

The Assortment Packing Problem: Multiperiod Assortment Planning for Short-Lived Products

Felipe Caro* Victor Martínez-de-Albéniz[†] Paat Rusmevichientong[‡]

January 16, 2012

Abstract

Motivated by retailers' frequent introduction of new items to refresh product lines and maintain their market share, we present the assortment packing problem in which a firm must decide, in advance, the release date of each product in a given collection over a selling season. Our formulation models the trade-offs among profit margins, preference weights, and limited life cycles. A key aspect of the problem is that each product is short-lived in the sense that, once introduced, its attractiveness lasts only a few periods and vanishes over time. The objective is to determine when to introduce each product to maximize the total profit over the selling season. Even for two periods, the corresponding optimization problem is shown to be NP-complete. As a result, we study a continuous relaxation of the problem that approximates the problem well when the number of products is large. When margins are identical and product preferences decay exponentially, its solution can be characterized: it is optimal to introduce products with slower decays earlier. The relaxation also helps us to develop several heuristics, for which we establish performance guarantees. Numerical experiments show that these heuristics perform very well, yielding profits within 1% of the optimal in most cases.

1. Introduction

Keeping customers interested is one of the challenges in industries with short-lived products. Firms launch products frequently in order to keep their presence in the market place and capture the attention of customers. This forces firms to plan their assortments over time. Indeed, carrying a static assortment – one that remains the same over time – becomes ineffective and possibly unprofitable because consumers are quickly “bored” with the choices within assortment and they divert their purchases to other consumption options. In other words, the customers' preference for a particular product in the assortment decays over time, as it ages on the shelf. Moreover, due to substitution effects, it might not be optimal for a firm to release all the products at once because it dilutes its market share, so timing product entry becomes relevant.

An industry where assortment renewal strategies have been massively adopted is apparel retailing. The industry traditionally used to launch two collections a year and push them to the stores at the beginning of the Spring and Autumn (specifically, once the discounting season is over). Over the last decade, powerful players such as the Swedish clothing retailer H&M has chosen to continuously release their products into the stores. H&M claims that it introduces new products into stores daily (H&M 2007). Interestingly, the timing of release is not necessarily linked to when a product is

*UCLA Anderson School of Management, Los Angeles, CA 90095, fcaro@anderson.ucla.edu

[†]IESE Business School, University of Navarra, Barcelona, Spain, valbeniz@iese.edu

[‡]USC Marshall School of Business, Los Angeles, CA 90089, rusmevic@marshall.usc.edu

designed or produced. Indeed, H&M is well known for selling mostly *basic* and *fashion-basic* items, which are respectively products with low or medium fashion content, e.g., neutral colors such as black and white, standard cuts and few embellishments such as ribbons (see Abernathy et al. 1999 for a more detailed discussion on basics and fashion-basics). These products are typically designed and ordered long in advance, but are only required at the store in the middle of the selling season. Similarly, Florida-based Chico's claims that "You'll find something new every day at Chico's.", and the new products introduced in store have a six-month lead-time, which means that the company decouples design decisions and introduction timing (Tiplady 2006).

Besides apparel, other industries exhibit similar dynamics. Book stores typically announce and promote "recent arrivals" but then move them to less prominent locations because the attractiveness of the product quickly drops after introduction. Restaurants also frequently change the items on their menu to avoid customer satiation, with some restaurants committing to completely change the menu every three months (Moskin 2011). Finally, motion picture distributors schedule the release dates of their movies over a season. Since customers and theaters tend to prefer new releases, it is important to avoid having too many movies in the theaters at the same time to prevent cannibalization of ticket sales (The Economist 2007).

While this trend seems to provide better value to consumers from a marketing standpoint, it leads to a more complex execution from an operations perspective because the firm needs to decide how to deploy a catalog of products in stores over a season. Traditionally, the entire catalog was pushed to the store in the first day of the season, and would not be changed over time. However, such a strategy is no longer appropriate when products are short-lived. Indeed, under this strategy, the store would have big spike in sales at the beginning of the season but then it would see its market share drastically shrink. In contrast, a strategy of continuously releasing a few products every period would keep some products' valuation high, balancing the low valuation of items that were introduced long ago, maintaining the firm's market share later in the season.

With a large number of products, determining each product's release date in order to maximize the total profit over a season is challenging. In practice, companies typically solve this problem manually, with a qualitative assessment of the value of the assortment for each week. For example, we have interacted with a Spanish accessories retailer that introduced a large number of items at the beginning of the season and then added a few items to the assortment each week. The decision process was completely manual, and it was based purely on price and visual characteristics of the items. How to arrange product introductions over time constitutes a difficult multiperiod assortment planning problem. Our objective in this paper is to shed some light on the issue, and

requires us (i) to develop a model to capture the trade-offs related to product introductions, (ii) to provide exact or approximate solutions to support introduction decisions, and (iii) to derive some managerial insights to guide retailers on these issues.

For this purpose, we present the assortment packing problem in which a planner must decide, in advance, when to introduce each product in a given collection over a selling season. The products have different profit margins, preference weights, and life cycle patterns. Once a product is introduced to the store, the product becomes stale and its attractiveness decays over time. The objective is to determine if and when to release each product to maximize the total profit over the selling season. We model demand in the form of market shares following an attraction model (Bell et al. 1975). The model takes as its inputs each product's profit margin, its initial preference weight (i.e., the value of the product in the attraction model when it is introduced), and its preference decay (i.e., the decay of the product's preference weight over time). As an output, the model determines the the release date of each product over the season that maximizes the firm's discounted market share.

When the season has a single period, the assortment packing problem reduces to a classical static assortment optimization problem under the attraction demand model, which can be solved efficiently (see, for example, Talluri and van Ryzin 2004, Kök et al. 2008). However, even with just two periods, the problem becomes NP-hard and the particular structure found in the one-period assortment problem is lost. As a result, we study a relaxation of the problem that corresponds to the situation where there are many different products of each type, and we show that it approximates the original problem well. We can characterize the optimal solution of this relaxation and find that, when margins are identical and product preferences decay exponentially, it is optimal to introduce products sequentially, in a decreasing order of preference decay, i.e., always introduce a product with slower decay first. We then build several heuristics to solve the original problem, for which a performance guarantee can be derived. As shown in our numerical experiments, these heuristics perform remarkably well, yielding profits within 1% of the optimal in most cases.

Our work thus contributes to the literature in three dimensions: modeling, methodological development and managerial insights. While there has been a number of papers on understanding the benefits of frequent assortment changes on demand learning (see, for example, Caro and Gallien 2007, Alptekinoglu et al. 2011), there has been less attention on multiperiod assortment planning when learning is not a predominant issue. In that context, we provide a formulation that explicitly models when each product is introduced, and determines the optimal assortment in each period taking into account the natural decay that occurs over time.

2. Literature Review

This paper is closely related to the literature on assortment planning and product roll-over in operations management. Our work extends the static assortment planning problem to a multiperiod setting. Kök et al. (2008) provide an excellent review of the literature in this area. The assortment optimization problem is typically driven by the trade-off between product margins and sales. Indeed, a larger assortment induces a higher sales probability, but at the same time, it also includes lower-margin items. Talluri and van Ryzin (2004) first show in a single-period revenue management setting that the optimal assortment is a set comprising of highest-margin products. In more general contexts, the assortment problem does not have a simple solution. For instance, Bront et al. (2009) show that when there is an heterogeneous population the single-period problem becomes NP-Hard.

An important stream of operations papers focus on the relationship between variety benefits and inventory costs, and they typically assume a certain product substitution structure when stockouts occur (see, for example, van Ryzin and Mahajan 1999, Smith and Agrawal 2000). Cachon et al. (2005) explicitly account for customer search in the assortment planning process and show that ignoring customer search in demand estimation can lead to sub-optimal assortment decisions and lower expected profits. Kök and Fisher (2007) also develop a method for estimating demand under substitution. In most of these papers, the inventory decision follows the assortment decision, in the sense that a newsvendor-type formula is used to compute the inventory cost for a given assortment. As a result, the evolution of the assortment (as some products “die”) is ignored. There are some exceptions: Mahajan and van Ryzin (2001) show that when substitution is stockout-based, the problem quickly becomes intractable and profits are no longer concave in inventory levels; Honhon et al. (2010) propose a dynamic programming formulation to find the optimal stocking levels when there is a fixed proportion of customer types and random demand; Bernstein et al. (2010) dynamically adjust the assortment over time, depending on each customer’s preferences and the remaining inventory, so as to “hide” products with low inventories to reserve them for future customers. Furthermore, pricing decisions have also been considered: Aydin and Porteus (2008) study the joint inventory and pricing decision for an assortment. The competitive aspect of these assortment and price decisions has been studied in Besbes and Sauré (2010), under the multinomial logit (MNL) model. In our paper, we consider the assortment decision only, in a multiperiod environment.

Dynamic assortments have been studied before, usually in a context where the underlying demand function is unknown and must be estimated, and inventory is ignored. Caro and Gallien (2007) apply a finite-horizon multi-armed bandit model with Bayesian learning to a dynamic as-

assortment problem for seasonal goods, and derive a closed-form dynamic index policy that captures the key exploration versus exploitation tradeoffs. Rusmevichientong et al. (2010) and Sauré and Zeevi (2008) consider the case of MNL demand model with a capacity constraint, and present an assortment planning algorithm that simultaneously learns the underlying MNL parameters and optimizes profit. Farias and Madan (2011) consider the setting where a product cannot be used once it is removed from the assortment. Farias et al. (2010) consider the problem of estimating a choice model that is defined as a distribution over all permutations, where each permutation defines a preference-ordered list of products. Alptekinoglu et al. (2011) also learn about customer preferences in a locational model by dynamically adjusting the assortment. Honhon and Kök (2011) study the impact of variety-seeking customers and identify cyclical patterns in the optimal assortment. All of the papers mentioned thus far assume that the attractiveness of each product remains the same over time. To our knowledge, we are the first to explicitly model the reduction of a product’s attractiveness over its life cycle: we propose an attraction model in which product preference weights decay over time. A similar approach has been used in Ainslie et al. (2005), where box-office sales are modeled as a MNL demand model where the market attractiveness of each movie follows a Gamma distribution of its age in the theaters (and hence eventually decays over time). That paper provides the empirical methodology that can potentially be used to calibrate the decay parameters in our assortment packing model. It does not study how to time movie releases, which is the focus of our work.

In addition, product introduction timing has also been studied by the product rollover literature. Most of the work relevant to this paper studies the trade-off between the market expansion associated with a new product and the saturation effect, i.e., how the older product is cannibalized by the new one. Lim and Tang (2006) study whether introducing a new product and removing an older one should be done simultaneously or not, and characterize the best product prices associated with these strategies. Demand diffusion models have also been used to analyze the timing decision. For instance, Kalish et al. (1995) study whether two markets need to be entered simultaneously or sequentially. With a similar diffusion pattern, Savin and Terwiesch (2005) study the timing of entry for a new product competing against an established one. Druehl et al. (2009) analyze the pace of product introduction over multiple product generations with market growth, cannibalization, margin decays and pace-dependent introduction costs.

3. The Assortment Packing Problem

In this section, we describe the assortment packing problem and discuss the features of our model in detail, with examples to demonstrate the qualitative nature of the solution. We then study its

computation complexity. Note that all the proofs are contained in the Appendix.

3.1 Model Formulation.

Consider a retailer that is planning the assortment for a season consisting of T periods (e.g., weeks). The retailer has n products available, and must decide when to introduce each product $i \in I := \{1, \dots, n\}$. For $i \in I$ and $t = 1, \dots, T$, let x_{it} be the binary variable that is equal to 1 if product i is introduced in period t , and 0 otherwise. Each product can only be introduced once; after it is introduced, it remains in the assortment for several consecutive periods. Let r_i denote the unit gross margin of product i . Without loss of generality, assume that the products are numbered so that $r_1 \geq r_2 \geq \dots \geq r_n > 0$. Period t has a discount or seasonality factor α_t .

