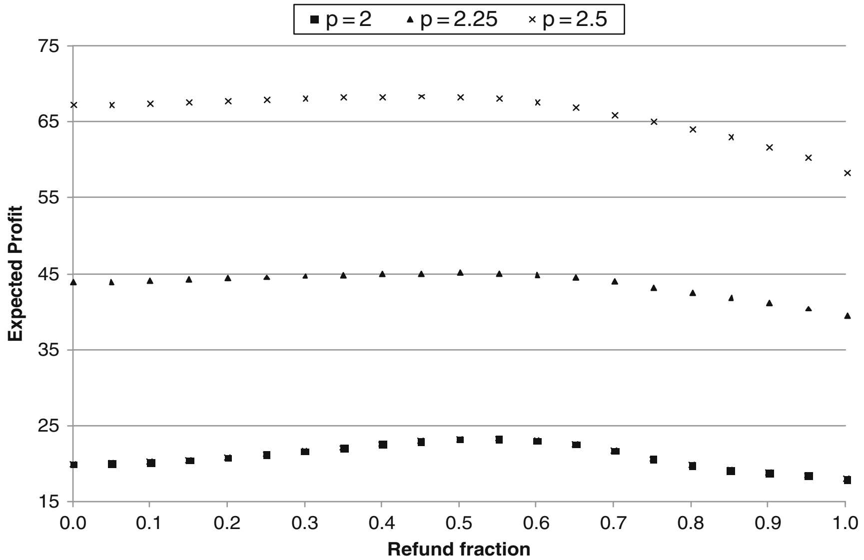

Fig. 8.4 Profit versus refund fraction: Expected profit as refund fraction (α) varies for different prices (p) under optimal assortment (S^*).

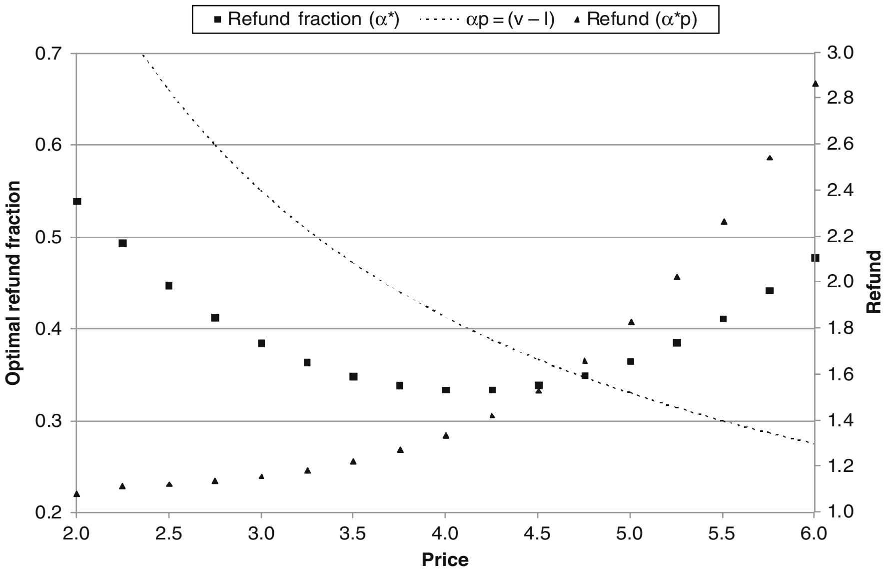

Fig. 8.5 Refund versus price: Optimal refund fraction (α^*) and refund ($\alpha^* p$) for different prices (p) under optimal assortment (S^*).

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, $p - c - (\alpha p + l - v)P_{\text{return}|i}$, 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.

8.4.4 Optimal Assortment for Multiple Periods in Extension 3

In this subsection, we extend the problem to a multiple-period setting. We assume that all returns are kept in 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 8.2, we compute the optimal assortment for different values of α as we did in Table 8.1. 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 Theorems 1 and 2). Tables 8.3 and 8.4 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 8.3 and 8.4 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 8.3 Optimal assortment S^*, composed of products that correspond to shaded cells, for a multiple-period problem with three-period planning horizon.

		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 8.4 Optimal assortment S^*, composed of products that correspond to shaded cells, for a multiple-period problem with ten-period planning horizon.

		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	■	■	■	■	■	■	■	■	■		

8.5 Concluding Remarks

Motivated with the question of whether retailers should consider product returns when merchandising (as they compose their product assortments), we first explore in our base model the interactions between product assortment decision and return policy of a price-taking retailer under two basic operational modes, make-to-order (MTO) and make-to-stock (MTS). We have a demand model grounded on individual consumer behavior. Consumers decide which product to buy in the first stage of a nested multinomial logit model, and then decide to keep or return the item in the second stage. In their purchase and keep/return decisions, consumers take both the assortment and refund fraction, the percentage of price refunded upon return, into account. The retailer, an expected profit maximizer, makes its assortment decision from an exogenous set of potential products that are horizontally differentiated. We call products with high (low) attractiveness *popular* (*eccentric*), because they are more (less) likely to be purchased by a typical consumer. In the MTS case, the retailer also makes an inventory decision for each product offered.

