

**When to Carry Eccentric Products?
Optimal Retail Assortment under
Consumer Returns**

**Aydin Alptekinoglu,
Alex Grasas**

**This version November 2012
(April 2011)**

Barcelona GSE Working Paper Series

Working Paper n° 680

When to Carry Eccentric Products? Optimal Retail Assortment under Consumer Returns

November 10, 2012

Aydın Alptekinoglu, Alex Grasas

Edwin L. Cox School of Business, Southern Methodist Univ., Dallas, Texas 75275, USA, aalp@cox.smu.edu
Dept. of Economics and Business, Univ. Pompeu Fabra, Barcelona, Spain 08005, alex.grasas@upf.edu

To understand whether retailers should consider consumer returns when *merchandising*, we study how the optimal assortment of a price-taking retailer is influenced by its return policy. The retailer selects its assortment from an exogenous set of horizontally differentiated products. Consumers make purchase and keep/return decisions in nested multinomial logit fashion. Our main finding is that the optimal assortment has a counterintuitive structure for relatively strict return policies: It is optimal to offer a mix of the most popular and most eccentric products when the refund amount is sufficiently low, which can be viewed as a form of risk sharing between the retailer and consumers. In contrast, if the refund is sufficiently high, or when returns are disallowed, optimal assortment is composed of only the most popular products (a common finding in the literature). We provide preliminary empirical evidence for one of the key drivers of our results: more eccentric products have higher probability of return – conditional on purchase. In light of our analytical findings and managerial insights, we conclude that retailers should take their return policies into account when merchandising.

1. Introduction

Consumer return policies and product assortment are typically considered as separate realms of the retailing business. Returns are often viewed as micro and more operational, and assortment as strategic and more marketing related. This state of affairs leads to decisions in each area to be made separately and independently of the other (Stock et al. 2006, Olavson and Fry 2006). We counter this conventional thinking by proving that optimal assortment decisions are fundamentally different when returns are taken into account.

Returns are financially important for retailers. The annual value of returned goods in the United States was \$194 billion in 2010, 8.12% of total retail industry sales (Anonymous 2010), and one estimate puts the industry’s annual spending on reverse logistics for processing and disposition of returns at more than \$40 billion (Enright 2003). Ever-increasing product variety only serves to

Table 1 Data for SONY Cyber-Shot DSC-W610 digital camera, as of May 3, 2012

	Sony	Target	Best Buy	Fry's	Amazon	eCost	Circuit City
Price	\$109.99	\$109.99	\$108.99	\$109.99	\$109.00	\$109.99	\$109.99
Restocking fee*	**	No fee	No fee	15%	15%	20%	25%
Refund fraction*	**	100%	100%	85%	85%	80%	75%
Refund	**	\$109.99	\$108.99	\$93.49	\$92.65	\$87.99	\$82.49
Average open-box price	\$96.59						
Net incremental value per unit returned	**	-\$13.40	-\$12.40	\$3.10	\$3.94	\$8.60	\$14.10

*Percentage of price. **Restocking fee not specified on the website.

make returns even more costly, as it increases the complexity of various processes used for returns management (see Guide et al. 2006, pp. 1202-1203, and Stock et al. 2006, pp. 59-61).

In this paper we investigate how the product assortment decision of a price-taking retailer is influenced by its return policy. We treat two basic operational environments, make-to-order (MTO) and make-to-stock (MTS), that allow us to draw a sharp distinction between retail settings that differ on the timing of supply decision. Under MTO (MTS), procurement of products happens after (before) consumers make their purchase decisions. Our demand model is grounded on a nested multinomial logit model of consumer choice behavior, in which consumers make purchase and keep/return decisions in two stages with random utilities. 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. Under MTS, the retailer also decides how much inventory to hold for each product. Price is exogenous: retailers often act as price-takers in numerous product categories or with particular brands, e.g., they use *manufacturer suggested retail price* (MSRP). Products differ only in terms of their attractiveness. We call products with high attractiveness *popular*, because a typical consumer is more likely to find them utility-maximizing, and call those with low attractiveness *eccentric*. We focus on a single aspect of return policies: refund amount. Both full and partial refunds are commonly observed in retailing (Shulman et al. 2009, p. 578). We term the percentage of price that the retailer refunds upon return *refund fraction*, and presume that it is exogenous (driven perhaps by category- or store-wide considerations, which are beyond the scope of our analysis).

To put our research in context, we collected some data on a digital camera recently released by Sony and sold by many online retailers and Sony itself (Table 1). The camera was available in several different colors and not all retailers carried all of them. Prices were almost identical across retailers, suggesting that they mostly follow Sony's MSRP. Their return policies included a

restocking fee, a common practice for consumer electronics products. The restocking fees as well as the net incremental value of returned units, estimated by deducting the refund amount from the average of open-box prices we were able to observe at the time (at Amazon and Best Buy), exhibit large variation among retailers. Note that a high enough restocking fee means that *returns can be a net benefit*. Examples of this phenomenon have also been documented by Shulman et al. (2009, p. 584). This leads us to the main research question we address in the paper: how does a retailer's return policy affect its assortment selection?

Our results establish that the retailer's optimal assortment is structurally different depending on how high or low the refund fraction is. Specifically, we show three results that apply to both MTO and MTS environments. With a strict return policy — meaning, a sufficiently low refund fraction — the retailer's optimal assortment has a counterintuitive structure: it contains a mix of the most popular and most eccentric products. On the other hand, with a lenient return policy, or when returns are not allowed, the retailer offers some of the most popular products. Carrying popular products in a retail assortment is consistent with common intuition, previous work (e.g., van Ryzin and Mahajan 1999, Aydın and Ryan 2000, Hopp and Xu 2005, Maddah and Bish 2007, Cachon and Kök 2007), and industry practice (e.g., Cargille et al. 2005, Olavson and Fry 2006). We find that considering returns in assortment planning can reverse this intuitive finding.

The main motive in offering an eccentric product as part of the optimal assortment is to profit from the resale of returned items (with positive net incremental values as in Table 1). The lower the refund fraction, as in strict return policies, the higher this benefit. This favors eccentric products, because they have higher probability of return, a consequence of our choice model for which we show preliminary empirical evidence in §3.2. Low refunds serve a risk-sharing purpose between the retailer and consumers; they make eccentric products economically viable. (We conjecture in §5 and §6 that our rationale for carrying eccentric products would likely extend to endogenous price, endogenous refund fraction, and multiple resale opportunities.) On the flip side, the reasons for carrying popular products include: they minimize return probability, which becomes desirable when every return is a net loss under a lenient return policy, and they have smaller demand variability (as measured by coefficient of variation), which reduces the inventory risk under MTS.

Our contribution also includes: (1) showing that a more lenient return policy may not imply less variety and (2) that the structure of optimal assortment sharply differs between MTO and MTS environments when variety has negligible fixed cost; and (3) taking steps to verify the robustness

of our findings to optimizing the refund fraction, to quantity-dependent salvage values, and to reselling as a post-purchase alternative to returning.

In light of our analytical findings (§4) and managerial insights (§5), we conclude that retailers should take their return policies into account when merchandising. We close the paper in §6 by suggesting ideas for future research and listing three empirically testable implications of our work.

2. Literature Review

Planning and management of retail assortment (product variety) have enjoyed substantial research attention from several different angles. A growing body of work explores relevant operational concerns such as inventory (van Ryzin and Mahajan 1999, Smith and Agrawal 2000, Aydın and Ryan 2000, Gaur and Honhon 2006), joint inventory-pricing decisions (Maddah and Bish 2007), substitution upon stockout (Honhon et al. 2010), delivery leadtime (Alptekinoğlu and Corbett 2010), and modularity in product design (Hopp and Xu 2005). Another line of research tackles strategic and competitive aspects such as entry deterrence (Bayus and Putsis 1999), channel coordination (Aydın and Hausman 2009), and competitive implications of consumer search (Cachon et al. 2008), basket-shopping behavior (Cachon and Kök 2007), brand preference (Kök and Xu 2011) and product satiation (Caro and Martinez-de-Albeniz 2012).

Consumers returning products is a widespread and well-accepted practice in retailing, and it clearly complicates assortment planning, yet it has never been addressed in this literature (Ramdas 2003, Kök et al. 2009). In a companion paper (Anonymous 2009) we present a special case of our base model and numerically explore three extensions (endogenous price, endogenous return fraction, and multiple periods) that are analytically intractable. In this paper, we conduct an analytical investigation of how product return policies influence optimal assortment selection. We prove that the presence of product returns not only introduces sizable difficulties to assortment decisions but also fundamentally changes them.

Product return policies were first studied in the context of retailer-to-supplier returns. For example, Pasternack (1985) and Emmons and Gilbert (1998) analyze the virtues of accepting returns for better channel coordination. We are interested in consumer-to-retailer returns. Under this umbrella, Moorthy and Srinivasan (1995) explore the quality signalling aspect of return policies. Su (2009) studies the impact of returns on supply chain coordination. Shulman et al. (2009) analyze the fundamental question of how much to refund. In our work, we also focus on the refund amount as

the key element of return policies, yet we are unique in exploring how return policies interact with assortment planning.

3. Model

We consider a retailer that takes product assortment decisions in a single-period setting. Our research objective is to explore the interactions between optimal assortment and return policy in make-to-order (MTO) and make-to-stock (MTS) environments.

3.1. Product Assortment and Return Policy

Let $N = \{1, 2, \dots, n\}$ denote a set of potential products that the retailer can choose from, and S denote the retailer's *assortment* — the subset of products that the retailer offers ($S \subseteq N$). The retailer incurs a fixed cost f for each product included in S (for a discussion of what this fixed cost may entail in retailing, see Smith and Agrawal 2000, p. 55).

Assortment decision S is for a horizontally differentiated product line with variants that differ only on a certain taste attribute, e.g., color or finish of a fashion good. Consistent with this, we assume that all products in N have identical unit procurement cost c , unit retail price p , and unit salvage values v_o and v_n for open-box (returned) and new-and-never-sold items (excess inventory that occurs only under MTS), respectively. Products differ only by their *attractiveness* (a 's defined in §3.2). To ensure that the retailer's quantity decision is worth investigating — that it is risky — we assume that returns and leftovers can be sold below cost in a secondary market, i.e., $v_o \leq v_n < c < p$. Salvage values are independent of quantity (§5.3.2 relaxes this).

