

A Model for Determining Retail Product Category Assortment and Shelf Space Allocation

Norm Borin

*Department of Business Administration, California Polytechnic State University,
San Luis Obispo, CA 93407*

Paul W. Farris and James R. Freeland

*Darden Graduate School of Business Administration, The University of Virginia, Box 6550,
Charlottesville, VA 22901-6550*

ABSTRACT

We develop a category management model to aid retailers in the space constrained decisions of which products to stock (assortment) and how much shelf space to allocate to those products. The model is formulated as a constrained optimization problem with two basic decision variables: assortment and allocation of space to the items in the assortment. The non-linearities in the objective function and the zero-one decision variables disallow a closed form solution. We develop a heuristic solution procedure based on simulated annealing and test it on a problem with a known optimum. We also apply the technique to a larger problem without a known optimum. Finally, the solution found by simulated annealing is compared against a solution produced using a shelf allocation rule based on share of sales.

Subject Areas: Heuristics, Marketing Management, Resource Allocation, and Simulation.

INTRODUCTION

Retail selling space is a fixed resource. Managing this space means making frequent decisions about which products to stock (assortment) and how much shelf-space to allocate those products. Although the average size of a store has continued to increase, it has not kept pace with the overwhelming number of new product introductions. In 1988, 5,694 new products and 10,558 new varieties were offered to supermarket retailers: many supermarket chains choose their assortment from a product list of over 60,000 stock-keeping units (SKUs) [51].

Retailers currently use a wide range of methods to choose assortments and allocate space. The widespread availability of scanner data and commercial shelf-management models allow retailers to quickly detect and eliminate unprofitable items. These commercial models typically use sales data and product and shelf dimensions to produce recommended shelf allocations. An increasing number of these models attempt to maximize retailer profits by allocating space according to item profitability and shelf-space elasticity [53]. In some models, there are options to allocate shelf space in order to minimize out-of-stocks.

We are not aware of any models which combine the effects of the assortment and space decisions. The lack of such models may be attributed to (1) the belief

that assortment decisions can be adequately modeled with space variables, that is, a decision not to include an item is modeled by allocating zero space, and (2) the opinion that including both space and assortment will create a complex model which cannot be optimized. While allocating zero space for an item is equivalent to eliminating it from the assortment, we argue that the sales elasticities for assortment decisions are apt to be much higher than for space decisions. This is especially true when minimum “pack-outs” (minimum shelf-space allocations) are often equal to two facings or more of a product.

An ideal shelf-management model provides assortment and shelf space recommendations that:

1. correspond to actual product dimensions, minimum economic pack-outs, and delivery cycles;
2. are based on differences in item profitabilities, including costs of selling, stocking, storing, and transportation;
3. incorporate shelf-space elasticities and cross-elasticities among brands in the same category;
4. consider the strength of consumer loyalty to an SKU in making assortment decisions.

This paper presents a model capable of incorporating all of these features. The model is formulated as a constrained optimization problem with two basic decision variables: product assortment and allocation of fixed space to the items in the assortment. The non-linearities in the objective function and the zero-one decision variables disallow a closed form solution. We develop a heuristic solution procedure based on simulated annealing and test it on a problem with a known optimum as well as on a larger problem without a known optimum. Finally, the solution found by simulated annealing is compared against a solution produced using shelf allocation rules based on share of sales. We conclude the paper by discussing potential applications of the model as well as some of its limitations.

REVIEW OF THE LITERATURE

The model which we propose is structured to correspond to an adaptation and extension of a “Push-Pull” model developed by Farris, Olver, and De Kluyver [23]. A key characteristic of this model is the division of a product’s market share into two components: uncompromised demand and compromised demand. Uncompromised demand stems from: (1) consumer preference for SKUs; (2) in-store merchandising support such as shelf-space, display, and advertising; and (3) product availability across retail stores. A brand may also capture compromised demand if some purchasers are willing to compromise choice (i.e., select alternatives) when the specific brands they desire are not available.

This paper models in-store support as a function of space—the space allocated to an item and the space allocated to its competitors. Resistance to compromise is used to account for both the long term and short term effects of missing items respectively due to assortment decisions and temporary stockouts.

Shelf Management Models and Experiments

If consumers are completely brand loyal and the product is available, the space allocated to an item has no effect on its sales [1]. An individual always purchases

the same product if it is present and either delays the purchase or proceeds to another location if it is not available. However, past work [10] [20] [22] [42] [53] shows that many consumers are willing to compromise their initial choice and switch to other products, either because their brand was not available or the shelf display changed their choice. In fact, many consumers brand choice decisions are made at the point of purchase [39] [43]. For these consumers the final choice may be influenced by one or more in-store merchandising factors, including space.

Shelf-Space Studies. Early studies in space management concentrated on establishing whether a relationship exists between the space allocated to an item and that item's sales [8] [15] [16] [24] [34] [35] [44]. These studies have consistently demonstrated a weak link between the two variables. Bultez and Naert [6] attribute these weak results to (1) poor experimental design, (2) low variation in space allocation, and (3) unreliable sales data.

Space Models. Typically, shelf-space models hypothesize that the ratio of sales/space decreases as space increases [1] [5] [6] [7] [9] [13] [14] [17] [18] [26]. We review four key studies below.

In Anderson's [1] work, space allocation decisions were based on marginal analysis of a logistic function modeling the relationship between a product's share of space and its market share. Brand loyal or nonswitchers were identified as those consumers whose purchases are not influenced by item shelf space. Cross-space effects were accounted for by using share instead of absolute category space for an SKU. Furthermore, although stockout costs were modeled as a function of item demand, the lost sales due to stockouts were not incorporated into either the demand function or the cost function.

Hansen and Heinsbroek's [26] solution to the assortment decision was to set an item's space to zero when it was not available. Furthermore, because their sales function neglects cross-space effects which measure demand interdependencies [6] [7] [13] [14], this model did not measure the positive effect on sales of other items or the category when a product is absent. The authors solve for a "near-optimal solution" using a generalized Lagrange multiplier approach.

Corstjens and Doyle [13] used geometric programming to optimize category space allocation with profit maximization being the objective. Their multiplicative model incorporated direct and cross-space elasticities, and it allowed for different product profit margins. Costs were modeled as a function of inventory investment and handling. The model's constraints included a store size limitation that prevented the sum of all category space from exceeding the store's size, upper and lower limits on space for each category, and a production or availability constraint. In 1983, Corstjens and Doyle extended their model to include conditions of widely varying product growth potentials [14].

Corstjens and Doyle's model was a major step towards a managerially useful space allocation model. However, their approach had a few drawbacks and areas of concern. First, the multiplicative model predicts zero sales for a given category, if the space of any of the stores' other categories is set to zero. Therefore, the model cannot be used for simultaneous assortment and space decisions. Second, the space solutions for 7 out of the 10 product categories were outside the model's constraints. The authors did not provide an explanation. Third, their optimization method cannot be used on non-linear functions different from the polynomial form.

Finally, metric space solutions are provided which may not match the product space requirements if the model was used for allocating space to SKUs within a category.

