

Chapter 11

Optimizing Retail Assortments for Diverse Customer Preferences

Stephen A. Smith

1 Introduction

Assortment selection is one of the most important and difficult decisions that retailers face. Assortments are typically chosen subjectively, often before any sales have been observed for some candidate products. Compared to pricing or advertising decisions, assortment decisions are more difficult to adjust later on. For multi-featured items such as consumer electronics and durable goods, the large number of product options, together with limited display space and financial constraints all contribute to the complexity of this decision. Consumer preferences for the various product attributes may also be heterogeneous, which requires assessing tradeoffs between the products that appeal to diverse customer segments. Because of these complexities, intuitively chosen retail assortments are likely to be suboptimal.

This paper develops an operational methodology for selecting optimal retail assortments based on an underlying multinomial logit (MNL) choice model for each customer's selection of product and retailer. A formulation is developed for optimizing the retailer's expected profit across customers with heterogeneous preferences. The formulation can also include a variety of additional merchandising constraints, such as display space, price point coverage or brand offerings.

Choice models have been successfully applied in consumer package goods to predict customers' response to assortment changes, based on observing repeat purchase behavior. The increased use of the Internet as a shopping guide for more complex, less frequently purchased products provides an opportunity to

S.A. Smith (✉)
Department of Operations Management and Information Systems,
Leavey School of Business, Santa Clara University, 500 El Camino Real,
Santa Clara, CA 95053, USA
e-mail: ssmith@scu.edu

obtain detailed preference information for broader classes of merchandise. A commercial data base of consumer preferences for attributes and features of DVD players, which was obtained through interactive Internet sessions, is used to illustrate the methodology. Consumer surveys or past buying behavior of individuals might also be used as alternative sources for the preference information needed for this assortment optimization methodology.

The methods in this paper provide a basis for several strategic retailer decisions including: (1) determining the optimal set of SKUs to offer and their estimated selling proportions; (2) how the retailer's relative market strength affects the contents of the optimal assortment; (3) how changing the contents of the assortment affects the probability that customers choose a given retailer and (4) how the customers' preference structure affects the optimal assortment and the corresponding expected profits. In analyzing our sample data set, it was found that accounting for preference heterogeneity and customers' use of consideration sets both had significant impacts on the retailer's expected profits.

1.1 Literature Review

Kok et al. (2015) provide a comprehensive survey of recent papers in retail assortment planning, and thus this paper's literature review will focus on a few papers that are particularly relevant for the optimization model developed here. Several recent papers have developed models for assortment optimization based on a newsvendor type model for inventory cost. van Ryzin and Mahajan (1999), Cachon and Gurhan Kok (2007) and Cachon et al. (2005) use a multinomial logit (MNL) model in which customers have homogeneous expected utilities. In Mahajan and van Ryzin (2001), customers are heterogeneous with regard to utility and their paper explicitly models the substituted demand that results from random stockouts of the retailer's inventory, but optimizing the assortment requires solution heuristics that are based on the set of possible inventory trajectories over the season. Guar and Honhon (2006) used a Lancaster type of model of substitution for products distributed along a single attribute dimension, and analyzed the impacts of static and dynamic substitution under this preference structure. Honhon et al. (2010) consider assortment optimization with stockout based substitution for more general deterministic preference structures. This leads to a dynamic programming formulation, for which they develop solution heuristics. Rusmevichientong and Topaloglu (2012) consider a generalized version of the MNL in which the model parameters are random, and show that the optimal assortments satisfy the nested set properties that hold for the MNL choice model with fixed parameters. Sauré and Zeevi (2013) develop a retail assortment optimization model that incorporates learning through experimentation with alternative assortments, and study the tradeoff between gaining information and maximizing current revenue. Smith and Agrawal (2000) used a probability of substitution matrix across products to optimize assortments in combination with an approximate

newsvendor inventory model. Miller et al. (2010) provide a method for optimizing assortments for infrequently purchased products and compare the results for several simple assortment selection heuristics. Kok and Fisher (2007) develop a heuristic for optimizing the allocation of shelf facings and inventory levels for a supermarket based on a particular substitution structure that also considers stockouts. Chong et al. (2001) developed a more general hierarchical market model for retail assortment planning for repeat purchase items, but due to the complexity of the resulting objective function, used a local improvement heuristic for optimization.

Only two of the above papers address the issue of retailer choice. Cachon et al. (2005) investigates how three different consumer models for the value of additional search at alternative retailers can affect the optimal assortment. Cachon and Gurhan Kok (2007) develop a more general category management model based on the retailer choice probabilities obtained from the nested logit model, but require mean utilities that are homogeneous across customers.

Product line optimization models have used mathematical programming formulations to solve a related problem. In this setting, a manufacturer decides which set of products to produce, where each potential product is viewed as a collection of adjustable product attributes. Chen and Hausman (2000) considered product line selection based on the MNL choice model, with homogenous customer preferences. Green and Krieger (1985), McBride and Zufryden (1988), Dobson and Kalish (1988, 1993) and Kohli and Sukumar (1990) consider heterogeneous customer utilities, but assume deterministic product choices. Green and Krieger treat discrete price options as product attributes, as is done in this paper, while Dobson and Kalish treat product prices as separate decision variables. With the exception of Chen and Hausman, these mathematical programming formulations are computationally difficult to solve, in part because they assume strict utility maximization by customers. Some product line selection papers developed solution heuristics (Kohli and Sukumar 1990; Dobson and Kalish 1993) or suggested clustering of customer preferences to reduce the problem size (Green and Krieger 1985) so that iterative search methods can be applied. These product line optimization methods do not model retailer choice, nor do they include inventory management costs.

1.2 *Summary of Results*

This paper provides an operational assortment optimization model that includes general heterogeneous consumer preferences as well as the customer's choice of retailer within the MNL framework. It is shown that the input parameters required for modeling product choice and retailer choice can be estimated separately, which facilitates their use in an operational model for assortment optimization. Assuming homogeneous mean utilities, van Ryzin and Mahajan (1999) showed that the optimal assortments form nested sets as the assortment size increases. Rusmevichientong and Topaloglu (2012) extend this result for the case of unknown MNL parameters. For heterogeneous customer utilities and competing retailers, this

chapter shows that this nested set property no longer holds, but that nested optimal assortment sets do occur for two limiting cases: (1) a monopoly retailer and (2) perfect competition among retailers. An optimization formulation is developed, which can include linear retailer constraints on the contents of the assortment, such as brand coverage and display space limitations. Finally, a commercial data base of preferences for DVD players is analyzed to illustrate the sensitivity of the expected profit and optimal assortment to the customer preference structure. The results for this data set illustrate the importance of including preference heterogeneity and customers' use of considerations sets in assortment optimization, as well as the sensitivity of the retailer's profit to assortment size.