We assume an attraction demand model, where each product's market share contribution is proportional to its preference weight or attractiveness in each period. Let v_i denote the weight of product i when it is first introduced. As the product remains in the store, however, its attractiveness changes over time. For instance, the product might become stale, in which case its attractiveness decays as time goes by. If product i is introduced in period t_i , then we assume that its weight in period $t \geq t_i$ is given by $\kappa_{i,t-t_i} v_i$, where $\kappa_{i,d} \in [0, 1]$ is a product-specific decay parameter.¹ In particular, if the item has a planned life cycle of ℓ_i periods, then $\kappa_{i,d} = 0$ for $d \geq \ell_i$. In any period, a customer can choose not to purchase from the assortment. This outside option has a weight v_0 , independent of time. As a result, the assortment packing problem (*APP*) can be formulated as the following nonlinear combinatorial optimization problem:

$$\begin{aligned}
 (APP) \quad V^* = \max \quad & \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \times \left(\frac{v_i \sum_{u=1}^t \kappa_{i,t-u} x_{iu}}{v_0 + \sum_{j=1}^n v_j \sum_{u=1}^t \kappa_{j,t-u} x_{ju}} \right) & (1) \\
 \text{s.t.} \quad & \sum_{t=1}^T x_{it} \leq 1 & \forall i \in I, & (2) \\
 & x_{it} \in \{0, 1\} & \forall i \in I, t = 1, \dots, T. & (3)
 \end{aligned}$$

The objective function (1) is the sum of the discounted (gross) profits over all products from each of the T periods in the season. The expression in the parentheses represents the market share of product i in period t . Constraint (2) ensures that each product is introduced at most once, and constraint (3) imposes the binary requirement on the decision variables.

For notational convenience, let $z_{it} := \sum_{u=1}^t \kappa_{i,t-u} x_{iu}$ denote the contribution of product i to the assortment's attractiveness or load in period t . Clearly, $z_{it} = 0$ if product i is introduced after period t . On the same lines, we define $z_t := \sum_{i=1}^n v_i z_{it}$ as the total load in period t and $\phi_t := z_t / (v_0 + z_t)$ as

¹We refer to $\kappa_{i,d}$ as a decay parameter, but in general the model allows for life cycles that are not necessarily decreasing.

the firm’s market share in period t . Throughout the paper, we say that two products i and j have the same *type* if they have exactly the same decay pattern, i.e., $\kappa_{i,d} = \kappa_{j,d}$, for $d \geq 0$. In particular, we refer to products with no decay as *basics*. In other words, a product i belongs to the basic type if $\kappa_{i,d} = 1$ for all $d \geq 0$. This represents products with an attractiveness that might be low but remains constant from the release date until the end of the season, which is characteristic of basic apparel items as pointed out in the introduction. Interestingly, it is easy to show that when a basic product i is such that $r_i \geq r_j$ for all j , then it is optimal to introduce it in the first period, and hence this product can be omitted from the analysis (although i ’s preference weight will still appear in the fractional terms of the objective functions).

3.2 Model Discussion. We review here the model’s features, justify our modeling assumptions, and describe settings where our model might be applicable. We believe that this model strikes a good balance between maintaining tractability and capturing important constraints faced by the retailers in planning their product introductions.

First, we can see that the problem formulation requires that an open-loop policy is used. In other words, one must decide in advance when to introduce the different products, and cannot adjust release dates during the selling season. This assumption is motivated by applications where products have a long production lead time and a relatively short selling season, making it impractical to re-order products during the season. As pointed out in the introduction, this is a reasonable assumption for basic and fashion-basic apparel, books, accessories or movies. As a result, since there is no possibility of adjusting the assortment for demand learning (in contrast to Caro and Gallien 2007), any uncertain demand can be translated into its deterministic counterpart, i.e., its expectation.

One key characteristic of our model is the formulation of the demand function. We use an attraction model where the sales of product j in time t is given by $\alpha_t \times \frac{v_j z_{jt}}{v_0 + \sum_{j=1}^n v_j z_{jt}}$. Such form is one of the most commonly used demand models in marketing and operations management to capture assortment-based substitution (Kök et al. 2008). In period t , a number α_t of consumers are ready to buy a product, and then choose product i with probability $\frac{v_i z_{it}}{v_0 + \sum_{j=1}^n v_j z_{jt}}$. There are well-established techniques for estimating the initial attractiveness v_i associated with each product (see e.g., Talluri and van Ryzin 2004). In addition, Ainslie et al. (2005) discuss the empirical toolbox for estimating the decay parameters $\kappa_{i,d}$; it is worth pointing out that, in their case, the weight of blockbuster movies decays exponentially, a functional form of $\kappa_{i,d}$ which we study in depth later.

In contrast with single-period models, where z_{it} is either one if the product is introduced, or

zero if it is not. In our model, the contribution of each product depends on *when the product was introduced*. Specifically, if the product has not been introduced yet at time t , then $z_{it} = 0$; if the product has just been introduced at t , then $z_{it} = 1$; otherwise, the product has been in the assortment for a few periods, since t_i , and $z_{it} = \kappa_{i,t-t_i} \leq 1$, which implies that the value of the item in the consumers' eyes is lower than at introduction. This feature captures the customers' preference for new items, which is prevalent for apparel, books, restaurant menus or movie consumptions. As a result, it is clearly sub-optimal to introduce all the items in the first period. Indeed, if that was the chosen solution, sales would be high in the beginning of the horizon, but limited due to strong product competition; in contrast, they would be low at the end of the horizon, since the load of later periods z_t would be much reduced compared to the outside option.

Furthermore, note that each product is introduced only once; it stays in the assortment until the end of the horizon. Limited product life cycles can be modeled by setting $\kappa_{i,d} = 0$ for $d \geq \ell_i$ where ℓ_i is the planned life of the item. This implies that the product introduction decision is irrevocable, similar to Farias and Madan (2011). Although sometimes it may be possible to show a product to the customers, then remove it from the assortment and then reintroduce it again, this is usually not the case in retail settings: it would be too complex and expensive to execute, involving handling, logistics and merchandizing costs. For instance, most apparel retailers introduce a product only once, and keep it in the stores for a number of periods.

Finally, to keep the problem tractable and focused on the product introduction question, we do not consider inventory decisions. In practice, this is a decision that is taken after assortment plans are finalized, and hence can be set sufficiently high later on so as to avoid losing sales. Similarly, we do not consider budget or shelf-space constraints, e.g., a maximum number of products to be introduced over the horizon or a maximum number of items per period. We discuss how these change our analysis and results in Section 7.

3.3 Computational Complexity. When there is no outside option ($v_0 = 0$) and $\alpha_t = 1$ for all t , the *APP* is trivial because it is optimal to introduce the most profitable product r_1 in the first period and to have it as the sole product in the assortment for the entire duration of its life cycle. The reason is that having an assortment with more products yields a per-period profit that is a weighted average of the individual margins r_1, \dots, r_n , which clearly cannot exceed r_1 . Once the attractiveness of the first product declines to zero, it is optimal to introduce the second most profitable product r_2 , and so on and so forth.

Another case that can be solved efficiently is when all products are basic, with $\kappa_{i,d} = 1$ for all i and d . In this case, the optimal solution is myopic, and the problem is equivalent to a single-

period setting with $T = 1$. Talluri and van Ryzin (2004) show that for a single-period setting it is optimal to introduce products with the highest margins, corresponding to a revenue-ordered subset $A_k = \{1, \dots, k\}$ for some product k . As shown in the following example, the nice structure is not present in the *APP* with decay parameters, and in general, it may be beneficial to introduce products with lower margins in the beginning of the season.

Example 1 (Lower Margin Products May Be Introduced First). Suppose that $\alpha_1 = \alpha_2 = 1, n = 2, T = 2, r_1 = \$10, r_2 = \$9, v_0 = 1, v_1 = 3, v_2 = 7$, and $\kappa_{1,d} = \kappa_{2,d} = 0.4^d$ for $d = 0, 1$. For this instances it can be verified that the optimal release schedule is to first introduce product 2 with lower margin in period 1 and only release the higher margin product (product 1) in period 2. Intuitively, by releasing product 1 later, it preserves its attractiveness, which allows for a higher profit from offering both products together.

The example above suggests that the general *APP* is computationally intractable, which is formally established below: Theorem 1 shows that, even when there is only two periods, no discount, all products have equal margin, and a single-period life cycle with $\kappa_{i,0} = 1$ and $\kappa_{i,d} = 0$ for all i and $d \geq 1$, the problem remains NP-hard. The proof follows from a reduction from the well known NP-complete PARTITION problem. We first formulate a decision-theoretic version of our two-period *APP*.

TWO-PERIOD EQUAL-MARGIN ASSORTMENT PACKING WITH ONE-PERIOD LIFE CYCLE

INPUTS: The set of products indexed by i ; the preference weights v_1, v_2, \dots, v_n , where $v_i \in \mathbb{Z}_+$ for $i = 1, \dots, n$; a no-purchase weight $v_0 \in \mathbb{Z}_+$; and the target profit $K \in \mathbb{Q}_+$;

QUESTION: Is there a partition S_1 and S_2 such that $S_1 \cap S_2 = \emptyset$ and $S_1 \cup S_2 = \{1, \dots, n\}$ such that

$$\frac{\sum_{j \in S_1} v_j}{v_0 + \sum_{j \in S_1} v_j} + \frac{\sum_{j \in S_2} v_j}{v_0 + \sum_{j \in S_2} v_j} \geq K.$$

Theorem 1. *For any $v_0 \in \mathbb{Z}_+$, the TWO-PERIOD EQUAL-MARGIN ASSORTMENT PACKING WITH ONE-PERIOD LIFE CYCLE is NP-complete.*

4. The Continuous Relaxation

In the previous section, we observe that the *APP* is computationally intractable even in the simplest setting, making it difficult to establish managerial insights and structural properties of the optimal solution, and to develop an algorithm for solving this problem. In this section, we introduce a continuous relaxation of the original problem, by replacing the binary constraints $x_{it} \in \{0, 1\}$ by $x_{it} \geq 0$ for all i and t . The corresponding continuous optimization problem is given by:

$$\bar{V} = \max \left\{ \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \times \left(\frac{v_i \sum_{u=1}^t \kappa_{i,t-u} x_{iu}}{v_0 + \sum_{j=1}^n v_j \sum_{u=1}^t \kappa_{j,t-u} x_{ju}} \right) \mid \sum_{t=1}^T x_{it} \leq 1 \forall i, \quad x_{it} \geq 0 \forall i, t \right\}, \quad (4)$$

and let $\{\bar{x}_{it} : i = 1, \dots, n, t = 1, \dots, T\}$ be an optimal solution to the above problem.

When all products have equal margins, the above objective function is concave because the market share $(\sum_{i=1}^n v_i \sum_{u=1}^t \kappa_{i,t-u} x_{iu}) / (v_0 + \sum_{j=1}^n v_j \sum_{u=1}^t \kappa_{j,t-u} x_{ju}) = \sum_{i=1}^n v_i z_{it} / (v_0 + \sum_{j=1}^n v_j z_{jt})$ in each period t is a concave function, and therefore the maximization can be carried out efficiently. When products have unequal margins, however, the continuous relaxation becomes a particular case of fractional programming known as the linear sum-of-ratios problem, see Schaible and Shi (2003), for which there is still no general polynomial optimization method but several branch-and-bound (b&b) schemes have been developed, e.g., Kuno (2002) and Benson (2007).

Note that the continuous relaxation provides an upper bound \bar{V} for the optimal value V^* of the original *APP*. However, it should be noted that there are instances where the objective value \bar{V} is far from V^* . This happens when there is a “fat” item that is significantly more attractive than the rest, i.e., its weight v_i is significantly larger than the weight of the other items (see Section 6 for some examples).

Another interesting feature of the relaxation problem is that all products of the same type with the same margin can be bundled together. Thus, if $r_i = r_j$ and $\kappa_{i,d} = \kappa_{j,d}$ for all $d \geq 0$, then we can replace combine products i and j and replace it with a new product having weight $v_i + v_j$, reducing the number of variables in the relaxation problem by one. Thus, the continuous relaxation can be seen as a strategic version of the original *APP* because it deals only with product types. The original *APP* is more of a tactical planning model, which might be used to determine when to release each specific product. The bundling of products with the same type and margin suggests that when we have a lot of products of each type, the continuous relaxation provides a good approximation to the original *APP*, as shown in the following section.