Bultez's and Naert's SHARP model [6] and SHARP II model [7] (Shelf Allocation for Retailers' Profit) optimized space allocation within a product class. The model was similar to that of Corstjens and Doyle in that the SHARP model incorporated both direct and cross-space effects and modeled costs as a function of sales per unit space. However, rather than develop an explicit sales-space function they performed marginal analysis on a general theoretical model and developed a search heuristic that was based on the convergence of the SHARP rule for each brand on the shelf. Space elasticities were estimated by using a symmetric attraction model in SHARP I and an asymmetric model in SHARP II. Although the model does not restrict itself to a specific formulation, the marginal analysis becomes impractical with non-linear models. This limits the number of variables as well as the functional form.

Current shelf management models focus on space responsiveness and neglect issues of assortment and stockouts. The existing models attempt to allocate space to shelves using only space elasticities which have been shown to be weak. As Lee [37] observed, using space elasticities to make assortment decisions invites problems which are compounded with the multiplicative model formulation used by many authors.

Assortment and Stockouts

Research on the effects of stockouts and limited assortments has been conducted in connection with diverse topics: retail classification [27]; variety seeking [40]; consumer store/product choice [49]; relationship of retail assortment to profitability [2] [31]; manufacturer assortment optimization projects [48]; and estimation of retail stockout costs [11] [21] [42] [50] [53] [54]. Studies evaluating consumer purchase effects of item stockouts are the most relevant for our approach.

In an in-store survey, Peckham [45] found that when a preferred item was missing 42 percent of the shoppers did not buy in the category. An intensive survey by Nielsen Marketing Research [46] [47] revealed that in some categories over 60 percent of the respondents refused to switch brands under an out-of-stock (OOS) scenario. Walter and Grabner [53] and Walter and La Londe [54] used a self-administered questionnaire to produce a frequency distribution of "intended" behaviors under a hypothesized stockout scenario. Liquor store patrons were asked their projected response if their desired item was not found. Although 82 percent would buy another item that day, 14 percent indicated that they would visit another store first. This number rose to 40 percent if the item was not available on a return trip. More recently, Emmelhainz, Stock, and Emmelhainz [22] measured the stockout behaviors within a store environment. An out-of-stock was created for the leading stock-keeping unit (SKU) in five categories. Twenty-seven percent would not switch to an alternate item if their preferred product was not available.

The literature available on consumer response to item stockouts seems to indicate that in many cases shoppers are resistant to compromise their original preferences. This implies that the assortment variable is a critical part of the shelf management decision. In addition, the simultaneity of retailer assortment and space

decisions, which is due to category space constraints, also demands that space and assortment be modeled in an interactive relationship. The shelf management model developed here accomplishes this goal and produces a model that is logically consistent and can be effectively optimized.

MODEL FORMULATION

Figure 1 presents a conceptual framework of our model. Item sales within a store are composed of four factors: unmodified demand, modified demand, acquired demand, and stockout demand. As the figure illustrates, unmodified, modified, and acquired demand are positive influences on an item's sales while stockout demand represents a loss of sales for that item. Before these factors are discussed individually, we will define our terms and assumptions.


Definitions

1. Space is the number of facings allocated to an SKU. We measure in units of a "standardized facing" and assume that a shelf has a standardized depth. These assumptions simplify the process of calculating equivalent facings and inventory when one item is dropped and another of unequal size is added. Increases in space are added in increments of a physical facing, but the addition is expressed in standardized facings.
2. Unmodified demand or preference, P_i , is the sales an SKU would receive if all SKUs were stocked and received identical retail support.
3. Modified demand is the sales an SKU receives from its unmodified demand and retail support such as promotions, space, pricing, and advertising.
 - a. In-store attractiveness, β_i , represents the incremental sales an SKU receives from retail support.
 - b. Space elasticity, γ , represents the change in demand of SKU_i per unit change in the space allocation of SKU_j .
4. Acquired demand is the sales each SKU receives from the items not selected for the assortment.
 - a. Alpha, α_i , represents the proportion of consumers who, finding their preferred item unavailable, are unwilling to switch to the available SKUs. This is used to calculate the opportunity cost of a stockout or of an item deleted from the assortment. (As pointed out by an anonymous reviewer, the loyalty that consumers display toward a brand is also a function of the particular retail stocking environment. Consumers can only be loyal to those brands which have a certain degree of availability so that they can eventually get the product to which they are loyal.)
5. Stockout demand is the sales each SKU receives from the items which have temporarily stocked out.

Assumptions

1. We assume that the retailer's objective is to maximize the category's return on inventory.

Figure 1: Shelf management model.



2. We assume that the retailer's inventory investment represents the retailer's purchase cost of a full shelf for all items in the assortment.
3. The model and objective function focuses on a given retailer's assortment and space decision.
4. Consistent with prior research and the assumption of substitutability, an SKU's direct space elasticity, γ_{ii} , is assumed to lie between 0 and 1, while its cross space elasticity, γ_{ij} , ranges from -1 to 0.
5. We assume that if an SKU is selected for the assortment, the minimum number of facings will hold one case of packages [30]. We assume that this represents the threshold for the minimum presence required for the consumer to realize the product is being carried.
6. We assume that, beyond the minimum number of facings, a retailer can stock additional items on a single package basis. This is a compromise between using a full direct product profitability model to maximize profits while accounting for use of capital and also assumes adequate backroom space to store partial cases.

Unmodified Demand

Unmodified demand represents the intrinsic preference for the SKU and is usually estimated using laboratory experiments conducted to gather brand or SKU choice. Consumers are asked to select from a group of SKUs that receive the same merchandising treatment, that is, demand for SKUs is unmodified by in-store support. Each SKU's unmodified demand is calculated using its share of choice. Unmodified demand is similar to Farris' et al. [23] unmodified preference which represented a brand's market strength exclusive of in-store support, Shugan's alpha [50] which measured a brand's absolute market potential, and Corstjens and Doyle's alpha [13] [14] which measured the effects of all marketing variables except space. Brand i 's unmodified demand (U) is presented in Equation 1.

$$U_i = P_i, \quad (1)$$

where

$$\begin{aligned} P_i &= \text{unmodified preference,} \\ \text{Space}_i &= \text{Space}_j \text{ for all } j. \end{aligned}$$

Modified Demand

The concept of modified demand reflects the differential in-store merchandising support each SKU receives. The in-store support an SKU may receive includes space, shelf location, special displays, shelf tags, backroom inventory, and window displays. A brand may also receive special support if the retailer stocks a relatively large number of its sizes or varieties. In our model we assume that differences in in-store merchandising support are solely a function of space allocation and other variables are held constant, therefore they are reflected in P_i . The “unmodified” demand of each SKU is “modified” to reflect its differential space allocation. Much of the reviewed literature supports decreasing marginal sales response to an SKU’s own space. In addition, space models typically incorporate the negative effects of space allocated to other SKUs. If some SKUs are more directly competitive than others, that is, have higher cross-elasticities, sales of an item can be affected by reallocation of space among competitive brands. The most parsimonious model that allows for competitive interactions, decreasing returns, and relative ease of estimation is the multiplicative model. Multiplicative models, including pricing [48], marketing mix decisions [52], and space models [6] [7] [13] [14] [26], are well represented in the marketing literature.