2 Model Description

This paper focuses on the assortment decision for a particular retailer r , whose objective is to maximize the expected profit over a fixed time period, e.g., the Fall season. It is assumed that other retailers do not react competitively to this retailer's decisions. The retailer's assortment is defined by a binary vector $y = y_1, y_2, \dots, y_n$, where $y_j = 1$ if the retailer's assortment includes product j and 0 otherwise. Then let


$$D_j(y) = \text{the random demand for product } j,$$

which depends on y as well as other factors that affect demand. We now develop a choice model that determines the probability distribution for $D_j(y)$.

2.1 Modeling the Consumer's Purchase Decision

First, suppose that customers are classified according to n distinct customer types indexed by $i = 1, \dots, n$. It is assumed that customers of the same type assign the same expected values to various choice alternatives, but their actual purchase decisions also reflect individual random variations.

Actual purchases are the result of a sequential process that can be diagrammed as follows:



The choice decisions in each of these steps can be described in terms of the iPACE model for retail shopping decisions that has been developed in the marketing literature, where iPACE stands for information, Price, Assortment, Convenience and Entertainment, (see e.g., Hanson and Kalyanam 2006, Chap. 13) By

becoming an active shopper, the customer is sufficiently interested in the product category to gather information. Using a variety of sources, which may include both Internet research and store visits, customers assess their utilities for the available products and the relative values of purchasing from the alternative retailers. This process allows a customer to narrow the set of choices to a “consideration set” of products. The customer selects a retailer based on the retailer’s assortment, as well as the assessed convenience and entertainment values of shopping at that retailer. Finally, the customer makes a product selection from the choice set, which is defined as the intersection of the consideration set and the chosen retailer’s assortment. Although this description is sequential, these decisions do not necessarily need to be made in any specific order. For example, the customer might choose the most preferred product first, and then select the retailer from which to purchase. The key assumption is that the combination of the utility of the retailer and utility of the chosen product jointly determine the customer’s decision. This chapter assumes that these decisions are made normatively by customers, based on maximizing expected utility.

From the perspective of a particular retailer r , the customer may also choose the “no purchase” option for two reasons: (1) no product in the consideration set has positive net value, i.e., the choice set is empty or (2) the combined value of shopping and purchasing from this particular retailer’s assortment either does not exceed the product’s price, or is less than the combined value obtained from another retailer.

The mathematical models for each of these steps can be summarized as follows. The assortment decision is made for a fixed period of time, e.g., one season, and the time dependent parameters correspond to the length of this season. A random number N_i of customers of type i will become “active shoppers,” i.e., they will gather information and make a purchase decision this season for this product category. We assume that N_i is a Poisson random variable with rate parameter λ_i . For the N_i shoppers, define

$$q_{ij}(y) = P\{\text{customer type } i \text{ chooses product } j \text{ from this retailer} \mid \text{assortment } y\}.$$

This implies that $D_j(y)$, the random demand for product j defined previously, has a Poisson distribution with mean

$$\mu_j(y) = \sum_i \lambda_i q_{ij}(y). \quad (11.1)$$

The remaining customer decisions, which determine $q_{ij}(y)$, are based on the following utility model. The underlying choice model is a multinomial logit (MNL) in which customer i ’s combined utility for product j and retailer r is a random variable of the form

$$U_{ij}^r = U_{ij} + V_{ir} + \varepsilon_{ijr}, \quad (11.2)$$

where ε_{ijr} = Gumbel distributed error terms with mean 0 and scale parameter ξ_i ,
 U_{ij} = the expected (net) utility obtained from purchasing product j ,
 V_{ir} = the additional utility obtained by purchasing from retailer r .

For this paper's analysis, the product price is included in U_{ij} as a fixed attribute, rather than a decision variable. For many retailers, this is justified based on operational practice. At the individual product level, tactical pricing decisions such as temporary markdowns are typically made by the retailer later on during the selling season, as part of promotional and advertising activities. Strategic pricing decisions, such as how to price relative to competitors, are typically made less frequently and at a higher level than just one product category. For assortment planning purposes, the product price is therefore the estimated average price for the season. A combined model that simultaneously optimizes product prices and the retail assortment is conceptually superior to separate decision models, but it cannot feasibly include all the other aspects of customers' purchasing decisions that are analyzed here.

Additive MNL models of the form (11.2) are frequently used for two dimensional choice decisions. (See, e.g., Ben Akiva and Lerman 1985 for further discussion.) In the context of this application, the error terms ε_{ijr} can capture both the customer's imprecise knowledge of his or her own utilities, as well as the retailer's imperfect knowledge of customers' utilities. It is common practice to rescale the utilities for each customer i so that the scale parameters $\xi_i = 1$ for all i . This is possible because dividing all utilities with subscript i by the same scalar ξ_i does not change which utility is the maximum for customer i . That is, probability statements about the maximum utility for customer i are not affected by this rescaling.

2.1.1 Narrowing the Product Choices

Narrowing the product choices is a "prescreening" step that does not change the fundamental structure of the underlying logit model. When there are many product alternatives to consider, marketing researchers have found that customers typically use some criteria to narrow their choices to a "consideration set" of products, which are then investigated in more detail. (See, e.g., Roberts and Lattin 1991; Andrews and Srinivasan 1995; Siddarth et al. 1995). In a normative framework, customer i would form a consideration set by eliminating all products with expected utility less than some threshold u_i , where the threshold is based on his or her cost of considering additional alternatives. Thus we define

u_i = customer i 's minimum acceptable expected utility for considering a product,
 $X_{ij} = 1$ if $U_{ij} \geq u_i$ and 0 otherwise, for all i, j .

Consideration sets can have a significant impact on the assortment optimization, as the numerical analysis in Sect. 3 illustrates.

2.1.2 Determining $q_{ij}(y)$

The definition of conditional probability implies that

$$\begin{aligned} q_{ij}(y) &= P\{\text{customer } i \text{ purchases product } j \text{ from retailer } r | y\} \\ &= P\{\text{customer } i \text{ purchases product } j | \text{purchases from retailer } r, y\} \\ &\quad * P\{\text{customer } i \text{ purchases from retailer } r | y\} \end{aligned} \quad (11.3)$$

This equation does not necessarily imply that the customer chooses the retailer first, but this decomposition allows a separable estimation of the required model parameters, as will be discussed later.