4.1 A Fluid Approximation. The continuous relaxation has a fundamental interpretation as a fluid approximation of the *APP*. To see this, consider a sequence of *APP* instances indexed by k , which we call APP_k , $k = 1, 2, \dots, \infty$, in which each original product i is subdivided into k subproducts with equal attractiveness v_i/k . Let V_k be the optimal value of instance k .

$$\begin{aligned}
 (APP_k) \quad V_k = \max \quad & \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \times \left(\frac{\frac{v_i}{k} \sum_{u=1}^t \kappa_{i,t-u} \sum_{q=1}^k \nu_{iuq}}{v_0 + \sum_{j=1}^n \frac{v_j}{k} \sum_{u=1}^t \kappa_{j,t-u} \sum_{q=1}^k \nu_{juq}} \right) \\
 \text{s.t.} \quad & \sum_{t=1}^T \nu_{itq} \leq 1 \quad \forall i \in I, q = 1, \dots, k, \\
 & \nu_{itq} \in \{0, 1\} \quad \forall i \in I, t = 1, \dots, T, q = 1, \dots, k.
 \end{aligned}$$

All instances in the sequence APP_k have the same continuous relaxation, so $V_k \leq \bar{V}$ for all k . As the (sub)products become infinitesimally small, i.e, as they become a fluid, the objective value V_k gets arbitrarily close to \bar{V} . The following theorem established the convergence rate to \bar{V} .

Theorem 2 (Fluid Approximation). *For any $k \geq 1$,*

$$0 \leq \bar{V} - V_k \leq \left(\frac{\max_{j=1, \dots, n} v_j}{k v_0 + \max_{j=1, \dots, n} v_j} \right) \bar{V},$$

which implies that $\lim_{k \rightarrow \infty} V_k / \bar{V} = 1$.

Theorem 2 tells us that when there are a lot of products of the same type with similar decay profiles, then we can obtain a good approximation to the APP by solving the continuous relaxation problem. Given the fluid interpretation, we proceed to analyze the continuous relaxation as a way to shed some light on the structure of the APP , which we later use as guideline to develop general heuristics.

4.2 Exponential Preference Decay. We can provide structural results when the preference decay parameters are given by $\kappa_{i,d} = \kappa_i^d$ with $0 < \kappa_i \leq 1$ and for all i and $d \geq 0$. As discussed in Section 3.2, this is one of the most prevalent decay patterns observed in practice. In what follows, we first assume that all products have equal margins with $r_i = 1$ for all i . Towards the end we discuss an extension with unequal margins.

Recall that for any product i and time period t , $z_{it} = \sum_{u=1}^t \kappa_{i,t-u} x_{iu}$ denote the load of product i in period t . For any t , let $D_t = \frac{\alpha_t v_0}{(v_0 + z_t)^2}$. Note that there is a one-to-one correspondence between $z_t = \sum_{i=1}^n v_i z_{it}$ and D_t . We can roughly interpret D_t as the marginal contribution to the total profit in period t . To see this, if $J : [0, 1]^{n \times T} \rightarrow \mathbb{R}_+$ denotes the objective function of the continuous relaxation problem, then

$$J_{i,t} := \frac{\partial J}{\partial x_{it}} = v_i \sum_{l=0}^{T-1} \kappa_i^l D_{t+l} = v_i D_t + \kappa_i J_{i,t+1}.$$

By studying the dual variables associated with the constraint of each product, we can determine which product will be introduced first. Since the continuous relaxation with equal margins is a concave optimization problem, we can use the Karush-Kuhn-Tucker (KKT) conditions with the dual variable $\lambda_i \geq 0$ associated with the constraint $\sum_{t=1}^T x_{it} \leq 1$. At optimality, we then have the following system of equations:

$$\lambda_i \geq J_{i,t} \quad i = 1, \dots, n, \quad t = 1, \dots, T, \quad \text{and} \quad 0 = \lambda_i \left(1 - \sum_{t=1}^T x_{it} \right) \quad i = 1, \dots, n,$$

along with the complementary slackness constraints that $x_{it}(\lambda_i - J_{i,t}) = 0$ for all i and t . This implies that $\bar{x}_{it} > 0$ only if $J_{i,t} = v_i \sum_{l=0}^{T-1} \kappa_i^l D_{t+l} = \lambda_i$, and $\bar{x}_{it} = 0$ if $v_i \sum_{l=0}^{T-1} \kappa_i^l D_{t+l} < \lambda_i$ (letting $D_t = 0$ if $t > T$). In other words, at optimality

$$\frac{\lambda_i}{v_i} = \max_{t \geq 1} \left\{ \sum_{l=0}^{T-1} \kappa_i^l D_{t+l} \right\}, \quad i = 1, \dots, n. \quad (5)$$

Suppose $\kappa_1 > \kappa_2$, e.g., type 1 are regular staple items while type 2 are fashion items such that their attractiveness is expected to fade away faster. Then Equation (5) implies that $\lambda_1/v_1 > \lambda_2/v_2$. This means that when there is a constraint on the total number of products that can be offered, the firm should choose those with the highest κ_i , which is intuitive. In addition, one finds that product 1 should be introduced earlier than product 2. This structure can be generalized to n types of products: products are introduced in decreasing order of κ_i over time, as shown in the following theorem.

Theorem 3 (Staggered Product Introduction). *With equal margins, if $\kappa_1 > \dots > \kappa_n$ and $\alpha_t = \alpha^{t-1}$ for all t , then for each product $i \in I$ there exists a time window $[S_i, E_i]$ such that in the optimal solution of the continuous relaxation: $\bar{x}_{it} > 0$ if $S_i \leq t < E_i$, $\bar{x}_{iE_i} \geq 0$, and $\bar{x}_{it} = 0$ for all $t \notin [S_i, E_i]$. Moreover, $S_1 = 1$, $E_i \leq S_{i+1}$ for all i , and optimality, D_t is strictly decreasing after E_n , i.e. $D_{E_n} > D_{E_n+1} > \dots > D_T$.*

This result demonstrates that, when products have equal margins and exponential decays, they are introduced earlier if they stay “fresh” for longer (i.e., they have a higher κ_i). Theorem 3 requires a strict ordering of the decay parameters, but when there are several products of the same type (i.e., that share the same decay factor κ_i), the staggered product introduction property still applies: there is one optimal solution with such structure. As shown in the following example, the intervals $[S_1, E_1], [S_2, E_2], \dots, [S_n, E_n]$ may be disjoint.

Example 2 (Non-overlapping intervals). Suppose that $n = 2, T = 4, r_1 = r_2 = \$1, \alpha = 0.95, v_0 = 1, v_1 = 10, v_2 = 1, \kappa_1 = 0.8$ and $\kappa_2 = 0.4$. The unique solution of the relaxation is to set $\bar{x}_{11} = 1, \bar{x}_{1t} = 0$ for $t \neq 1, \bar{x}_{23} = 1, \bar{x}_{2t} = 0$ for $t \neq 3$. Thus, $[S_1, E_1] = [1, 1]$ and $[S_2, E_2] = [3, 3]$.

When there is a single product type, the solution of the continuous relaxation can be characterized even further, as shown in the next two results. We know from Theorem 3 that there is a cut-off period E_1 after which no product is introduced. Let us denote this cut-off period by $E_1(\kappa)$ to emphasize its dependence on κ , since the cut-off period will change with κ . The following theorem shows that this cut-off period is non-increasing in κ .

Proposition 4 (More Fashionable Products are Introduced over Longer Periods). *Assume a single product type and $\alpha_t = \alpha^{t-1}$ for all t . Then, $E_1(\kappa)$ is non-increasing in κ .*

Note that when $\kappa = 1$, we have a truly basic product and there is no decay. In this case, it is easy to verify that it is optimal to introduce everything in the first period, that is, $E_1(1) = 1$. Proposition 4 shows that as the product becomes more fashionable, with κ decreasing to zero, it would gradually be introduced over a longer horizon. In fact, with $\kappa = 0$, we have a single-period life cycle, and in this case, $E_1(0) = T$. As a consequence of the above proposition, we see that basic products are always introduced earlier than fashion products.

The next result shows that the market share under the optimal product introduction is declining over time. The optimal market share in period t is given by $\bar{z}_t/(v_0 + \bar{z}_t)$. The following proposition shows that \bar{z}_t decreases with t .

Proposition 5 (Decreasing Optimal Market Share). *Assume a single product type and $\alpha_t = \alpha^{t-1}$ for all t . In the optimal solution of the continuous relaxation the market share decreases over time, that is, $\bar{z}_1 \geq \bar{z}_2 \geq \dots \geq \bar{z}_T$.*

The previous results assume that all products have equal margins. With unequal margins the problem is hard to analyze because the continuous relaxation is not even concave. An exception is when there is a basic item j with a margin $r_j \geq 1$ and all other items have equal margins normalized to 1. If $v_0 \geq v_j(r_j - 1)$, then the problem remains concave and the results in Theorem 3 and Proposition 5 continue to hold. In other words, the optimal solution of the continuous relaxation will follow the staggered introduction property; and if all items $i \neq j$ have the same type, then the optimal market share will be decreasing.

5. Approximation Algorithms for the APP

In this section we introduce three approximation algorithms to solve the APP. The first two assume equal margins while the last one is for the general case and allows for products with different margins. Though these are approximate methods, we are able to establish performance guarantees.

5.1 Two Heuristics under Equal Margins. Here we consider the case when all products have the same gross margin, which as before we normalize to one, i.e., $r_i = 1, \forall i = 1, \dots, n$. The objective function of the APP becomes $J := \sum_{t=1}^T \frac{\alpha_t z_t}{v_0 + z_t}$. Since each fractional term $\frac{\alpha_t z_t}{v_0 + z_t}$ is increasing in z_t , it follows that with equal margins it is optimal to introduce all products (this is not necessarily the case with unequal margins). The two heuristics that we introduce next are based on the structure of the continuous relaxation (4) when margins are equal.

5.1.1 GREEDY Heuristic. To solve the *APP* with equal margins, consider the partial derivative of the objective function J with respect to x_{it} :

$$\frac{\partial J}{\partial x_{it}} = v_0 \sum_{s=t}^T \frac{\alpha_s v_i \kappa_{i,s-t}}{(v_0 + z_s)^2}. \quad (6)$$

Motivated by the concavity of the objective function J , we propose the following GREEDY heuristic: sequentially assign products to periods (bins), each time selecting the product-period pair with the highest partial derivative (6) among those products that have not been previously assigned. Intuitively, we use the index (6) because this is a measure of the marginal increase in profits, and products are introduced according to this index in a greedy fashion. Note that this heuristic does not require solving the continuous relaxation. Its computation time is $O(n^2 T^2)$ for general decay patterns but can be computed more efficiently for specific structures such as exponential decay.

While the performance of GREEDY is generally hard to evaluate analytically, when all products have a single-period life cycle, a performance guarantee can be developed. This corresponds to the case where $\kappa_{i,d} = 0$ when $d \geq 1$, i.e., products completely lose their attractiveness one period after being introduced. This can be seen as an extreme case of exponential preference decay in which the decay parameter tends to zero for all products, or it can also represent the situation when all products have the same planned life cycle, i.e., $\ell_i = \ell$, $\forall i = 1, \dots, n$ (in this case the horizon T should be a multiple of ℓ).

Without loss of generality, as all products have the same margin, we sort the items by decreasing sizes, $v_1 \geq \dots \geq v_n$, and we assume that $\alpha_1 \geq \dots \geq \alpha_T \geq 0$. Let V^* be the optimal solution to the *APP*, and let $z_t^* = \sum_{i=1}^n v_i x_{it}^*$ be the load for period t in an optimal solution. Note that because α_t is decreasing in t , it must be true that z_t^* is also decreasing in t . Otherwise, if $z_{t_1}^* < z_{t_2}^*$ with $t_1 \leq t_2$, the products allocated to t_1 and t_2 could be exchanged and profits would be increased. Hence, we observe that the decreasing market share property that was shown in Proposition 5 for the continuous relaxation holds here for the optimal integral solution of the *APP*.

