

Pricing Promotional Products Under Upselling

Goker Aydin

Department of Industrial and Operations Engineering, University of Michigan,
Ann Arbor, Michigan 48109, ayding@umich.edu

Serhan Ziya

Department of Statistics and Operations Research, University of North Carolina at Chapel Hill,
Chapel Hill, North Carolina 27599, ziya@unc.edu

Upselling is offering an additional product to a customer who just made a purchase. Most catalogers and online sellers, in addition to some traditional retailers, use upselling often to clear inventories of slow-moving items. We investigate the pricing and discounting questions for such an item, which we call the promotional product. In our model, an arriving customer may purchase this promotional product or one of the other products that the firm sells. If the customer purchases one of the other products, the promotional product is offered to the customer, possibly with a discount. While deciding whether to offer a discount and, if so, how big a discount to offer, the firm uses the information that the customer has just bought a certain product with a certain price. We investigate how discounting decisions depend on the inventory levels, time, type of pricing policy in use, and the relationship between the customers' reservation prices for the promotional product and the other products (negatively or positively correlated). In particular, we find that if the firm sets prices and discounts dynamically and the customers' reservation prices for the promotional product are negatively correlated with their reservation prices for the product they purchased, then customers are always offered a discount regardless of the inventory levels and time. On the other hand, if the customers' reservation prices for the promotional product are positively correlated with their reservation prices for the product they purchased, then the customer may or may not be offered a discount, depending on the inventory levels and time. Our numerical study shows that the benefit to the firm from using customer purchase information is high when the firm uses a static price, but chooses discounts dynamically. We also find that although dynamic discounting decisions bring modest improvements, setting the price dynamically seems to have a more significant effect on the firm's profits.

Key words: dynamic pricing; pricing of limited inventories; cross-selling; bundling

History: Received: June 27, 2006; accepted: June 1, 2007. Published online in *Articles in Advance* January 4, 2008.

1. Introduction

Every time one buys a burger from a fast-food restaurant, one will be asked the famous question: "Do you want fries with that?" Those who avoid burgers in favor of subs will face a similar question after ordering a sandwich: "Do you want to add any chips or drinks?" Upselling refers to this practice of offering an additional product to a customer who just made a purchase.¹ Many catalog retailers such as Lillian Vernon, based in Rye, NY (Andrews

1999), Corona Cigar Co., based in Ocoee, FL (Ferriolo 2003), toy and novelty cataloger Oriental Trading Co., and apparel cataloger DM Management (Miller 1995) use upselling. Online retailers such as Amazon and Buy.com also use upselling by recommending additional products to customers who just purchased an item. Traditional retailers are no strangers to the practice either. Most people can remember being offered a tie that will go with their new shirt or a blouse that will be perfect with their new pants.

One goal of upselling is to increase revenue per customer visit. Retailers typically try to sell each customer an additional product related to the one the customer just bought, e.g., offering a battery charger to a customer who just bought a digital camera or offering a carrying case to a customer who just ordered a cell phone. Another popular reason for the

¹ The practice is also called cross-selling. Because the Federal Trade Commission (FTC) defines upselling in the context of telemarketing as "soliciting the purchase of goods or services following an initial transaction during a single telephone call," we use the term upselling instead of cross-selling. (See FTC's Telemarketing Sales Rule in Code of Federal Regulations Title 16, Part 310.)

use of upselling is liquidation of excess inventories. In fact, according to a study by *Catalog Age*, 42% of all catalog retailers and 63% of apparel catalogers use upselling to reduce excess merchandise (Del Franco et al. 1999). The retailer identifies slow-moving items with excess inventory and tries to liquidate that inventory by offering the items, possibly at a discount, to customers who purchase other items. For example, the cataloger DM Management is reported to use a computerized method to identify overstock items, which are then promoted by sales representatives (Miller 1995). In this paper, we focus on upselling practices aimed at liquidation of excess inventories. In particular, we assume that the firm has already chosen an item for upselling (hereafter, the *promotional product*), and we focus on the pricing of this promotional product. Some natural choices for the promotional product would be a discontinued product or a seasonal product with excess inventory.

Whereas the inventory of the promotional product is liquidated through upselling, the products that are not subject to upselling, which we refer to as *regular products*, are likely to be staple items sold over an extended period of time, or fast-moving seasonal products. Staple items are replenished regularly, and we expect that lost sales due to stockouts will be rare for such products. On the other hand, when the regular product is a fast-selling seasonal product with no replenishment opportunities, there is a chance that the regular product will run out of stock before the promotional product does, which needs to be taken into account when upselling the promotional product. We first consider the case in which the regular product has unlimited supply, so that it is available whenever a customer demands it. We then consider what happens when the regular product has limited availability.

When a firm is upselling a product, in addition to those customers buying the product on an upsell offer there will be others who arrive with the intention of purchasing the product in the first place. Such customers will observe an announced price for the product. If the firm is trying to clear the excess inventory of the product, then it would be natural for the firm to adjust the announced price over time as long as it is feasible to do so. This may be quite possible for an online retailer, whereas a cataloger may

find it harder to do because the initial price is already published in the catalog. Regardless of whether the firm engages in such dynamic pricing the firm can choose to use upselling. In cases where dynamic pricing is used, upselling can be seen as a promotional tool that complements dynamic pricing as a means of clearing inventory. We consider both the case in which the firm is dynamically adjusting the announced price over time and the case in which there is a fixed price that cannot be adjusted.

When making an upsell offer, the firm must decide whether the offer should be accompanied by a discount from the announced price. Whether the upsell offer will include a discount may depend on the remaining inventory of the promotional product and the time remaining until the end of the selling season. In addition, the firm may be able to tailor the discounting decisions at the individual customer level. A firm using upselling has one valuable piece of information about the customer who is about to receive the upsell offer: The firm knows that this customer has just purchased a certain product at a certain price. This information sets this particular customer apart from a random customer in the population and allows the firm to make a better-informed decision on whether to offer a discount. Here we analyze the benefits from using such information.

In some cases, the firm may be able to decide not only whether the upsell offer to an individual customer should include a discount, but also how deep the discount should be. For example, an online retailer could send customized coupons to individual customers. Likewise, a cataloger can make different discount offers to different customers while the customer is making a purchase over the phone. On the other hand, a firm may choose to fix the discount level (e.g., using a coupon with a given face value) while deciding whether to offer the discount at the individual customer level (e.g., whether to offer the coupon to a given customer.) We consider both the case in which the discount level can differ across customers and the case in which the discount level needs to be the same for all customers who get a discount.²

² Offering different discounts to different customers may raise questions of legality. In the form of upselling that we investigate, the discount is akin to offering a discount on a bundle of products.

We first analyze the dynamic price, dynamic discount case where the firm can dynamically adjust the announced price of the promotional product as well as customize the discount that comes with an upsell offer. We show that an upsell offer will include a discount if and only if the promotional product and the regular product are *dissimilar*, irrespective of the inventory status and time left until the deadline. The two products being dissimilar refers to a set of conditions that imply that a customer's reservation prices for the two products are negatively correlated. Likewise, we say that two products are *similar* if a customer's reservation prices for the two products are positively correlated. When the products are similar, we show that the firm may or may not offer a discount, depending on the inventory levels of the regular and promotional products.

We also consider a number of extensions to the base model: the case of multiple regular products; the static price, dynamic discount case (where the announced price is fixed at the beginning but the discount can be adjusted); and the static price, static discount case (where the announced price and the discount are both fixed at the beginning, and the only dynamic decision is whether an upsell offer will include the discount). In addition, we investigate, by a numerical study, the benefits of using purchase information, dynamic pricing, and dynamic discounting.

In the next section, we position our work with respect to the existing literature. We then describe the model in §3 for the case where the announced price is dynamically adjusted and the discount level is customized for each customer. Section 4 presents a discussion of our results for that case. In §5, we consider the static price, dynamic discount and static price, static discount cases. We discuss the results of our numerical study in §6. We conclude with a discussion of future research directions in §7.

2. Literature Review

We discuss the relevant prior work under three headings: bundling, dynamic pricing, and cross-selling and upselling.

Stremersch and Tellis (2002) discuss the legality of bundling practices and claim that any kind of bundling is legal as long as the customer has the option to purchase the products separately, which is the case in this paper.