We assume that price is exogenous. Endogenizing the price would be clearly interesting, yet analytically quite challenging (see Maddah and Bish (2007) who model price as a decision variable in an MNL-choice-based assortment problem, in which inventory is also a decision, but returns are ignored). van Ryzin and Mahajan (1999) note that there are “realistic cases in which a retailer's pricing flexibility is quite limited” (p. 1498). Following their lead, we also limit our analysis to a setting where the retailer has minimal or no control over prices and simply uses MSRP pricing. (We discuss potential consequences of endogenizing price in §6.)

We consider an exogenous return policy, characterized by the fraction of price refunded by the retailer. Let α denote the *refund fraction* ($0 \leq \alpha \leq 1$). We assume that the same refund fraction is valid for all products in S , so the refund amount for a return is αp . Note that this assumption reflects the typical practice in retailing, particularly for a given product line. If the retailer incurs

a unit reverse logistics cost (for sorting, repackaging, and restocking activities), this figure can be subtracted from v_o without loss of generality, except the following minor caveat. We still need $v_o > 0$ for returns to exist at all. Otherwise, if salvage value was smaller than the reverse logistics cost, the retailer would be better off asking customers who want to return a product to just keep it, which is hardly ever the case in practice.

Also, we disregard the option of product exchange. This is in alignment with current practice in online retailing. Customers wanting to exchange a product with another item in the same product line are typically asked to place a new order. Leaving exchanges beyond the scope of our analysis still limits the generality of our results. Because, those new orders would be associated with subsequent periods, which we do not explicitly model. (We state the consequences of extending our model to multiple periods in §6.) Besides, some retail stores do allow exchanges. Analytically speaking, however, exchanges are analogous to dynamic substitution: consumers may eventually buy a product other than their first choice, which poses substantial analytical difficulties even if product returns were ignored (Honhon et al. 2010).

Disposition of returned products in retailing takes many forms: (1) *restocking*, in which case the item is resold by the retailer either “as-new” at full price or as “open-box” at a discount, and the latter may occur in a different channel (e.g., Nordstrom Rack); (2) *liquidation*, in which case the item is sold to a third party, sometimes through intermediaries such as Genco that handles many aspects of returns management for their retailer clients (Hindo 2007); (3) *disposal*, which tends to happen more with defective or fraudulent returns (Shlachter 2010). We assume in our model a simple version of either the first or the second form, and ignore the third form altogether. In §5.3.2 we explore an alternative disposition scenario.

3.2. Individual Consumer Choice Behavior and Aggregate Demand

When making a purchase decision, a consumer has in mind all the products in S along with the *outside option* of not buying any of those products, denoted by 0. We use a nested multinomial logit (N-MNL) model to represent the consumer’s choice among $S \cup \{0\}$ and her post-purchase decision to either keep or return the product.

In the N-MNL framework, choice is viewed as a sequential process in which the consumer first chooses a product in S or the outside option with probability P_i^S , $i \in S \cup \{0\}$. Then, conditional on this first choice, she chooses to keep or return her purchase (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. At this stage consumers make their post-purchase decisions to keep or return the product given their decision to purchase a product $i \in S$ in the first stage. Let a_i be the *attractiveness* of product i . Without loss of generality label products such that $a_1 \geq a_2 \geq \dots \geq a_n$. Let $u_{i,keep} = a_i - p + \varepsilon_i + \epsilon_{i,keep}$ and $u_{i,return} = -d - (1 - \alpha)p + \varepsilon_i + \epsilon_{i,return}$ be the utilities associated with keeping and returning product i , respectively. Here, d is the disutility of returning an item (e.g., making an extra trip to the store), ε_i are independent and identically distributed (*iid*) Gumbel random variables with zero mean and $1/\mu_1$ scale parameter ($\mu_1 > 0$) that represent consumers' pre-purchase (stage 1) uncertainty, and $\epsilon_{i,keep}$ and $\epsilon_{i,return}$ are *iid* Gumbel random variables with zero mean and $1/\mu_2$ scale parameter ($\mu_2 > 0$) that represent consumers' post-purchase (stage 2) uncertainty. Note that the expected utility of buying and keeping product i , $E[u_{i,keep}]$, is the attractiveness net of price, and the expected utility of buying and returning product i , $E[u_{i,return}]$, is the disutility of returning and foregoing the non-refundable portion of the price. (If shopping itself involved a fixed cost or disutility for the consumer, subtracting a positive constant from both $u_{i,keep}$ and $u_{i,return}$ would account for it, and this would not change any of our results.)

At stage 2 consumers choose between keep and return options, whichever maximizes their idiosyncratic utility, having observed the realizations of all random terms. Therefore, the conditional probability of a typical consumer returning product i given purchase is

$$P_{return|i} \equiv \Pr \{u_{i,return} > u_{i,keep} | i\} = \left[1 + \exp \left(\frac{a_i + d - \alpha p}{\mu_2} \right) \right]^{-1}$$

(See Anderson et al. (1992) for generic derivations of this and other N-MNL formulas.) In addition, $P_{keep|i} = 1 - P_{return|i}$. If the consumer chooses the outside option at the first stage, then there is no second-stage decision.

Stage 1. At this stage consumers make their purchase decisions not knowing for certain if they would return the product or not. Typical of N-MNL models, this is captured by the consumers' inability to observe post-purchase random utility terms associated with stage 2 ($\epsilon_{i,keep}$ and $\epsilon_{i,return}$). Each consumer wanting to make a purchase decision assesses the utility of choosing nest $i \in S \cup \{0\}$ as $U_i = A_i + \varepsilon_i$, where the *expected utility* of buying product $i \in S$ is $A_i \equiv E[\max(a_i - p + \epsilon_{i,keep}, -d - (1 - \alpha)p + \epsilon_{i,return})] = \mu_2 \ln[\exp(a_i/\mu_2) + \exp((\alpha p - d)/\mu_2)] - p$, and

the expected utility of the outside option is $A_0 = 0$ (without loss of generality). Then the probability of a typical consumer choosing nest $i \in S \cup \{0\}$ in the first stage is given by

$$P_i^S \equiv \Pr \left\{ U_i = \max_{j \in S \cup \{0\}} U_j \right\} = \frac{\exp(A_i/\mu_1)}{1 + \sum_{j \in S} \exp(A_j/\mu_1)} = \frac{\omega_i}{1 + \sum_{j \in S} \omega_j}$$

where $\omega_i \equiv \exp(A_i/\mu_1)$ denotes the *preference* of product i (a term adopted from van Ryzin and Mahajan 1999), and P_0^S is the probability of not purchasing (preferring the outside option). By assumption, $\omega_0 = 1$.

Let us now review some basic properties of this choice model. More generous return policies with higher refunds (i.e., higher refund fraction, α) result in higher conditional return probability ($P_{return|i}$), higher expected utility (A_i) and higher purchase probability (P_i^S) for all products. More variety (i.e., larger assortment, S) implies lower demand for an existing product and higher total demand. Whether it implies higher or lower (unconditional) return probability depends on how the assortment is expanded. Finally, a product with higher attractiveness (a_i) experiences higher probability of purchase (P_i^S) and lower conditional probability of return ($P_{return|i}$).

This is a two-stage random utility model. From an empirical standpoint, random terms capture consumer heterogeneity. Specifically, ε_i reflect consumers' pre-purchase preferences; heterogeneity in stage 1 stems from consumers' varying preferences for products and return policies, reasons for needing the product, information states, etc. Similarly, $\epsilon_{i,keep}$ and $\epsilon_{i,return}$ reflect consumers' post-purchase preferences; heterogeneity in stage 2 comes from how consumers consider the keep and return options after a purchase decision. For example, customers may be swayed one way or another depending on feedback from their spouses on a piece of apparel. Larger μ_1 and μ_2 indicate higher variance, hence more heterogeneity. We need $\mu_1 > \mu_2$ for technical integrity of the N-MNL model (Anderson et al. 1992). This is reasonable in our modeling framework; consumers' pre-purchase heterogeneity is usually higher than their post-purchase heterogeneity. Consumers who buy the same product will differ less from each other due to their shared experience with, and understanding of, that one particular product.

While independence of the random terms may at first appear like a major limitation, many economists – taking an empirical stance – argue otherwise in sequential choice settings similar to ours. Train (2009), for example, views it as “the ideal rather than a restriction” (pp. 35-36). A well-specified empirical model, he reasons, must capture all observed utilities in deterministic components and what remains must be purely idiosyncratic differences from one consumer to another

at each stage. To the extent that the two choice stages involve distinct sources of heterogeneity, independence across stages also seems plausible. We consider pre-purchase heterogeneity as informational, driven by marketing mix variables like advertising and packaging; and post-purchase heterogeneity as hedonic, driven by experiential factors such as aesthetic fit and feel.

We presume that unhappy consumers are captive; they have no option but to return unwanted products to the retailer. This is not a real limitation, as all of our analytical results can be extended to a model with consumers having an option to resell, say on eBay (see §5.3.3).

Product Popularity versus Returns. High and low values of attractiveness point to a semantic distinction between products. A product with higher attractiveness has a higher expected utility for the ‘purchase and keep’ option, hence a higher probability of purchase. By the principle of utility maximization every consumer buys what she considers to be the “best” product. That is, a_i does not really determine the ‘attractiveness’ of product i in colloquial sense, but rather the purchase probability of product i . Hence, we refer to products with high attractiveness as *popular*, and those with low attractiveness as *eccentric*. Recall that we sort products in N in decreasing order of attractiveness, which makes lower-indexed products more popular and higher-indexed products more eccentric.