Equation 2 presents the model for modified demand (M) which incorporates both direct and cross-space elasticities. The total number of parameters in the modified demand model is $n \times n + n$ where n is the number of SKUs in the category.

$$\text{Modified Demand}_i = M_i = P_i \beta_i, \quad (2)$$

where

$$\begin{aligned} \beta_i &= \text{in-store attractiveness,} \\ &= \prod_{j=1}^n s_j^{z_j}, \\ z_i &= \text{space allocation for SKU}_i \text{ (number of facings),} \\ s_j &= 1 \text{ if } z_j = 0, \\ &= z_j \text{ if } z_j > 0, \\ s_i &= 0 \text{ if } z_i = 0. \end{aligned}$$

Acquired Demand

Acquired demand is the portion of the shelf management model that captures the effects of assortment decisions. To formulate acquired demand, assume there exists a market with n SKUs, n_1 of which are part of the category assortment, and $n - n_1$ of which have been excluded from the retailer’s shelf. Each of the n_1 SKUs stocked captures a portion of the available sales that each of the $n - n_1$ SKUs would have obtained, if they had been stocked. The available sales from each missing

SKU will be determined by its potential modified demand and the consumers willingness to switch to the n_1 available SKUs.

SKU_i 's acquired demand will consist of two parts in a multiplicative relationship.

1. SKU_i 's relative sales strength. Prior research [22] indicates that large market share SKUs receive a greater portion of the sales from those SKUs which are absent. To ensure that all of the $n-n_1$ item's available sales are distributed amongst the remaining n_1 items, an attraction model is:

$$\frac{\gamma_{ij} P_i \beta_i}{n_1} \cdot \sum_{i=1} \gamma_{ij} P_i \beta_i \quad (3)$$

The numerator represents the modified demand for SKU_i while the denominator sums the modified demand of all of the stocked SKUs. This fraction will sum to one across all stocked SKUs. The degree of substitutability with the missing SKUs will also affect the level of acquired demand. The higher the cross-elasticity value, the stronger the effects of change in the space allocation of SKU_j will have on the sales of another SKU. Combining the information available in the cross-elasticities with the relative SKU strength produces SKU_i 's share of the available sales from the non-stocked SKU_j .

2. The amount of sales available from j . The potential sales from j will depend upon its modified demand. However, some of SKU_j 's sales will be lost to the store because some fraction of buyers, represented by α_j , will be resistant to compromising their original purchase choice. Therefore, $1-\alpha_j$ represents the fraction of j 's sales that will be distributed amongst the n_1 stocked items. Equation 4 presents the demand from SKU_j (SKU_j was dropped from the assortment) that is available to be distributed to the stocked SKUs:

$$P_j \beta_j [1 - \alpha_j]. \quad (4)$$

Incorporating these factors into the model produces SKU_i 's acquired demand presented in Equation 5. To summarize, Part 1 of the equation represents SKU_i 's proportion of the available demand from non-stocked SKUs. Part 2 presents the demand from the non-stocked SKUs.

$$A_i = \sum_{j=n_1+1}^n \left[\begin{array}{cc} \text{Part 1} & \text{Part 2} \\ \frac{\gamma_{ij} P_i \beta_i}{n_1} & P_j \beta_j [1 - \alpha_j] \end{array} \right], \quad (5)$$

A = Acquired demand,

α_j = Resistance to compromise.

Stockout Demand

Finally, the complete model must include a correction for the possibility that the predicted sales from unmodified, modified, and acquired sources exceeds the shelf inventory for a given SKU. If the sum of the unmodified, modified, and acquired demand for an item is larger than the item's inventory (a stockout), then this difference is potentially available to the SKUs that are in stock. By applying the loyalty factors (α), we can determine exactly how much of a stockout is available for the other SKUs. This amount is then allocated to other items that are in stock in the same proportion as used in (4) for acquired demand. The amount allocated to an item is a stockout gain.

This procedure is iterative because when the allocation is made to other items, it may cause the inventory to be exceeded for other items. If so, a new stockout loss is calculated, and $(1-\alpha)$ of this amount is allocated to those remaining SKUs that are not stocked out. This process continues until either all items have stocked out or the sum of the unmodified, modified, and acquired demand and the total stockout gain does not exceed the inventory.

We use the same loyalty factors for both temporary out-of-stock situations and permanent adjustments to product assortment. In practice, one might argue that a temporary out-of-stock would be associated with different loyalties than one which the consumer recognizes as part of a store's assortment. We know of no published empirical studies which address this issue.

The question of whether consumer reactions to temporary "out-of-stocks" is the same as the reaction to permanent changes in assortment is also relevant to the illegal practice of "bait-and-switch" (purposely stocking out of some items in order to get consumers to buy other items which presumably have higher profit margins). An anonymous reviewer points out that this model might be relevant to the analysis of such practices if we knew more about these differences. Work by Moinzadeh and Ingene [41] is relevant to this issue.

Measurement of the Parameters

Although there are several published examples of studies reporting measures of direct and cross-elasticities to shelf space [8] [15] [16] [24] [34] [35] [44], the difficulties of obtaining good estimates should not be minimized. Lilien, Kotler, and Moorthy [38] provide an overview of these measurement problems. In practice three techniques have been used to arrive at estimates: experiments, time-series data, and cross-sectional data. Other methods are emerging which use surveys and new media technologies.

Presently, in-store experiments (manipulating shelf space and measure sales changes) are probably the most reliable ways to estimate the direct and cross-elasticities to shelf space. The major drawbacks to these studies are the time and labor needed to collect the data, and the difficulty of obtaining store approval for often disruptive shelf manipulations. Because of these problems Bultez and Naert [6] [7] used time series data and an attraction model to estimate direct space elasticities. Cortjens and Doyle [13] [14] assumed that space allocation rules collected cross-sectional data from ice cream stores to estimate direct (range -.01 to .19) and cross space elasticities (range -.11 to .10).

Although resistance to compromise has not formally been defined in the literature, consumer response to stockouts provides an approximate measure [22] [46] [47] [53] [54]. IRI and Bishop Consulting [29] deleted items and asked shoppers whether they “noticed a difference.” Other studies have either questioned individuals on their probable behavior if a desired item is missing, or have removed items from the shelf, and used scanner data to measure the sales effects of these stockouts. Depending on the category, the percentage of consumers who would purchase elsewhere if their preferred brand was missing, ranged from 6 percent to 83 percent. These numbers rose if the item was missing on a second occasion.

The potential gains from implementing the shelf management model cannot be estimated without “guessing” initial values of the parameters. These guesses can be informed by previous studies as well as cross-sectional and time-series analyses on variations in sales and shelf space. As with the IRI study [29] the most promising shelf arrangements can then be selected for in-store tests, surveys, and sales tracking studies for verification. Virtual reality technology may be creating new opportunities: consumers walking through a store on a simulated trip select items from the shelf, read labels, and place items in their shopping basket. In spite of measurement problems, there are several alternative techniques and others are emerging.