2.1. Bundling

The practice of upselling is closely related to bundling, i.e., selling two or more products as one package. One can think of upselling as a special form of bundling in which the customer is given the option to purchase the bundle only after she has committed to buying a product. There is a rich literature in marketing and economics on the practice of bundling. For a review, see Stremersch and Tellis (2002). The research on bundling in economics literature starts with the work of Adams and Yellen (1976) and generally focuses on the question of whether bundling is beneficial to the seller or the consumer population. The marketing research on bundling has considered the same question in addition to how consumers evaluate bundles (Yadav and Monroe 1993) and the optimal pricing of bundles (Hanson and Martin 1990). Taking a more operational perspective, given exogenously fixed bundling decisions, Ernst and Kouvelis (1999) consider the question of setting the inventory levels for individual products and bundles. More recently, Gurler et al. (2005) investigate the dynamic pricing of two products with limited inventories, which are sold separately or as a bundle.

2.2. Dynamic Pricing

Our work is also closely related to the stream of research dealing with dynamic pricing of limited inventories pioneered by Gallego and van Ryzin (1994) and by Bitran and Mondschein (1997). (In particular, our model follows the framework used by Bitran and Mondschein in that we divide the selling season into short discrete time intervals.) For recent reviews of the literature on dynamic pricing of limited inventories, see Bitran and Caldentey (2003) and Elmaghraby and Keskinocak (2003). As dynamic pricing has become more popular in the retail industry, it has been receiving increased attention. Many interesting variations of the problem have been considered recently, such as dynamic pricing for multiple products (Zhang and Cooper 2005 and Maglaras and Meissner 2005); dynamic pricing in the presence of strategic consumers (Elmaghraby et al. 2008, Aviv and Pazgal 2005a); and the choice of initial stock level under dynamic pricing (Monahan et al. 2004).

In their review paper, Elmaghraby and Keskinocak (2003) note that a promising direction for future

research is the possibility of using customer-specific information to make customized offers in the presence of inventory considerations. We consider this in the context of upselling. In our model, the use of customer purchase information is rather simple in that the firm uses only the fact that the customer has just bought a certain product with a certain price, ignoring any past information it might have. The firm could start with some beliefs on the arrival rate of customers and their reservation price distributions, and could update those beliefs on the basis of observations made throughout the selling season. For an analysis of such sophisticated uses of learning by a firm with limited inventories, see Aviv and Pazgal (2005b).

2.3. Cross-Selling and Upselling

A recent work that is closely related to ours is by Netessine et al. (2006), who refer to the practice as cross selling. Netessine et al. (2006) address the problem of what additional product to offer to a customer who made a purchase, i.e., the packaging question. In contrast, we eliminate the packaging question by assuming that the firm has already decided what product to offer. This simplification allows us to obtain additional insights as follows: (i) Netessine et al. (2006) assume that a customer's reservation price for a two-product package has an exogenously given distribution. In contrast, we explicitly model the reservation prices of a customer for individual products. We find that whether the correlation between those reservation prices is negative or positive has a substantial effect on pricing decisions. (ii) In addition to the case of static announced prices considered by Netessine et al. (2006), we consider the case where the promotional product's announced price is dynamically adjusted. This allows us to gain insights into the interaction between upselling and dynamic pricing. (iii) In upselling, knowing that the customer bought a product at a certain price has informational value to the firm when offering the promotional product. We are able to investigate the value of such information because we explicitly model the relationship between the reservation prices for different products.

3. Model Description

Initially, we assume that the firm is offering two products. One of these is the *promotional product*, the product the firm is upselling. We refer to the other product

as the *regular product*. Later, we extend our results to the case where there are multiple regular products.


In this section, we describe our model assuming that both the announced price of the promotional product and the discount that comes with an upsell offer can be dynamically adjusted over time. Later, in §6, we will modify this model to analyze the cases where the firm has limited flexibility in adjusting the price and/or the discount. We assume that the regular product's price is exogenously fixed at r . This would certainly be true in the case of a cataloger because the initial prices published in the catalog cannot be changed during the relatively short selling season of the promotional product. In addition, certain retailers (e.g., retailers following the everyday-low-price format) may avoid frequent price changes on regular products, while using dynamic pricing to clear the inventory of a promotional product.

In keeping with our focus on the use of upselling to liquidate excess inventories, we assume that the promotional item has a fixed initial inventory that needs to be sold by a certain deadline, and that there are no replenishment opportunities for the promotional product. Initially, we assume that the regular product is always available when a customer demands it; we relax this assumption later to consider the limited availability of the regular product.

3.1. Reservation Prices and Consumer Population

For each product, we assume that the customer population is clustered into two segments: the *target* segment and the *nontarget* segment. The target customers are the primary consumers of the product, whereas the nontarget customers are all the others in the population. A customer may belong to the target segment of one product and the nontarget segment of the other product. We use 1 as the index for the target segment and 2 as the index for the nontarget segment. We let q_{Ri} , $i = 1, 2$, denote the probability that a customer belongs to segment i of the regular item. Similarly, q_{Pi} , $i = 1, 2$, denotes the probability that a customer belongs to segment i of the promotional item. Let δ_{ij} , $i, j = 1, 2$, denote the probability that a customer belongs to segment j of the promotional item given that she belongs to segment i of the regular item. (Note that, by definition, $\sum_{j=1}^2 \delta_{ij} = 1$, $i = 1, 2$.) For clarification, consider the following example. Suppose the promotional product is a slow-selling

Figure 1 Let a, b, c, d Denote the Fraction of Customers that Belong to the Groups Depicted in the Venn Diagram Below



Notes. Then: $q_{p1} = b + c$, $q_{R1} = a + b$, $\delta_{11} = b/(a + b)$, $\delta_{21} = c/(c + d)$.

new album by the Rolling Stones and the regular product is a classic album of the Rolling Stones that has had steady sales over the years. Then, one would expect δ_{11} to be large. On the other hand, if the regular product is an album by a pop star, say Britney Spears, then we would expect δ_{11} to be smaller. See Figure 1 for a visual depiction of the problem parameters pertaining to the customer population. Note that for the model parameters to be consistent with each other, the parameter values need to satisfy the following

$$q_{pi} = q_{R1}\delta_{1i} + q_{R2}\delta_{2i}, \quad i = 1, 2. \quad (1)$$

Let $F_{Ri}(\cdot)$, $i = 1, 2$, be the cumulative distribution function (cdf) of the reservation price of segment i customers for the regular item. Similarly, $F_{Pi}(\cdot)$, $i = 1, 2$, is the cdf of the reservation price of segment i customers for the promotional item. We let $\bar{F}_{Ri}(x) := 1 - F_{Ri}(x)$ and $\bar{F}_{Pi}(x) := 1 - F_{Pi}(x)$. For two cdfs Φ_1 and Φ_2 (with corresponding pdfs ϕ_1 and ϕ_2), if the failure rate of Φ_1 is strictly less than that of Φ_2 , i.e., if $\phi_1(x)/\bar{\Phi}_1(x) < \phi_2(x)/\bar{\Phi}_2(x)$, $\forall x$, then we say Φ_1 strictly dominates Φ_2 in failure rate ordering, and we write $\Phi_1 >_{fr} \Phi_2$. (See, for example, Müller and Stoyan 2002.) For any cdf Φ and the corresponding pdf ϕ , generalized failure rate is defined to be $x\phi(x)/(1 - \Phi(x))$. Throughout the paper, we use increasing/decreasing and positive/negative in the weak sense, unless specifically qualified as strictly increasing/decreasing or nonpositive/negative. We make the following assumptions on the reservation price distributions.

ASSUMPTION 1(A1). $F_{Ri}(\cdot)$, $F_{Pi}(\cdot)$, $i = 1, 2$, are twice continuously differentiable, strictly increasing functions, and they all have the same nonnegative support.

ASSUMPTION 2(A2). $F_{Ri}(\cdot)$, $F_{Pi}(\cdot)$, $i = 1, 2$, have strictly increasing generalized failure rates.

ASSUMPTION 3(A3). $F_{R1}(\cdot) >_{fr} F_{R2}(\cdot)$ and $F_{P1}(\cdot) >_{fr} F_{P2}(\cdot)$.