An interesting implication of our consumer choice model and – it turns out – a key driver of our results is that eccentric products have higher return probability (i.e., $P_{return|i}$ is higher for products with lower attractiveness, a_i). Datasets that we obtained from a major department store and Sena Cases (a maker and retailer of leather cases for mobile devices) offer preliminary empirical evidence for this. See Figures 1 and 2 for plots of relative sales volume (as % of total sales) versus return rate (as % of SKU sales) in the online channel for 21 bath towels from the department store and 18 iPhone cases from Sena. Both product lines are horizontally differentiated; SKUs within them differ by color only. Consistent with our choice model, SKUs with lower sales volume – those that are more eccentric – tend to have a higher return rate. Each data plot exhibits a statistically significant fit (with p-values 3.58% and 2.33% for towels and cases, respectively) to a downward-sloping power function, our theoretical prediction based on the N-MNL model. A more detailed description of the data and our statistical analysis is in Appendix A.

Aggregate Demand. Next, we state how the individual choice behavior results in an aggregate demand for each product $i \in S$. Let λ be the average number of consumers. Following suit with the existing literature (e.g., van Ryzin and Mahajan 1999), we assume that consumer choice depends




Figure 1 Bath Towels from a Major Department Store, Sales versus Returns for 21 SKUs, June–November 2010




Figure 2 Ultraslim for iPhone 3G/3GS from Sena Cases, Sales versus Returns for 18 SKUs, March 2009 – November 2010

solely on the set S , and is insensitive to the specifics of the retailer’s operational context such as MTO/MTS, supply leadtime, inventory policy, etc. We model the aggregate product i demand as a normal random variable D_i with mean λP_i^S and standard deviation $\sigma (\lambda P_i^S)^\beta$, where $\sigma > 0$ and $0 \leq \beta < 1$. Secondly, we model the aggregate product i returns as a normal random variable R_i with mean $\lambda P_{i,return}^S$ and standard deviation $\sigma (\lambda P_{i,return}^S)^\beta$. In a realistically calibrated model, the probability of realizations with $R_i > D_i$ will be so low that they can be safely ignored. In fact, this possibility is eliminated completely by Poisson demands and returns, which is a special case of our aggregate demand model (with $\sigma = 1$, $\beta = 1/2$, and normal approximation to Poisson).

3.3. Supply Process and the Timing of Events

We consider two operational environments – MTO and MTS – and 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 a supplier once an order is placed by a consumer. In the case of MTS, the quantity decision is made *before* the realization of demand. In both cases, we disregard any capacity limitations, assuming no limits on order quantities for the products in S .

MTO Environment. Orders are placed after demand is observed. So matching supply with demand is simple: the order quantity for each product must be equal to the observed demand, as any excess inventory would be salvaged incurring a unit loss of $(c - v_n)$. In this setting, the only

quantity risk associated with the supply decision is due to returns: Some products will be returned and will have to be salvaged at a value below cost ($v_o < c$).

Hence, the expected profit for the MTO case is given by:

$$\Pi_{MTO}(S) = \sum_{j \in S} E[(p - c)D_j - (\alpha p - v_o)R_j] - f|S| = \sum_{j \in S} [p - c - (\alpha p - v_o)P_{return|j}] \lambda P_j^S - f|S|$$

The terms within expectation are the sales revenue, procurement cost, and net cost of returns, respectively. The last term is the fixed cost of the assortment. Note that returns can only be salvaged. A more general model could allow reselling of returned items in the store, which would require a multiple-period modeling approach (more on that in §6).

MTS Environment. Orders are placed prior to observing the demands. The supply decision here is riskier as the retailer might over-stock or under-stock each product in S . Let x_j denote product j inventory ordered and received in advance of demand realizations.

When a stockout occurs, (1) the retailer can receive an emergency replenishment at a unit cost of e ($c < e < p$), and (2) the consumer agrees to wait for her most preferred product in the assortment rather than substituting for another one that may be currently in stock. In retailing, emergency replenishments are prevalent. For example, Famous Footwear promises a free delivery, if it turns out that a store does not have the desired size or color of a particular product. Not every consumer would go along with this obviously. This assumption lets us simplify the problem and concentrate on the influence of return policies on optimal assortment. Furthermore, if we were to model dynamic substitution in a stockout situation (where a consumer may switch from her most preferred product that is not in stock to a different one in stock), our model would be substantially more complex and most likely intractable. In fact, assortment and inventory optimization with an explicit consideration of stockout based substitution is a difficult problem in its own right, even without considering returns (Gaur and Honhon 2006, Honhon et al. 2010).

The expected profit under MTS is given by

$$\Pi_{MTS}(S) = \sum_{j \in S} \max_{x_j \geq 0} \{E[pD_j - cx_j - e(D_j - x_j)^+ - (\alpha p - v_o)R_j + v_n(x_j - D_j)^+]\} - f|S|$$

where $(y)^+ = y$ if $y > 0$, and 0 otherwise, for any real number y . The terms within expectation are revenue from sales, cost of regular supply, cost of emergency replenishment, net cost of returns, and salvage revenue from excess inventory, respectively. If the retailer incurs a holding cost for each

unit of excess inventory, this figure can be deducted from v_n without loss of generality. Again, the last term is the fixed cost of the assortment.

Given normally distributed demands, and using a well-known newsvendor result, the optimal order quantity for product j is $x_j^* = \lambda P_j^S + z^* \sigma (\lambda P_j^S)^\beta$ for all $j \in S$, where $z^* = \Phi^{-1}((e-c)/(e-v_n))$, and $\Phi(\cdot)$ denotes the cumulative distribution function of a standard normal random variable. Plugging the optimal order quantities back into the above profit expression, we obtain

$$\Pi_{MTS}(S) = \sum_{j \in S} [p - c - (\alpha p - v_o) P_{return|j}] \lambda P_j^S - (e - v_n) \sigma \phi(z^*) (\lambda P_j^S)^\beta - f|S| \quad (1)$$

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

Backlogging of excess demand through emergency orders, a standard assumption in the inventory management literature to gain analytical tractability, is a crucial compromise. Lost sales case, with consumers walking away upon stockout, is much more difficult: the newsvendor critical fractile (z^*) would then depend on P_j^S , and thus on the assortment decision.

Timing of Events. Having a return policy with an exogenous refund fraction α , the retailer seeks for an optimal assortment S that maximizes the expected profit, $\Pi_{MTO}(S)$ or $\Pi_{MTS}(S)$. Under MTO, the retailer observes the demand first, and then orders the demand quantity for all $j \in S$. Under MTS, the retailer orders x_j^* for all $j \in S$ first, and then observes the demand. The purchase decisions consumers take in stage 1 of the N-MNL model drives the demand process (§3.2). If MTS is in effect, emergency orders are placed and delivered to those consumers whose most preferred item is not in stock. Then, the return process is driven by the keep-return decisions taken in stage 2 of the N-MNL model by those consumers who make a purchase in stage 1. At the end the retailer salvages returns along with excess inventory (if any).

Determining the optimal refund fraction is an interesting problem in its own right. It inevitably requires a store-wide 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. Refund fractions in practice are usually round numbers, so one can always compute the expected profit for $\alpha = 0\%$, 1% , 2% , \dots , 100% , to find the near-optimal refund fraction for a given assortment (see §5.3.1). What is difficult is to find the optimal assortment because there are $2^n - 1$ different possibilities. We present structural results in the next section that can save significant time and effort in search of optimality.

4. Structure of the Optimal Assortment

In this section, we provide an analytical characterization of the optimal assortment S^* (used with a subscript, where appropriate, to indicate the operational environment). Recall that, without loss of generality we sort products in N in decreasing order of attractiveness (a_i); lower indices refer to more popular products, and higher indices to more eccentric products. Since ω_i is increasing in A_i , and A_i is increasing in a_i , this same ordering applies to preference values as well, i.e., $\omega_1 \geq \omega_2 \geq \dots \geq \omega_n$. (All proofs are in Appendix B.)

4.1. MTO Environment with Returns

Consider a retailer that operates in MTO environment and allows returns. In an effort to formally characterize the optimal assortment for such a retailer, we first consider an intermediate question: To an existing assortment $S \subset N$, which product (if any) would be the most profitable to add? To this end, we need to (i) investigate how the expected profit for the new assortment behaves as a function of δ , the preference of a hypothetical product, so that we can identify the best product to add to S , and (ii) determine if adding a particular product to S improves the expected profit or not. (We note that Theorem 1 of van Ryzin and Mahajan (1999) uses a similar proof technique.) Lemmas 1 and 2 below address (i) and (ii), respectively. Let $P_{return|\delta}$ be the conditional return probability for a product with preference δ (allowing a slight abuse of notation).

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

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

If $\alpha \geq v_o/p$, then $h_{MTO}(\delta)$ is quasiconcave and non-decreasing. Else, if $\alpha < v_o/p$, 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 $h_{MTO}(\cdot)$ lets us establish a local optimality result on which product to add, if we must, to an existing assortment S . Lemma 1 essentially says that: for a sufficiently

lenient return policy with refund fraction $\alpha \geq v_o/p$, the best product to add is the most popular of the remaining products in $N \setminus S$; whereas, for a strict return policy with $\alpha < v_o/p$, it must be either the most popular or the most eccentric product in $N \setminus S$.

We now take step (ii). Let $M_j \equiv p - c - (\alpha p - v_o)P_{return|j}$ be the expected profit margin per unit sales of product j . Note that, since $P_{return|j}$ is decreasing in a_j , $\alpha \geq v_o/p$ implies $M_j \geq M_{j+1}$, and $\alpha < v_o/p$ implies $M_j \leq M_{j+1}$ for $j = 1, \dots, n-1$.

LEMMA 2. *Adding product $i \in N \setminus S$ to an existing assortment $S \subset N$ improves the expected profit under MTO, i.e., $\Pi_{MTO}(S \cup \{i\}) \geq \Pi_{MTO}(S)$, if and only if*

$$M_i \geq \sum_{j \in S} P_j^S M_j + \frac{f}{\lambda P_i^{S \cup \{i\}}}$$

Lemma 2 shows that, for product i to be included in an existing assortment S , its expected profit margin must be greater than or equal to the expected profit margin of the current set plus a fixed cost term. Rules of thumb similar in nature to this result have been documented in practice (e.g., Cargille et al. 2005, Olavson and Fry 2006).