SOLVING THE SHELF MANAGEMENT PROBLEM

The shelf management model is expressed as a constrained optimization problem. The decision variable, z_i , represents the number of facings allocated to SKU_i . The objective function in (6) represents the category’s return on the retailer’s cost of total shelf inventory. The space constraint in (7) states that the sum of the space allocated to the SKUs must be equal to the stipulated category space level. Space constraints in (8) place lower and upper limits on an individual SKU’s space and establish the desired relationship between the space variables and the zero-one indicator variables.

Given the non-linearities and zero-one decision variables, it is not possible to get a closed form solution. However, the nature of the problem seems suited for simulated annealing (SA) because: (1) for the typical number of items in a category, the number of possible combinations is too large for complete enumeration; (2) functions are highly non-linear; and (3) simulated annealing provides a number of alternative solutions that can be evaluated on criteria not included in the model.

Simulated Annealing

Simulated annealing is a combinatorial optimization algorithm, which finds near-optimal solutions for different kinds of problems. It is based on the method of cooling/annealing metals. At high temperatures, the molecules in a metal can rearrange themselves fairly easily. But at low temperatures, only very limited motion is possible. With a slow and gradual decline in the temperature, and thus the energy of the system, the end product is normally a very uniform block of metal at a minimal energy state [19] [25].

Shelf Management Problem

Find $Space_i \forall_i$ in order to:

$$\text{Maximize } \Pi = \frac{\sum_{i=1}^n G_i(\text{Price}_i(M_i + A_i + B_i + L_i))}{n \sum_{i=1}^n (1 - G_i)\text{Price}_i\text{Inventory}_i},$$

subject to

$$\sum_{i=1}^n Space_i = \text{Total Category Space},$$

$$Space_i \leq \text{Total Category Space},$$

$$Space_i \geq \text{Casepack}_i, \\ = \text{if item } i \text{ is stocked.}$$

where

- Π = Category Return on Inventory,
- n = Number of SKUs in category,
- G_i = Gross Margin of SKU_i ,
- M_i = Modified Demand of SKU_i ,
- A_i = Acquired Demand of SKU_i ,
- B_i = Stockout Benefit of SKU_i ,
- L_i = Stockout Loss of SKU_i ,
- Inventory_i = Units of i on hand at beginning of period.

Our objective for the shelf management problem is to find an optimal allocation of space and assortment of SKUs so as to maximize the return on inventory. When applying simulated annealing to this optimization problem, we need to change (or randomly mutate) the set of SKUs in the assortment and the space allocated to each, and then accept those sets that improve the objective function value. However, if we do not allow for the acceptance of a few poor solutions, which may lead us to even better solutions, we may get stuck at a local maximum (like the top of a small hill, rather than the peak in a mountain range). Simulated annealing allows for this kind of a decision process.

Like the annealing of metals, earlier on in the search (cooling process) we can allow for the acceptance of fairly poor solutions. We accept a poor solution with a certain probability. This probability is fairly high earlier in the process (so as to examine an entire surface of possible solution sets or metal forms), but is steadily decreased as we iterate the process over and over again. An analogy to this would be the search for the summit of a mountain range (the optimum). Because we are fresher earlier in the day, we may be more likely to go downhill in the hope of

finding a higher mountain. But as we lose energy towards the end of the day, we would be less likely to go downhill, and climb another mountain, but rather we would climb to the summit of the mountain we have found ourselves on. However, if we kept a record of the highest mountain we had seen thus far, then we would at least have a knowledge of where we had found the highest known mountain. Similarly, we can save the best shelf arrangement during the simulated annealing process.

Simulated annealing works best for problems that have many local peaks because it avoids getting trapped at local maximums. Heuristics like SA are often applied to optimization problems where the “running time for any algorithms currently known to guarantee an optimal solution is an exponential function of the size of the problem” [19, p. 271]. The shelf management problem is an example of a problem where the number of shelf combinations (unique space allocation and assortment solution) increases exponentially as the number of items increases. For example, a category with only six products, each of which can take on the value of 0 to 12 facings, will have approximately 4.8 million (126) possible combinations. For a small sized category of 12 items the number of combinations is well over a trillion.

SA versus Alternative Heuristics

Simulated annealing (SA) has been applied in a number of diverse applications, from the travelling salesperson problem [36] to pollution control to graph partitioning [32] [33] (see [12] for an annotated bibliography of the technique). A number of others have suggested modifications of the approach including storing the best solution so far, sampling the neighborhood without replacement, and alternative acceptance probabilities. Others have attempted to improve the results by combining SA with other methods, for example, using an alternative method to provide an initial solution.

Because of the many modifications available to the basic SA algorithm, conclusions about the relative results of SA versus other heuristics are difficult to make. In addition, the results are often dependent upon the type of problem solved and the setup of the neighborhood structure selected. However, many authors have tested the algorithm against other heuristics. Eglese [19] provides an excellent review of this topic. Johnson, Aragon, McGeoch, and Schevon [32] [33] and van Laarhoven [36] demonstrated that SA had significantly better results than repeating a descent algorithm using different random starting positions. Johnson et al. [32] [33] found that: (1) in the graph partitioning problem SA outperformed traditional algorithms on random graphs but was beaten on graphs with built-in-geometric structures; (2) in the graph coloring problem, when large amounts of computing time are available, SA dominates traditional techniques and; (3) for number partitioning, SA is inferior to the differencing algorithm of N. Karmarkar and R.M. Karp. Finally, Hertz and de Werra’s [28] tabu search technique was superior to SA on the graph coloring problem while Bland and Dawson’s [3] use of SA proved better than the tabu search for layout optimization problems.

Because the shelf management problem and its feasible solutions can be clearly formulated, and a neighborhood structure explicitly defined, we felt that SA

could be used to find a good solution. The programming language, C++, was used to develop the SA algorithm.

Before we detail the SA procedure for the shelf management problem we define a few terms.

1. Objective function (Equation (6))—The total return on inventory generated by a give space allocation.
2. Trial—The process of evaluating a specific shelf allocation against the specified constraints, calculating the category return on inventory using the shelf allocation and model, and calculating the value of the acceptance function.
3. Delta (δ)—The difference in the value of the objective function in two successive trials. A poor shelf allocation will produce a negative delta.
4. Annealing schedule—A set of SA parameters that control the rate at which the probability of accepting a poor shelf allocation declines and the conditions for terminating the search.
 - a. Control parameter (T)—An annealing schedule parameter selected by the researcher to control the probability of accepting a poor shelf allocation.
5. Acceptance function—The specific function whose value represents the probability of accepting a poor shelf allocation. Shelf allocations with positive deltas are always accepted. The probability of accepting poor shelf allocations is a function of the annealing schedule.
6. Run—The sequence of trials from an initial random shelf allocation to the termination of the SA search heuristic.

Figure 2 details the simulated annealing process. The search procedure starts with a randomly selected allocation (a feasible allocation of space to the SKUs). The return on inventory for this shelf is calculated. A new shelf allocation is randomly selected within the “neighborhood” of the original one. Consistent with retailer practice, a neighborhood move is an exchange between items. This may represent simply an exchange of one facing of one item for another, or if the items are different sizes (package widths), multiple items may be involved in one neighborhood exchange. As this process continues new items are continually added or deleted from the assortment.