Our main finding from the base model is that the structure of the optimal assortment critically depends on the refund fraction and whether the products are supplied on an MTO or MTS basis. More specifically, we have two major analytical 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 confirm this. Including only the most popular products in an assortment agrees with common intuition, previous results in the literature (e.g., 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 (e.g., 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. (We argue by numerical examples in this chapter that this logic would likely survive extensions of our base model to endogenous price, endogenous refund fraction, and multiple resale opportunities.) The case for popular products, on the other hand, is twofold. 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.

Our analytical and numerical results so far 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. Nevertheless, price and refund fraction simultaneously influence the purchase and return probabilities in a complex way, which proves to be quite challenging to investigate analytically. Our numerical study in this chapter demonstrates that the strict return policy region, where most of the interesting interactions occur, is still prominent after endogenizing either of these variables.

Another extension would be to consider multiple periods, which may allow richer inventory management issues and more sophisticated return behavior. Even in the simplest possible case, if assortment decision was to be made for once at the beginning of a finite planning horizon, the question of optimal assortment becomes analytically intractable. On a positive note, extending our model to multiple periods can only strengthen our result about strict return policies. Having multiple resale opportunities can only increase the resale value of returned products; therefore, the basic rationale for carrying eccentric products would actually be even more prominent over a multiple-period planning horizon. We indeed demonstrate this in our numerical studies: the strict return policy region remains to be highly salient.

In light of our analytical results and numerical observations, we conclude that retailers should not only carefully consider their return policy when merchandising, but also take their basic operational mode (MTO versus MTS) into account.

References

Alptekinoğlu A, Corbett CJ (2008a) Leadtime – variety tradeoff in product differentiation. Working paper, SMU, Dallas, Texas

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

Anderson SP, de Palma A, Thisse JF (1992) Discrete choice theory of product differentiation. MIT Press, Cambridge, Massachusetts

Aydin G, Ryan JK (2000) Product line selection and pricing under the multinomial logit choice model. Working paper, Purdue University, West Lafayette, Indiana

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

Borle S, Boatwright P, Kadane JB, Nunes JC, Shmueli G (2005) The effect of product assortment changes on customer retention. Marketing Science 24(4):616–622

Cachon GP, Kök AG (2007) Category management and coordination in retail assortment planning in the presence of basket shopping consumers. Management Science 53(6):934–951

Cachon GP, Terwiesch C, Xu Y (2005) Retail assortment planning in the presence of consumer search. Manufacturing Service Operating Management 7(4):330–346

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

Dekker R, Fleischmann M, Inderfurth K, van Wassenhove LN (2004) Reverse logistics: Quantitative models for closed-loop supply chains. Springer-Verlag, New York

Emmons H, Gilbert SM (1998) The role of returns policies in pricing and inventory decisions for catalogue goods. Management Science 44(2):276–283

Enright T (2003) Post-holiday logistics. trafficWORLD, January 20

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

Grasas A, Alptekinoğlu A, Akçalı E (2008) When to carry eccentric products? Optimal assortment under product returns. Working paper, University of Florida, Gainesville, Florida

Guide VDR, Jr, Souza GC, van Wassenhove LN, Blackburn JD (2006) Time value of commercial product returns. Management Science 52(8):1200–1214

Hopp WJ, Xu X (2005) Product line selection and pricing with modularity in design. Manufacturing Service Operating Management 7(3):172–187

Kim J, Allenby GM, Rossi PE (2002) Modeling consumer demand for variety. Marketing Science 21(3):229–250

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

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

McFadden D (1978) Modelling the choice of residential location. North-Holland Publishing Company, Amsterdam

Mostard J, Teunter R (2006) The newsboy problem with resalable returns: a single period model and case study. European Journal of Operational Research 169(1):81–96

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

Pasternack BA (1985) Optimal pricing and return policies for perishable commodities. Marketing Science 4(2):166–176

Robert CP, Casella G (1999) Monte Carlo statistical methods. Springer-Verlag, New York

Rogers DS, Tibben-Lembke RS (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

Shulman JD, Coughlan AT, Savaskan RC (2008) Optimal restocking fees and information provision in an integrated demand-supply model of product returns. Working paper, University of Washington, Seattle, Washington

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

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

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

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

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

Yalabik B, Petruzzi NC, Chhajed D (2005) An integrated product returns model with logistics and marketing coordination. European Journal of Operational Research 161(1):162–182