Given that customer i selects retailer r 's assortment for a purchase, his or her choice set is defined as the intersection of the consideration set and retailer r 's assortment, i.e.,

$$S_{ri} = \{j | y_j X_{ij} = 1\}, \quad \text{for all } i.$$

Given any choice set S_{ri} , the probability of selecting item $j \in S_{ri}$ is the standard MNL probability, which in this case is

$$P\{\text{customer } i \text{ purchases product } j | \text{chooses retailer } r, y\} = \frac{e^{U_{ij}}}{\sum_{k \in S_{ri}} e^{U_{ik}}}. \quad (11.4)$$

Ben Akiva and Lerman (1985, p. 282) show that the maximum utility that customer i obtains from the choice set S_{ri} has a Gumbel distribution, with mean

$$V_{ir}^* = \ln \left(\sum_{j \in S_{ri}} e^{U_{ij}} \right),$$

and the same scale parameter as the individual utilities. Thus, the total utility of purchasing from retailer r 's assortment is Gumbel distributed with mean $v_{ir} = V_{ir} + V_{ir}^*$. The analogous result holds for all other retailers' assortments, which we index by ρ . Therefore, the maximum utility that customer i could obtain from shopping at other retailers also has a Gumbel distribution with mean

$$v_{i0} = \ln \left(\sum_{\rho \neq r} e^{V_{i\rho} + V_{i\rho}^*} \right),$$

and the same scale parameter $\xi_i = 1$ as the individual utilities. This allows the retailer choice probability to be written as a binary logit probability

$$\begin{aligned}
 f_i &= P\{\text{customer } i \text{ selects retailer } r | y\} = \frac{e^{v_{ir}}}{e^{v_{ir}} + e^{v_{io}}} \\
 &= \frac{\sum_{j \in S_{ri}} e^{U_{ij}}}{e^{a_{ir}} + \sum_{j \in S_{ri}} e^{U_{ij}}} \text{ with } a_{ir} = v_{io} - V_{ir}.
 \end{aligned} \tag{11.5}$$

The second fraction results if we multiply top and bottom by $\exp\{-V_{ir}\}$. From this point onward, we focus on the particular retailer r and simply write a_i for a_{ir} .

Combining the two probabilities in (11.3) using the assortment y for retailer r and the X_{ij} for customer i to define the choice set S_{ri} , we obtain the formula

$$q_{ij}(y) = \frac{y_j X_{ij} e^{U_{ij}}}{e^{a_i} + \sum_k y_k X_{ik} e^{U_{ik}}}, \tag{11.6}$$

after cancelling the term $\sum_{j \in S_{ri}} e^{U_{ij}}$. A key result in (11.6) is that a_i is a constant that is

independent of retailer r 's assortment decision y .

The size of a_i indicates the relative strength of retailer r 's competitors for customer type i . The value of a_i can be obtained in various ways. One method is to assume that customer i knows the contents of all the retailers' assortments and chooses the best retailer by maximizing the total utility as described above. Alternatively, the customer might simply decide whether to continue shopping at other retailers based on an estimated value a_i , which corresponds to the estimated maximum utility improvement obtained from other retailers' products, plus the improvement in value obtained by buying from an alternative retailer versus buying from retailer r . For assortment optimization using (11.6), retailer r does not need to know which behavioral model applies to customer i , since a_i is simply a parameter to be estimated, as discussed below.

Kahn and Lehmann (1991) and others have suggested adding terms to V_{ir} to capture the additional customer value associated with properties of the assortment that increase its "breadth," such as the total number of products or the number of brands offered. The structure of the optimization model in this chapter does not allow these additional variables to be included in the retailer's objective function. But features such as the total number of products or the number of brands in the assortment can be included as constraints for the assortment optimization model, with their corresponding values being added as constant terms to V_{ir} . This allows a sensitivity analysis to be done with respect to these assortment parameters.

2.1.3 Estimation and Empirical Testing

An estimate for a_i can be obtained using (11.5) from the observed fraction f_i of customers of type i who choose retailer r for *any particular given* assortment. Assuming that f_i can be obtained approximately from market research data for the current assortment, we can solve for the corresponding a_i as follows

$$a_i = \left(\frac{1}{f_i} - 1 \right) \sum_{j \in S_{ri}} e^{U_{ij}}.$$

This formula requires utility estimates for each product, which can be obtained from (11.4) as discussed previously. Thus, (11.4) and (11.5) allow the $\{U_{ij}\}$ and $\{a_i\}$ to be estimated separately, and they could in fact be obtained from different assortments.

Purchasing behavior for consumer package goods based on multi-stage logit models has been studied empirically for a variety of model forms. For example, Cintagunta (1993) provides a summary of articles that include empirical studies of three stages of consumer purchase decision making: (a) whether or not to purchase from this retailer (b) item choice from a retailer and (c) purchase quantity. See also Roberts and Lattin (1997) for a literature review. In forecasting demand for consumer package goods, “purchase incidence,” which is defined as the probability that the customer makes a shopping trip to a given retailer that results in a purchase from the category, plays a role that is similar to retailer choice in this paper. [See, e.g., Bucklin and Lattin 1991 for a discussion of using the binary logit model for purchase incidence.]

2.1.4 Elasticity Comparisons

Formula (11.6) shows that adding another product to the assortment *increases* the probability that customer i purchases from this retailer, but *decreases* the probability that each of the original products in the assortment is selected. The magnitudes of these effects depend on a_i , as shown below.

Let $Q_i(y) = P\{\text{customer } i \text{ purchases from this retailer}\}$, where

$$Q_i(y) = \sum_j q_{ij}(y) = \frac{P_i(y)}{e^{a_i} + P_i(y)}, \quad \text{with } P_i(y) = \sum_j y_j X_{ij} e^{U_{ij}}.$$

Interpreting partial derivatives as changes in y_k from 0 to 1, we can define the two elasticities

$$\frac{1}{Q_i(y)} \frac{\partial Q_i(y)}{\partial y_k} = \frac{e^{a_i} X_{ik} e^{U_{ik}}}{P_i(y)[e^{a_i} + P_i(y)]} \quad \text{and} \quad \frac{1}{q_{ij}(y)} \frac{\partial q_{ij}(y)}{\partial y_k} = - \frac{X_{ij} e^{U_{ij}}}{[e^{a_i} + P_i(y)]^2}.$$

The first and second elasticities show, respectively, that:

1. The percentage increase in total sales to customer i from adding product k is greater when a_i is larger.
2. The percentage of cannibalization of product j 's sales due to adding product k is smaller when a_i is larger.

Taken together, these results imply that including additional products in the assortment is more advantageous to the retailer when the retailer's competition is stronger.