The category return on inventory of this new allocation is compared against the return on inventory from the previously accepted shelf set, and if it is larger, the new allocation is automatically accepted. However, if it is smaller, it may still be accepted based on the following acceptance function: $e^{(k/\delta T)}$ where: (1) δ is the difference between the value of the current and prior objective function (the smaller the value of δ the higher the probability of accepting a poorer shelf allocation); (2) T is a control parameter representing the number of trials (the higher the value of T the lower the probability of accepting a poorer shelf allocation) and; (3) k is a scaling factor. The process ends when the stopping criteria is reached.

The programming language, C++, was used to develop the SA algorithm.

A Small Category Test of the SA Heuristic

To evaluate the SA heuristic a six-SKU category example was generated, and the category return on inventory calculated using both SA and complete enumeration

Figure 2: Modified demand.

SALE	A	A	A	B	B	C	C	D
	A	A	A	B	B	C	C	D
	A	A	A	B	B	C	D	D
	A	A	A	B	B	C	D	D

of all possible shelf allocations. The total category space was 24 facings with each SKUs space constrained to lie between 0 and 12 facings. Since the focus here was an evaluation of the heuristic, the package width, retail price, and gross margin of each SKU was assumed to be identical. Using these constraints, there are 92,547 feasible shelf combinations from the 4.8 million total shelf combinations. The effectiveness of the SA algorithm can be evaluated by comparing its results against the global optimum achieved through complete enumeration.

Six-SKU Category Parameters. The reviewed literature provides an indication of the range of parameter values for the shelf management model variables. Space elasticities have been reported in [1] [7] [13] [17]. Though the term resistance to compromise has not been formally defined in the literature, outside of the push-pull model [23], the concept is often represented by out-of-stock purchase behavior [22] [46] [47] [49] [54]. The selected parameters are presented in Table 1. A more complete description of the distributions used to derive the parameters is presented later.

Six-SKU Results. Table 2 presents the shelf allocations which produced the 10 highest and 10 lowest category return on inventory levels. The mean and standard deviation across all feasible combinations were 11.48 percent and .49 percent, respectively. As revealed in Table 2, there were a number of shelf allocations that produced category return on inventory levels near the global optimum. This ability to approximate an optimal solution indicates that a search algorithm that approaches the global maximum, but does not necessarily achieve it, is a useful alternative to a complete enumeration. The bottom half of the table also reveals that making arbitrary assortment and space decisions may lead to poor returns. The approximately 3 percent difference in return on inventory between the best and poorest shelf arrangements translates into approximately 40 units of product. Given the assumption that return on inventory was during one restocking period, and the large number of categories within a typical store, the loss in profits over a year would be substantial.

In order to test the SA heuristic on the shelf management problem a range of annealing schedules was selected, and 25 runs were made for each schedule. The three parameters that constituted the annealing schedule were: (1) stopping criteria (i.e., when the search is terminated); (2) number of trials until the control parameter, T , is increased; and (3) the scaling factor, k . Each run began at a randomly selected shelf allocation.

To illustrate the search process, the annealing schedule parameters 15, 5, and 2 were selected; 15 represents the stopping criteria (e.g., after 15 consecutive trials

Table 1: Parameters for the 6-item category.

Unmodified Demand and Resistance to Compromise		
Item	P_i	Alpha
1	28.53	.45
2	23.62	.40
3	25.59	.35
4	22.40	.28
5	15.62	.19
6	10.50	.10

SKU Space Elasticities

Item	1	2	3	4	5	6
1	.1532	-.0630	-.0100	-.0089	-.0101	-.025
2	-.048	.2273	-.0159	-.0303	-.0101	-.001
3	-.0232	-.0463	.2089	-.0504	-.028	-.012
4	-.0242	-.0606	-.0628	.2143	-.030	-.024
5	-.0130	-.0571	-.0165	-.0296	.2955	-.058
6	-.0125	-.0543	-.0221	-.0239	-.074	.3104

which are not “accepted,” the run terminates; 5 is the number of trials until T is increased; and 1 is the scaling factor). A randomly assigned initial shelf allocation was selected for each SA run.

Figure 3 displays the search process for runs 3 and 5. These runs demonstrate that the search algorithm explores a number of hills and valleys as it progresses. In addition, despite the fact that the runs started at significantly different shelf allocations, they are converging to the same category return on inventory level and leveled out when the search was terminated. The global optimum (return on inventory=14.193 percent) was found in all 25 runs. The mean number of trials was 726 which represents only .7 percent of the total number of shelf allocations (92,547).

This process was repeated for each of the 48 different annealing schedules selected. Each annealing schedule was evaluated on (1) the average number of trials until the search was terminated, (2) the average maximum return on inventory reached, (3) the standard deviation of these maximums, and (4) the number of global maxima found during the 25 runs. Based on these results, the annealing schedule of stopping criteria 15, control parameter 15, and scaling factor 2 was selected. (Details of the process of selecting the best cooling schedule can be found in Borin [4].)

Conclusion for the Six-SKU Category. Our analysis suggests that SA can be used to obtain “good” shelf allocations using our shelf management model. In addition, the number of trials required to find a “good” shelf allocation is substantially below the total number of possible shelf arrangements. However, a larger category will have many more items to make alternative exchanges and may require much longer search times than an annealing schedule of 15-15-2 would produce. A

Table 2: Results from a complete enumeration of the 6-item category.

Category ROI (%)	Space					
	Item 1	Item 2	Item 3	Item 4	Item 5	Item 6
10 Best Shelves						
12.42	4	2	7	2	9	0
12.419	4	2	8	2	8	0
12.413	4	2	6	2	10	0
12.413	3	2	8	2	9	0
12.413	5	2	7	2	8	0
12.412	4	2	9	2	7	0
12.41	3	2	7	2	10	0
12.41	5	2	6	2	9	0
12.409	3	2	9	2	8	0
12.408	5	2	8	2	7	0
10 Poorest Shelves						
9.033	0	1	0	0	11	12
9.053	0	1	0	0	12	11
9.081	0	0	0	1	11	12
9.096	0	0	1	0	11	12
9.1	0	0	0	1	12	11
9.115	0	0	1	0	12	11
9.233	0	2	0	0	10	12
9.249	1	0	0	0	11	12
9.254	0	2	0	0	11	11
9.268	1	0	0	0	12	11
9.272	0	2	0	0	12	10

conservative approach would involve increasing the number of trials until a stable shelf arrangement has been reached for a repeated number of trials.

Simulated Annealing Results for a Ketchup Category

In order to further test the performance of SA on the shelf management problem, data from a ketchup category were collected from a local supermarket. The ketchup category was chosen because it has a relatively large number of items and has clearly defined boundaries.