With single-period life cycles, the GREEDY heuristic simplifies to the following procedure. First, sort the items in decreasing size. Second, allocate items to periods (bins) in a greedy fashion: place item i in the period with the highest current index $\frac{\alpha_t}{(v_0 + z_t)^2}$ where z_t is the current load. Clearly the heuristic runs in polynomial time. Moreover, we can provide the following performance guarantee.

Theorem 6. *When all products have a single-period life cycle and $\alpha_t = 1$ for $t = 1, \dots, T$, let V^{Greedy} be the objective value of the GREEDY solution. Then $1 \leq \frac{V^*}{V^{Greedy}} \leq \frac{9}{8}$.*

Remarkably, when there is no seasonality or discounting, Theorem 6 provides a 9/8-approximation for the *APP*. This result can then be used for more general situations.

Corollary 7. *When all products have a single-period life cycle, compute the solution provided by the GREEDY heuristic applied to the problem with $\alpha_t = 1$ for $t = 1, \dots, T$. Then this heuristic is a $\frac{9}{8}\Lambda$ -approximation for the *APP*, where $\Lambda = \frac{\max_t \alpha_t}{\min_t \alpha_t}$.*

5.1.2 EARLYENTRY Heuristic. The second heuristic, which we call EARLYENTRY, is motivated by the staggered entry property shown in Theorem 3 for exponential preference decay. The procedure is simple: solve the continuous relaxation and determine the time window in which each product is released. Then, in the *APP*, introduce each product at the beginning of its time window. In other words, introduce product i in the earliest period t such that $\bar{x}_{it} > 0$, where \bar{x}_{it} represents the optimal solution of the continuous relaxation. When the values of κ_i are large, the EARLYENTRY heuristic generates solutions that are close to the optimum of the *APP*.

Theorem 8. *When all products have exponential preference decay and $\alpha_t = \alpha^{t-1}$ for all t , let $[S_i, E_i]$ denote the release time window in the optimal solution of the continuous relaxation for product $i \in I$. Let $V^{EarlyEntry}$ be the objective value of the EARLYENTRY solution. Then $1 \leq \frac{V^*}{V^{EarlyEntry}} \leq \max_{i=1, \dots, n} \left\{ \frac{1}{\kappa_i^{E_i - S_i}} \right\}$, where κ_i is the decay parameter of product $i \in I$.*

This performance bound depends on the actual solution to the continuous relaxation, through the values S_i and E_i . Hence, if $E_i = S_i$ for all i , then the heuristic leads to the optimum, because the relaxation solution is integral. When the number of products is large, the time windows are compressed, so $E_i - S_i$ is small for all i , which leads to a performance bound that only depends on the decay parameters κ_i . The performance bound is tighter when all κ_i are high. Hence, the EARLYENTRY heuristic is an appropriate complement to GREEDY, which performs well when all κ_i tend to zero.

5.2 A Randomized Heuristic under General Margins. The two previous heuristics were developed for the case of equal margins. Note that these heuristics provide a feasible solution and an upper bound for the general case if one replaces the margins r_i by its maximal value. For a heuristic that can work with different margins directly, we resort to a randomized approach (Raghavan and Thompson 1987). Here we use the fact that the constraint (2) in the *APP* allows giving a probabilistic interpretation to the solution of the continuous relaxation.

As before, let $(\bar{x}_{i,t} : i \in I, t = 1, \dots, T)$ denote the solution to the continuous relaxation problem and let \bar{V} be its objective value. Let Q_i denote a random variable that represents the time period

when product $i \in I$ is introduced. The probability distribution of Q_i is given by the solution to the continuous relaxation. To be precise, let $\Pr(Q_i = t) = \bar{x}_{i,t}$, $\forall i \in I, t = 1, \dots, T$. Since with unequal margins constraint (2) might not be binding, we consider a fictitious period $T + 1$ to represent the non-release option with probability mass $\Pr(Q_i = T + 1) = 1 - \sum_{t=1, \dots, T} \bar{x}_{i,t}$, $\forall i \in I$. Let $f(Q_1, \dots, Q_n)$ denote the (random) profit associated with the random vector (Q_1, \dots, Q_n) , i.e.,

$$f(Q_1, \dots, Q_n) := \sum_{t=1}^T \alpha_t \frac{\sum_{u=1}^t \sum_{i=1}^n r_i \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + \sum_{u=1}^t \sum_{i=1}^n \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}, \quad (7)$$

where $\mathbb{1}_A$ is the indicator function for event A . The following theorem gives a performance guarantee associated with the random vector (Q_1, \dots, Q_n) .

Theorem 9. *Let $\rho := \frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j}$. Then, $1 \leq \frac{V^*}{\mathbb{E}[f(Q_1, \dots, Q_n)]} \leq \frac{\bar{V}}{\mathbb{E}[f(Q_1, \dots, Q_n)]} \leq \frac{1}{\rho}$.*

Based on the probabilistic interpretation described above, we define a RANDOMIZED heuristic that consists in sampling a fixed number of solutions for the APP according to the probability distribution of (Q_1, \dots, Q_n) and keeping the best one. Note that the inequality $\mathbb{E}[f(Q_1, \dots, Q_n)] \geq \rho \bar{V}$ in Theorem 9 ensures that sampling the distribution of (Q_1, \dots, Q_n) will eventually yield a solution with an objective value greater than $\rho \bar{V}$, where $\rho = \frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j}$. This solution is guaranteed to be within percentage $1 - \rho$ of the optimal value V^* . Clearly, the performance guarantee for the RANDOMIZED heuristic is tighter when ρ is larger, which occurs when the outside option has significant weight v_0 vis-a-vis the attractiveness of the products the firm can offer. This situation usually occurs when the market is competitive and customers have many other alternatives where to purchase. For instance, if $v_0 \geq v_i$, $\forall i \in I$, which is a common assumption found in the literature (e.g., Allon et al. 2010), then the RANDOMIZED solution is guaranteed to be within 50% of the optimum.

As mentioned in Section 4, solving the continuous relaxation with unequal margins is challenging and one of the prevalent methods used is b&b. Given the substantial computational effort that is required, in practice the b&b procedure is usually terminated when a feasible solution is found that satisfies a sub-optimality tolerance. If that is the case, the performance guarantee of the RANDOMIZED heuristic can be modified according to the following corollary.

Corollary 10. *Let $(\tilde{x}_{i,t} : i \in I, t = 1, \dots, T)$ represent the incumbent solution in the branch-and-bound procedure and let \tilde{J} be its objective value. Let \tilde{V} be the current upper bound and let $\Gamma := (\tilde{V} - \tilde{J})/\tilde{V}$ be the sub-optimality gap. Then, the RANDOMIZED heuristic using probability weights $\tilde{x}_{i,t}$ guarantees that $1 \leq \frac{V^*}{\mathbb{E}[f(Q_1, \dots, Q_n)]} \leq \frac{1}{(1 - \Gamma)\rho}$, where ρ is defined as in Theorem 9.*

6. Numerical Study

Our numerical study has two parts. First, in Section 6.1 we describe in detail a small illustrative example to shed further light on the problem, its structure, and the merits of the approximation algorithms. In Section 6.2 we consider larger instances trying to replicate real situations that firms face.

Throughout this section we compare the performance of the three approximation algorithms. As a benchmark, we also include a fourth heuristic that consists in just rounding the solution of the continuous relaxation, which is the most common approach used in practice to solve problems with integer variables. Specifically, for each product i , identify the period t such that \bar{x}_{it} is the largest in the optimal solution of the continuous relaxation, and then introduce product i in period t in the *APP*, which amounts to setting $x_{it} = 1$. Following the probabilistic interpretation discussed in Section 5.2, this rounding procedure is equivalent to introducing product i in the period that represents the mode of the probability distribution $(\bar{x}_{i,t} : i \in I, t = 1, \dots, T)$. Hence, we refer to it as *MODEROUNDING*.

6.1 Small Illustrative Example. Here we consider a small instance with four products and ten periods. The products have equal margins ($r_i = 1, i = 1, 2, 3, 4$) and attractiveness $v_1 = v_2 = v_3 = 1, v_4 = 100$, so it corresponds to a situation with three minor items and one “fat” product. We assume exponential decay with parameters: $\kappa_1 = 0.9, \kappa_2 = 0.6, \kappa_3 = 0.5, \kappa_4 = 0.4$. The weight of the outside option is $v_0 = 1$.

Table 1 show the results for this small instance. The objective value of the different solutions is reported on left hand side of the table, while the actual solution is shown to the right. The first entry of the table is the optimal solution to the continuous relaxation, which has a objective value of 9.457. This value is a very loose upper bound for the *APP* due to the presence of the fat item. Therefore, we did ten b&b steps — branching on the release period for product 4 — which reduced the upper bound to 8.269. This result is shown as the second entry in Table 1. The third entry is the optimal integer solution, which we were able to find by exhausting the b&b tree for this small example. Though there is a 0.54% gap between the optimal integer solution and the partial b&b upper bound, we assess the sub-optimality of the approximation algorithms with respect to the latter emulating the fact that the true integer optimum would be computationally prohibitive to find for any real-size instance. The remaining three entries in Table 1 report the performance of the heuristics described in Section 5. Note that for this example, *MODEROUNDING* provided the same solution as *EARLYENTRY* so it is not reported.

As expected, the continuous relaxation solution has the staggered entry property shown in

Algorithm	Objective (% gap)	Solution x_{it}											
		i, t	1	2	3	4	5	6	7	8	9	10	
Continuous relaxation	9.457	1	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		3	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		4	0.120	0.096	0.099	0.102	0.103	0.104	0.104	0.105	0.105	0.063	
Partial b&b upper bound	8.269	1	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	0.000	0.000	0.000	0.000	0.000	0.685	0.315	0.000	0.000	0.000	0.000
		3	0.000	0.000	0.000	0.000	0.000	0.000	0.176	0.391	0.413	0.021	
		4	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
OPTIMAL	8.224 (0.54%)	1	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000
		3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000
		4	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GREEDY	8.198 (0.85%)	1	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	0.000	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000
		3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000
		4	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
EARLYENTRY	6.836 (17.33%)	1	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		3	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		4	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RANDOMIZED	8.067 (2.50%)	1	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		2	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		3	1.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
		4	0.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 1: Small numerical example with 4 products and 10 periods. The gaps are with respect to the upper bound (8.269).

Theorem 3, with the first three items released only in the first period while the fat item is released continuously throughout the horizon, which makes the solution very fractional. In general, we observed this behavior whenever there was a fat item, i.e., a product that was significantly more attractive than the rest. In those cases, the optimal continuous solution has a large integrality gap. For the purpose of this example, we reduced the gap through b&b, first partially and then exhaustively. The idea of doing partial b&b is to show that typically branching on the fat item(s) is enough to narrow the integrality gap to a reasonable level, which can be very useful in practice.

In contrast to the continuous relaxation, the optimal integral solution has the fat item as the sole product introduced in the first period despite the fact that it has the fastest decay. The remaining products are introduced only once the attractiveness of the fat item has vanished. The GREEDY solution has this same structure — i.e., the fat item first and the rest later — and it achieves the best performance among the solutions considered. In fact, except for product 2, the GREEDY solution introduces all items optimally, which shows that the partial derivative (6) is indeed a good optimality indicator. The EARLYENTRY and RANDOMIZED solutions are based on the continuous relaxation, which we know has a large integrality gap. In particular, EARLYENTRY solution is misguided by the fractional values of the fat item, so in this case it has very poor performance. The

RANDOMIZED algorithm is able to detect that it would be better to introduce the fat item later, so it has a much better performance, though in terms of structure it is exactly the opposite from the optimal integer solution. Note that once the release of the fat item is fixed by doing partial b&b, the continuous solution for the remaining items also has the staggered entry property. Hence, the performance of EARLYENTRY and RANDOMIZED usually can be improved by using the partial b&b solution as the input instead of the original continuous solution.