The first assumption is technical in nature. The second assumption is needed to provide some regularity to the revenue function, and it is satisfied by a variety of probability distributions, including the positive-valued part of the normal distribution and all Weibull distributions. (For a comparison of various assumptions on reservation prices used in revenue management problems, see Ziya et al. 2004.) The ordering stated in (A3) can be interpreted to mean that the absolute price elasticity of demand of a target customer is less than that of a nontarget customer; i.e., customers in the target segment are less price sensitive. A further implication of (A3) is that, for each product, the reservation price of a customer in the target segment stochastically dominates that of a customer in the nontarget segment, i.e., $F_{R1}(x) < F_{R2}(x)$, and $F_{P1}(x) < F_{P2}(x)$, $\forall x$.

With these assumptions, correlation between the reservation prices for the two products and δ_{ij} values are related as described in Proposition 1. (See Appendix A in the online supplement for a proof.³)

PROPOSITION 1. Let X and Y denote the reservation prices of a randomly chosen customer for the regular and the promotional products, respectively. These two reservation prices are negatively (positively) correlated (i.e., $\text{Cov}(X, Y) < (>)0$) if and only if $\delta_{11} + \delta_{22} < (>)1$.

Throughout the remainder of the paper, we will use the following definition for ease of exposition.

DEFINITION 1. If $\delta_{11} + \delta_{22} < 1$, we say that the promotional and the regular products are dissimilar. Otherwise, we say that the promotional and the regular products are similar.

Hence, in this paper, two products are dissimilar if a randomly chosen customer's reservation prices for the two products are negatively correlated. Otherwise, the two products are similar.

³ An online supplement to this paper is available on the *Manufacturing & Service Operations Management* website (<http://msom.pubs.informs.org/ecompanion.html>).

3.2. Timing of Events

We now focus on the time horizon during which the promotional item will be sold. We assume that each arriving customer has only one product in her consideration set, either the regular product or the promotional product.⁴ Henceforth, we refer to a customer as a *potential customer for the regular product* if the customer has the regular product in her consideration set. Likewise, a *potential customer for the promotional product* is a customer who has the promotional product in her consideration set. We assume that the time horizon is divided into T periods, each of which is short enough that either one customer arrives in each period or none at all. (This assumption is common in dynamic pricing models. See, for example, Bitran and Mondschein 1997.) If an arriving consumer is a potential customer for the regular product and ends up buying it, then the firm offers her the promotional product. To model this process, we divide each period into two stages: initial stage and upsell stage. In the initial stage, the following is the sequence of events.

1. At the beginning of the initial stage, the seller announces p_t , the price for the promotional item that will be in effect during the initial stage of period t . (Recall that the regular item's price is fixed at r .)

2. During the initial stage, one of three events happens: A potential customer for the regular item arrives with probability λ_R ; a potential customer for the promotional item arrives with probability λ_P ; and no customer arrives with probability $1 - \lambda_R - \lambda_P$.

3. A potential customer for the regular product will purchase it if r is less than or equal to her reservation price for the regular item. Likewise, a potential customer for the promotional product will purchase it if p_t is less than or equal to her reservation price for the promotional item.

The upsell stage occurs only if a customer arriving in the initial stage purchases the regular item. The following is the sequence of events for the upsell stage in a period.

1. The customer who purchases the regular item in the initial stage receives an offer from the firm for

the promotional item. Let d_t denote the absolute discount offered to the customer in the upsell stage of period t . Note that the firm may choose to offer no discount; i.e., the firm may offer the promotional item at price p_t .

2. The customer will purchase the promotional item on the upsell offer if $p_t - d_t$ is less than or equal to her reservation price for the promotional item.

3.3. Firm's Revenue-Maximization Problem

In the initial stage of period t , given that a potential customer for the regular item arrived, let $\beta_R(r)$ denote the probability that the customer will buy the regular item. Note that the customer belongs to the target segment with probability q_{R1} and the nontarget segment with probability q_{R2} , and will buy the regular product with probability $\bar{F}_{R1}(r)$ if she is in the target segment and with probability $\bar{F}_{R2}(r)$ if she is in the nontarget segment. Therefore, we can write $\beta_R(r)$ as follows

$$\beta_R(r) = q_{R1}\bar{F}_{R1}(r) + q_{R2}\bar{F}_{R2}(r).$$

Henceforth, we write $\beta_R(r)$ as β_R , because the regular product's price, r , is exogenous to our model. Let $\beta_P(p_t)$ denote the probability that a customer will buy the promotional item, given that a potential customer for the promotional item arrived in the initial stage of period t to observe the announced price of p_t . Then we can write $\beta_P(p_t)$ as follows

$$\beta_P(p_t) = q_{P1}\bar{F}_{P1}(p_t) + q_{P2}\bar{F}_{P2}(p_t). \quad (2)$$

If a customer buys the regular product, this provides the firm with additional information about the customer, namely, that the customer's reservation price for the regular product is above r . When making an upsell offer, the firm can use this new information to update the probabilities with which the customer belongs to the target and nontarget segments of the promotional item. Let \hat{q}_{P1} denote the probability that a customer belongs to the target segment of the promotional item given that she purchased the regular item in the initial stage, and \hat{q}_{P2} the probability that she belongs to the nontarget segment. Using Bayes' rule, we can write these updated probabilities as follows

$$\begin{aligned} \hat{q}_{P1} &= \frac{q_{R1}\delta_{11}\bar{F}_{R1}(r) + q_{R2}\delta_{21}\bar{F}_{R2}(r)}{\beta_R}, \\ \hat{q}_{P2} &= \frac{q_{R1}\delta_{12}\bar{F}_{R1}(r) + q_{R2}\delta_{22}\bar{F}_{R2}(r)}{\beta_R}. \end{aligned} \quad (3)$$

⁴ A consideration set is the set of products the consumer is considering for purchase. It would not be too difficult to extend the model to allow a third class of customers who have both products in their consideration set.

Let $\alpha(x)$ denote the probability that a customer will buy the promotional item on an upsell offer when the firm offers her the promotional product at a price of x . If the firm offers a discount d_t from the announced price p_t , the customer will be facing a price of $p_t - d_t$ for the promotional product. The following gives the probability that such a customer will buy the promotional item

$$\alpha(p_t - d_t) = \hat{q}_{P1} \bar{F}_{P1}(p_t - d_t) + \hat{q}_{P2} \bar{F}_{P2}(p_t - d_t). \quad (4)$$

Now we are ready to formulate the dynamic program for the firm's revenue-maximization problem over the T -period horizon. (We will denote the first period as period T and the last period as period 1.) Let $V_t(y)$ denote the firm's optimal expected revenue from the promotional product when starting period t with y units of inventory. The optimality equations are given by

$$V_t(y) = \max_{p_t, d_t: p_t \geq d_t \geq 0} \{ \lambda_R \beta_R [\alpha(p_t - d_t)(p_t - d_t + V_{t-1}(y - 1))] + \lambda_P \beta_P(p_t)(p_t + V_{t-1}(y - 1)) + (1 - \lambda_R \beta_R \alpha(p_t - d_t) - \lambda_P \beta_P(p_t)) V_{t-1}(y) \},$$

$$y > 0, t = 1, \dots, T,$$

$$V_t(0) = 0, \quad t = 1, \dots, T, \quad \text{and} \quad V_0(\cdot) = 0.$$

The first term of $V_t(y)$ is the revenue-to-go in the event that a customer purchases the promotional product on an upsell offer, the second term is the revenue-to-go if a customer arrives to purchase the promotional product in the first place, and the third term is the revenue-to-go if no promotional product is sold in that period. The terminal conditions indicate that the salvage value of the item is zero, an assumption that could be relaxed. After some algebraic manipulation, one can write $V_t(y)$ as follows

$$V_t(y) = V_{t-1}(y) + \max_{p_t, d_t: p_t \geq d_t \geq 0} \{ \lambda_P \beta_P(p_t)(p_t - \Delta(y, t)) + \lambda_R \beta_R \alpha(p_t - d_t)(p_t - d_t - \Delta(y, t)) \},$$