These two lemmas provide ammunition for finding the globally optimal assortment. Define $\mathcal{A}_0 = \mathcal{Z}_0 = \emptyset$, $\mathcal{A}_i = \{1, \dots, i\}$ as the set of i most popular products in N , and $\mathcal{Z}_j = \{n-j+1, \dots, n\}$ as the set of j most eccentric products in N , for all positive integers i and j between 1 and n .

THEOREM 1. *For a retailer operating in MTO environment, the optimal assortment is composed of (a) some number of most popular products from N , i.e., $S_{MTO}^* = \mathcal{A}_k$ for some $k \in \{0, 1, \dots, n\}$, if the return policy is sufficiently lenient so that $\alpha \geq v_o/p$, or (b) some number of most popular and some number of most eccentric products from N , i.e., $S_{MTO}^* = \mathcal{A}_j \cup \mathcal{Z}_{k-j}$ for some $j \in \{0, \dots, k\}$ and $k \in \{0, 1, \dots, n\}$, if the return policy is sufficiently strict so that $\alpha < v_o/p$.*

Return policy has a fundamental impact on optimal assortment. If the retailer adopts a lenient return policy, by using a sufficiently large refund fraction, the optimal assortment includes only the most popular products. This intuitive result concurs with those of van Ryzin and Mahajan (1999) and others cited in §1. A larger refund fraction renders product returns more costly to the retailer. Consequently, the retailer becomes more averse to returns and lines up the assortment with popular products, which are less likely to be returned.

In contrast, if the retailer adopts a strict return policy, by using a relatively small refund fraction ($\alpha < v_o/p$), the optimal assortment is a combination of the most popular and the most eccentric

products. Inclusion of eccentric products enables higher profits, mainly because an item that is sold and returned nets a profit of $(p - c - \alpha p + v_o)$, but an item sold and kept only $(p - c)$. This favors eccentric products, because they are more likely to be returned. On the other hand, having a higher probability of purchase favors popular products, because there is a per-product fixed cost for offering variety. Depending on whether the net incremental benefit from returns dominates the fixed cost effect, the retailer chooses one extreme (the most popular product) or the other (the most eccentric) when constructing its optimal assortment.

Hence the retailer may offer eccentric products if returned items have a positive net incremental value. We argue that this occurs often in practice. See Table 1 for a consumer electronics example. Shulman et al. (2009, p. 584) and Xie and Gerstner (2007, pp. 18–19) cite other examples ranging from antiques, to jewelry, and to service cancelations. Typically retailers set return policies for entire stores or product categories for ease of implementation and consumer relations reasons. Once a retailer decides to charge a restocking fee for a certain category, some product lines in that category will likely fall within the strict policy region. What we show is that whenever the return policy for a specific product line is in the strict region, the retailer has incentive to offer some of the most eccentric products in that product line. Our paper takes return policy as exogenous and does not delve into reasons for why retailers select strict return policies. But once they do, a strict return policy allows them to share the risk of carrying an eccentric product with consumers.

4.2. MTS Environment with Returns

Consider a retailer that operates in MTS environment and allows returns. The analysis starts with a thought experiment similar to the one employed in §4.1. We first find a necessary condition for when adding a product to an existing assortment S is profit-improving (Lemma 4), and then consider the question of which product (if any) should be added (Lemma 5).

Let $\widetilde{M}_j^S \equiv p - c - (\alpha p - v_o)P_{return|j} - (e - v_n)\sigma\phi(z^*)(\lambda P_j^S)^{\beta-1}$ be the expected profit margin per unit sales of product j under MTS. Note that, unlike the MTO case, the expected profit margin now depends on S ; it decreases as the assortment becomes broader with the addition of more products, because $(P_j^S)^{\beta-1} < (P_j^{S \cup \{i\}})^{\beta-1}$, which follows from $P_j^S > P_j^{S \cup \{i\}}$ and $\beta < 1$. Consequently, the inventory cost term, $(e - v_n)\sigma\phi(z^*)(\lambda P_j^S)^{\beta-1}$, increases as demand is dispersed among more products – weakening the effect of inventory pooling.

LEMMA 3. Adding product $i \in N \setminus S$ to an existing assortment $S \subset N$ improves the expected profit under MTS, i.e., $\Pi_{MTS}(S \cup \{i\}) \geq \Pi_{MTS}(S)$, if and only if

$$\widetilde{M}_i^{S \cup \{i\}} \geq \sum_{j \in S} P_j^S \widetilde{M}_j^{S \cup \{i\}} + \sum_{j \in S} [\widetilde{M}_j^S - \widetilde{M}_j^{S \cup \{i\}}] \frac{P_j^S}{P_i^{S \cup \{i\}}} + \frac{f}{\lambda P_i^{S \cup \{i\}}}$$

Lemma 3 gives a necessary and sufficient condition for identifying a profitable addition to an existing assortment. It seems difficult to interpret, yet leads us to a simpler and more intuitive necessary condition that proves useful in determining the structure of the optimal assortment.

LEMMA 4. A necessary condition for product $i \in N \setminus S$ to improve the expected profit when added to an existing assortment $S \subset N$ is given by:

$$\widetilde{M}_i^{S \cup \{i\}} > \sum_{j \in S} \frac{P_j^{S \cup \{i\}}}{1 - P_i^{S \cup \{i\}}} \widetilde{M}_j^{S \cup \{i\}}$$

Therefore, if adding product i to an existing assortment S is profitable, then its expected profit margin must exceed a threshold. That threshold, the right-hand-side above, is the new expected profit margin of existing products in S , conditional on the consumer not buying product i . A weaker necessary condition would be $\widetilde{M}_i^{S \cup \{i\}} > \sum_{j \in S} P_j^{S \cup \{i\}} \widetilde{M}_j^{S \cup \{i\}}$.

We now turn to the question of which product to add to S .

LEMMA 5. Let $h_{MTS}(\delta)$ for $\delta \in [\min_{j \in N \setminus S}(\omega_j), \max_{j \in N \setminus S}(\omega_j)]$ be the expected profit function when a product with preference δ is added to an existing assortment $S \subset N$ under MTS. That is, $h_{MTS}(\delta) = [g(\delta) + \tilde{g}(\delta)] / \gamma(\delta)$, where $\gamma(\delta) = 1 + \sum_{j \in S} \omega_j + \delta$ and

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

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

If $\alpha \geq v_o/p$, then $h_{MTS}(\delta)$ is increasing in δ for all δ that is profit-improving. Else, if $\alpha < v_o/p$, then $h_{MTS}(\delta)$ is quasiconvex.

In the lenient case the result is subject to the existence of a profit-improving product. This is innocuous; if no profit-improving product existed, the firm obviously would not add any more products to the current assortment, which would render this entire thought experiment unnecessary.

THEOREM 2. For a retailer operating in MTS environment, the optimal assortment is composed of (\mathbf{a}) some number of most popular products from N , i.e., $S_{MTS}^* = \mathcal{A}_k$ for some $k \in \{0, 1, \dots, n\}$, if

the return policy is sufficiently lenient so that $\alpha \geq v_o/p$, or **(b)** some number of most popular and some number of most eccentric products from N , i.e., $S_{MTS}^* = \mathcal{A}_j \cup \mathcal{Z}_{k-j}$ for some $j \in \{0, \dots, k\}$ and $k \in \{0, 1, \dots, n\}$, if the return policy is sufficiently strict so that $\alpha < v_o/p$.

The operational environment amplifies our counterintuitive result for strict return policies; the result continues to hold despite the fact that having to carry inventory favors popular products. Under MTS, the ordering decision for each product runs an inventory risk, which is proportional to the standard deviation of demand for the product (see equation 1). The operational risk of including a product in the assortment can thus be gauged by the coefficient of variation, standard deviation / mean demand. In our modeling framework, more attractive products are more likely to be purchased and have lower coefficients of variation. Hence, a retailer operating in MTS environment with a strict return policy has conflicting preferences (1) for popular products to capitalize on their lower demand variability and lower overhead that stems from the fixed cost of variety, and (2) for eccentric products to profit from their resale. The optimal assortment structure is a product of this tension. When a lenient return policy is in effect, however, the retailer clearly prefers popular products; they have lower demand variability *and* they are less likely to be returned.

It is apparent that eccentric (popular) products having higher (lower) return probability is a key driver of our results under both environments. We offer preliminary empirical evidence for this implication of our choice model in §3.2.

4.3. MTO and MTS Environments with No Returns

The retailer accepting no returns is a special case. By imposing infinite hassle on returns ($d = +\infty$), both of our models can capture the case of consumers never returning products in stage 2 of the choice process. This renders the N-MNL model equivalent to the standard MNL model.

THEOREM 3. *When the retailer accepts no returns, the optimal assortment in both operational environments – MTO and MTS – is composed of some number of most popular products from N , i.e., $S^* = \mathcal{A}_k$ for some $k \in \{0, 1, \dots, n\}$.*

We skip the proof for this special case. Theorem 3 echoes a structural result by van Ryzin and Mahajan (1999) in an MTS setting with lost sales and no returns. In conjunction with Theorems 1 and 2, it reveals that retailers need to explicitly take product returns into account when making assortment decisions, particularly if they adopt a strict return policy.

4.4. Solution Procedure and a Numerical Example

Our analytical results contain all the necessary ingredients for efficiently constructing the optimal assortment in MTO and MTS environments with or without returns. The greedy algorithm that starts with the empty set and adds one product at a time (in the lenient return policy and no-returns cases, the most popular of the remaining products, and in the strict return policy case, either the most popular or the most eccentric of the remaining products) until the expected profit stops improving (Lemmas 2 and 3) need not find the globally optimal assortment. Except for two special cases, finding the optimal assortment generally involves checking all subsets that has the structure shown in Theorems 1, 2, and 3. This means evaluating n , $n(n+1)/2$, and n subsets (ignoring the empty set) in the lenient return policy, strict return policy, and no-returns cases, respectively. Still, from an algorithmic standpoint, the value of our structural results is immense: they reduce the worst-case computational complexity of the assortment planning problem from exponential-time (complete enumeration of $2^n - 1$ possible assortments) to polynomial-time.