6.2 Large Instances. The previous small instance serves as an illustrative example of the *APP* and the solutions provided by the approximation algorithms. In order to test our methods further on real-size instances, we constructed a problem consisting of 200 products. To use realistic parameters, we use a data set of DVD titles from Rusmevichientong et al. (2010) which provide the values for attractiveness (v_i 's). The attractiveness of the DVDs ranged from 0.005 to 0.03, with a total sum equal to 1.92, and the weight of the outside option was $v_0 = 1$. We considered a 52-week horizon with discount $\alpha_t = \alpha^{t-1}, \forall t \geq 1$, and we assumed equal margins ($r_i = 1, \forall i \in I$) in order to make the continuous relaxation tractable. For the life cycle pattern we assumed exponential decay, which we believe is reasonable since the DVDs were movies or series that had already been released in theaters or on TV so they were past their peak in popularity. We did not have data as in Ainslie et al. (2005) to estimate the decay parameters, so instead, we generated the decay parameters κ_i randomly to ensure that there was no correlation with the weights v_i . The decay parameters ranged from 0.6 to 0.999, with an average value of 0.82. We also considered the cases when the decay κ_i and the weights v_i were positively and negatively correlated, but the result were the same so we do not report them in the paper. The continuous relaxation was coded in GAMS and solved using the commercial non-linear solver MINOS5, while the approximation algorithms were coded in Perl. The running times in all cases did not surpass a few minutes.

We first explore the properties of the continuous relaxation (cf. Section 4.2). Figure 1 shows the product release time windows and the market share trajectory ϕ_t/ϕ_1 for three different parameter sets (recall that $\phi_t = \frac{z_t}{v_0 + z_t}$). The DVDs are sorted in decreasing order of the decay parameter κ_i , and therefore we observe a staircase shape that follows from the staggered product introduction property in Theorem 3. Plot (a) in Figure 1 is the base case with $v_0 = 1$ and $\alpha = 1$. In plot (b) the weight of the outside option is ten times higher so the firm prefers to introduce more items early to boost its initial market share. Plot (c) has a more pronounced discount rate that reduces the value of later-period profits, which again pushes the items to be introduced earlier. It is interesting to see that with multiple products the market share is also decreasing in the same way it was shown in Proposition 5 for a single product type.




Figure 1: Market share trajectory ϕ_t/ϕ_1 (top curve) and product release time windows (bottom curve) for the instance with 200 DVDs, 52 periods and exponential decay.

Table 2 shows the suboptimality gap of the solutions from the approximation algorithms and the MODEROUNDING heuristic. The suboptimality is measured with respect to the upper bound given by the optimal value of the continuous relaxation. (Here there is no fat item so b&b was not necessary.). The top part of the table shows the sensitivity of the gaps with respect to the discount parameter α and the bottom part shows the sensitivity with respect to the attractiveness of the outside good v_0 . For each set of parameters, the algorithm with the best performance is shown in boldface. The first remark is that the all the algorithms perform very well, with all gaps below 2%. In contrast with the small example described in the previous section, here GREEDY has a slightly worse performance, whereas EARLYENTRY, RANDOMIZED and MODEROUNDING have gaps below 0.005%. The main reason is that in this case the number of products (200) is large with respect to the number of periods (52). From the staggered entry property in Theorem 3, packing many products in a few periods forces the time windows $[S_i, E_i]$ to be very narrow, which means that the optimal solution of the continuous relaxation is not too fractional. This observation is also in line with the fluid interpretation of the continuous relaxation discussed in Section 4.1.

The instance with DVD data has a large number of products so the approximation algorithms based on the continuous relaxation perform very well. To show what happens when there are fewer products, we generated a random instance in which a planner has to introduce a small number of products each period. The planner could represent a movie distributor or a category manager

Discount α	v_0	Relaxation	GREEDY	EARLYENTRY	RANDOMIZED	MODEROUNDING
1.0	1	13.27	0.980%	0.003%	0.001%	0.001%
0.999	1	12.96	0.859%	0.003%	0.001%	0.001%
0.99	1	10.65	0.225%	0.003%	0.001%	0.001%
0.95	1	6.16	0.008%	0.001%	0.000%	0.000%
0.9	1	4.15	0.003%	0.000%	0.000%	0.000%
1.0	10	1.74	0.143%	0.000%	0.000%	0.000%
1.0	2	7.63	0.698%	0.001%	0.001%	0.001%
1.0	1	13.27	0.980%	0.003%	0.001%	0.001%
1.0	0.3	27.66	1.125%	0.003%	0.002%	0.002%
1.0	0.03	47.78	0.296%	0.001%	0.001%	0.001%
1.0	0.005	51.24	0.056%	0.000%	0.000%	0.000%

Table 2: Sub-optimality gaps when varying the discount parameter α (top) and the weight of the outside good v_0 (bottom) for an instance consisting of 200 DVDs and 52 weeks. The gaps are measured with respect to the optimal value of the continuous relaxation.

in apparel or accessories, as those described in Section 1. For that, we considered 52 products and the same number of periods. We assumed no discounting or seasonality, and we generated the attractiveness v_i from a uniform $[0,1]$ distribution for all items except one, which had an attractiveness of 10. The idea was to represent the situation in which there was a big disparity among products and one was substantially more appealing. For the outside option, we considered three different values of v_0 (0.5, 1 and 101). The results are shown in Table 3. The first three rows assume exponential decay — with the κ_i parameters again generated randomly, between 0 and 1 — and the last three rows correspond to the case when all products have a single-period life cycle, i.e., when $\kappa_i \rightarrow 0, \forall i \in I$. To tighten the upper bound, we did partial b&b branching on the “fat” item, as in the small example described in the previous section. With exponential decay, all the approximation algorithms continue to perform very well. In contrast, when all products have a single-period life cycle and the weight of the outside good is low, the performance of the algorithms is very poor except for GREEDY, which as expected achieves sub-optimality gaps within 1/8 given by the performance guarantee in Theorem 6.

v_0	Life cycle	Relaxation	Upper bound	GREEDY	EARLYENTRY	RANDOMIZED	MODEROUNDING
0.5	52	44.47	44.06	0.265%	1.420%	0.795%	1.603%
1	52	38.84	38.25	0.385%	1.971%	1.108%	2.202%
101	52	1.48	1.47	0.067%	0.274%	0.098%	0.207%
0.5	1	30.95	27.29	9.569%	25.306%	15.417%	15.417%
1	1	22.03	18.67	7.648%	16.856%	9.144%	9.144%
101	1	0.38	0.37	0.123%	0.208%	0.123%	0.123%

Table 3: Sub-optimality gaps for a randomly generated instance with 52 periods and 52 products including a “fat” item. Here there is no discounting or seasonality ($\alpha_t = 1$ for $t = 1, \dots, T$).

7. Discussion and Conclusions

We have presented the *APP* for a firm that has to decide when to release a set of products throughout a season. Our formulation assumes that the attractiveness and decay pattern of each product are known in advance. This would be the case if the products have already been designed and produced, like basics and fashion-basics in apparel, and their attractiveness has been estimated, for instance, through marketing techniques, such conjoint analysis, regression models or focus groups. Though the assortment decision for these types of products is usually made only a few times a year, planners are constantly evaluating different scenarios. Hence, it pays to have efficient methods to solve the *APP*. We have presented here three approximation algorithms for the *APP* and we have shown analytically as well as numerically that the overall performance is close to optimal. In particular, the *EARLYENTRY* and *RANDOMIZED* algorithms perform very well when the solution of the continuous relaxation is not too fractional, which is the case whenever there is a large number of products or the decay pattern is exponential. In contrast, when the continuous relaxation is highly fractional, the *GREEDY* algorithm tends to do well.

We conclude the paper by discussing a few possible extensions to the *APP*. One of the implicit assumptions in the original model formulation is that the number of product introduced in a given period is only restricted by how many are available. In practice, it could occur that there is limitation that prevents the firm from releasing more than a certain number of new items. For instance, category managers in apparel some times are given a budget and there is an operational cost associated to each new introduction, so they have to be more selective and not all products are released. This situation can be included in the model by considering the constraint $\sum_{t=1}^T \sum_{i=1}^n x_{it} \leq N$ or a variation of it, where N represents the total number of products that can be introduced (if the budget is B and the introduction cost is c , then $N = \lfloor B/c \rfloor$). Another reason that could limit the number of products is some sort of shelf space constraint, which could be incorporated in the model by imposing the inequality $\sum_{i=1}^n \sum_{u=1}^t S_{i,t-u} x_{iu} \leq S, \forall t \geq 1$, where S is the available space for display and $S_{i,l}$ is the space that product i takes l periods after its introduction.

Adding a budget or shelf space constraint is likely to affect the structure of the problem. However, it is interesting to notice that the fluid interpretation of the continuous relaxation and Theorem 2 remain valid. With equal margins, the *GREEDY* and *EARLYENTRY* can be modified to accommodate the additional constraint but the performance guarantees in Section 5.1 are lost, or at least they would require a different proof. In contrast, with general margins, the performance bound for the *RANDOMIZED* algorithm in Theorem 9 continues to hold. Generating the randomized solution would have to incorporate the product introduction limitation, but this can be achieved

by sampling the variables sequentially and resolving the *APP* each time a variable is fixed to update the probability distributions, or alternatively one can try to adapt the dependent rounding algorithms described in Bertsimas et al. (1999).

Finally, our model assumes that all parameters are well-known to the firm. In particular, the attractiveness and decay patterns are taken as given. If the parameters are highly uncertain and there is an opportunity to gain additional information between decision periods, then a dynamic programming (DP) or closed-loop approach might be more suitable. However, DPs are rarely solvable in practice so researchers resort to approximate methods. One of the most common techniques is the certainty-equivalence approximation that replaces the uncertain parameters by their expected values and resolves the deterministic problem in a rolling horizon fashion (see Bertsimas and Popescu 2003). Hence, this method relies on repeatedly solving the open-loop formulation, for which our model and algorithms would come in handy.

Appendix: Proofs

Proof of Theorem 1

Proof. Let v_0 be given. We show that 2-PERIOD EQUAL-MARGIN ASSORTMENT PACKING is in NP, because we can transform an arbitrary instance of PARTITION, which is a well-known NP-complete problem, to an equivalent 2-PERIOD EQUAL-MARGIN ASSORTMENT PACKING problem. The PARTITION problem is defined as follows.

PARTITION

INPUTS: The set of items indexed by $1, \dots, n$ and the size $c_i \in \mathbb{Z}_+$ associated with each item.

QUESTION: Is there a subset $S \subseteq \{1, \dots, n\}$ such $\sum_{i \in S} c_i = \sum_{i \notin S} c_i$?

Let $C = \frac{1}{2} \sum_{i=1}^n c_i \in \mathbb{Z}_+$. Since $\sum_{i \in S} c_i = \sum_{i \notin S} c_i$ if and only if $\sum_{i \in S} c_i = C$, we may assume without loss of generality that $C \in \mathbb{Z}_+$. An instance of PARTITION can be solved by solving the following instance of 2-PERIOD EQUAL-MARGIN ASSORTMENT PACKING: define the preference weights as $v_i = c_i$ for $i = 1, 2, \dots, n$, and set the target profit as $K = 2C/(v_0 + C)$. The PARTITION problem indeed has a solution *if and only if* there exists a partition X_1 and X_2 such that $\frac{\sum_{j \in X_1} v_j}{v_0 + \sum_{j \in X_1} v_j} + \frac{\sum_{j \in X_2} v_j}{v_0 + \sum_{j \in X_2} v_j} \leq \max_{y \in [0, 2C]} \left\{ \frac{y}{v_0 + y} + \frac{2C - y}{v_0 + 2C - y} \right\} = \frac{2C}{v_0 + C} = K$. This is true because $G(y) = \frac{y}{v_0 + y} + \frac{2C - y}{v_0 + 2C - y}$ is concave in y over the interval $[0, 2C]$ and achieves a unique maximum at $y = C$. ■

Proof of Theorem 2

Proof. This proof makes use of a later result, Theorem 9, which is proved independently of this theorem. Fix an arbitrary k . Clearly the relaxations of APP_k and APP are identical: $\overline{V}_k = \overline{V}$.