To facilitate the exchange of different sized packages the space of each item was standardized, with the smallest bottle occupying one shelf facing. For example, each spot occupied by Heinz 64 oz. is equivalent to approximately three bottles of Heinz 14 oz. Using this measure, the ketchup category occupied 241 standard facings. Furthermore, if an item is selected as part of the category's assortment, it must have a minimum number of facings to allow all the bottles in one case to fit on the shelf. For example, Heinz 32 oz. has a case pack of 24 items and has 5 items per facing. Therefore, it must be allocated a minimum of 5 actual or 10 standard facings across. Although each item could conceivably receive all of the

Figure 3: Pseudocode for the simulated annealing algorithm.

```
Generate an initial Shelf I
Calculate ROI(I)
Repeat
Set T increment counter to 0
Repeat
Generate J a Neighbor of I (Exchange 1 facing of 2 Random Items)
Calculate ROI (J)
Delta = ROI(J) - ROI(I)
If Delta > 0
    I = J
    If ROI(J) > Max ROI
        Max ROI = ROI J
        Max Shelf = J
    Stopping Criteria = 0;
Else
    If random(0,1) > exp(k/Delta*T)
        I = J
        Stopping Criteria = 0;
    Else
        Stopping Criteria + 1;
T counter + 1;
until T counter = Increment Level
until stopping criteria is reached
```

category space, it was decided to limit the maximum number of standard facings any item could receive to be 25 percent of the total category facings, that is, 60.

The parameters found in the literature again served as a guide to the selection of the shelf management model parameters. The unmodified preferences were first selected. It was assumed that the larger the current space allocation, the higher the items unmodified preference. Therefore, P_i was estimated as each SKU's share of the current total category space. This is demonstrated in Table 3.

The selected values for the alpha, resistance to compromise, ranged from 5 percent to 45 percent and are presented in Table 3. Items with stronger preferences and larger space allocations were assumed to have higher levels of brand loyalty measured by resistance to compromise.

Consistent with the work of Bultez and Naert [6] each item's own space elasticity, γ_{ii} , varied inversely with its unmodified preference. Thus, those items with higher consumer preferences would be less responsive to increases in space allocation and therefore would have lower values of elasticity. A truncated normal distribution (minimum=0, maximum=1.0) was selected to assign each item a space elasticity. A single sample from the truncated normal distribution was drawn and is presented in the final column in Table 3. These values will be used as a measure of each item's direct space elasticity.

To select a range of values for the cross-space elasticities two assumptions were made: (1) the sales that any specific SKU will gain from an increase in its

Table 3: Unmodified preference, resistance to compromise and space elasticity for the ketchup category (each SKU's cross space elasticity, γ_{ij} , is available from the authors).

Product	Size (oz.)	Facings Across	$P_i =$ Facings/241	Alpha α_i	γ_{ii}
Private Brand	32	46	.194	.450	.109
Heinz	32	32	.109	.450	.118
Hunt	40	16	.048	.380	.156
Del Monte	44	21	.081	.420	.088
Heinz	14	7	.028	.420	.221
Heinz	40	16	.059	.440	.196
Cost Cutter	32	18	.060	.400	.178
Private Brand	40	12	.059	.250	.181
Del Monte	32	12	.041	.270	.263
Heinz	64	12	.063	.230	.134
Heinz Hot	14	4	.018	.050	.296
Hunt	44	12	.053	.260	.202
Private Brand	64	12	.045	.220	.211
Hunt	32	6	.015	.150	.264
Hunt No Salt	14	3	.008	.100	.342
Featherweight	14	2	.011	.050	.358
Heinz	28	8	.019	.180	.171
Heinz Lite	13	2	.008	.100	.304

own space cannot exceed the sales given up from the remaining SKUs, that is, there is no category sales effect from space changes; and (2) the elasticities are calculated when the sales of all items are identical. Although these two assumptions are quite strong, it permits a basis for setting the cross-space elasticities. When these two assumptions hold, the cross-space elasticities must sum to the direct elasticity in any column of the space elasticity matrix. Using these relationships, a space elasticity matrix was formed for the ketchup category. Each cross-space elastic value, within a column, was calculated by sampling from a truncated normal distribution (minimum=-1, maximum=0) with a mean equal to $\gamma_{ii}/17$ (every column had 17 cross-space elasticities). A small standard deviation (.003) was introduced to provide some variability among the cross-space elasticities. (Cross-space elasticities are available upon request from the authors.)

Ketchup Category Results. Twenty-five randomly selected shelf allocations served as the starting point for the SA run. Using the annealing schedule of 15, 15, and 2 the results were: (1) return on inventory ranged from 13.4 percent to 16.1 percent with a standard deviation of .7 percent; and (2) number of trials averaged 523. We attempted to reduce the range of return on inventory by increasing the number of trials through a new annealing schedule. Based on these results an annealing schedule of stopping criteria 25, 25 trials until T is changed, and a scaling factor of 1 was selected. Figure 4 presents the category return on inventory from the explored shelf allocations of runs 15 and 18. The graphs indicate that the search has leveled off.

Figure 4: Simulated annealing search for the six-item optimum shelf arrangement
annealing schedule 15-15-2.





Table 4: Simulated annealing results for the ketchup category.

Trials Until Exit	Maximum Category ROI (%)	Standard Facings Items																	
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
502	15.13385	44	46	4	51	0	0	26	6	0	21	0	18	24	0	0	1	0	0
1274	15.4749	58	22	4	60	0	0	10	0	0	15	0	57	15	0	0	0	0	0
1196	15.60505	42	20	0	60	0	6	32	0	0	15	0	54	12	0	0	0	0	0
725	15.65921	58	20	0	51	0	6	22	0	0	15	0	57	12	0	0	0	0	0
828	15.69959	48	22	0	51	0	0	36	0	0	15	0	57	12	0	0	0	0	0
887	15.75707	44	20	0	57	0	6	42	0	0	15	0	42	15	0	0	0	0	0
958	15.76892	38	20	0	60	0	6	48	0	0	15	0	42	12	0	0	0	0	0
825	15.78741	34	18	4	54	0	4	54	4	0	15	0	42	12	0	0	0	0	0
1264	15.78875	40	20	0	48	0	6	46	0	0	15	0	54	12	0	0	0	0	0
1441	15.7917	38	18	0	39	0	6	50	0	0	15	0	60	15	0	0	0	0	0
591	15.81682	36	18	0	54	0	6	54	4	0	15	0	42	12	0	0	0	0	0
1099	15.8415	42	20	0	51	0	0	50	0	0	15	0	48	15	0	0	0	0	0
809	15.87207	40	20	4	36	0	0	54	0	0	15	0	60	12	0	0	0	0	0
939	15.89147	52	18	0	57	0	6	48	0	0	15	0	30	15	0	0	0	0	0
1092	15.95811	40	20	0	45	0	6	58	0	0	15	0	45	12	0	0	0	0	0
1189	15.96582	52	18	4	30	0	6	50	0	0	15	0	54	12	0	0	0	0	0
967	16.02314	44	22	4	27	0	0	60	0	0	15	0	57	12	0	0	0	0	0
840	16.04101	46	26	0	15	0	6	60	4	0	15	0	57	12	0	0	0	0	0
750	16.05536	44	28	4	45	0	6	60	0	0	15	0	27	12	0	0	0	0	0
858	16.0989	52	22	4	21	0	6	60	4	0	12	0	48	12	0	0	0	0	0
541	16.10292	52	20	4	33	0	0	60	0	0	15	0	45	12	0	0	0	0	0
829	16.1031	50	20	0	33	0	0	60	0	0	15	0	48	15	0	0	0	0	0
599	16.12344	46	50	0	21	0	6	60	4	0	15	0	24	15	0	0	0	0	0
881	16.26441	60	28	0	33	0	6	60	0	0	15	0	27	12	0	0	0	0	0
841	16.34664	58	54	0	18	0	6	60	0	0	15	0	18	12	0	0	0	0	0
Mean Trials	Mean ROI	St. Dev.																	
		(%)																	
909	158.7885	.2528																	

Table 5: Comparative performance of the simulated annealing heuristic and the rule of thumb—Share of shelf = Share of sales.