2.2 Retailer's Assortment Optimization

The profit function $\Pi_j(D_j(y))$ for each product j is based on a newsvendor type model. In general, a fixed cost

$$F_j = \text{the fixed cost of stocking product } j$$

should also be included. The expected profit $\Pi(y)$ for the planning period as a function of y can therefore be written as the sum of the expected profits for the various products

$$\Pi(y) = \sum_j \{E[\Pi_j(D_j(y))] - F_j\}.$$

[It should be noted that even though the random variables $\Pi_j(D_j(y))$ are not independent, their expectations are still additive.] A more general optimization problem can be defined if there are nonlinear cost interactions between the products, but that formulation will not be developed in this chapter.

The newsvendor expected profit for product j for a fixed time period as a function of the assortment y can be written as

$$E[\Pi_j(D_j(y))] = \max_{s_j} \left\{ m_j \mu_j(y) - c_{uj} E[D_j(y) - s_j]^+ - c_{oj} E[s_j - D_j(y)]^+ \right\} \quad (11.7)$$

where $E[x]^+$ denotes the expected value of $\max\{0, x\}$ and

s_j = the base stock level for product j for the time period

m_j = unit profit margin for product j

$\mu_j(y)$ = expected demand during the time period = $E[D_j(y)]$

c_{uj} = "understock" cost per unit c_{oj} = "overstock" cost per unit

The financial input quantities can be calculated in the usual way, *i.e.*,

m_j = selling price – unit cost,

c_{uj} = shortage loss – unit cost

c_{oj} = unit cost – salvage value.

From (11.6), we see that $y_j=0$ implies $D_j(y)=0$ with probability 1, which implies that the expected profit is 0. That is, there is no specific shortage cost c_{uj} that results from not including a given item in y , but there is a net loss of expected utility for the retailer's assortment, which increases the likelihood that the customer will choose another retailer. This is because when a customer's most preferred item is missing, the customer either substitutes another item from this retailer's assortment or chooses another retailer. The demand that results from substitutions for items not in the retailer's assortment is captured in $\mu_j(y)$. Substitutions from stockouts are ignored, as discussed below. From the standpoint of this retailer r , the probability of choosing another retailer is lumped together with the "no purchase" option.

Using the newsvendor critical ratio formula, the optimal base stock level s_j^* satisfies

$$s_j^* = \arg \min_s \left\{ s | P\{D_j(y) \leq s\} \geq \alpha_j = \frac{c_{uj}}{c_{uj} + c_{oj}} \right\}.$$

The overstock cost c_{oj} above can have a variety of interpretations. For continuing products that will be offered in subsequent seasons, it is the unit holding cost for the season, while for "seasonal" products, it is the unit cost minus the expected salvage value per unit for any excess inventory at the end of the season.

There are various fixed costs F_j that can be associated with stocking items in a product category. For larger items such as furniture, it is common to display one unit in the store and hold additional inventory elsewhere, for example. In this case, F_j would include the required floor space for display. For smaller items, there may be a shelf facing with one item viewable, and the remaining items stored behind it. In both these cases, F_j would include the fixed cost of the required display space in the store when the item is in the assortment.

2.2.1 Incremental Demand Arising from Substitution

Kok et al. (2015) define two kinds of substitution-based demand: (1) assortment based substitution in which a customer switches to another product when a more preferred product is not carried in the assortment and (2) stockout-based substitution in which the customer substitutes another product if a more preferred alternative is in the assortment, but it is out of stock. This chapter captures assortment-based substitution through the MNL choice model discussed previously, but it ignores stockout-based substitution. Some recent papers have modeled stockout-based substitutions, but this generally leads to complex optimizations, and thus solution heuristics are required. Mahajan and van Ryzin (2001) and Guar and Honhon (2006) and Honhon et al. (2010) assume that customers maximize utility over the items that are currently available, i.e., they treat the retailer's assortment as dynamic. These approaches are quite general, but require heuristic solutions for

most customer preference structures. The other assortment optimization models discussed previously in the literature review have either not treated this stockout-based substitution or have bounded its effects.

This chapter assumes that the customer chooses the retailer based on the complete assortment y , and that product demands which encounter stockouts of products in the retailer's assortment become lost sales. Thus the demand arising from stockout-based substitutions is ignored. Smith and Agrawal (2000) argue using bounds, that the absolute percentage error in expected demand that results from ignoring demand from stockout based substitutions is bounded by $(1 - \alpha)(1 - L)$, where α is the target service level and L is the probability that the customer is unwilling to substitute. For retailers that set high service level targets for most products during the normal selling season, this bound implies that stockout substitutions will rarely occur. The stockout based demand needs to be counted only if the customer is willing to switch to another product from the same retailer. Given that alternative retailers exist for many items, customers who choose another retailer instead of substituting a different product will be correctly captured by the lost sales assumption. Also, when the service level is defined as the probability of no stockout using the normal distribution, the fraction of demand served is typically larger than the service level. For example, for the normal distribution with $P\{\text{no stockout}\} = \alpha = 0.9$, approximately 96 % of demand will be served, and with $\alpha = 0.95$ approximately 98 % of demand will be served before a stockout occurs. The error corresponding to the unserved demand is further reduced by eliminating those customers who do not substitute another product from this retailer. Thus, for retail products that have high service levels such as 0.9 or 0.95, it seems reasonable to ignore substitution demand arising from stockouts.

2.2.2 Two Variants of the Objective Function

Products that may have purchase quantities larger than one can be handled in a variety of ways. One method is to use a compound Poisson distribution for demand, in which customers arrive according to a Poisson process and then select their purchase quantities randomly. For example, Poisson arrivals with a purchase quantity selected from a logarithmic distribution result in a negative binomial distribution for total demand during any fixed period. Smith and Agrawal (2000) used the negative binomial distribution and found that a linear approximation to the newsvendor objective function worked well in that case. Other papers on assortment optimization (e.g., van Ryzin and Mahajan 1999; Mahajan and van Ryzin 2001; Guar and Honhon 2006) have used a normal approximation for demand to obtain a newsvendor expected profit function.

When there are time based holding costs, it may be advantageous for retailers to restock more frequently than once per season. This feature can be added to the newsvendor model (11.7), provided that the assortment does not change in midseason. If there is an additional cost $h = \text{unit holding cost for one restocking}$

period, a cost term of the form $0.5h \left[s + |s - D_j(y)|^+ \right]$ is subtracted from the objective function. The critical ratio stock level formula still holds, where c_{oj} is replaced by $c_{oj} + h$ and c_{uj} is replaced by $c_{uj} - 0.5h$. The costs c_{oj} and c_{uj} may also be allowed to vary by time period.