Given the optimal solution $(\bar{v}_{itq} : i = 1, \dots, n, t = 1, \dots, T, q = 1, \dots, k)$ associated with the continuous relaxation of APP_k , let $(Q_{itq} : i = 1, \dots, n, t = 1, \dots, T, q = 1, \dots, k)$ denote the randomized rounding solution as defined by Theorem 9. It follows that

$$\left(\frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j/k} \right) \bar{V}_k \leq \mathbb{E}[f(Q_{itq} : i = 1, \dots, n, t = 1, \dots, T, q = 1, \dots, k)] \leq V_k$$

where the first inequality follows from the fact that in APP_k , the weight of each product i is subdivided into k equal parts and each subproduct has a weight of v_i/k . Then, we have the following series of inequalities:

$$\bar{V} = \bar{V}_k \geq V_k \geq \left(\frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j/k} \right) \bar{V}_k = \left(\frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j/k} \right) \bar{V},$$

which gives the desired result. \blacksquare

Proof of Theorem 3

Proof. Without loss of generality, we can use in this proof $v_\ell = 1$ for all ℓ (otherwise, we can redefine λ_ℓ below as λ_ℓ/v_ℓ). To avoid confusion with the subindices, in this proof we write x_{it} as $x_{i,t}$. Recalling that $z_{i,t} = \sum_{s=1}^t \kappa_i^{t-s} x_{i,s}$, we have $z_{i,t} = x_{i,t} + \kappa_i z_{i,t-1}$ and $z_{i,0} = 0$. As J denotes the objective function of the continuous relaxation, with preference weights equal to one, we have

$$J_{i,t} := \frac{\partial J}{\partial x_{i,t}} = \sum_{l=0}^{T-1} \kappa_i^l D_{t+l} = D_t + \kappa_i J_{i,t+1}.$$

We know from the KKT conditions that there exists nonnegative dual variables $\lambda_1, \dots, \lambda_n$ such that $x_{i,t} \geq 0$ when $J_{i,t} = \lambda_i$ and $x_{i,t} = 0$ when $J_{i,t} < \lambda_i$. We will hence prove that for any $\lambda_1, \dots, \lambda_n$, the optimal solution to the relaxed problem has the desired property.

Claim 1: If all n products are introduced, then we can restrict our attention to the case where $\lambda_1 > \dots > \lambda_n$ and $\lambda_1(1 - \kappa_1) < \dots < \lambda_n(1 - \kappa_n)$.

Proof of Claim 1: Indeed, suppose that $\lambda_i \leq \lambda_{i+1}$ for some i . Then it is easy to see that for all i and all values of $D_1, \dots, D_T > 0$ $J_{i,t} > J_{i+1,t}$ for $t < T$. Since $J_{i,t} \leq \lambda_i$, this implies that $J_{i+1,t} < \lambda_{i+1}$ and hence $x_{i+1,t} = 0$, for all t , i.e., product $i+1$ will never be introduced (except possibly in period T , in which case it can be replaced by product i without loss of optimality). Hence $\lambda_1 > \dots > \lambda_n$.

Similarly, suppose that $\lambda_i(1 - \kappa_i) \geq \lambda_{i+1}(1 - \kappa_{i+1})$ for some i . Then we prove by induction that for all t , $\lambda_i - J_{i,t} > \lambda_{i+1} - J_{i+1,t}$. This is indeed true for $t = T$ because $\lambda_i > \lambda_{i+1}$ and $J_{i,T} = J_{i+1,T} = D_T$. If the result is true for $t+1$, then it must also be true for t because $J_{i,t} = D_t + \kappa_i J_{i,t+1}$, which implies that $\lambda_i - J_{i,t} = \lambda_i(1 - \kappa_i) - D_t + \kappa_i(\lambda_i - J_{i,t+1}) > \lambda_{i+1}(1 - \kappa_{i+1}) - D_t + \kappa_{i+1}(\lambda_{i+1} - J_{i+1,t+1}) = \lambda_{i+1} - J_{i+1,t}$, where the inequality follows from the induction hypothesis. This completes the induction and thus, $\lambda_i - J_{i,t} > 0$ for all t , and product i will never be introduced. \blacksquare

Claim 2: If $x_{i,t} > 0$ for some i and t , then $D_{t-1} \leq \lambda_i(1 - \kappa_i) \leq D_t$.

Proof of Claim 2: Since $J_{i,t} = D_t + \kappa_i J_{i,t+1}$, $x_{i,t} > 0$ implies that $J_{i,t} = \lambda_i$ and hence $D_t = \lambda_i - \kappa_i J_{i,t+1} \geq \lambda_i(1 - \kappa_i)$ and $D_{t-1} = J_{i,t-1} - \kappa_i \lambda_i \leq \lambda_i(1 - \kappa_i)$. ■

Claim 3: If $D_t > D_{t+1}$ for some t , then $D_t > D_{t+1} > \dots > D_T$ and no product is introduced after period t .

Proof of Claim 3: It suffices to show that $D_{t+1} > D_{t+2}$ because the proof for other cases is similar. Suppose on the contrary that $D_{t+2} \geq D_{t+1}$. Then, we have that $D_t > D_{t+1}$ and $D_{t+2} \geq D_{t+1}$, which implies that $D_{t+1}^2 < D_t D_{t+2}$, or equivalently,

$$\left(\frac{\alpha^t}{(v_0 + \sum_{i=1}^n z_{i,t+1})^2} \right)^2 < \frac{\alpha^{t-1}}{(v_0 + \sum_{i=1}^n z_{i,t})^2} \cdot \frac{\alpha^{t+1}}{(v_0 + \sum_{i=1}^n z_{i,t+2})^2} \leq \frac{\alpha^{t-1}}{(v_0 + \sum_{i=1}^n z_{i,t})^2} \cdot \frac{\alpha^{t+1}}{(v_0 + \sum_{i=1}^n \kappa_i z_{i,t+1})^2},$$

where the last inequality follows from the fact that $z_{i,t+2} = x_{i,t+2} + \kappa_i z_{i,t+1} \geq \kappa_i z_{i,t+1}$ for all i . Since $D_t > D_{t+1}$, it follows that $x_{i,t+1} = 0$ for all i by Claim 2. Thus, $z_{i,t+1} = \kappa_i z_i$ for all i , and the above

inequality implies that $\left(\frac{\alpha^t}{(v_0 + \sum_{i=1}^n \kappa_i z_{i,t})^2} \right)^2 < \frac{\alpha^{t-1}}{(v_0 + \sum_{i=1}^n z_{i,t})^2} \frac{\alpha^{t+1}}{(v_0 + \sum_{i=1}^n \kappa_i^2 z_{i,t})^2}$, which implies that $\left(v_0 + \sum_{i=1}^n \kappa_i z_{i,t} \right)^2 > \left(v_0 + \sum_{i=1}^n z_{i,t} \right) \left(v_0 + \sum_{i=1}^n \kappa_i^2 z_{i,t} \right)$ or in other words $\left(\sum_{i=1}^n \kappa_i z_{i,t} \right)^2 > \left(\sum_{i=1}^n z_{i,t} \right) \left(\sum_{i=1}^n \kappa_i^2 z_{i,t} \right) + v_0 \left(\sum_{i=1}^n (1 - \kappa_i)^2 z_{i,t} \right) > \left(\sum_{i=1}^n z_{i,t} \right) \left(\sum_{i=1}^n \kappa_i^2 z_{i,t} \right)$.

However, an application of Cauchy-Schwarz inequality to to vectors $(\sqrt{z_{1,t}}, \dots, \sqrt{z_{n,t}})$ and $(\kappa_1 \sqrt{z_{1,t}}, \dots, \kappa_n \sqrt{z_{n,t}})$ implies the opposite, which is a contradiction. Thus, D_t is strictly decreasing after t . Moreover, it follows immediately that no product is introduced after period t . ■

To complete the proof of Theorem 3, we consider the periods when product i is introduced. Define E_i the latest period in which $J_{i,E_i} = \lambda_i$. Note that $x_{i,E_i} \geq 0$. Since $J_{i,E_i} = D_{E_i} + \kappa_i J_{i,E_i+1}$, it follows that $D_{E_i} = \lambda_i - \kappa_i J_{i,E_i+1} \geq \lambda_i(1 - \kappa_i)$. Let \tilde{S}_i be the highest t earlier than E_i such that $D_t \geq \lambda_i(1 - \kappa_i)$, i.e., $\tilde{S}_i = \min \{t : t \leq E_i \text{ and } D_t \geq \lambda_i(1 - \kappa_i)\}$. By definition, for any $t < \tilde{S}_i$, $D_t < \lambda_i(1 - \kappa_i)$, which implies that $x_{i,t} = 0$ by Claim 2. Similarly, for $\tilde{S}_i \leq t < E_i$, $D_t = \lambda_i(1 - \kappa_i)$, and $x_{i,t} \geq 0$. Since $\lambda_i(1 - \kappa_i) < \lambda_{i+1}(1 - \kappa_{i+1})$, then we must have that $\tilde{S}_{i+1} \geq E_i$.

Furthermore, we now prove that we can define $S_i \geq \tilde{S}_i$ such that $x_{i,t} > 0$ for $S_i \leq t < E_i$ and $x_{i,t} = 0$ for $\tilde{S}_i \leq t < S_i$. Indeed, suppose that there is a period t such that $\tilde{S}_i < t < E_i$ and $x_{i,t} = 0$. Since only product i could be introduced at t , we have $x_{j,t} = 0$ for all $j = 1, \dots, n$. Since $D_{t-1} = D_t = \lambda_i(1 - \kappa_i) \leq D_{t+1}$, we have that $D_t^2 \leq D_{t-1} D_{t+1}$. Using the fact that $z_{i,t} = \kappa_i z_{i,t-1}$ and $z_{i,t+1} \geq \kappa_i z_{i,t} = \kappa_i^2 z_{i,t-1}$ for all i , we have that $\left(\frac{\alpha^{t-1}}{(v_0 + \sum_{i=1}^n \kappa_i z_{i,t-1})^2} \right)^2 \leq$

$\frac{\alpha^{t-2}}{(v_0 + \sum_{i=1}^n z_{i,t-1})^2} \frac{\alpha^t}{(v_0 + \sum_{i=1}^n \kappa_i^2 z_{i,t-1})^2}$ which generates a contradiction again unless $z_{j,t} = 0$ for all $j = 1, \dots, n$, using the same argument as in Claim 3 above. As a result, it must be true that product i is not introduced before time period t . Hence we can define S_i as the first period after \tilde{S}_i and before E_i such that $x_{i,t} > 0$. Finally, note that $S_1 = 1$ because otherwise $D_1 = \frac{1}{v_0^2} > D_t$ for all $t \geq 2$, which would mean that no product is introduced at all.

We have thus shown that each product i is introduced during the interval $[S_i, E_i]$ with $\bar{x}_{it} > 0$ for $S_i \leq t < E_i$, $\bar{x}_{i,E_i} \geq 0$, and $\bar{x}_{it} = 0$ for all $t \notin [S_i, E_i]$. Moreover, we have shown that $S_1 = 1$, $E_i \leq S_{i+1}$ for all i . The last part of Theorem 3 follows from Claims 2 and 3, and the fact that $x_{n,E_{n-1}} > 0$, $x_{n,E_n} \geq 0$, and $x_{j,t} = 0$ for all j and $t > E_n$. ■

Proof of Proposition 4

Proof. Fix an arbitrary κ . To simplify our notation, we will refer to $E_1(\kappa)$ as simply E . We make use of an equivalent continuous relaxation problem given in terms of $\mathbf{z} = (z_1, \dots, z_T)$ where $z_t = \sum_{u=1}^t \kappa^{t-u} x_u = x_t + \kappa z_{t-1}$, for any $t \geq 1$. Note that there is a one-to-one correspondence between \mathbf{x} and \mathbf{z} . The optimization problem is given by:

$$\max \left\{ H(z_1, \dots, z_T) := \sum_{t=1}^T \frac{\alpha_t z_t}{v_0 + z_t} \mid \sum_{t=1}^T (z_t - \kappa z_{t-1}) \leq v_1 \text{ and } z_t - \kappa z_{t-1} \geq 0, t = 1, \dots, T \right\}$$

which in turn is equivalent to the optimization problem:

$$\max \left\{ H(z_1, \dots, z_T) \mid z_T + (1 - \kappa) \sum_{t=1}^{T-1} z_t \leq v_1 \quad \text{and} \quad \kappa z_{t-1} - z_t \leq 0 \quad t = 1, \dots, T \right\}$$

Let $\bar{\mathbf{z}} = (\bar{z}_1, \bar{z}_2, \dots, \bar{z}_T)$ denote an optimal solution associated with the above optimization problem, and let $\bar{\mathbf{x}} = (\bar{x}_1, \bar{x}_2, \dots, \bar{x}_T)$ denote the corresponding optimal product schedule.