$$y > 0, t = 1, \dots, T, \quad (5)$$

where $\Delta(y, t) = V_{t-1}(y) - V_{t-1}(y - 1)$. Here, $\Delta(y, t)$ can be interpreted as the expected future benefit from carrying one unit of the promotional item into

period $t - 1$ when there are y units in inventory, i.e., the marginal value of one unit of inventory. Define

$$\Pi_1(p, \Delta) = \beta_P(p)(p - \Delta), \quad (6)$$

$$\Pi_2(p, \Delta) = \alpha(p)(p - \Delta). \quad (7)$$

Here, $\Pi_1(p, \Delta)$ can be interpreted as the expected marginal contribution from one unit of the promotional item in the initial stage of a period when the announced price is set at p and the marginal value of one unit of the promotional item is Δ . Likewise, $\Pi_2(p, \Delta)$ is the expected marginal contribution from one unit of the promotional item in the upsell stage, when the upsell price is p . Given $\Pi_1(p, \Delta)$ and $\Pi_2(p, \Delta)$ as defined by (6) and (7) and the optimality equations in (5), the maximization problem the firm is solving for a given t and y can be written as

$$\max_{p_t, d_t: p_t \geq d_t \geq 0} \{ \lambda_P \Pi_1(p_t, \Delta(y, t)) + \lambda_R \beta_R \Pi_2(p_t - d_t, \Delta(y, t)) \}. \quad (8)$$

Define

$$z_1^*(y, t) = \inf \{ p^*: \Pi_1(p^*, \Delta(y, t)) \geq \Pi_1(p, \Delta(y, t)), \forall p \}; \quad (9)$$

$$z_2^*(y, t) = \inf \{ p^*: \Pi_2(p^*, \Delta(y, t)) \geq \Pi_2(p, \Delta(y, t)), \forall p \}; \quad (10)$$

$$z^*(y, t) = \inf \{ p^*: \lambda_P \Pi_1(p^*, \Delta(y, t)) + \lambda_R \beta_R \Pi_2(p^*, \Delta(y, t)) \geq \lambda_P \Pi_1(p, \Delta(y, t)) + \lambda_R \beta_R \Pi_2(p, \Delta(y, t)), \forall p \}. \quad (11)$$

In light of the interpretations of $\Pi_1(p, \Delta)$ and $\Pi_2(p, \Delta)$ discussed earlier, we can interpret $z_1^*(y, t)$ and $z_2^*(y, t)$ as follows. With y units in inventory and t periods to go, $z_1^*(y, t)$ is the announced price that maximizes the profit in the initial stage of period t and $z_2^*(y, t)$ is the upsell price that maximizes the profit in the upsell stage of period t . Likewise, $z^*(y, t)$ would be the price the firm would choose if it had to use the same price in both stages. (When $\Pi_i(p, \Delta)$ has multiple optima, $z_i^*(y, t)$ and $z^*(y, t)$ are taken to be the smallest maximizers.) It is easy to construct numerical examples to demonstrate that $\Pi_1(p, \Delta)$ and $\Pi_2(p, \Delta)$ are not necessarily unimodal in p . Nevertheless, the

possibly complicated behavior of these functions does not render the problem intractable insofar as one can obtain a number of results on the optimal announced price and discount levels, which we discuss next.

4. Dynamic Prices and Discounts

Here, we first discuss our analytical results for the base model presented in the previous section. We then discuss how these results change with changes in the model (under limited availability of the regular product and when there are multiple regular products).

4.1. Results for the Base Model

One question of interest is whether the firm should accompany upsell offers with discounts on the promotional item and, if so, how deep those discounts should be. The answer to this question is closely related to the following comparison. Consider a customer randomly drawn from the population and another customer who is known to have purchased the regular item. Which customer is more likely to be in the target segment of the promotional item? In other words, using the notation defined earlier, which probability is larger: q_{P1} or \hat{q}_{P1} ? As the following proposition states, this comparison is decided simply on the basis of the values of δ_{ij} s. (The proofs of all the results in this section are given in Appendix A of the online supplement.)

PROPOSITION 2. *In comparison to a random customer, a customer who is known to have bought the regular item is less likely to be in the target segment of the promotional item if and only if the regular and the promotional items are dissimilar, i.e., $\hat{q}_{P1} < q_{P1}$ if and only if $\delta_{11} + \delta_{22} < 1$.*

The larger δ_{11} is, the more likely it is that a target customer for the regular item is a target customer for the promotional item as well. Likewise, the larger δ_{22} is, the more likely it is that a nontarget customer for the regular item is also a nontarget customer for the promotional item. Proposition 2 states that if $\delta_{11} + \delta_{22} < 1$ (with the minimum possible sum being zero and the maximum possible sum being two), then a customer who bought the regular product is less likely to be a target customer for the promotional product, compared to a customer chosen from the population at random.

Given the result of Proposition 2, intuition would suggest that if the two products are dissimilar, then a customer who bought the regular product should be given a discount on the promotional item as part of the upsell offer (because a customer who bought the regular item, compared to a random customer, is less likely to be in the target segment of the promotional item and, hence, more likely to have a smaller reservation price for the promotional item). On the other hand, if the two products are similar, then a customer who bought the regular item is more likely to be in the target segment of the promotional item, so the firm will not offer a discount from the announced price. If anything, the firm would like to charge an upsell price in excess of the announced price (had it been allowed), because a customer buying the regular product signals a higher willingness to pay for the promotional product as well. Theorem 1 formalizes this result.

THEOREM 1. *Let $z_1^*(y, t)$, $z_2^*(y, t)$, and $z^*(y, t)$ be as defined by (9), (10), and (11). Given y units in inventory and t periods to go*

(a) *If the two products are dissimilar, then the firm offers a discounted upsell price. In particular, we have $z_1^*(y, t) > z_2^*(y, t)$, and it is optimal for the firm to set $p_t = z_1^*(y, t)$ and $d_t = z_1^*(y, t) - z_2^*(y, t) > 0$.*

(b) *If the two products are similar, then the firm does not offer a discount as part of the upsell offer. In particular, we have $z_1^*(y, t) \leq z^*(y, t) \leq z_2^*(y, t)$, and it is optimal to set $p_t = z^*(y, t)$ and $d_t = 0$.*

Hereafter, we let $p_t^*(y)$ and $d_t^*(y)$ denote the optimal price and discount pair prescribed by Theorem 1, when the firm has y units in inventory with t periods to go. We see from Theorem 1 that when the two products are similar the firm chooses an announced price higher than $z_1^*(y, t)$, the price that maximizes the revenue in the initial stage of period t . By doing so, the firm sacrifices some of its initial-stage revenues but improves its upsell-stage revenues by charging an upsell price in excess of $z_1^*(y, t)$. When the two products are dissimilar, no such trade-off occurs, because the firm would prefer charging an upsell price lower than the announced price.

One important implication of Theorem 1 is that whether the firm will offer a discount has nothing to do with time or inventory, but depends entirely on

the similarity of the products (i.e., correlation between the reservation prices for the two products). However, in cases where the firm chooses to offer a discount, the size of the discount, $d_t^*(y) = z_1^*(y, t) - z_2^*(y, t)$, does depend on t and y .

Next, we state a result that characterizes how $V_t(y)$ and $\Delta(y, t) = V_{t-1}(y) - V_{t-1}(y - 1)$ depend on t and y . This result is important in that it helps us understand how the optimal announced prices and upsell prices depend on time and inventory level.

PROPOSITION 3. *Given $y \geq 1$ units in inventory and $t \geq 1$ periods to go, we have*

- (a) $\Delta(y, t + 1) \geq \Delta(y, t)$,
- (b) $\Delta(y, t) \geq \Delta(y + 1, t)$,
- (c) $V_{t+1}(y) - V_t(y) \geq V_{t+2}(y) - V_{t+1}(y)$.

The first part of the proposition states that the marginal value of one unit of the promotional item is decreasing as we are getting closer to the end of the horizon. The second part implies that the same marginal value is larger when there are fewer items in inventory. Finally, part (c) of the proposition says that, at a fixed y , the revenue-to-go function is convex in t . Given the results of Proposition 3, the following results that characterize the behavior of the optimal announced price and the optimal upsell price in t and y are to be expected.