Heuristic	ROI (%)	Decrease in ROI from Maximum (%)
Simulated Annealing	16.47	NA
Share of Shelf = Share of Sales (Identical Assortment)	12.70	23
Share of Shelf = Share of Sales (Total Assortment)	11.21	32

Table 4 lists the results for each of the 25 runs. The return on inventory range is now only 1.2 percent, with 24 of the runs within 1 percent of each other, demonstrating a convergence towards the same return on inventory level. The maximum return on inventory for the category was found to be 16.46 percent return on inventory with the shelf allocation indicated in Table 4. Eleven of the 18 items were not chosen as part of the final assortment.

Figure 4 and Table 5 clearly demonstrate that a randomly selected shelf arrangement can produce significantly poorer results than that achieved with the simulated annealing algorithm. The simulated annealing search improved the ketchup category's return on inventory by approximately 7.5 percent from the initial starting point. However, alternative shelf management methods exist that also assist retailers with space allocation. We now turn to one such method.



COMPARATIVE ANALYSIS WITH MANAGEMENT RULES OF THUMB

The performance of the SA heuristic is compared to the results achieved using a proportional shelf allocation rule. One of the most frequently used methods a retailer uses to determine space for each SKU is to allocate space to an SKU in approximate proportion to its historical share of category unit sales. Those SKUs with more sales will earn more shelf space than those SKUs which do not sell well. In addition, if there is a minimum stocking requirement (at least one case), the best selling items will always have proportionately less space than slower selling items. This method suffers from the following disadvantages.

1. The causal direction is assumed to flow from sales to space.
2. It is only a space allocation rule. A retailer must have a priori decided on the assortment before using this rule.
3. The method assumes a linear relationship between space and sales.

The shelf management model was used to calculate the item and category sales when used with a rule that specifies an SKU's space allocation equal to its share of unit sales. The rule was used on both the total assortment of items within the ketchup category and on the best assortment found with the SA algorithm. Standard facings and minimum and maximum facing constraints were observed. Given a shelf allocation for each stocked SKU, the sales for each SKU were calculated

Figure 5: Simulated annealing search for the eighteen item ketchup category annealing schedule 25-25-1.



using the shelf management model. The resulting shelf allocation was calculated using the rule of thumb, the shelf management model was rerun, and a new set of SKU and category sales was calculated. This process was repeated until the same shelf configuration appeared on successive occasions. Nine different starting shelf allocations were used. Using this scheme a stable shelf allocation was found in under 20 iterations, and the final allocation reached was identical in eight out of nine cases. The one shelf location that differed varied only in the space allocation between two SKUs.

The return on inventory produced from the final shelf arrangement using the share of shelf decision rule is displayed in Table 5. Even when applied to an optimal assortment, allocating space according to a proportionality rule led to a return on inventory reduction of 23 percent. Aggregated across time periods and categories, this would lead to a significant drop in profits for the retailer. One of the major strengths of the SA algorithm is its ability to escape from a local optimum. Since the rule of thumb is not a search routine, it cannot perform this function and will unlikely lead to the best shelf allocation.

CONCLUSION

The results demonstrate a shelf management model that incorporates both space and assortment effects. More importantly, the SA search heuristic finds a "good" shelf allocation within a relatively small number of trials. The algorithm is flexible enough to allow for alternative category and package sizes and different restocking practices such as by item or by case. The comparative analysis between different shelf stocking rules and algorithms reveal that retailers who ignore assortment/stockout effects, or base their space allocation on proportionality rules, may be losing substantial amounts of sales.

Future research in shelf management might investigate how optimal shelf arrangements change depending upon whether the objective function is sales, dollar profit, gross margin return on inventory, or another measure which would include costs of restocking, such as direct product profitability.

Extensions of the model could include multiple objectives, such as minimizing stockouts or maximizing assortment within a category. The authors are presently working on an analysis of the sensitivity of the model to errors in the parameter estimates to help determine the circumstances under which its application is best suited. This work will help address the concerns about the level of measurement accuracy required to support a given model formulation. [Received: January 28, 1993. Accepted: May 16, 1994.]

REFERENCES

- [1] Anderson, E. E. An analysis of retail display space: Theory and methods. *Journal of Business*, 1979, 52(1), 103-118.
- [2] Bawa, K., Landwehr, J. T., & Krishna, A. Consumer response to retailers' marketing environments: An analysis of coffee purchase data. *Journal of Retailing*, 1989, 65(Winter), 471-495.
- [3] Bland, J. A., & Dawson, G. P. *Tabu search applied to layout optimization*. Report, Department of Maths, Stats and O.R., Trent Polytechnic, UK, presented at CO89, Leeds, July 1989.