2.2.3 A Linear Approximation for the Objective Function

It can be verified by numerical calculation that for common ratios of profit margin to overstock and understock costs, the newsvendor expected profit function (11.7) is approximately linear in the expected demand $\mu_j(y)$ for the Poisson distribution. That is, when the various costs are held fixed and expected demand increases, the target service level remains constant and the safety stock increases in such a way that the sum of the terms in (11.7) increases approximately linearly as a function of the mean.

For the Poisson demand distribution, this approximation is illustrated for a range of parameter ratios in Fig. 11.1. To simplify the graph, all $F_j = 0$ and all profits have been divided by c_u . That is, when all cost parameters are expressed as multiples of c_u , the graphs can be expressed as (expected profit)/ c_u , which implies that the only required variables are the service level α and the mean demand. Using linear regression, the R^2 values for all the linear fits to the points in this figure are at least 0.998.




Fig. 11.1 Maximum profit vs. mean demand $m = cu, co = cu(1/\alpha - 1)$

The linear approximation implies that are constants π_j and b_j derived from the slope and intercept of the regression line for product j such that the expected profit can be approximated as follows

$$E[\Pi_j(D_j(y))] \approx \pi_j \mu_j(y) - y_j(b_j + F_j).$$

In general, it appears that the quality of the fit improves as the mean increases and as the service level α increases. When $F_j = 0$, Fig. 11.1 shows that the b_j values are positive. This is because the expected profit becomes negative for low enough mean demand, but in these cases $y_j = 0$ will be optimal.

Using $C_j = b_j + F_j$ to combine the constants b_j with the fixed costs F_j , and recalling that $\mu_j(y) = \sum_i \lambda_i q_{ij}(y)$, the retailer’s approximate objective function can therefore be written as

$$\Pi^*(y) = \sum_j \pi_j \sum_i \lambda_i q_{ij}(y) - \sum_j C_j y_j. \tag{11.8}$$

This objective can be maximized with respect to y , subject to various constraints such as display space or brand representation in the assortment.

2.3 Properties of the Optimal Assortment

When customers’ utilities are Gumbel distributed with homogeneous means, van Ryzin and Mahajan (1999) showed that the optimal assortments form nested sets. This case corresponds to $U_{ij} = U_j$ for all i in this paper’s notation. That is, if S^K is the best assortment of size K , then $S^K \subseteq S^{K+1}$ for all K . With nonhomogeneous means U_{ij} , however, this property no longer holds, as demonstrated by the following counterexample. Let $\lambda_i = 1$ and $\exp(a_i) = 10$ for all i and consider the following matrix of $\exp(U_{ij})$ values

		Products		
		1	2	3
Customers	1	1,000	2	1,000
	2	1,000	1,000	2
	3	2	1,000	2
	4	2	2	1,000

Let the unit profits for the three products be 10, 9, 9 respectively. Clearly, the best single product is Product 1. But it can be seen from the table of expected profits below that the best two products are 2 and 3.

	y		Expected Profit
0	1	1	35.6
1	1	0	31.0
1	0	1	31.0

Thus, although Product 1 is the best single product, it is not part of the best set of two products.

Nested set properties do hold for two limiting cases, however. First, let us consider the case in which $a_i = a$ for all i and a is very large. Then rewrite $\Pi^*(y)$ as

$$\Pi^*(y) = e^{-a} \sum_{i,j} \lambda_i \pi_j \left(\frac{y_j X_{ij} e^{U_{ij}}}{1 + e^{-a} P_i(y)} \right) - \sum_j C_j y_j. \tag{11.9}$$

As a becomes sufficiently large, the term in parenthesis approaches $y_j X_{ij} e^{U_{ij}}$. Thus, if the products are ordered so that

$$\pi_1 \sum_i \lambda_i X_{i1} e^{U_{i1}} - C_1 \geq \pi_2 \sum_i \lambda_i X_{i2} e^{U_{i2}} - C_2 \geq \dots, \tag{11.10}$$

then the optimal assortments will be $\{1\}$, $\{1, 2\}$, \dots for a sufficiently large. This implies that there is an optimal product ordering for the assortment, if the retailer’s competition is sufficiently strong, even when consumer preferences are heterogeneous. In microeconomic terms, this might be called the “perfectly competitive” case.

A second special case arises when $\exp(a_i)$ approaches 0 for all i . In this case, the retailer is effectively a monopolist, since any consumer who purchases will choose this retailer. For the case in which $X_{ij} = 1$ for all i, j , every product in the retailer’s assortment is in every customer i ’s choice set. Thus, the optimal strategy for a monopoly retailer is to rank products in order of profitability, based on ranking the expected profits as follows

$$\pi_1 \sum_i \lambda_i - C_1 \geq \pi_2 \sum_i \lambda_i - C_2 \geq \dots \tag{11.11}$$

But if some $X_{ij} = 0$, this property may not hold, because some customers may not consider the retailer’s most profitable product and thus would not choose it. *Thus, with considerations sets, there may be no specific nested set property when $\exp(a_i)$ approaches 0, for all i .*

2.3.1 Sensitivity to the Retailer’s Market Strength

To illustrate the difference in the two rankings (11.10) and (11.11), let us consider an example with 5 customer types and 20 products, where the utilities U_{ij} were generated by taking samples from a uniform distribution on $[0, 2]$. Let all $X_{ij} = 1$




Fig. 11.2 Profitability calculations for 20 randomly generated products

and all $\lambda_i = 1$ in this example. The 20 products are assigned gradually decreasing unit profits π_j : \$10.00, \$9.90, \$9.80, . . . , \$8.10 and fixed costs $C_j = 0$ for all j . Thus, for the case in which the retailer’s competitive position is very strong, the products’ ranking is based on (11.11), which implies that the products would be ranked in order of the unit profits, 1, 2, 3, Therefore, for a retailer with a dominant market position, the optimal assortment of K products is $\{1, 2, \dots, K\}$.

On the other hand, for a retailer in a weak competitive position, the product rankings are based on the rankings in (11.10). The calculated results for (11.10) are illustrated in Fig. 11.2.

The height of the bars in Fig. 11.2 shows that the expected values for this case are quite different from those that would produce the ranking of 1, 2, 3, . . . determined by (11.11). For example, the top 5 products based on ranking the values in Fig. 11.2 are: {9, 11, 20, 13, 7}.