PROPOSITION 4. *Both the optimal announced price (i.e., $p_t^*(y)$) and the optimal upsell price (i.e., $p_t^*(y) - d_t^*(y)$) are decreasing in inventory (y) and increasing in time-to-go (t).*

The proposition states that both the optimal announced price and the optimal upsell price will be lower if there is more inventory of the promotional item or if there is less time until the end of the horizon. However, the difference between the two, the optimal discount, turns out to be nonmonotonic in time at a fixed inventory level and in inventory level at a fixed time, as shown in Figure 2. As inventory level increases, the discount may increase first, but once the inventory level becomes large enough, the announced price starts dropping so fast that the discount needed to get to the optimal upsell price starts to decrease. In addition, there are examples showing that not only the nominal discount, but also the percentage discount, is not monotone with respect to the inventory level and time.

4.2. Limited Availability of the Regular Product

Our original model introduced in §3 assumes that the regular product is always available when demanded. However, if the regular product is a fast-selling seasonal product, then there is a chance that it might go out of stock, which would in turn affect the pricing decisions for the promotional product. If it is profitable for the firm to make *emergency replenishments* whenever the regular product is demanded but is out of stock, then the promotional product would be offered to the customer regardless of the on-hand inventory level of the regular product. As a result, the limited availability of the regular product would not have any effect on the pricing of the promotional product and all

Figure 2 Discount Is Not Monotonic in Time-to-Go at a Fixed Inventory Level and in Inventory Level at a Fixed Time-to-Go



Notes. In this example, $\delta_{11} = 0, \delta_{22} = 0, q_{R1} = 0.3, \lambda_P = 0.2, \lambda_R = 0.5, r = 65$ and $F_{P1}, F_{P2}, F_{R1}, F_{R2}$ are all Weibull distributions with shape and scale parameters given by, respectively, (2, 90), (2, 50), (2, 100), and (2, 50).

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

of our results for the unlimited availability case would go through. (For more on the use of emergency replenishment in a similar context, see Netessine et al. 2006.) Here, we investigate how the pricing and discounting decisions are affected if the emergency replenishment option is not available to the firm.

Let $V_t(x, y)$ denote the firm's optimal expected revenue from the promotional product when starting period t with x units of the regular product and y units of the promotional product. Then, the optimality equations are given by

$$\begin{aligned}
 V_t(x, y) &= \max_{p_t, d_t; p_t \geq d_t \geq 0} \{ \lambda_R \beta_R [\alpha(p_t - d_t)(p_t - d_t + V_{t-1}(x-1, y-1)) \\
 &\quad + (1 - \alpha(p_t - d_t))(V_{t-1}(x-1, y))] \\
 &\quad + \lambda_P [\beta_P(p_t)(p_t + V_{t-1}(x, y-1)) \\
 &\quad + (1 - \beta_P(p_t))V_{t-1}(x, y)] \\
 &\quad + (1 - \lambda_R \beta_R - \lambda_P)V_{t-1}(x, y) \}, \\
 &\quad x > 0, y > 0, t = 1, \dots, T,
 \end{aligned}$$

$$\begin{aligned}
 V_t(0, y) &= \max_{p_t; p_t \geq 0} \{ \lambda_P [\beta_P(p_t)(p_t + V_{t-1}(0, y-1)) \\
 &\quad + (1 - \beta_P(p_t))V_{t-1}(0, y)] + (1 - \lambda_P)V_{t-1}(0, y) \}, \\
 &\quad y > 0, t = 1, \dots, T,
 \end{aligned}$$

$$V_t(x, 0) = 0, \quad x \geq 0, t = 1, \dots, T, \quad \text{and} \quad V_0(\cdot) = 0. \quad (12)$$

For any given state with $t > 0$, $x > 0$, and $y > 0$, the single-period optimization problem can be written as

$$\begin{aligned}
 \max_{p_t \geq d_t \geq 0} \{ &\lambda_R \beta_R \alpha(p_t - d_t)(p_t - d_t - \Delta(x-1, y, t)) \\
 &+ \lambda_P \beta_P(p_t)(p_t - \Delta(x, y, t)) \}, \quad (13)
 \end{aligned}$$

where

$$\Delta(x, y, t) = V_{t-1}(x, y) - V_{t-1}(x, y-1).$$

Note that $\Delta(x, y, t)$ can be interpreted as the marginal value of one unit of the promotional product in period t when there are x units of the regular product and y units of the promotional product. The following proposition describes how $\Delta(x, y, t)$ changes with x, y , and t . (See Appendix B in the online supplement for the proofs of the results in this section.)

PROPOSITION 5. *Suppose that the regular product and the promotional product are dissimilar. Then, for $x \geq 0$, $y \geq 0$, $t \geq 1$, we have*

- (a) $\Delta(x+1, y+1, t) \geq \Delta(x, y+1, t)$,
- (b) $\Delta(x, y+2, t) \leq \Delta(x, y+1, t)$,
- (c) $\Delta(x, y+1, t+1) \geq \Delta(x, y+1, t)$.

Parts (b) and (c) of Proposition 5 extend Proposition 3 to the case where the regular product has limited availability. In addition, Proposition 5(a) states that the marginal value of one unit of the promotional product increases with the inventory level of the regular product. This result is not surprising because higher inventory levels of the regular product imply more opportunities for upselling (and thus selling the promotional product) until the end of the season. From Proposition 5(a), it follows that compared with the case where the regular product is always available when demanded, in the limited availability case, the marginal value of the promotional product will be smaller. Therefore, it is reasonable to expect that one would be more inclined to offer discounts. Using Proposition 5(a), we can prove the following theorem:

THEOREM 2. *Suppose that there is a finite amount of the regular product at time T and its inventory is never replenished. Furthermore, suppose that the regular and promotional products are dissimilar. Then, the firm offers a discounted upsell price.*

Theorem 2 states that the firm always offers discounts when the regular and promotional products are dissimilar. For such products, the following proposition further describes how the announced price and the upsell price depend on the inventory levels of the regular and promotional products and time until the end of the season.

PROPOSITION 6. *Suppose the regular and promotional products are dissimilar. Then, both the optimal announced price and the optimal upsell price are decreasing in inventory of the promotional product (y), increasing in the inventory of the regular product (x), and increasing in time-to-go (t).*

Theorem 2 and Proposition 6 concern products that are dissimilar. As for similar products, recall from Theorem 1 that when the regular product is always

available, customers are never offered discounts. In contrast, when the regular product has limited availability, the firm may choose to offer discounts. Consider the example where $q_{R1} = 0.7$, $q_{P1} = 0.7$, $\delta_{11} = \delta_{22} = 1$, $\lambda_R = 0.25$, $\lambda_P = 0.25$, $r = 85$, F_{R1} , F_{R2} , F_{P1} , and F_{P2} are all Weibull distributions with shape and scale parameters (2, 100), (2, 50), (3, 190), and (3, 150), respectively. For this example, when there is one unit of promotional product and one unit of regular product in inventory, and there are 12 time periods left, the firm will offer a discount of approximately 6.3 from the announced price.

When the two products are similar, one might expect certain monotonicity structures on discounting decisions to hold. For example, one would expect Proposition 6 to hold even when the two products are similar, and indeed our numerical study supports that claim. On the other hand, however, certain monotonicity properties that might be expected to be true do not hold. For example, in some cases, although discount is being offered at a certain inventory level (of the regular or promotional product), it is not offered at higher and lower levels of inventory, indicating that optimal discount policy does not have a simple monotonic structure with respect to inventory levels. (See Figure 3 for an example.)