The following result spells out the two happy exceptions.

PROPOSITION 1. The greedy algorithm is optimal in MTO environment with a sufficiently lenient return policy ($\alpha \geq v_o/p$) and in MTO environment with no returns.

Hence, the greedy algorithm may terminate in less than n iterations; one need not always evaluate all n possibilities to find the global optimum in these two environments. In contrast, van Ryzin and Mahajan (1999) discover that it generally fails to find the optimal assortment in their model.

Due to shelf-space or storage limitations retailers may sometimes place a hard constraint on the size of an assortment. Suppose a retailer can carry at most K products and wishes to find the best possible assortment. All of our structural results would still apply. Let k^* be the optimal number of products to offer. By the same reasoning used in the proofs of Theorems 1 and 2, it would be optimal to either offer the $\min(K, k^*)$ most popular products under a lenient return policy or a mix of $\min(K, k^*)$ popular and eccentric products under a strict return policy.

We close our analysis by illustrating the array of solutions that one may see in MTO and MTS environments with a strict or a lenient return policy. Table 2 displays an example set of parameter values, along with the attractiveness levels for 10 potential products, and the optimal assortment under MTO and MTS with refund fraction α set to various values. In this example the critical refund fraction that differentiates strict and lenient return policies is $v_o/p = 0.8$.

Table 2 Base parameter values and the corresponding optimal assortment (indicated by shaded cells) for a 10-product problem instance at various refund fraction values.

Product, i	a_i	Parameter	Value
1	6.00	λ	50
2	5.44	p	4
3	4.89	e	3.8
4	4.33	c	3.6
5	3.78	v_n	3.4
6	3.22	v_o	3.2
7	2.67	d	0.2
8	2.11	μ_1	1
9	1.56	μ_2	0.5
10	1.00	σ	1
		β	0.5
		f	0.1

		α									
		i	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00
MTO	1										
	2										
	3										
	4										
	5										
	6										
	7										
	8										
	9										
	10										
MTS	1										
	2										
	3										
	4										
	5										
	6										
	7										
	8										
	9										
	10										

As Table 2 illustrates, the results in Theorems 1b and 2b are tight. There exist problem instances in the strict region for which the optimal assortment is composed of most eccentric products only, some most eccentric and some most popular products, or most popular products only ($\{1, 2, 3, 4\}$ and $\{1, 2\}$ are optimal under MTO and MTS, respectively, for $\alpha = 0.79 < 0.80 = v_o/p$).

To exemplify how the greedy algorithm fails to find the optimal assortment, take the MTO case with $\alpha = 0.70$. Starting with an empty set, the greedy algorithm first adds product 1 for an expected profit of 17.54 (adding product 10 instead would yield 7.78). The second step adds product 2 (the remaining most popular item), and the next five steps add products 10, 9, 8, 7 and 6. Then the algorithm stops, as adding one more product would not increase the expected profit. The resulting set $\{1, 2, 6, 7, 8, 9, 10\}$ is a local optimum with an expected profit of 18.29, lower than 19.64 obtained from the globally optimal solution (seven most eccentric products) reported in Table 2.

5. Insights and Discussion

In this section we provide managerial insights and robustness checks by a mix of analytical and computational means. All numerical observations reported in this section appear to be robust; equivalent experiments with different sets of parameters yield qualitatively similar results.

5.1. 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. In fact, in MTO environment with negligible variety cost, we can provide a mild sufficient condition for a more lenient return policy to always result in a reduction in optimal variety.

PROPOSITION 2. *Suppose the retailer is operating under MTO environment with negligible fixed cost for variety ($f \approx 0$) 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 $\alpha \geq v_o/p$, the cardinality of the optimal assortment $|S_{MTO}^*|$ is decreasing in refund fraction α .*

In order to probe this further, we plot in Figure 3 the cardinality of the optimal assortment under both operational environments as a function of α (ranging from 0 to 1 with 0.01 increments). Base parameter values in Table 2 apply except that we set $f = 0$ to better isolate the effect of α . Numbers in brackets next to each decimal data point show the composition of the optimal assortment, e.g., [2,3] indicates the assortment with two most popular and three most eccentric products.

Higher refund fraction leads to less variety for a sufficiently high α . 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 the highest and lowest α values within the strict return policy region ($0 \leq \alpha \leq v_o/p$). As α goes from 0 to v_o/p , there is first a decrease and then an increase in MTO variety. Whereas for MTS, variety increases with α in most of the strict policy region.

The managerial take-away from this experiment is that more lenient return policies *can* require assortments with higher variety. This effect occurs when the refund fraction is below a critical level (v_o/p). The typical reason is the incentive to offer eccentric products under strict return policies. For product categories with strong secondary markets ($v_o \approx p$) then, this insight is more salient.

5.2. What if variety was free? (The case of negligible fixed cost for variety)

If the fixed cost for variety is zero, or negligible, our analysis produces a sharper prescription on which products to offer under MTO. Further, this prescription differs from the one for MTS, so the operational environment moderates the fundamental impact of returns on optimal assortment.

PROPOSITION 3. *Suppose the retailer incurs no fixed cost for variety, i.e., $f = 0$.*

(a) *For a sufficiently lenient return policy ($\alpha \geq v_o/p$), the optimal assortment under both MTO and MTS environments is composed of some number of most popular products from N , i.e., $S_{MTO}^* = \mathcal{A}_k$ and $S_{MTS}^* = \mathcal{A}_{k'}$ for some $k, k' \in \{0, 1, \dots, n\}$.*

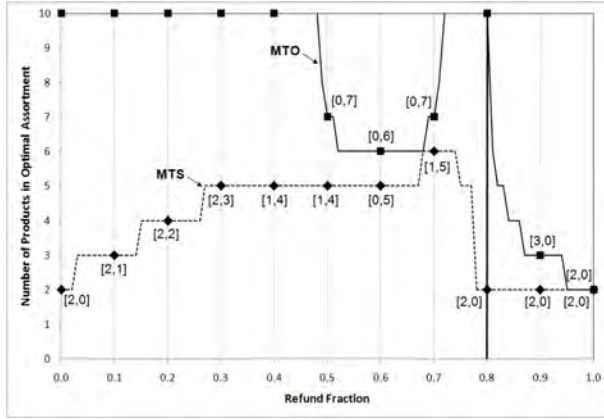


Figure 3 Optimal assortment versus refund fraction (α) in MTO and MTS environments with zero fixed cost per product ($f = 0$)




Figure 4 Optimal expected profit versus refund fraction (α) in MTO and MTS environments (NR indicates no-returns)

(b) For a sufficiently strict return policy ($\alpha < v_o/p$), the optimal assortment under MTO is composed of some number of most eccentric products from N , i.e., $S_{MTO}^* = \mathcal{Z}_j$ for some $j \in \{0, 1, \dots, n\}$. Whereas, under MTS, it is optimal to carry some number of most popular and some number of most eccentric products from N , i.e., $S_{MTS}^* = \mathcal{A}_j \cup \mathcal{Z}_{k-j}$ for some $j \in \{0, \dots, k\}$ and $k \in \{0, 1, \dots, n\}$.

We omit the proof as it exactly mirrors our previous analyses. Proposition 3 establishes that a retailer's optimal assortment critically depends on the operational environment when variety is free and the return policy is strict. In particular, the retailer carries only the most eccentric products in the MTO case, but a mix of popular and eccentric products in the MTS case (see Figure 3).

Note that, even when variety is free under MTO, it is not necessarily optimal for the retailer to offer all products in N . Returns are the central reason. The more popular a product, the less the relative rate of returns. This makes expected profit margins unequal despite uniform prices, e.g., if the return policy is lenient, more popular products have higher expected profit margins. Thus, fearing cannibalization, the retailer may not offer all potential products in N even if there is no fixed cost per product. Secondly, the eccentric-product result under MTO must be interpreted with caution. One of the assumptions we make is that all returned items and excess inventories can be cleared at a certain salvage value in the secondary market. This puts a limitation on the kind and breadth of products we can admit in N . Proposition 3b is not to suggest that a retailer operating in MTO environment must offer the most objectionable products it can find. Rather, it recommends eccentricity in retail assortments within the bounds of our modeling assumptions.

5.3. Robustness of our main finding

We now demonstrate the robustness of our main insight – sufficiently strict return policies call for hybrid popular-eccentric assortments – to three factors that may possibly work against it.

5.3.1. What if the retailer optimized the refund fraction? We present numerical examples for which the optimal refund fraction falls in the strict policy region, i.e., $\alpha^*p \leq v_o$. Using the base parameter values (Table 2), we plot in Figure 4 the optimal expected profit for all practically relevant values of refund fraction, from 0% to 100% with 1% increments. In the process, we take advantage of our structural results in §4 for optimizing the assortment at each data point. The peak of each curve represents the globally optimal refund fraction for the corresponding operational environment. We also include the no-returns cases (denoted by NR) as benchmark. Figure 4 suggests that our strict return policy results are not an artifact of refund fraction being exogenous. (Anonymous (2009) present a similar example with a different parameter set and numerically explore other interesting aspects of endogenizing the refund fraction.)