2.4 Solving the Optimization Problem

If the total number of products is small, optimal assortments can be obtained by an exhaustive search, but this becomes more difficult for larger numbers of products. Based on the structure of the problem, certain products may be eliminated from the assortment *a priori*, which reduces the problem size. Substituting the definition (11.6) of $q_{ij}(y)$ into (11.8), the objective function can be written as

$$\begin{aligned} \text{Max } \Pi^*(y) &= \sum_{j \geq 1} y_j \{ \pi_j r_j(y) - C_j \} \quad \text{with } y_j = 0, 1 \text{ for all } j \geq 1, \\ \text{where } r_j(y) &= \sum_i \lambda_i \left(\frac{y_j X_{ij} e^{U_{ij}}}{e^{a_i} + \sum_{k \geq 1} y_k X_{ik} e^{U_{ik}}} \right) \end{aligned} \tag{11.12}$$

For any y such that $y_k = 0$, define

$$\Delta_k r_j(y) = r_j(y + e_k) - r_j(y), \quad \text{where } e_k = \text{the unit vector with } k\text{th element} = 1.$$

It can be verified that

$$\text{If } y_k = 0, \text{ then } \Delta_k y_j [r_j(y) - C_j] \leq 0 \text{ for all } j \neq k.$$

This has the implication that if $\pi_k r_k(e_k) - C_k \leq 0$ for any k , then $y_k = 0$ must hold. That is, $y_k = 1$ cannot be optimal since y_k could be changed to 0 and all terms in the objective function will improve or stay the same. This observation can be used to eliminate some products before searching on y . However, it appears that an exhaustive search over the remaining 0,1 variables is required to optimize the assortment.

Retailer imposed constraints, such as the number of products must be at least K , or at least one product of Brand B must be included, can be added as linear constraints on y . For example, if the assortment must include at least one product of Brand B, define the logical inputs

$$I_{Bj} = 1 \text{ if product } j \text{ is of brand } B, \text{ and } 0 \text{ otherwise.}$$

Then the brand constraint is of the form

$$\sum_j y_j I_{Bj} \geq 1 \quad \text{for brand } B.$$

We can also include a display space constraint of the form

$$\sum_j d_j y_j \leq D, \text{ where}$$

d_j = the space required for product j

D = total available display space for this category.

These additional constraints also reduce the number of alternatives to be searched.

3 Illustrative Application for a DVD Player Data Base

This section illustrates the application of the optimization model to a set of customer utilities derived from a conjoint analysis of actual consumer Internet responses. The preference data were collected through the *Active Decisions' Active Buyers Guide Sales Assistant* website. [See www.activedecisions.com. This company has been acquired by Knova Systems, who plan to offer conjoint utility encoding as a

consulting service.] Visitors to *activebuyersguide.com*, *yahoo.com* and other e-commerce sites completed an interactive survey to elicit their preference tradeoffs for product attributes. These preferences are defined so as to be *independent* of the specific set of products in the market. Product utilities were then derived from additive conjoint analysis of 2,213 customer responses for the DVD player category. That is, each customer's net utility for a particular product was calculated as the sum of his or her "part worths" for the attributes of that product, including the price. (See Green and Srinivasan 1978; Cattin and Wittink 1982; Wittink and Cattin 1989 for discussions of conjoint analysis methods. The conjoint analysis of this data was performed by *Active Decisions* and the author is indebted to them for sharing their results).

The utility values were then normalized by dividing each utility U_{ij} by customer i 's maximum utility to obtain

$$S_{ij} = \frac{U_{ij}}{\max_{k \in \Omega} U_{ik}} \text{ for all } i, j.$$

After this normalization, it was assumed that $\xi_i = 1$ for all i . Consideration sets based on utility thresholds can then be defined as a fixed fraction θ of each customer's maximum utility over all products. That is,

$$u_i = \theta \max_{j \in \Omega} U_{ij}, \text{ where } \Omega = \text{the set of all products in the market.}$$

Thus, $X_{ij} = 1$ if and only if $S_{ij} > \theta$.

Assortment optimization for this example was done for the case of "large" a_i , i.e., the retailer's competitive position is weak. Thus, the optimal assortments will form nested sets according to (11.10), as discussed previously. Because of the highly competitive nature of the DVD player market and because this retailer was not a dominant player in consumer electronics, this assumption seemed appropriate. However, the database had no data available on retailer preference so this assumption could not be tested.

3.1 Comparing the Model's Predictions to a Retailer's Sales Data

In order to test the predictive accuracy of utilities in the data base and the MNL choice models, we obtained data on the observed selling proportions for an assortment of 30 DVD players offered by a major retail chain. These selling proportions were compared to those predicted by the MNL choice model fitted to the product attribute utilities in the DVD Player data base. The actual selling proportions of the products ranged from 0.2 to 16 %. [There were 117 different DVD player products at the time the data set was collected, and the retailer data was obtained for the same time period.] A variety of θ values were tested to obtain the correlations and the R-square values are shown below in the table below.

3.1.1 Actual vs. Predicted Selling Proportions for 30 Products

θ	Correlation	R-square
0	69 %	47 %
0.9	78 %	60 %
0.95	79 %	62 %
1.0	72 %	53 %

This table indicates that the fit is reasonably good for all θ values, but the accuracy improves somewhat when customers are assumed to use moderately restrictive consideration sets. Further investigation also revealed that most of the error in these predictions resulted from over-predictions for three products, which the retailer reported were unavailable in some stores. This test supports the use of the utilities in the data base, and also suggests a fairly high θ value such as $\theta = 0.9$ or 0.95 may be appropriate for this data set.

3.2 Comparing the Expected Revenue of the Retailer's Assortment vs. the Optimal Assortment

The objective function in (11.12) was then applied to the set of 117 DVD player products available at that point in time to determine the optimal assortment of 30 products. For the optimization, it was assumed that each of the 2,213 respondents to the online survey represents a customer segment of equal size, *i.e.*, the λ_i were assumed to be equal for all i . The fixed costs C_j were set to zero and the product prices from the DVD Player data base were used to compute the expected revenue from a given assortment. Since the revenue comparisons will be done on a percentage basis, it is not necessary to know the actual number of buyers per segment. For percentage calculations with $\lambda_i = \lambda$ for all i , the λ will cancel out of the profit comparisons. Therefore, for the case of "very large" a_i , the objective function in (11.12) can be maximized by substituting a linear objective function that is similar to the ranking calculation in (11.10),

$$\begin{aligned} \text{Max } \Pi_0(y) &= \sum_{j \geq 1} y_j \pi_j X_{ij} e^{U_{ij}}, \\ \text{subject to } y_j &= 0, 1 \text{ for all } j \geq 1 \text{ and } \sum_{j=1}^n y_j = 30. \end{aligned} \quad (11.13)$$

The optimal assortments were then determined for various values of $\theta = 0.9$, 0.95 and 1.0 , which are captured by changes in the X_{ij} . The table below compares the percentage improvements achieved by the optimal assortment over the retailer's current assortment, for the various θ choices.