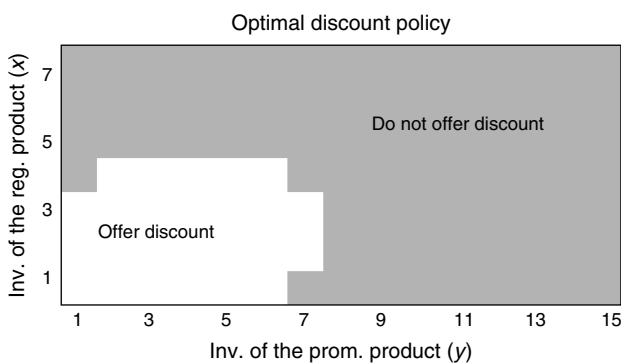
4.3. Multiple Regular Products

In §4.1, we derived all our results for the case where there is one product in addition to the promotional one. Suppose now that there are $n \geq 2$ regular products instead of just one. In this section, we extend the

results of the dynamic price, dynamic discount model to this case. A detailed development of the model along with the additional notation needed, is given in Appendix C in the online supplement. Here, we will provide only the notation needed to summarize the results. With n regular products, in each period, one of $n + 2$ events happens in the initial stage: A potential customer for regular product $k \in \{1, \dots, n\}$ may arrive, or a potential customer for the promotional product may arrive, or no customer arrives at all. In the upsell stage, when making an upsell offer to a customer who bought one of the regular items, the firm can offer different discounts to different customers based on which regular product the customer bought.

Let \hat{q}_{Pik} denote the probability that a customer belongs to segment i of the promotional item (segment 1 being the target segment, and segment 2 the nontarget), given that the customer bought regular product k , $k = 1, \dots, n$, in the initial stage of a period. In the rest of this section, it will be convenient to assume that the regular items are indexed in ascending order of \hat{q}_{P1k} s; i.e., they are indexed so that $\hat{q}_{P11} \leq \hat{q}_{P12} \leq \dots \leq \hat{q}_{P1n}$. Then a customer who bought regular item i in the initial stage is less likely to be in the target segment of the promotional item, compared to customers who bought regular items $i + 1$ or higher. Given this indexing of the regular products, one would expect the firm to be more willing to offer a discount to purchasers of lower-indexed regular items compared to purchasers of higher-indexed regular items. The next theorem states this result, among others. (See Appendix C in the online supplement for a proof.)

Figure 3 Optimal Discount Policy for $t = 20$



Notes. In this example, $\delta_{11} = 0.3$, $\delta_{22} = 1$, $q_{R1} = 0.7$, $\lambda_P = 0.25$, $\lambda_R = 0.5$, $r = 85$ and F_{P1} , F_{P2} , F_{R1} , F_{R2} are all Weibull distributions with shape and scale parameters given by, respectively, (3, 90), (3, 150), (2, 100), and (2, 50).

THEOREM 3. Suppose the regular products are indexed so that $\hat{q}_{P11} \leq \hat{q}_{P12} \leq \dots \leq \hat{q}_{P1n}$. Given y units of the promotional item in stock and t periods to go, let p_t and d_{tk} denote, respectively, the announced price of the promotional item in period t and the discount offered on the promotional item to a customer who purchased regular item k in the initial stage of period t . Then

(a) If regular product $k \in \{1, \dots, n\}$ is dissimilar to the promotional one, i.e., $\delta_{11k} + \delta_{22k} < 1$, then the firm offers a nonzero discount when making upsell offers to customers who buy regular product k ; i.e., it is optimal for the firm to set $d_{tk} > 0$.

(b) Suppose there exists at least one regular product that is dissimilar to the promotional product. Then, there

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

exists $m \in \{1, \dots, n\}$ such that the firm offers a nonzero discount when making upsell offers to customers who buy regular product k for $k \leq m$; i.e., it is optimal to set $d_{ik} > 0$ for $k \leq m$. Furthermore, for any two products $i < j \leq m$, it is optimal to have $d_{ii} \geq d_{ij}$.

The first part of the theorem states that, regardless of time and inventory, purchasers of any regular item that is dissimilar to the promotional item will receive discounted upsell offers. This is similar to the result in the case of a single regular product. In addition, as stated in part (b), all customers who purchased regular items 1 through m will receive discounted upsell offers for some $m \in \{1, \dots, k\}$. What is implicit in the statement of part (b) is that some of these m products may be similar to the promotional product, but customers who purchase these regular products will nevertheless receive discounted upsell offers. This is in contrast to the single regular product case, where a customer never receives a discount if the regular product is similar to the promotional product. Here is the intuition behind this result. Among all the regular products, there may exist some whose purchasers are very likely to be in the target segment of the promotional item, i.e., products for which \hat{q}_{P1k} is particularly large. Purchasers of such regular products are willing to pay relatively high prices for the promotional item. In order to be able to charge high upsell prices to such customers, the firm keeps the announced price high, and offers a discount to purchasers of regular products for which \hat{q}_{P1k} is not particularly large, even though those products may be similar to the promotional one.

5. Static Price and/or Discount

So far we have assumed that the firm can adjust both the announced price and the discounts dynamically. In this section, we consider the firm's revenue-maximization problem when the firm has less flexibility in adjusting the announced price and the discount.⁵

⁵ Here, we provide only an overview of the results for the case of static price and/or discounts. A more detailed version of this section and formal proofs of the results are available from the authors upon request.

5.1. Static Price and Dynamic Discounts

First, consider the case where the firm charges announced price p over the entire horizon, but is free to choose the discount d_t that will accompany an upsell offer in period t with y units of promotional product and x units of regular product in inventory. For this static price/dynamic discount problem, one can show a result analogous to Propositions 3 and 5. The marginal value of one unit of the promotional product increases in the remaining time (t) and inventory level of the regular product (x) and decreases in the inventory level of the promotional product (y). Given such a result on the marginal value, one can establish the following proposition, which states the effects of inventory levels and remaining time on the decision to offer a discount.

PROPOSITION 7. *Suppose the announced price is fixed at some p .*

(a) *At any fixed inventory levels y for the promotional product and x for the regular product, if it is optimal to pick some $d_t > 0$ at any time $t = t_0$, then it is also optimal to pick some $d_t > 0$ for any $t < t_0$.*

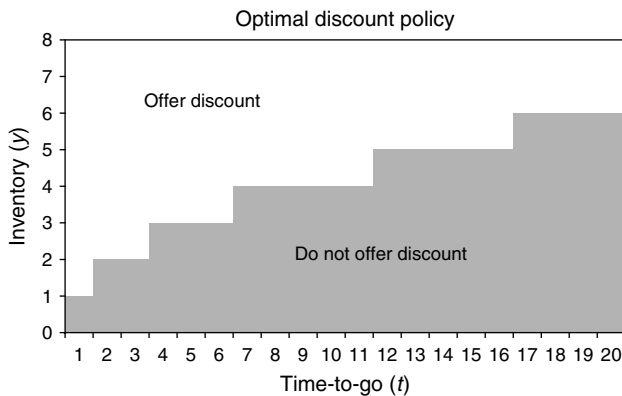
(b) *At any fixed time t and fixed inventory level x for the regular product, if it is optimal to pick some $d_t > 0$ when there are $y = y_0$ units in inventory, then it is also optimal to pick some $d_t > 0$ when there are $y > y_0$ units in inventory.*

(c) *At any fixed time t and fixed inventory level y for the promotional product, if it is optimal to pick some $d_t > 0$ when there are $x = x_0$ units in inventory, then it is also optimal to pick some $d_t > 0$ when there are $x < x_0$ units in inventory.*

According to Proposition 7, when price is static, the discount policy is the switching-curve type, as shown in Figure 4. (For purposes of this figure, we assumed that the regular product is always available when demanded; i.e., x is larger than the number of periods to go, 20.) In addition, we note that even when the regular product is always available and the two products are similar, it is possible that the firm will offer a discount along with the upsell offer, which would never happen in the case in which the announced price is set dynamically.

The pricing problem that arises in this case is akin to that considered by Netessine et al. (2006), who fix the prices of the products, but allow the price of a

Figure 4 When the Announced Price Is Fixed at the Beginning, There Is a Boundary of Inventory Level and Time-to-Go Pairs Above Which It Is Optimal to Offer a Discount



Notes. In this example, $T = 20$, $\delta_{11} = 0$, $\delta_{22} = 0$, $q_{R1} = 0.3$, $\lambda_p = 0.2$, $\lambda_R = 0.5$, $r = 65$ and F_{P1} , F_{P2} , F_{R1} , F_{R2} are all Weibull distributions with shape and scale parameters given by, respectively, (2, 90), (2, 50), (2, 100), and (2, 50).

two-product bundle to be determined optimally every time a consumer is offered a bundle. In fact, we follow their approach in establishing some of our results.