5.3.2. What if salvage values were quantity-dependent? During clearance sales, retailers often charge lower prices for items with higher inventory. Such quantity-dependent salvage values would weaken the rationale for carrying eccentric products, because while they are more likely to be returned, they may also be less valuable once returned. We take our single-period MTO scenario and incorporate a salvage value $v_i(R_i)$ for each product that depends on the realized number of returns. For simplicity, we model the demand for returned products during clearance using an additive demand model of the form $D_i(v_i) = \varphi - \psi v_i$ ($\varphi > 0$, $\psi > 0$) (*à la* Petruzzi and Dada 1999), and assume that the retailer clears all returned items by setting $v_i = (\varphi - R_i)/\psi$. The expected profit using the new salvage value is $\Pi_{MTO}(S) = \sum_{j \in S} E \left[(p - c)D_j - \left(\alpha p - \frac{\varphi - R_j}{\psi} \right) R_j \right] - f|S|$, or

$$\sum_{j \in S} \left\{ \left[p - c - \left(\alpha p - \frac{\varphi}{\psi} \right) P_{return|j} \right] \lambda P_j^S - \frac{1}{\psi} \left[(\lambda P_{j,return}^S)^2 + \sigma^2 (\lambda P_{j,return}^S)^{2\beta} \right] \right\} - f|S|$$

The optimal assortment is analytically intractable, so we proceed by numerical means. Using the same base parameter values (Table 2), and setting $\varphi = 18$ and $\psi = 5$, which ensure positive salvage values strictly less than $c = 3.6$, we compute the optimal assortment by complete enumeration for different values of α . We observe that the optimal assortment is most-eccentric only for $\alpha = 0.6$, a mix of most popular and most eccentric for $\alpha = 0.7$, and most-popular only for $\alpha = 0.8$ and above. This numerical example (and many others we experimented with using different sets of parameters)

suggest that the structure of the optimal assortment coincides with our main result: a sufficiently strict return policy results in a hybrid popular-eccentric assortment.

5.3.3. What if consumers could resell instead of return? One valid criticism of our choice model is that unhappy consumers are captive: if they are not happy with a product, they have no option but to return it to the retailer. All of our structural results extend to the case with consumers having an option to resell to a third party, say to another consumer through eBay.

The new choice model functions as follows. Besides *keep* and *return*, the second stage of the N-MNL model has a third option, *sell*, with utility $u_{i,sell} = v_s - d_s - p + \epsilon_{i,sell}$. The deterministic portion is the revenue v_s that the consumer obtains by reselling, minus the cost or disutility d_s of reselling, minus the original selling price p that the consumer foregoes. The random component $\epsilon_{i,sell}$ is again a Gumbel random variable with mean zero and scale $1/\mu_2$, independent of $\epsilon_{i,keep}$ and $\epsilon_{i,return}$. Reselling gives consumers an extra post-purchase option, which increases the expected utility of each product: $A_i \equiv E[\max(u_{i,keep}, u_{i,return}, u_{i,sell})] = \mu_2 \ln \left[\exp\left(\frac{a_i}{\mu_2}\right) + \exp\left(\frac{\alpha p - d}{\mu_2}\right) + \exp\left(\frac{v_s - d_s}{\mu_2}\right) \right] - p$. Possibility of reselling diminishes both the probability of returning and keeping. The new probabilities are: $P_{return|i} = \left[1 + \exp\left(\frac{a_i + d - \alpha p}{\mu_2}\right) + \exp\left(\frac{v_s - d_s + d - \alpha p}{\mu_2}\right) \right]^{-1}$, $P_{sell|i} = \left[1 + \exp\left(\frac{\alpha p - d - v_s + d_s}{\mu_2}\right) + \exp\left(\frac{a_i - v_s + d_s}{\mu_2}\right) \right]^{-1}$, and $P_{keep|i} = 1 - P_{return|i} - P_{sell|i}$.

Despite these changes our structural results apply identically. In particular, our counterintuitive result for strict return policies continue to hold. The reason is quite intuitive. Reselling option enhances all products; their expected utility increases by a positive constant. Although this constant may differ among products (due to the logarithmic expression in the expected utility), the popularity ranking of products remains unaffected. In other words, these ‘product improvements’ are technically equivalent to higher attractiveness values for each product.

6. Concluding Remarks

We believe that this paper highlights an interesting interaction between product assortment and return policy. The optimal assortment is composed of the most popular products if the return policy is lenient (if it imposes a sufficiently low restocking fee, meaning a high refund fraction). On the other hand, the optimal assortment has a distinct and counterintuitive structure if the return policy is sufficiently strict: a mix of the most popular and the most eccentric products. These results hold for pure MTO and pure MTS environments; they can also be extended to a hybrid environment where the firm is able to choose between MTO and MTS for every product that it

offers (we omit this extension). When the fixed cost for variety is negligible, MTO environment differs from MTS in the strict return policy case; the optimal assortment is then composed of most eccentric products only. Our results underscore the need to consider consumer returns and operational environment in making retail assortment decisions.

Evident from the many unexpected implications we observe of returns on optimal assortment, endogenizing the refund fraction (α) would be a useful extension. Likewise, endogenizing the price (p) could also reveal some interesting insights. We note that either of these directions is likely to bolster the counterintuitive results we obtain for strict return policies. Lower refund fractions are likely to require lower prices, which together make the condition $\alpha < v_o/p$ more likely to be satisfied. Anonymous (2009) provide numerical experiments that indeed support this idea.

Another important direction could involve multiple periods, allowing for a more comprehensive treatment of inventory control and return management issues. For a finite planning horizon, optimizing the assortment appears analytically intractable even if the assortment decision was static (i.e., made just once prior to the beginning of the planning horizon). However, in multiple-period settings, inclusion of eccentric products in the optimal assortment would be further justified as there would be multiple resale opportunities for returned items. So multiple periods can only reinforce our result on strict return policies, a point supported numerically by Anonymous (2009).

We close with three implications of our theory that seem empirically testable:

- *H1: Stricter return policies in the form of higher restocking fees imply a larger propensity of retail assortments including eccentric products.* This follows from Theorems 1 and 2, and holds for both operational environments.
- *H2: If a retailer allows returns and actively considers them when taking merchandising decisions, then the retailer is more likely to offer eccentric products.* Ignoring returns in assortment decisions (extant literature), or disallowing them altogether (Theorem 3), results in optimal assortments with popular products only. In contrast, by Theorems 1 and 2, offering a mix of popular and eccentric products can be optimal when returns are taken into account.
- *H3: If offering variety is cheap, in that it imposes on retailers a negligibly small fixed cost per product, then H1 holds more strongly in MTO environments than in MTS environments.* This is due to Proposition 3; for sufficiently strict return policies, it is optimal to offer only eccentric products under MTO, but a mix of popular and eccentric products under MTS.

Acknowledgments

We benefitted from insightful discussions with and invaluable feedback from Elif Akçalı of University of Florida and numerous colleagues at the ITOM and Marketing departments of SMU Cox School of Business. Also, we would like to express our deep gratitude to Sena Cases for sharing their data and fielding our endless questions on their assortment planning and returns management processes.

Appendix A: Empirical Evidence on How Returns are Linked to Popularity

Here we provide empirical evidence for the following aspect of our choice model: eccentric products have a higher conditional probability of return, i.e., $P_{return|i}$ is decreasing in a_i . We obtained data from online channels of two retailers on product lines that approximately fit our research context.

The first dataset is from a major department store. It contains SKU-level monthly sales and returns (in units and in dollars, June-November, 2010) for a bath towel in 21 different colors. Except for their color all SKUs were identical, including prices (any change in prices applied to all SKUs uniformly), and the assortment stayed the same. Our base analysis was to aggregate the sales and returns over 6 months for each SKU, calculate their relative sales volume (units sold as % of total sales) and return rate (units returned as % of units sold), and then run a nonlinear regression to estimate return rates from relative sales volumes. Our theoretical prediction based on the N-MNL model is that the return rate must be a power function of the relative sales volume. In particular, the product of conditional return probability and a power of normalized purchase probability is a constant, i.e., $P_{return|i} \left[\omega_i / (\sum_{j \in S} \omega_j) \right]^B = C$ for all $i \in S$, where B and C are positive constants for a given assortment S . It would be consistent with the N-MNL model then to observe a relationship of the sort: $\log y = -B \log x + \log C$, with y representing the return rate, and x the relative sales volume. This is indeed born by the data; significance level of the nonlinear regression model is 3.58%. Figure 1 shows a scatter plot and the power function that best fits the data.

The second dataset is from Sena Cases, a maker and retailer of leather cases for mobile devices. It contains transaction-level data for a popular case purposely designed for iPhone 3G/3GS, called *Ultraslim*. The current Ultraslim offering has 18 colors (as of January 22, 2011), available since March 11, 2009. As with the previous analysis, we first aggregated the data for each SKU. We considered the sales transactions from March 11, 2009 to November 30, 2010, and the corresponding returns, if any, up until December 31, 2010 (this allows enough time for returns to occur, as Sena has a 30-day limit). We then conducted a similar regression analysis; the theoretical model explains the data at 2.33% significance level. Figure 2 shows a scatter plot of the data and the power function that best fits it.

The results provide sufficient evidence for product popularity and return rates to be inversely related: eccentric products, which capture a lower market share, tend to have a higher return rate. One possible complication is the effect of variation in prices over time on return rates. In both cases, however, controlling for price does not change the basic conclusion.

Appendix B: Proofs

[Throughout the appendix we 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 $\alpha \geq v_o/p$, and quasiconvex when $\alpha < v_o/p$. We use the following result from Mangasarian (*Nonlinear Programming*, 1969. McGraw-Hill, New York, p.148): *The function $h(\cdot) = \frac{g(\cdot)}{\gamma(\cdot)}$ is quasiconcave (quasiconvex) on a set $\Gamma \subset \mathbb{R}^n$ if $g(\cdot)$ is concave (convex) on Γ , $\gamma(\cdot) > 0$ on Γ , and $\gamma(\cdot)$ is linear on \mathbb{R}^n .*

We show concavity (convexity) of $g(\delta)$ by examining its second derivative: $g''(\delta) = -\lambda(\alpha p - v_o) [2P'_{r|\delta} + \delta P''_{r|\delta}]$ 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 simplifies to $2P'_{r|\delta} + \delta P''_{r|\delta} = \frac{\mu_1}{\mu_2} \left(\frac{\mu_1}{\mu_2} - 1 \right) P_{r|\delta} \delta^{-1} > 0$ (since $\mu_1 \geq \mu_2$). We can therefore state that $g(\delta)$ is concave when $\alpha \geq v_o/p$, and convex when $\alpha < v_o/p$. Since $\gamma(\delta)$ is strictly positive and linear in δ , it follows that $h_{MTO}(\delta)$ is quasiconcave when $\alpha \geq v_o/p$, and quasiconvex when $\alpha < v_o/p$.