θ	Revenue improvement	Common products
0.9	169 %	11 (37 %)
0.95	185 %	9 (30 %)
1.0	208 %	7 (23 %)

The revenue improvements in this table are optimistic because they assume that each customer i 's buying behavior exactly matches the MNL model. However, even recognizing this, it appears that using the MNL-based optimal assortment with consideration sets has substantial potential to improve this retailer's revenues.

3.3 *The Impact of Customer Preference Structure*

The analysis above is based on the use of both consideration sets and heterogeneous customer market segments. To test the impact of these structural assumptions, we focus on three sensitivity questions:

1. *What is the impact of including customer preference heterogeneity in determining optimal assortments?*
2. *How does customers' use of consideration sets impact the optimal assortments and expected profits?*
3. *How does the expected profit increase with assortment size, i.e., how does the optimal assortment size depend on the fixed costs of offering additional products?*

3.3.1 Customer Heterogeneity

To examine the role of customer preference heterogeneity in developing the optimal assortment, optimal assortments for homogeneous preferences were generated by replacing the S_{ij} with "average" values S_j , which equal the average S_{ij} value over all customer types i . The expected profits for these optimal assortments were then compared to the profits for the optimal assortment with heterogeneous preferences S_{ij} in Fig. 11.3.

The potential revenues of the two optimal assortments converge when essentially all positive utility products are carried by the retailer. However, for assortment sizes 10–30 that are relevant to most retailers, the optimal assortments for heterogeneous preferences result in profits almost twice as large. Examining the contents of the assortments produced by the two methods found only about 5 % common items in the assortments of sizes 5–30. Thus, for this data set, ignoring customer heterogeneity has significant financial consequences and major impacts on the optimal assortment.




Fig. 11.3 Including preference heterogeneity in assortment optimization




Fig. 11.4 The effect of consideration sets on expected profit

3.3.2 The Use of Consideration Sets

To analyze the impact of consideration sets, the optimal assortments for $\theta = 0, 0.9$ and 1.0 are compared in Fig. 11.4, where $\theta = 0$ corresponds to “no consideration sets.” Figure 11.4 shows that when customers use consideration sets and the retailer uses this information correctly in developing the optimal assortments, a substantial increase in expected profit results for typical assortment sizes. For assortments in the 5–10 item range, the $\theta = 0.9$ or 1.0 cases yielded two to three times the profit of the optimal assortment without consideration sets.

Consideration sets allow the retailer to use a more focused assortment. When customers use consideration sets, the retailer can achieve 80–90 % of the maximum possible profit with an assortment sizes of only about 30 items, while these assortment sizes can achieve only about 50 % of the maximum without consideration sets. For $\theta = 1$, all customers can receive their first choice product with an assortment size of 66, but for $\theta = 0$ additional products always increase expected sales.

The shape of the curves in Fig. 11.4 also determines the impact of the fixed costs C_j on the optimal assortment size. For an assortment of size 30, for example, the slopes of the lines are approximately, \$3,500, \$5,000 and \$6,000, respectively, which correspond to the marginal benefits of an additional product. [These dollar figures correspond to one purchase by each of the 2,213 active shoppers in the category. This level of sales would correspond to an aggregate across multiple stores.] Thus, consideration sets allow high fixed costs to be justified for small numbers of products, but tend to limit the optimal assortment size as the number of products increases.

It was assumed in Fig. 11.4 that the optimal assortment was determined for the correct θ value in each case. But since customers' behavior with regard to consideration sets may be difficult to predict, it is interesting to consider the impact of incorrect assumptions about consideration sets. This calculation is illustrated in Fig. 11.5, where the optimal assortment for $\theta = 0$ was used when the correct value was $\theta = 0.9$, and vice versa.

This shows that if customers form consideration sets based on $\theta = 0.9$, the optimal assortment for $\theta = 0$ results in a reduction in expected profit of 12–50 % for assortments in the range of 10–30. On the other hand, if customers do not use consideration sets to prescreen the products, i.e., $\theta = 0$ is correct, the optimal assortment for $\theta = 0.9$ results in a 10–20 % reduction in expected profit. Thus, for this data set, the less risky alternative is to assume that customers do use considerations sets.

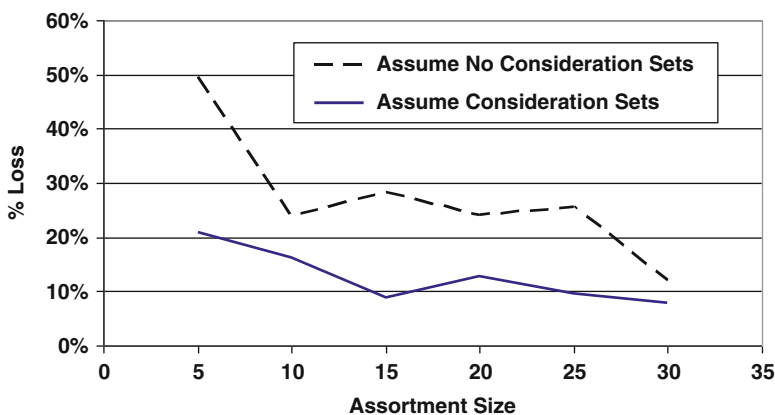


Fig. 11.5 Percentage profit loss for incorrect consideration set assumption with $\theta = 0, 0.9$

4 Summary and Conclusion

This chapter has developed an operational model for assortment optimization, based a multinomial logit choice model with general heterogeneous customer preferences. The structure of the model allows the required input parameters for product choice and retailer choice to be estimated separately from product sales and retailer market shares. These estimates can be based on observed consumer choices for previous assortments, which need not be optimal. The linear approximation of the newsvendor cost function assumes that temporary stockouts result in lost sales, which restricts the model's use to retailers or categories of products with relatively high service levels. However, this assumption leads to a closed form objective function that captures the impact of the assortment on both retailer choice and product choice. While the optimal assortments may no longer form nested sets for heterogeneous preferences, it is shown that the special cases of perfect competition and retailer monopoly do lead to different sequences of optimal nested sets, and it is illustrated how the optimal assortment transitions between these two extremes as the retailer's market share increases.