By associating the dual variable $\lambda \geq 0$ with the constraint $z_T + (1 - \kappa) \sum_{t=1}^{T-1} z_t \leq v_1$, and the variable $\eta_t \geq 0$ with the constraint $\kappa z_{t-1} - z_t \leq 0$, it follows from the KKT condition that $\bar{\mathbf{z}}$ satisfies the following set of equations:

$$\begin{aligned} 0 &= \frac{\partial H}{\partial z_1}(\bar{\mathbf{z}}) - \lambda(1 - \kappa) + \eta_1 - \kappa\eta_2 && \Leftrightarrow && \lambda(1 - \kappa) + \kappa\eta_2 = \eta_1 + D_1 \\ &\vdots && && \\ 0 &= \frac{\partial H}{\partial z_{T-1}}(\bar{\mathbf{z}}) - \lambda(1 - \kappa) + \eta_{T-1} - \kappa\eta_T && \Leftrightarrow && \lambda(1 - \kappa) + \kappa\eta_T = \eta_{T-1} + D_{T-1} \\ 0 &= \frac{\partial H}{\partial z_T}(\bar{\mathbf{z}}) - \lambda + \eta_{T-1} && \Leftrightarrow && \lambda = \eta_{T-1} + D_T \end{aligned}$$

and we have the following complementary slackness condition: $\lambda \left(z_T + (1 - \kappa) \sum_{t=1}^{T-1} \bar{z}_t - v \right) = 0$ and $\eta_t(\bar{z}_t - \kappa \bar{z}_{t-1}) = 0$ for all t .

From Theorem 3, we know that $\bar{x}_t > 0$ for $t = 1, \dots, E$ and $\bar{x}_t = 0$ for $t > E$. We claim that $0 = \eta_1 = \eta_2 = \dots = \eta_E < \eta_{E+1} < \eta_{E+2} < \dots < \eta_T$. Indeed, it follows immediately from the complementary slackness condition that $\eta_1 = \eta_2 = \dots = \eta_E = 0$. To prove the strict inequalities, we start showing that $0 = \eta_E < \eta_{E+1}$. Indeed, if this was not true and $\eta_{E+1} = 0$, then $\eta_{E+1} + D_{E+1} < \eta_E + D_E$, because, from Theorem 3 we know that $D_E > D_{E+1} > \dots > D_T$. But we know from the KKT condition that $\eta_{E+1} + D_{E+1} = \lambda(1 - \kappa) + \kappa\eta_{E+2} \geq \lambda(1 - \kappa) = \lambda(1 - \kappa) + \eta_{E+1} = \eta_E + D_E$, which would be a contradiction. We then prove that $\eta_{E+1} < \eta_{E+2} < \dots < \eta_T$. Suppose on the contrary that $\eta_{E+2} \leq \eta_{E+1}$. This means that $\eta_{E+2} + D_{E+2} \leq \eta_{E+1} + D_{E+1}$. By the KKT condition once again, $\lambda(1 - \kappa) + \kappa\eta_{E+3} = \eta_{E+2} + D_{E+2} \leq \eta_{E+1} + D_{E+1} = \lambda(1 - \kappa) + \kappa\eta_{E+2}$, which implies that $\eta_{E+3} \leq \eta_{E+2}$. Continuing in this fashion, it follows that $\eta_T \leq \eta_{T-1} \leq \dots \leq \eta_{E+2} \leq \eta_{E+1}$. However, the last two equations in the KKT condition imply that $\eta_T = \eta_{T-1} + D_{T-1} - D_T > \eta_{T-1}$ because $D_T < D_{T-1}$. This is again a contradiction. Therefore, it must be the case that $\eta_{E+2} > \eta_{E+1}$. The same argument shows that $\eta_{E+3} > \eta_{E+2}$, and so on.

Since the above result holds for an arbitrary κ , it follows that

$$E(\kappa) = \begin{cases} \min \{t : \eta_t(\kappa) > 0\} - 1, & \text{if } \eta_t(\kappa) > 0 \text{ for some } t, \\ T, & \text{otherwise.} \end{cases}$$

It is easy to verify that if $\kappa = 0$, then $\bar{x}_t(\kappa) > 0$ for all t , which implies that $E(\kappa) = T$. Since the dual variables are monotone, as κ increases, the first dual variable that will become positive is $\eta_T(\kappa)$, corresponding to $E(\kappa) = T - 1$. Then, the next dual variable that will become positive is $\eta_{T-1}(\kappa)$, corresponding to $E(\kappa) = T - 2$. Continuing in this fashion, we observe that as κ increases, the cut-off period $E(\kappa)$ will gradually decrease from T to $T - 1$ to $T - 2$, and so on. \blacksquare

Proof of Proposition 5

Proof. We proceed as in the proof of Proposition 4 to derive the KKT conditions:

$$\begin{aligned} 0 = \frac{\partial H}{\partial z_1}(\bar{z}) - \lambda(1 - \kappa) + \eta_1 - \kappa\eta_2 & \Leftrightarrow \lambda(1 - \kappa) + \kappa\eta_2 = \eta_1 + \frac{v_0\alpha_1}{(v_0 + \bar{z}_1)^2} \\ \vdots & \\ 0 = \frac{\partial H}{\partial z_{T-1}}(\bar{z}) - \lambda(1 - \kappa) + \eta_{T-1} - \kappa\eta_T & \Leftrightarrow \lambda(1 - \kappa) + \kappa\eta_T = \eta_{T-1} + \frac{v_0\alpha_{T-1}}{(v_0 + \bar{z}_{T-1})^2} \\ 0 = \frac{\partial H}{\partial z_T}(\bar{z}) - \lambda + \eta_{T-1} & \Leftrightarrow \lambda = \eta_T + \frac{v_0\alpha_T}{(v_0 + \bar{z}_T)^2} \end{aligned}$$

with the following complementary slackness conditions: $\lambda \left(z_T + (1 - \kappa) \sum_{t=1}^{T-1} \bar{z}_t - v \right) = 0$ and $\eta_t (\bar{z}_t - \kappa \bar{z}_{t-1}) = 0$ for all t .

Since we have a single product, we know from Theorem 3 that there exists E such that $\bar{x}_t > 0$ for $t \leq E$ and $\bar{x}_t = 0$ for $t > E$. It follows that for any $k \geq 1$, $\bar{z}_{E+k} = \kappa^k \bar{z}_E$, which implies that $\bar{z}_E \geq \bar{z}_{E+1} \geq \bar{z}_{E+2} \geq \dots \geq \bar{z}_T$. So, it suffices to show that $\bar{z}_1 \geq \bar{z}_2 \geq \dots \geq \bar{z}_E$. We will prove the desired result by contradiction. Suppose on the contrary that $\bar{z}_{k-1} < \bar{z}_k$ for some $k \in \{2, 3, \dots, E\}$. There are two cases to consider: $k < T$ and $k = T$.

If $k < T$, since $\bar{x}_{k-1} > 0$ and $\bar{x}_k > 0$, we have $\eta_{k-1} = \eta_k = 0$, which implies that $\lambda(1 - \kappa) = \frac{v_0 \alpha_{k-1}}{(v_0 + \bar{z}_{k-1})^2}$ and $\lambda(1 - \kappa) + \kappa \eta_{k+1} = \frac{v_0 \alpha_k}{(v_0 + \bar{z}_k)^2}$. Subtracting the first equation from the second yields $0 \leq \kappa \eta_{k+1} = \frac{v_0 \alpha_k}{(v_0 + \bar{z}_k)^2} - \frac{v_0 \alpha_{k-1}}{(v_0 + \bar{z}_{k-1})^2} = v_0 \alpha_{k-1} \left(\frac{\alpha}{(v_0 + \bar{z}_k)^2} - \frac{1}{(v_0 + \bar{z}_{k-1})^2} \right) < 0$, where the last inequality follows from the fact that $\bar{z}_k > \bar{z}_{k-1}$. This is a contradiction.

If $k = T$, then again $\eta_{T-1} = \eta_T = 0$. A similar argument as above implies that $0 \leq \kappa \lambda = \frac{v_0 \alpha_T}{(v_0 + \bar{z}_T)^2} - \frac{v_0 \alpha_{T-1}}{(v_0 + \bar{z}_{T-1})^2} = v_0 \alpha_{T-1} \left(\frac{\alpha}{(v_0 + \bar{z}_T)^2} - \frac{1}{(v_0 + \bar{z}_{T-1})^2} \right) < 0$, yielding again a contradiction. \blacksquare

Proof of Theorem 6

Proof. Assuming $\alpha_t = 1$ for all $t = 1, \dots, T$, the continuous relaxation of the problem can be recast as an equivalent convex minimization problem:

$$\min \sum_{t=1}^T \frac{1}{v_0 + z_t} \quad \text{s.t.} \quad \sum_{t=1}^T z_t = W, z_t \geq 0, \forall t,$$

where $W = \sum_{i=1}^n v_i$. The minimum is achieved by setting $\bar{z}_t = W/T$.

Denote V^h and z_1^h, \dots, z_T^h the value and loads in each period provided by the heuristic. Note that when $\alpha_t = 1$ and products last for one period only, items are introduced in decreasing order of v_i . Without loss of generality, we can assume that no periods are such that they contain the same items in the optimal solution and the heuristic solution, because if so, the ratio V^*/V^h would become $(V^* + a)/(V^h + a) \geq V^*/V^h$ where a is the profit obtained in the periods where the two solutions coincide.

Since there are T periods with total load of $W = \sum_{i=1}^n v_i$ (more generally W would be the total load of the items where optimum and heuristic solution do not coincide), there must be at least one with load equal or larger than W/T , and at least one with load equal or smaller than W/T .

$$\text{First, we have a lower bound of } V^* \leq \bar{V} = T - \frac{v_0 T^2}{v_0 T + W} = \frac{TW}{v_0 T + W}.$$

Second, we show that $z_t^h \in [W/(2T), 2W/T]$. If $z_t^h > 2W/T$ then it is not possible that period t only contains one item, because if so this would also be part of the optimal solution (indeed,

because if it had two items in the optimal solution, then it would be possible to increase V^* by moving one item to another bin with load smaller than W/T , contradicting the fact that no bins coincide in the heuristic and optimal solution. Hence, period t has at least two items. Since there is at least a bin with $z_t^h \leq W/T$, then one should not have added the last item to that bin because there was a bin less full. Hence $z_t^h \geq 2W/T$. If $z_t^h < W/(2T)$ then we can examine all the bins such that $z_t^h \geq W/T$ (there is at least one like that). If none of them has more than one item, then again it must be true that all those are part of the optimal solution, which contradicts our initial assumption. We can then find one with at least two items. Select the smallest item there, $v_i \leq W/(2T)$. This item was introduced into the bin when the size of the bin was larger than $W/(2T)$, which contradicts the way the heuristic inserts items into periods.