Next, we show that $h_{MTO}(\delta)$ is non-decreasing when $\alpha \geq v_o/p$. To do so, we evaluate the first derivative of $h_{MTO}(\delta)$. For convenience, let $\widehat{S} = S \cup \{\delta\}$.

$$h'_{MTO}(\delta) = -\lambda P_0^{\widehat{S}} \sum_{j \in S} M_j P_j^{\widehat{S}} + \lambda P_0^{\widehat{S}} M_\delta \left(\sum_{j \in S} P_j^{\widehat{S}} + P_0^{\widehat{S}} \right) - \lambda(\alpha p - v_o) P_\delta^{\widehat{S}} P'_{r|\delta} \quad (2)$$

It is easy to see that $M_\delta = p - c - (\alpha p - v_o) P_{r|\delta}$ is increasing in δ since $P_{r|\delta}$ is decreasing in δ . Therefore, there must exist some $\widehat{\delta}$ such that $M_{\widehat{\delta}} \geq M_j$, $\forall j \in S$. Then, for all $\delta \geq \widehat{\delta}$, the first term in $h'_{MTO}(\delta)$ is always less than the second term, and their sum is therefore positive. The last term is always nonnegative since $P'_{r|\delta} < 0$. Hence, we can conclude that in the interval $[\widehat{\delta}, \infty)$, $h'_{MTO}(\delta)$ is positive, or $h_{MTO}(\delta)$ is increasing when $\alpha \geq v_o/p$.

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, $h_{MTO}(\delta_L) > h_{MTO}(\widehat{\delta})$, and $h_{MTO}(\widehat{\delta}) < h_{MTO}(\delta_H)$.

Furthermore, by the definition of quasiconcavity: $h_{MTO}(\Delta \delta_L + (1 - \Delta) \delta_H) \geq \min\{h_{MTO}(\delta_L), h_{MTO}(\delta_H)\}$, with $\Delta \in [0, 1]$. Since $\delta_L < \widehat{\delta} < \delta_H$, there exists $\widehat{\Delta} \in [0, 1]$ such that $\widehat{\Delta} \delta_L + (1 - \widehat{\Delta}) \delta_H = \widehat{\delta}$. Using the definition of

quasiconcavity for $\Delta = \widehat{\Delta}$, we obtain $h_{MTO}(\widehat{\delta}) \geq \min\{h_{MTO}(\delta_L), h_{MTO}(\delta_H)\}$, which contradicts $h_{MTO}(\delta_L) > h_{MTO}(\widehat{\delta})$, or $h_{MTO}(\widehat{\delta}) < h_{MTO}(\delta_H)$. Hence, when $\alpha \geq v_o/p$, we conclude that $h_{MTO}(\delta)$ is non-decreasing for $\delta < \widehat{\delta}$ as well.

Proof of Lemma 2

Product $i \in N \setminus S$ should be added to the current assortment S iff $\Pi_{MTO}(\widehat{S}) \geq \Pi_{MTO}(S)$, where $\widehat{S} = S \cup \{i\}$, or $\lambda \left(P_0^S - P_0^{\widehat{S}} \right) M_i \geq \sum_{j \in S} \lambda \left(P_j^S - P_j^{\widehat{S}} \right) (M_j - M_i) + f$. 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 plus the fixed cost. Using the N-MNL purchase probabilities, we rewrite it as follows:

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

The result follows from this inequality by further algebraic manipulations.

Proof of Theorem 1

Part (a), the lenient return policy case. The proof is by construction. Suppose the optimal assortment S_{MTO}^* has cardinality k , with $k \in \{1, \dots, n-1\}$. (The theorem holds trivially for $|S_{MTO}^*| = 0$ and $|S_{MTO}^*| = n$.) Take *any* subset S of N with cardinality k . Let $n_a = \max\{i \mid \mathcal{A}_i \subseteq S, i \in \{0, 1, \dots, k\}\}$ 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 $j \in S$ such that $j > n_a + 1$. Since $h_{MTO}(\delta)$ is non-decreasing due to Lemma 1, product j 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 $S_{MTO}^* = \mathcal{A}_k$.

Part (b), the strict return policy case. The proof is by construction. Suppose the optimal assortment S_{MTO}^* has cardinality $k \in \{1, \dots, n-2\}$. (The theorem holds trivially for $|S_{MTO}^*| = 0$, $|S_{MTO}^*| = n-1$ and $|S_{MTO}^*| = n$.) Take *any* subset S of N with cardinality k . Let $n_a = \max\{i \mid \mathcal{A}_i \subseteq S, i \in \{0, 1, \dots, k\}\}$ be the number of most popular products of N that belong to S ; and $n_z = \max\{j \mid \mathcal{Z}_j \subseteq S, j \in \{0, 1, \dots, k\}\}$ 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 $h_{MTO}(\delta)$ 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 $S_{MTO}^* = \mathcal{A}_j \cup \mathcal{Z}_{k-j}$ for some $j \in \{0, 1, \dots, k\}$.

Proof of Lemma 3

Product $i \in N \setminus S$ should be added to the current assortment S iff $\Pi_{MTS}(\widehat{S}) \geq \Pi_{MTS}(S)$, where $\widehat{S} = S \cup \{i\}$, or $\sum_{j \in S} \lambda P_j^{\widehat{S}} \widetilde{M}_j^{\widehat{S}} + \lambda P_i^{\widehat{S}} \widetilde{M}_i^{\widehat{S}} - f(|S| + 1) \geq \sum_{j \in S} \lambda P_j^S \widetilde{M}_j^S - f|S|$. Rearranging the terms and dividing by λ , we get $P_i^{\widehat{S}} \widetilde{M}_i^{\widehat{S}} \geq \sum_{j \in S} \left[P_j^S \widetilde{M}_j^S - P_j^{\widehat{S}} \widetilde{M}_j^{\widehat{S}} \right] + f/\lambda$. This is equivalent to $P_i^{\widehat{S}} \widetilde{M}_i^{\widehat{S}} \geq \sum_{j \in S} \widetilde{M}_j^{\widehat{S}} \left[P_j^S - P_j^{\widehat{S}} \right] + \sum_{j \in S} \left[\widetilde{M}_j^S - \widetilde{M}_j^{\widehat{S}} \right] P_j^S + f/\lambda$. Substituting the purchase probabilities into the first term of the right-hand-side, and dividing both sides by $P_i^{\widehat{S}}$, we obtain the desired inequality.

Proof of Lemma 4

We begin by noting that the last two terms of the condition given in Lemma 3 are strictly positive ($\widetilde{M}_j^S > \widetilde{M}_j^{\widehat{S}}$), because $P_j^S > P_j^{\widehat{S}}$ and $(P_j^S)^{\beta-1} < (P_j^{\widehat{S}})^{\beta-1}$ for all $j \in S$. Therefore, if product i improves the expected profit when added to S , it must satisfy $\widetilde{M}_i^{\widehat{S}} > \sum_{j \in S} P_j^S \widetilde{M}_j^{\widehat{S}}$. Now, multiplying the left-hand-side of this inequality by $\left(\sum_{j \in S} P_j^{\widehat{S}} + P_0^{\widehat{S}} \right)$ and the right-hand-side by an equivalent term $\left(\frac{1 + \sum_{j \in S} \omega_j}{1 + \omega_i + \sum_{j \in S} \omega_j} \right)$, and applying further algebraic manipulations, we obtain the result.

Proof of Lemma 5

We first show that $h_{MTS}(\delta)$ is quasiconvex when $\alpha < v_o/p$ using the same result due to Mangasarian (1969) stated in the first part of the proof of Lemma 1. We already know from that proof that $g(\delta)$ is convex when $\alpha < v_o/p$. We need to show that $\tilde{g}(\delta)$ is also convex when $\alpha < v_o/p$. Taking the first and second derivatives, we verify below that this is indeed the case for any α .

$$\begin{aligned} \tilde{g}'(\delta) &= -(e - v_n) \sigma \lambda^\beta \phi(z^*) \left[\beta \delta^{\beta-1} \gamma(\delta)^{1-\beta} + \left(\sum_{j \in S} \omega_j^\beta + \delta^\beta \right) (1 - \beta) \gamma(\delta)^{-\beta} \right] \\ \tilde{g}''(\delta) &= (e - v_n) \frac{\beta(1 - \beta) \sigma \lambda^\beta \phi(z^*)}{\gamma(\delta)^\beta} \left[\frac{\delta^{\beta-2} \left(1 + \sum_{j \in S} \omega_j \right)^2 + \sum_{j \in S} \omega_j^\beta}{\gamma(\delta)} \right] > 0 \end{aligned}$$

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

Next, we show that $h_{MTS}(\delta)$ is increasing when $\alpha \geq v_o/p$ as long as δ is profit-improving. To do so, we let $\widehat{S} = S \cup \{\delta\}$, and evaluate the derivative of $h_{MTS}(\delta)$:

$$\begin{aligned} h'_{MTS}(\delta) &= -\lambda(\alpha p - v_o) P_\delta^{\widehat{S}} P'_{r|\delta} - \lambda P_0^{\widehat{S}} \sum_{j \in S} \left[p - c - (\alpha p - v_o) P_{r|j} - (e - v_n) \sigma \phi(z^*) \beta \left(\lambda P_j^{\widehat{S}} \right)^{\beta-1} \right] P_j^{\widehat{S}} \\ &\quad + \lambda P_0^{\widehat{S}} \left[p - c - (\alpha p - v_o) P_{r|\delta} - (e - v_n) \sigma \phi(z^*) \beta \left(\lambda P_\delta^{\widehat{S}} \right)^{\beta-1} \right] \left(\sum_{j \in S} P_j^{\widehat{S}} + P_0^{\widehat{S}} \right) \end{aligned}$$