The optimization model can accommodate a variety of additional retailer constraints. For example, it may be important to: (1) require that certain top brands be represented in the assortment; (2) provide some level of assortment stability across time for customers; (3) stay within a given display space constraint; or (4) carry products with the full range of price points to promote the image of a category killer. The analysis of the DVD player data base illustrated the decreasing marginal benefits associated with increasing assortment size and also the sensitivity of the optimal assortment to the input assumptions regarding the customer choice process. Including customer heterogeneity had significant impacts on both the optimal assortments and the expected profits. Consideration sets, which have been studied in the context of modeling customer choice, but have not previously been included in assortment optimization, were found to strongly influence the optimal assortment for the DVD player data base. This analysis supports the importance of using a consumer choice model that includes heterogeneous preferences and consideration sets in obtaining optimal assortments. The sensitivity analysis also illustrates the potential profit improvement for additional selling effort designed to influence customers' product choices.

There are a number of promising avenues for future research. Clustering customers into fewer classes can reduce the problem size and lead to shorter computation times for the general competitive case. Analytical methods for choosing the best customer clusters for a given database of utilities could therefore extend the applicability of the optimization model. Clusters based on customers' preferences for product attributes, as opposed to individual product utilities, may lead to clusters that are more stable over time. Better optimization approaches that exploit the specific structure of the assortment problem may also exist. It is hoped that this chapter will also lead to additional research on the development of decision support systems for assortment planning that implement this optimization model for choosing assortments, taking into account both product choice and retailer choice.

Acknowledgement The author is grateful to Dale Achabal, Kirthi Kalyanam, Shelby McIntyre and Chris Miller for many valuable discussions and to Active Decisions, Inc. for providing the data base that was used for testing the optimization model. This research was partially supported by the Retail Workbench Research and Education Center at Santa Clara University.

References

- Andrews, R. L., & Srinivasan, T. C. (1995, February). Studying consideration effects in empirical choice models. *Journal of Marketing Research*, 32, 30–41.
- Ben Akiva, M., & Lerman, S. (1985). *Discrete choice analysis*. Cambridge: MIT Press.
- Bucklin, R., & Lattin, J. (1991, Winter). A two state model of purchase incidence and brand choice. *Marketing Science*, 10, 24–39.
- Cachon, G., Terwiesch, C., & Yi, X. (2005). Assortment planning in the presence of consumer search. *Manufacturing and Service Operations Management*, 7(4), 330–346.
- Cachon, G., & Gurhan Kok, A. (2007). Category management and coordination in retail assortment planning in the presence of basket shopping consumers. *Management Science*, 53(6), 934–951.
- Cattin, P., & Wittink, D. R. (1982, Summer). Commercial use of conjoint analysis: A survey. *Journal of Marketing*, 46, 44–53.
- Chen, K. D., & Hausman, W. H. (2000). Technical note - mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. *Management Science*, 46(2), 327–332.
- Cintagunta, P. K. (1993). Investigating purchase incidence, brand choice, and purchase quantity decisions of households. *Marketing Science*, 12, 184–208.
- Chong, J.-K., Ho, T.-H., & Tang, C. (2001). A modeling framework for category assortment planning. *Manufacturing and Service Operations Management*, 3(3), 191–210.
- Dobson, G., & Kalish, S. (1988). Positioning and pricing a product line. *Marketing Science*, 7(2), 107–125.
- Dobson, G., & Kalish, S. (1993). Heuristics for positioning and pricing a product line using conjoint and cost data. *Management Science*, 39(2), 160–175.
- Green, P. E., & Krieger, A. M. (1985). Models and heuristics for product line selection. *Marketing Science*, 4(1), 1–19.
- Green, P. E., & Srinivasan, V. (1978). Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research*, 5(2), 103–123.
- Guar, V., & Honhon, D. (2006). Assortment planning and inventory decisions under a locational choice model. *Management Science*, 52(10), 1528–1543.
- Hanson, W., & Kalyanam, K. (2006). *Internet marketing and e-commerce*. Cincinnati, OH: Southwestern College Publishing.
- Honhon, D., Guar, V., & Seshadri, S. (2010). Assortment planning and inventory management under stockout based substitution. *Operational Research*, 58(5), 1364–1379.
- Kahn, B. E., & Lehmann, D. R. (1991, Fall). Modeling choice among assortments. *Journal of Retailing*, 67, 274–299.
- Kohli, R., & Sukumar, R. (1990). Heuristics for product-line design using conjoint analysis. *Management Science*, 36(12), 1464–1478.
- Kok, A. G., & Fisher, M. (2007). Demand estimation and assortment optimization under substitution: Methodology and application. *Operational Research*, 55(6), 1001–1021.
- Kok, A. G., Fisher, M., & Vaidyanathan, R. (2015). Assortment planning: review of literature and industry practice. In N. Agrawal & S. Smith (Eds.), *Retail supply chain management* (2nd ed.). New York: Kluwer Academic Publisher.

- Mahajan, S., & van Ryzin, G. (2001). Stocking retail assortments under dynamic substitution. *Operational Research*, 49(3), 334–351.
- McBride, R. D., & Zufryden, F. S. (1988). An integer programming approach to the optimal product line selection problem. *Marketing Science*, 7(2), 126–140.
- Miller, C. M., Smith, S. A., McIntyre, S., & Achabal, D. (2010). Optimizing and evaluating retail assortments for infrequently purchased products. *Journal of Retailing*, 86(2), 159–171.
- Rusmevichientong, P., & Topaloglu, H. (2012). Robust assortment optimization in revenue management under the multinomial logit choice model. *Operational Research*, 60(4), 865–882.
- Roberts, J. H., & Lattin, J. M. (1991, November). Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28, 429–440.
- Roberts, J. H., & Lattin, J. M. (1997, August). Consideration: Review of research prospects and future insights. *Journal of Marketing Research*, 34, 406–410.
- Sauré, D., & Zeevi, A. (2013). Optimal dynamic assortment planning with demand learning. *Manufacturing and Service Operations Management*, 15(3), 387–404.
- Siddarth, S., Bucklin, R., & Morrison, D. (1995, August). Making the cut: Modeling and analyzing choice set restriction in scanner panel data. *Journal of Marketing Research*, 32, 255–266.
- Smith, S. A., & Agrawal, N. (2000). Management of multi-item retail inventories systems with demand substitution. *Operational Research*, 48, 50–64.
- van Ryzin, G., & Mahajan, S. (1999). On the relationship between inventory costs and variety benefits in retail assortments. *Management Science*, 45, 1496–1509.
- Wittink, D. R., & Cattin, P. (1989). Commercial use of conjoint analysis: An update. *Journal of Marketing*, 53(3), 91–96.