As a result, we have that

$$V^h \geq T - \max_{z_t} \left\{ \sum_{t=1}^T \frac{v_0}{v_0 + z_t} \mid \frac{W}{2T} \leq z_t \leq \frac{2W}{T} \text{ and } \sum_{t=1}^T z_t = W \right\}. \quad (8)$$

Since the objective is convex, the maximum is achieved by setting some $z_t = \frac{W}{2T}$ and some others to $z_t = \frac{2W}{T}$ and the last one to another value. Letting $Z_t = (2Tz_t/W - 1)/3$, we can transform the maximization above into

$$\begin{aligned} & \max \left\{ \frac{2T}{3W} \sum_{t=1}^T \frac{v_0}{\frac{2Tv_0/W+1}{3} + Z_t} \mid 0 \leq Z_t \leq 1 \text{ and } \sum_{t=1}^T Z_t = T/3 \right\} \\ & \leq \frac{Tv_0}{3W} \left(\frac{\frac{T}{3}}{\frac{2Tv_0/W+1}{3} + 1} + \frac{\frac{2T}{3}}{\frac{2Tv_0/W+1}{3}} \right) = \frac{T^2v_0(2Tv_0 + 3W)}{(Tv_0 + 2W)(2Tv_0 + W)}. \end{aligned}$$

Hence,

$$1 \leq \frac{V^*}{V^h} \leq \frac{\frac{TW}{Tv_0+W}}{T - \frac{T^2v_0(2Tv_0+3W)}{(Tv_0+2W)(2Tv_0+W)}} = \frac{(\frac{Tv_0}{W} + 2)(\frac{Tv_0}{W} + \frac{1}{2})}{(\frac{Tv_0}{W} + 1)^2} \leq \frac{9}{8}$$

because the function $\phi(z) = \frac{(z+2)(z+1/2)}{(z+1)^2} = 1 + \frac{z}{2(z+1)^2}$ is first increasing until $z = 1$ and then decreasing. \blacksquare

Proof of Theorem 8

Proof. Let $\bar{z}_{it} = \bar{x}_{it} + \kappa_i \bar{z}_{it-1}$ and $z_{it}^h = x_{it}^h + \kappa_i z_{it-1}^h$ where \bar{x}_{it} and x_{it}^h are the optimal solution to the relaxation and the EARLYENTRY decision respectively. From Theorem 3, we know that for each product i there exists S_i, E_i where the product is introduced during the interval $[S_i, E_i]$. The heuristic sets $x_{it}^h = 1$ for $t = S_i$ and zero otherwise. Hence, for $t < S_i$, $z_{it} - z_{it}^h = 0$; for $S_i \leq t \leq E_i$, $\bar{z}_{it} - z_{it}^h \leq 1 - \kappa_i^{t-S_i}$; for $t > E_i$, $\bar{z}_{it} - z_{it}^h \leq (1 - \kappa_i^{E_i-S_i})\kappa_i^{t-E_i}$. As a result,

$$\frac{\bar{z}_{it} - z_{it}^h}{z_{it}^h} \leq 1_{t > S_i} \left(\frac{1}{\kappa_i^{E_i - S_i}} - 1 \right) \text{ and hence}$$

$$0 \leq \bar{V} - V^h = \sum_{t=1}^T \left(\frac{\alpha_t \sum_{i=1}^n v_i z_{it}^h}{v_0 + \sum_{i=1}^n v_i z_{it}^h} \right) \left(\frac{v_0}{v_0 + \sum_{i=1}^n v_i z_{it}} \frac{\sum_{i=1}^n v_i (z_{it} - z_{it}^h)}{\sum_{i=1}^n v_i z_{it}^h} \right) \leq V^h \max_{i=1, \dots, n} \left\{ \frac{1}{\kappa_i^{E_i - S_i}} - 1 \right\}.$$

which leads to the final result. \blacksquare

Proof of Theorem 9

Proof. We only need to show that $\mathbb{E}[f(Q_1, \dots, Q_n)] \geq \rho \bar{V}$. For this purpose, note that

$$\begin{aligned} \mathbb{E}[f(Q_1, \dots, Q_n)] &= \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \mathbb{E} \left[\frac{\sum_{u=1}^t \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + v_i \sum_{u=1}^t \kappa_{i,t-u} \mathbb{1}_{[Q_i=u]} + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \mathbb{1}_{[Q_j=u]}} \right] \\ &\geq \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \mathbb{E}_{Q_i} \left[\frac{\sum_{u=1}^t \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + v_i \sum_{u=1}^t \kappa_{i,t-u} \mathbb{1}_{[Q_i=u]} + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \mathbb{E}_{Q_j} [\mathbb{1}_{[Q_j=u]}]} \right] \\ &= \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \mathbb{E}_{Q_i} \left[\frac{\sum_{u=1}^t \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + v_i \sum_{u=1}^t \kappa_{i,t-u} \mathbb{1}_{[Q_i=u]} + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \bar{x}_{ju}} \right] \\ &\geq \sum_{t=1}^T \alpha_t \sum_{i=1}^n r_i \mathbb{E}_{Q_i} \left[\frac{\sum_{u=1}^t \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + v_i + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \bar{x}_{ju}} \right]. \end{aligned}$$

The first inequality follows from the fact that the random variables Q_j 's are independent, and from applying Jensen's Inequality where we use the fact that the function $x \rightarrow \frac{a}{b+x}$ is convex in x . The second equality follows from the fact that $\mathbb{E}_{Q_j} [\mathbb{1}_{[Q_j=u]}] = \Pr(Q_j = u) = \bar{x}_{ju}$, and the last inequality follows because $\sum_{u=1}^t \kappa_{i,t-u} \mathbb{1}_{[Q_i=u]} \leq 1$ almost surely.

To complete the proof, note that for each product i , we have

$$\begin{aligned} \mathbb{E}_{Q_i} \left[\frac{\sum_{u=1}^t \kappa_{i,t-u} v_i \mathbb{1}_{[Q_i=u]}}{v_0 + v_i + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \bar{x}_{ju}} \right] &= \frac{\sum_{s=1}^t \bar{x}_{is} \kappa_{i,t-s} v_i}{v_0 \left(1 + \frac{v_i}{v_0} \right) + \sum_{j \neq i} v_j \sum_{u=1}^t \kappa_{j,t-u} \bar{x}_{ju}} \\ &\geq \left(\frac{v_0}{v_0 + v_i} \right) \left(\frac{\sum_{s=1}^t \kappa_{i,t-s} v_i \bar{x}_{is}}{v_0 + \sum_{j=1}^n \sum_{s=1}^t \kappa_{j,t-s} v_j \bar{x}_{js}} \right) \\ &\geq \left(\frac{v_0}{v_0 + \max_{j=1, \dots, n} v_j} \right) \left(\frac{\sum_{s=1}^t \kappa_{i,t-s} v_i \bar{x}_{is}}{v_0 + \sum_{j=1}^n \sum_{s=1}^t \kappa_{j,t-s} v_j \bar{x}_{js}} \right), \end{aligned}$$

which implies the desired result. \blacksquare

References

- Abernathy, F. H., J. T. Dunlop, J. H. Hammond, and D. Weil. 1999. *A stitch in time: Lean retailing and the transformation of manufacturing—lessons from the apparel and textile industries*. Oxford University Press, USA.
- Ainslie, A., X. Drèze, and F. Zufryden. 2005. Modeling Movie Life Cycles and Market Share. *Marketing Science* 24 (3): 508–517.

- Allon, G., A. Federgruen, and M. Pierson. 2010. Price Competition Under MNL Demand Functions with Random Coefficients. Working paper, Northwestern University, Kellogg School of Management.
- Alptekinoglu, A., D. Honhon, and C. Ulu. 2011. Learning Consumer Tastes Through Dynamic Assortments. Working paper, University of Texas at Austin.
- Aydin, G., and E. Porteus. 2008. Joint inventory and pricing decisions for an assortment. *Operations Research* 56 (5): 1247–1255.
- Bell, D. E., R. L. Keeney, and J. D. C. Little. 1975. A Market Share Theorem. *Journal of Marketing Research* 12 (2): pp. 136–141.
- Benson, H. P. 2007. A simplicial branch and bound duality-bounds algorithm for the linear sum-of-ratios problem. *European journal of operational research* 182 (2): 597–611.
- Bernstein, F., A. G. Kök, and L. Xie. 2010. Dynamic assortment customization with limited inventories. Working paper, Duke University.
- Bertsimas, D., and I. Popescu. 2003. Revenue management in a dynamic network environment. *Transportation Science* 37 (3): 257–277.
- Bertsimas, D., C. Teo, and R. Vohra. 1999. On dependent randomized rounding algorithms. *Operations Research Letters* 24:105–114.
- Besbes, O., and D. Sauré. 2010. Product Assortment and Price Competition Under Multinomial Logit Demand. Working paper, Columbia Graduate School of Business.
- Bront, J. M., I. Mendez Diaz, and G. Vulcano. 2009. A column generation algorithm for choice-based network revenue management. *Operations Research* 57 (3): 769–784.
- Cachon, G., C. Terwiesch, and Y. Xu. 2005. Assortment Planning in the Presence of Consumer Search. *Manufacturing and Service Operations Management* 7:330–346.
- Caro, F., and J. Gallien. 2007. Dynamic assortment with demand learning for seasonal consumer goods. *Management Science* 53 (2): 276–292.
- Druehl, C. T., G. M. Schmidt, and G. C. Souza. 2009. The optimal pace of product updates. *European Journal of Operational Research* 192 (2): 621 – 633.
- Farias, V. F., S. Jagabathula, and D. Shah. 2010. A New Approach to Modeling Choice with Limited Data. Working paper, MIT.
- Farias, V. F., and R. Madan. 2011. Irrevocable Multi-Armed Bandit Policies. *Operations Research* 59 (2): 383–399.
- H&M 2007. Annual report. Technical report, H&M.
- Honhon, D., V. Gaur, and S. Seshadri. 2010. Assortment planning and inventory decisions under stockout-based substitution. *Operations Research* 58 (5): 1364–1379.
- Honhon, D., and A. G. Kök. 2011. Assortment planning with variety-seeking customers. Presented at the 2010 INFORMS Annual Meeting, Austin, TX.
- Kalish, S., V. Mahajan, and E. Muller. 1995. Waterfall and sprinkler new-product strategies in competitive global markets. *International Journal of Research in Marketing* 12 (2): 105–119.
- Kök, A. G., and M. Fisher. 2007. Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application. *Operations Research* 55:1001–1021.
- Kök, A. G., M. Fisher, and R. Vaidyanathan. 2008. Assortment planning: Review of literature and industry practice. In *Retail Supply Chain Management*: Springer.
- Kuno, T. 2002. A branch-and-bound algorithm for maximizing the sum of several linear ratios. *Journal of Global Optimization* 22 (1): 155–174.
- Lim, W. S., and C. S. Tang. 2006. Optimal product rollover strategies. *European journal of operational research* 174 (2): 905–922.
- Mahajan, S., and G. J. van Ryzin. 2001. Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49 (3): 334–351.
- Moskin, J. 2011. The Perfect Menu. Now Change It. *The New York Times* February 15.
- Raghavan, P., and C. Thompson. 1987. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7:365–374.

- Rusmevichientong, P., Z.-J. M. Shen, and D. B. Shmoys. 2010. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* 58 (6): 1666–1680.
- Sauré, D., and A. Zeevi. 2008. Optimal dynamic assortment planning. Working Paper, Columbia Graduate School of Business.
- Savin, S., and C. Terwiesch. 2005. Optimal Product Launch Times in a Duopoly: Balancing Life-Cycle Revenues with Product Cost. *Operations Research* 53 (1): 26–47.
- Schaible, S., and J. Shi. 2003. Fractional programming: the sum-of-ratios case. *Optimization Methods and Software* 18 (2): 219–229.
- Smith, S. A., and N. Agrawal. 2000. Management of Multi-item Retail Inventories Systems with Demand Substitution. *Operations Research* 48:50–64.
- Talluri, K., and G. J. van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50 (1): 15–33.
- The Economist 2007. Endless summer. How the business of making blockbusters has changed. *The Economist* April 26:online.
- Tiplady, R. 2006. Zara: Taking the lead in fast-fashion. *Business Week* April 4.
- van Ryzin, G., and S. Mahajan. 1999. On the Relationship Between Inventory Costs and Variety Benefits in Retail Assortments. *Management Science* 45:1496–1509.