Note that the terms in square brackets are similar to the expected profit margin per unit sales, except for the value of β multiplying the inventory cost term. Using the identity $\widetilde{M}_i^{\widehat{S}} = M_i - (e - v_n) \sigma \phi(z^*) \left(\lambda P_i^{\widehat{S}} \right)^{\beta-1}$ and applying a few algebraic manipulations:

$$\begin{aligned}
h'_{MTS}(\delta) = & \lambda\beta \left[-(\alpha p - v_o)P_{\delta}^{\hat{S}}P'_{r|\delta} - P_0^{\hat{S}} \sum_{j \in S} \widetilde{M}_j^{\hat{S}}P_j^{\hat{S}} + P_0^{\hat{S}}\widetilde{M}_{\delta}^{\hat{S}} \left(\sum_{j \in S} P_j^{\hat{S}} + P_0^{\hat{S}} \right) \right] \\
& + \lambda(1-\beta) \left[-(\alpha p - v_o)P_{\delta}^{\hat{S}}P'_{r|\delta} - P_0^{\hat{S}} \sum_{j \in S} M_jP_j^{\hat{S}} + P_0^{\hat{S}}M_{\delta} \left(\sum_{j \in S} P_j^{\hat{S}} + P_0^{\hat{S}} \right) \right]
\end{aligned}$$

Using the derivative of h_{MTO} given in (2) and rearranging the terms, $h'_{MTS}(\delta)$ is:

$$h'_{MTS}(\delta) = \lambda\beta P_0^{\hat{S}} \left[\widetilde{M}_{\delta}^{\hat{S}} \left(\sum_{j \in S} P_j^{\hat{S}} + P_0^{\hat{S}} \right) - \sum_{j \in S} \widetilde{M}_j^{\hat{S}}P_j^{\hat{S}} \right] - \lambda\beta(\alpha p - v_o)P_{\delta}^{\hat{S}}P'_{r|\delta} + (1-\beta)h'_{MTO}(\delta)$$

We know from Lemma 1 that $h'_{MTO}(\delta)$ is non-negative. The second term is also non-negative since $P'_{r|\delta} < 0$ and $\alpha \geq v_o/p$. Assuming that the new product with preference δ is profit-improving, thus using the necessary condition stated in Lemma 4, the remaining term is also positive.

Proof of Theorem 2

Part (a), the lenient return policy case. The proof is by construction and analogous to the proof of Theorem 1a. We only comment on the technical assumption of Lemma 5 that requires the new product to be profit-improving for $h_{MTS}(\delta)$ to be increasing. If no such product exists, then the current assortment must be optimal. Else, it is easy to show that the product with maximum attractiveness (hence, the maximum value of δ) also improves the expected profit (in fact, maximally so). Assume product $i \in N \setminus S$ with preference $\delta = \omega_i$ improves the profit when added. The derivative of $h_{MTS}(\delta)$ evaluated at ω_i is positive (by Lemma 5). Then a product with preference $\delta = \omega_i + \Delta$, for any $\Delta > 0$, would also be profit-improving because $h_{MTS}(\omega_i + \Delta) > h_{MTS}(\omega_i) > \Pi_{MTS}(S)$. We could repeat the same reasoning and keep increasing Δ until $\omega_i + \Delta = \max_{j \in N \setminus S}(\omega_j)$.

Part (b), the strict return policy case. Since the proof is analogous to the proof of Theorem 1b, we only give a sketch of the argument. Due to the quasiconvexity of $h_{MTS}(\delta)$ when $\alpha < v_o/p$, a fact provided by Lemma 5, any assortment that does not conform to the structure stated in the theorem can be improved by pairwise interchanges, and therefore cannot be optimal.

Proof of Proposition 1

Take an MTO environment with a lenient return policy, and suppose the most profitable assortment S has cardinality $k < n$. By Theorem 1, $S = \mathcal{A}_k$. So, for the greedy algorithm to produce the globally optimal assortment, we need to establish that once the most popular of the remaining products fails to satisfy the necessary condition in Lemma 2, no more products can be added and increase the expected profit. Assume that product $k+1$ violates the necessary condition in Lemma 2, and therefore cannot be included

in S to improve the profit. By Lemma 1, no product i with $k+1 \leq i \leq n$ improves the profit either, i.e., $M_i < \sum_{j \in S} P_j^S M_j + f / (\lambda P_i^{S \cup \{i\}})$. Is it possible then to add two or more products at once and increase the profit? Let $B \subseteq N \setminus S$ be this subset of products to add. If $S \cup B$ improves the profit, then

$$\begin{aligned} \sum_{j \in S} \lambda P_j^{S \cup B} M_j - f|S| + \sum_{k \in B} \lambda P_k^{S \cup B} M_k - f|B| &> \sum_{j \in S} \lambda P_j^S M_j - f|S| \\ \sum_{k \in B} \lambda P_k^{S \cup B} M_k - f|B| &> \sum_{j \in S} \lambda (P_j^S - P_j^{S \cup B}) M_j = \left(\sum_{k \in B} P_k^{S \cup B} \right) \sum_{j \in S} \lambda P_j^S M_j \end{aligned}$$

Let i be the most popular product in B . The above inequality implies $M_i > \sum_{j \in S} P_j^S M_j + f|B| / (\lambda \sum_{k \in B} P_k^{S \cup B})$. Also considering $\sum_{k \in B} P_k^{S \cup B} \leq |B| P_i^{S \cup B}$ and $P_i^{S \cup \{i\}} > P_i^{S \cup B}$, we obtain

$$M_i > \sum_{j \in S} P_j^S M_j + \frac{f|B|}{\lambda \sum_{k \in B} P_k^{S \cup B}} \geq \sum_{j \in S} P_j^S M_j + \frac{f}{\lambda P_i^{S \cup B}} > \sum_{j \in S} P_j^S M_j + \frac{f}{\lambda P_i^{S \cup \{i\}}}$$

which contradicts the initial premise that product i could not be added to S and improve the profit. MTO environment with no returns is a special case with $d = +\infty$.

Proof of Proposition 2

By Theorem 1a, $S_{MTO}^* = \mathcal{A}_k$ for some $k \in \{0, 1, \dots, n\}$. In addition, by Lemma 2, when $f = 0$ the firm adds product $(i+1)$ to an existing assortment $S = \mathcal{A}_i$ iff $\left[1 + \sum_{j \in S} \omega_j \right] M_{i+1} \geq \sum_{j \in S} \omega_j M_j$. Note that the expected profit margin M_j is decreasing and preference ω_j is increasing in α . Therefore, if M_{i+1} decreases at least as fast as M_j for all $j \in S$ as α increases, the left-hand-side will be relatively smaller compared to the right-hand-side. Equivalently: the above inequality will be less likely to be satisfied; a more lenient return policy (higher α) within the lenient region ($\alpha \geq v_o/p$) will lead to lower variety. 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 $j \in S$. We have

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

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

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

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 $a_{i+1} \geq \alpha p - d$. 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.

References

- Alptekinoğlu, A., C. J. Corbett. 2010. Leadtime - variety tradeoff in product differentiation. *Manufacturing Service Oper. Management* **12**(4) 569–582.
- Anderson, S. P., A. de Palma, J. F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
- Anonymous. 2009. Is assortment selection a popularity contest? A study of assortment, return policy, and pricing decisions of a retailer. S. Netessine, C. Tang, eds., *Consumer-Driven Demand and Operations Management Models*. Springer, New York, NY, 205–228.
- Anonymous. 2010. Customer returns in the retail industry 2010. White paper, The Retail Equation and National Retail Federation, available at <http://www.theretailequation.com/>.
- Aydin, G., W. H. Hausman. 2009. The role of slotting fees in the coordination of assortment decisions. *Production Oper. Management* **18**(6) 635–652.
- 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.
- 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. 2008. On the effects of consumer search and firm entry in a multiproduct competitive market. *Marketing Sci.* **27**(3) 461–473.
- Cargille, B., C. Fry, A. Raphael. 2005. Managing product line complexity. *OR/MS Today* **32**(3) 34–41.
- Caro, F., V. Martinez-de-Albeniz. 2012. Product and price competition with satiation effects. Forthcoming in *Management Science*.
- 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.

- Hindo, B. 2007. What happens to that scarf that you really hated. *Business Week* (Jan 15, 2007).
- Honhon, D., V. Gaur, S. Seshadri. 2010. Assortment planning and inventory decisions under stockout-based substitution. *Oper. Res.* **58**(5) 1364–1379.
- Hopp, W. J., X. Xu. 2005. Product line selection and pricing with modularity in design. *Manufacturing Service Oper. Management* **7**(3) 172–187.
- Kök, A. G., M. L. Fisher, R. Vaidyanathan. 2009. Assortment planning: Review of literature and industry practice. N. Agrawal, S. Smith, eds., *Retail Supply Chain Management*. Springer, New York, NY, 99–153.
- Kök, A. G., Y. Xu. 2011. Optimal and competitive assortments with endogenous pricing under hierarchical consumer choice models. *Management Sci.* **57**(9) 1546–1563.
- 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.
- Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Sci.* **14**(4) 442–466.
- Olavson, T., C. Fry. 2006. Understanding the dynamics of value-driven variety management. *MIT Sloan Management Rev.* **48**(1) 63–69.
- Pasternack, B. A. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Sci.* **4**(2) 166–176.
- Petruzzi, N. C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* **47**(2) 183–194.
- Ramdas, K. 2003. Managing product variety: An integrative review and research directions. *Production Oper. Management* **12**(1) 79–101.
- Shlachter, B. 2010. Returns or ripoffs? It’s a fine line for retailers. *Star-Telegram* (Dec 10, 2010).
- Shulman, J. D., A. T. Coughlan, R. C. Savaskan. 2009. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. *Manufacturing Service Oper. Management* **11**(4) 577–594.
- 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. 2009. Consumer returns policies and supply chain performance. *Manufacturing Service Oper. Management* **11**(4) 595–612.
- Train, K. E. 2009. *Discrete Choice Methods with Simulation*. 2nd ed. Cambridge University Press, New York, NY.
- van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45**(11) 1496–1509.
- Xie, J., E. Gerstner. 2007. Service escape: Profiting from customer cancellations. *Marketing Sci.* **26**(1) 18–30.