- [4] Borin, N. An analysis of space and assortment tradeoffs in the retail environment. Unpublished doctoral dissertation. The Darden Graduate School of Business Administration, University of Virginia, 1992.
- [5] Brown, W. M., & Tucker, W. T. Vanishing shelf space. *Atlantic Economic Review*, 1961, 11(10), 9-13.
- [6] Bultez, A., & Naert, P. S.H.A.R.P.: Shelf allocation for retailers' profit. *Marketing Science*, 1988, 7, 211-231.
- [7] Bultez, A., Naert, P., Gijsbrechts, E., & Abelle, P. V. Asymmetric cannibalism in retail assortments. *Journal of Retailing*, 1989, 65(2), 153-192.
- [8] Burgoyne, D. G., & Johnston, C. B. Are shelf space and shelf location really important? *The Business Quarterly*, 1968, 33, 56-60.
- [9] Cairns, J. P. Suppliers, retailers, and shelf space. *Journal of Marketing*, 1962, 26, 34-36.
- [10] Carpenter, G. S., & Lehmann, D. R. A model of marketing mix, brand switching and competition. *Journal of Marketing Research*, 1985, 22, 318-329.
- [11] Chang, Y. S., & Niland, P. A model for measuring stock depletion costs. *Operations Research*, 1967, 15, 427-447.
- [12] Collins, N. E., Eglese, R. W., & Golden, B. L. Simulated annealing: An annotated bibliography. *American Journal of Mathematical Management Science*, 1988, 8, 209-307.
- [13] Corstjens, M., & Doyle, P. A model for optimizing retail space allocations. *Journal of Marketing*, 1981, 27, 822-833.
- [14] Corstjens, M., & Doyle, P. A dynamic model for strategically allocating retail space. *Journal of the Operational Research Society*, 1983, 34(10), 943-951.
- [15] Cox, K. K. The responsiveness of food sales to shelf space changes in supermarkets. *Journal of Marketing Research*, 1964, 1, 63-67.
- [16] Cox, K. K. The effect of shelf space upon sales of branded products. *Journal of Marketing Research*, 1970, 7, 55-58.
- [17] Curhan, R. C. The relationship between shelf space and unit sales in supermarkets. *Journal of Marketing Research*, 1972, 9, 406-412.
- [18] Curhan, R. C. Shelf space allocation and profit maximization in mass retailing. *Journal of Marketing*, 1973, 37, 54-60.
- [19] Eglese, R. W. Simulated annealing: A tool for operational research. *European Journal of Operational Research*, 1990, 46, 271-281.
- [20] Ehrenberg, A. S. C. An appraisal of Markov brand-switching models. *Journal of Marketing Research*, 1965, 2, 347-362.
- [21] Elton, M. C. J., & Mercer, A. Estimating the effect of variety on sales. *Operational Research Quarterly*, 1969, 20, 351-360.
- [22] Emmelhainz, M. A., Stock, J. R., & Emmelhainz, L. W. Consumer responses to retail stock-outs. *Journal of Retailing*, 1981, 67, 138-147.
- [23] Farris, P. W., Olver, J., & De Kluyver, C. The relationship between market share and distribution. *Marketing Science*, 1989, 8(2), 107-128.
- [24] Frank, R. E., & Massy, W. F. Shelf position and space effects on sales. *Journal of Marketing Research*, 1970, 7, 59-66.
- [25] Glover, F., & Greenberg, H. J. New approaches for heuristic search: A bilateral linkage with artificial intelligence. *European Journal of Operational Research*, 1989, 39, 119-130.
- [26] Hansen, P., & Heinsbroek, H. Product selection and space allocation in supermarkets. *European Journal of Operational Research*, 1979, 3, 474-484.
- [27] Hartley, R. F. *Retailing challenge and opportunity*. Boston, MA: Houghton Mifflin Company, 1984.
- [28] Hertz, A., & de Werran, D. Using tabu search techniques for graph colouring. *Computing*, 1987, 39, 345-351.
- [29] Information Resources, Inc. & Willard Bishop Consulting, Ltd. *Variety or duplication: A process to know where you stand*. Washington, D.C.: Food Marketing Institute, 1993.
- [30] Ireland, J. J. Estimating direct product profit from inexpensive data. Working Paper. Universitat Pompeu Fabra.
- [31] Judd, L. L., & Vaught, B. C. Three differential variables and their relation to retail strategy and profitability. *Journal of the Academy of Marketing Science*, 1988, 16, 30-37.
- [32] Johnson, D. S., Aragon, C. R., McGeoch, L. A., & Schevon, C. Optimization by simulated annealing: An experimental evaluation; Part I, Graph partitioning. *Operations Research*, 1989, 37, 865-892.

- [33] Johnson, D. S., Aragon, C. R., McGeoch, L. A., & Schevon, C. Optimization by simulated annealing: An experimental evaluation; Part II, Graph coloring and number partitioning. *Operations Research*, 1991, 39, 378-406.
- [34] Kotzan, J. A., & Evanson, R. V. Responsiveness of drug store sales to shelf space allocations. *Journal of Marketing Research*, 1969, 6, 465-469.
- [35] Krueckeberg, H. F., & Davis, P. C. Consumer response to space allocation in the supermarket dairy department. Bulletin #374. Agricultural Experiment Station, University of Delaware, Newark, Delaware, 1966.
- [36] Larrhoven, P. J. M. van. Theoretical and computational aspects of simulated annealing. Unpublished doctoral dissertation, Erasmus University, Rotterdam, 1988.
- [37] Lee, W. Space management in retail stores and implications to agriculture. In W.K. Dolva (Ed.), *Marketing keys to profits in the 1960's*. Chicago, IL: American Marketing Association, 1961, 523-537.
- [38] Lilien, G. L., Kotler, P., & Moorthy, K. S. *Marketing models*. NJ: Prentice Hall, 1992.
- [39] *Marketing News*. Pilot study finds final product choice usually made in store. 1982, August 6, 5.
- [40] McAlister, L., & Pessemier, E. Variety-seeking behavior: An interdisciplinary review. *Journal of Consumer Research*, 1982, 9(3), 311-322.
- [41] Moinzadeh, K., & Ingene, C. An inventory model of immediate and delayed deliver. *Management Science*, 1993, 39(5), 536-548.
- [42] Motes, W. H., & Castleberry, S. B. A longitudinal field test of stockout effects on multi-brand inventories. *Journal of the Academy of Marketing Science*, 1985, 13, 54-68.
- [43] Nielsen Marketing Research. *Category management. Positioning your organization to win*. Chicago, IL: NTC Business Books, 1992.
- [44] Pauli, H., & Hoecker, R. W. Better utilization of selling space in food stores. U.S. Department of Agriculture, Marketing Research Report No. 30, November 1952.
- [45] Peckham, J. O. The consumer speaks. *Journal of Marketing*, 1963, 27, 21-26.
- [46] *Progressive Grocer*. The out of stock study: Part I. 1968, 47(October), S1-S16.
- [47] *Progressive Grocer*. The out of stock study: Part II. 1968, 47(November), S17-S32.
- [48] Reibstein, D. J., & Gatignon, H. Optimal product line pricing: The influence of elasticities and cross-elasticities. *Journal of Marketing Research*, 1984, 21, 259-267.
- [49] Schary, P. B., & Christopher, M. The anatomy of a stock-out. *Journal of Retailing*, 1979, 55(2), 59-70.
- [50] Shugan, S. M. Product assortment in a triopoly. *Management Science*, 1989, 15, 304-320.
- [51] Swasy, A. Firms grow more cautious about new-product plans. *Wall Street Journal*, 1989, March 9, B1.
- [52] Urban, G. L. A mathematical modeling approach to product line decisions. *Journal of Marketing Research*, 1969, 6, 40-47.
- [53] Walter, C. K., & Grabner, J. R. Stockout cost models: Empirical tests in a retail situation. *Journal of Marketing*, 1975, 39, 56-68.
- [54] Walter, C. K., & La Londe, B. J. Development and tests of two stockout cost models. *International Journal of Physical Distribution*, 1975, 5(3), 121-132.

Norm Borin is Associate Professor of Marketing at California Polytechnic State University in San Luis Obispo, California. He received his Ph.D. from the University of Virginia, his M.B.A. from the California State University, Sacramento, and a B.S. from the University of California, Davis. His current research interests include channel management with a focus on shelf management models and category management. His papers have been published in the *Journal of Retailing*, *Journal of Direct Marketing*, and the *Journal of Education for Business*. He is a member of AMA, TIMS, and AAA.

Paul W. Farris is Landmark Communications Professor of Business Administration at the Darden School, University of Virginia. He has a B.S. in business economics from the University of Missouri, an M.B.A. from the University of Washington, and a D.B.A. from Harvard University. His research has focused on the assessment of marketing costs and productivity.

James R. Freeland is the Sponsors Professor of Business Administration and the Associate Dean for Faculty at the Darden Graduate School of Business Administration at

the University of Virginia and was formerly at Stanford University. He received his M.S. and Ph.D. in industrial and systems engineering from Georgia Institute of Technology and his B.S. in industrial engineering from Bradley University. He is the author of numerous technical and managerial papers in operations management and management science.