5.2. Static Price and Discount

Here, both the announced price, p , and the discount level, d , are fixed at the beginning of the horizon. In this case, the only dynamic decision available to the firm is whether to offer the discount in a given period with given inventory levels. As in the static price-dynamic discount case, it can be shown that the optimal discount policy is of switching-curve type. In addition, it is easy to show that for a fixed initial price p , at any given time t and inventory levels x and y , if it is optimal to offer a discount in the static discount model, it is also optimal to offer a discount in the dynamic discount model. This is not surprising, because one would be more inclined to offer discounts when there is the flexibility of choosing the discount level depending on the system state.

5.3. Multiple Regular Products

Proposition 7 can be generalized to the case where there are $n > 1$ regular products. In addition, regardless of whether discounts are dynamic or static, one can establish a result that orders the n regular products such that the purchaser of a lower-ranked regular product will be offered a discount only if the purchaser of a higher-ranked regular product is offered a

discount. This ordering depends on \hat{q}_{P1k} , $k = 1, \dots, n$, the probability that a customer who bought regular product k in the initial stage belongs to the target segment of the promotional product. For any two regular products i and j , the fact that $\hat{q}_{P1j} \geq \hat{q}_{P1i}$ indicates that purchasers of product j are willing to pay more than the purchasers of product i for the promotional product. Therefore, for any fixed inventory levels and time, customers who purchase product j are going to get a discount only if the purchasers of product i get a discount.

6. Numerical Study

So far, we have considered three different types of pricing and discounting policies: dynamic pricing and dynamic discounting policy (DPDD), static pricing and dynamic discounting policy (SPDD), and static pricing and static discounting policy (SPSD). In a DPDD policy, all decisions related to the promotional product are dynamic. Prices are set dynamically, the decision of whether to offer a discount is made dynamically, and the level of the discount is determined dynamically depending on the inventory level and time. In an SPDD policy, the advertised price of the product is kept the same throughout the time horizon, whereas discount decisions are still dynamic. Finally, in an SPSD policy, the only dynamic decision is whether to offer a discount to a customer, whereas the advertised price and the discount that will be offered are fixed throughout the season. We note that under the SPSD policy, the static price and discount are chosen optimally at the beginning of the horizon. Likewise, under the SPDD policy the static price is chosen optimally.

We generated 24,300 different problem instances, and for each instance we identified the optimal prices and discounts over a 20-period horizon under each pricing-discounting strategy (DPDD, SPDD, and SPSD). Finding the optimal solution for each problem instance requires solving a dynamic program to optimality, each of which takes a negligible amount of time. Initially, we identified several key model parameters and determined a value set for each parameter. (Parameters were assigned one of the values from this set.) Then, instances were generated by considering all possible combinations within these sets and among the parameters. Table 1 lists these parameters

Table 1 Parameter Value Sets

Parameter	Value set
q_{R1}	{0.3, 0.5, 0.7}
δ_{11}	{0.0, 0.3, 0.5, 0.7, 1.0}
δ_{22}	{0.0, 0.3, 0.5, 0.7, 1.0}
$\lambda_R + \lambda_P$	{0.3, 0.5, 0.7}
$\lambda_R / (\lambda_R + \lambda_P)$	{0.3, 0.5, 0.7}
r	{low, medium, high}
Initial inventory	{low, medium, high}
(W_{R1}, W_{R2})	{(100, 50), (80, 70)}
(W_{P1}, W_{P2})	{(90, 50), (75, 65)}
T	{20}

and their value sets. We assumed Weibull distribution for F_{R1} , F_{R2} , F_{P1} , and F_{P2} . In Table 1, W_{R1} and W_{R2} denote the Weibull scale parameter for F_{R1} and F_{R2} , respectively. (The shape parameters are two for both.) Similarly, W_{P1} and W_{P2} denote the Weibull scale parameter for F_{P1} and F_{P2} . (The shape parameters are three for both.) The “low,” “medium,” and “high” values for r and the initial inventory are not predetermined. For r , low corresponds to the value of the scale parameter of the reservation price distribution for the nontarget segment, high corresponds to the value of the scale parameter of the reservation price distribution for the target segment, whereas medium corresponds to the weighted average of the two values with respect to segment probabilities, q_{R1} and q_{R2} . For the initial inventory level, we first calculate an upper bound on the total demand over the whole horizon by $\bar{I} = (\lambda_R + \lambda_P)T$. The low inventory level corresponds to $(0.2)\bar{I}$, the “medium” level corresponds to $(0.5)\bar{I}$, and the high level corresponds to $(0.8)\bar{I}$.

6.1. Benefits of Using Customer Purchase Information

We first investigate the improvements in profits brought by using the customer purchase information while making upsell offers. Firms may simply choose not to use this information or they may not know how to use it. (If the firm does not know δ_{ij} values, which relate the two products, then the purchase information is useless.) In such a case, there is no information that would distinguish the purchaser of a regular product from a random customer. Hence, in the case of DPDD policy, there is also no need to change the price for the upsell offer because the announced price has already been adjusted according to time and the

inventory level of the promotional product. The optimality equations will be given by

$$V_t(y) = \max_{p_t: p_t \geq 0} \{ \lambda_R \beta_R [\beta_P(p_t)(p_t + V_{t-1}(y - 1))] + \lambda_P \beta_P(p_t)(p_t + V_{t-1}(y - 1)) + (1 - \lambda_R \beta_R \beta_P(p_t) - \lambda_P \beta_P(p_t)) V_{t-1}(y) \},$$

$$y > 0, t = 1, \dots, T,$$

$$V_t(0) = 0, t = 1, \dots, T, \text{ and } V_0(\cdot) = 0.$$

Likewise, one could write the optimality equations for the SPDD and SPSD policies when purchase information is not used. Of course, regardless of the pricing and discounting strategy, the use of purchase information will never worsen the firm’s expected profit. Here, we investigate whether the improvement due to purchase information differs significantly across pricing and discounting strategies. Table 2 gives percentage improvements in profits under DPDD, SPDD, and SPSD policies.

Numbers suggest that, overall, improvements under the SPDD policy are significantly larger than the improvements under the DPDD and SPSD policies. Under any pricing/discounting policy, the use of purchase information improves the firm’s revenues by helping the firm make better discounting decisions. By using the DPDD policy, the firm is already capturing large revenues due to the flexibility of

Table 2 Improvements in Profits when Using Customer Purchase Information

Percent of improvement	DPDD	SPSD	SPDD
[9, 10)	0	0	1
[8, 9)	0	0	1
[7, 8)	0	0	2
[6, 7)	0	0	4
[5, 6)	0	0	5
[4, 5)	10	6	23
[3, 4)	5	7	71
[2, 3)	41	40	314
[1, 2)	154	144	2,401
[0, 1)	24,090	24,103	21,479
Average	0.047	0.049	0.404
Min	0.000	0.000	0.000
Max	4.556	4.506	9.639

Notes. Numbers represent the number of instances, DPDD: dynamic price, dynamic discount, SPSP: static price, static discount, SPDD: static price, dynamic discount.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

adjusting the price, hence the use of purchase information leads to only modest improvements. Likewise, when the firm is using the SPSD policy, the improvement due to the use of purchase information is small because the firm's flexibility is limited to deciding whether to offer the discount, which limits the firm's ability to exploit purchase information to make better discounting decisions. However, when the firm is using the SPDD policy, setting the right discount level is more important because dynamic discounts can partially capture revenue management benefits that are not realized due to the static price. Hence, the use of purchase information leads to the highest improvements under this policy. This observation suggests that firms who do not have the flexibility to change prices dynamically but can adjust the discount levels (e.g., catalog retailers) will find using the customer purchase information more valuable.

6.2. Benefits of Dynamic Pricing and Discounting Decisions

Despite the clear advantage of dynamic policies, they are not always preferred, for various reasons. Some firms simply do not want to upset their customers by changing prices frequently, whereas some others do not have the technical capabilities required for the implementation of such policies. Therefore, it is of interest to understand the potential benefits that would be gained through dynamic pricing and discounting decisions. Here, we report our findings based on our numerical study. The pricing and discounting policies we consider can be ordered as DPDD, SPDD, and SPSD from more sophisticated to less sophisticated. However, some firms may not even have the means to implement SPSD policies. They may determine a discount level for the product (as in an SPSD policy) and offer this discount to all the purchasers of the regular products (unlike in an SPSD policy, where the firm decides optimally whether to offer the discount.) We call this a *Full Static* (FS) policy. In an FS policy, there is no dynamic decision at all, but as in the SPSD policy, the static price and the discount level are determined optimally at the beginning of the horizon.

For the 24,300 different instances generated, we computed the profits under each policy. We then, computed the percentage improvements that would

Table 3 Percentage Improvements in Profits Obtained by Switching from Less Sophisticated Policies to More Sophisticated Policies

	FS vs. SPSD	FS vs. SPDD	FS vs. DPDD	SPSD vs. SPDD	SPSD vs. DPDD	SPDD vs. DPDD
Max	3.0447	3.8301	7.1513	1.0830	6.5409	6.5038
Min	0.0000	0.0000	0.0001	0.0000	0.0001	0.0001
Average	0.2905	0.3489	1.9607	0.0579	1.6624	1.6030

be obtained by switching from less sophisticated policies to more sophisticated ones and determined the average, maximum, and minimum improvements over the 24,300 instances. The results are given in Table 3.

Table 3 clearly demonstrates the benefits of using dynamic policies—in particular, the fully dynamic policy DPDD over the fully static FS with an average improvement close to 2%. From the table, one can make the following observations.

6.2.1. Dynamic Pricing vs. Dynamic Discounting.

The average revenue improvement that would be obtained by switching from SPDD to DPDD is around 1.6%, whereas the average improvement that would be obtained by switching from FS to SPDD is around 0.35%. Therefore, although dynamic discounting decisions bring modest improvements, setting the price dynamically seems to have a much more significant effect.

6.2.2. Dynamically Changing Discount Levels vs. Dynamically Deciding to Offer a Fixed Discount.

Consider a firm currently using the FS policy. Furthermore, suppose this firm is committed to using a static announced price; for example, the firm may be a catalog retailer. There are two improvements such a firm can make: switching to SPSD policy or the SPDD policy. We observe from the table that switching from the FS policy to SPDD leads to an average revenue improvement of 0.35%, whereas switching from the FS policy to SPSD leads to an average revenue improvement of 0.29%. Therefore, most of the revenue improvement a firm can realize by switching to SPDD can be realized by switching to SPSD. Hence, our numerical analysis suggests that the firms would reasonably benefit from having the capability of deciding when it is best to offer a discount, but given that capability, having the additional flexibility

of changing the discount level is beneficial to a lesser degree.

7. Conclusion

In this paper, we investigated the interactions among upselling, the use of customer purchase information, and dynamic pricing. We found that if the price of the promotional product is dynamically adjusted and the regular product is always available when demanded, all customers who buy a regular product that is dissimilar to the promotional one (i.e., a product whose reservation price is negatively correlated with that of the promotional product) will be offered a discount along with the upsell offer, regardless of the inventory of the promotional product and the time until the end of the selling season. However, if the firm uses a fixed price for the promotional product and/or the regular product's availability is limited, then whether a customer will get a discounted upsell offer does depend on the inventory levels and the time until the end of the season.

Our numerical results indicate that the benefit of using customer purchase information is low when the firm can dynamically adjust the price and the discount. Likewise, when the firm needs to use a static price and a static discount throughout, the benefit to the firm from using customer purchase information is low. On the other hand, if the firm needs to use a static price but can adjust the discount levels dynamically, the use of purchase information is highly beneficial. Our numerical analysis also indicates that even though dynamic discounting decisions bring modest improvements over static decisions, setting prices dynamically seems to have a more significant effect on profits.

Given the prevalence of upselling among catalog and online retailers and the wealth of customer data available to such retailers, data-driven methods that help firms make better upselling decisions are likely to have an impact in practice. In this study, we focused mainly on obtaining insights into the issue of upselling. A further line of research would be to develop models with more relaxed assumptions (e.g., allowing multiple promotional products, and pricing flexibility for regular products) and to develop and evaluate heuristic methods amenable for use in practice. For such application-oriented research it is

important to have a model in which the parameters can be reliably estimated. Our consumer model may prove to be advantageous in this regard, because it does not require estimation of joint distributions for reservation prices of multiple products. Instead, the correlation across products is captured through cross-segment probabilities (e.g., the probability that a customer who belongs to the target segment of a product belongs to the target segment of the other product as well). Such cross-segment probabilities can be estimated while conducting the initial market segmentation. For example, conjoint analysis is widely used for market segmentation, and there is a large body of research on this. (See, for example, Green and Krieger 1991 for an overview of using conjoint analysis for segmentation, and Vriens et al. 1996 for a comparison of several approaches to segmentation through conjoint analysis.) For our model, one can use conjoint analysis to estimate the reservation prices of target and nontarget segments, and cross-segment probabilities can be estimated in the process as well.

Acknowledgments

The authors thank the associate editor and three anonymous referees for their comments, which helped to improve the paper.

References

- Adams, W. J., J. L. Yellen. 1976. Commodity bundling and the burden of monopoly. *Quart. J. Econom.* **90** 475–498.
- Andrews, K. 1999. Extending upselling and cross-selling efforts. *Target Marketing* **22**(November) 36–40, 94.
- Aviv, Y., A. Pazgal. 2005a. Optimal pricing of seasonal products in the presence of forward-looking consumers. Working paper, Olin School of Business, Washington University, St. Louis, MO.
- Aviv, Y., A. Pazgal. 2005b. A partially observed Markov decision process for dynamic pricing. *Management Sci.* **51** 1400–1416.
- Bitran, G. R., R. Caldentey. 2003. Commissioned paper: An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5** 203–229.
- Bitran, G. R., S. V. Mondschein. 1993. Pricing perishable products: An application to the retail industry. Working paper, Sloan School of Management, MIT, Boston, MA.
- Bitran, G. R., S. V. Mondschein. 1997. Periodic pricing of seasonal products in retail. *Management Sci.* **43** 64–79.
- Del Franco, M., P. Girard, L. Santo. 1999. Operations benchmark '99. *Catalog Age* **16**(May) 77, 78, 80, 82, 105.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Management Sci.* **49** 1287–1309.

- Elmaghraby, W., A. Gülcü, P. Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *Manufacturing Service Oper. Management*. **10**(1) 126–148.
- Ernst, R., P. Kouvelis. 1999. The effect of selling packaged goods on inventory decisions. *Management Sci.* **45** 1142–1155.
- Ferriolo, S. 2003. When less isn't more. *Catalog Age* **20**(January) 38–39.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40** 999–1020.
- Green, P. E., A. M. Krieger. 1991. Segmenting markets with conjoint analysis. *J. Marketing* **55** 20–31.
- Gurler, U., Z. Bulut, A. Sen. 2005. Bundle pricing of inventories with stochastic demand. Extended abstract, *2005 M&SOM Conf.*, Northwestern University, Evanston, IL.
- Hanson, W., K. Martin. 1990. Optimal bundle pricing. *Management Sci.* **36** 155–174.
- Lariviere, M. A., E. L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing Service Oper. Management* **3** 293–305.
- Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multi-product revenue management problems. *Manufacturing Service Oper. Management* **8** 136–148.
- Miller, P. 1995. Fine-tuning upsell technique. *Catalog Age* **12**(July) 61.
- Monahan, G. E., N. C. Petruzzi, W. Zhao. 2004. The dynamic pricing problem from a newsvendor's perspective. *Manufacturing Service Oper. Management* **6** 73–91.
- Muller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. Wiley, New York.
- Netessine, S., S. Savin, W. Xiao. 2006. Revenue management through dynamic cross-selling in e-commerce retailing. *Oper. Res.* **54** 893–915.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Stremersch, S., G. J. Tellis. 2002. Strategic bundling of products and prices: A new synthesis for marketing. *J. Marketing* **66** 55–72.
- Vriens, M., M. Wedel, T. Wilms. 1996. Metric conjoint segmentation methods: A Monte Carlo comparison. *J. Marketing Res.* **33** 73–85.
- Yadav, M. S., K. B. Monroe. 1993. How buyers perceive savings in a bundle price: An examination of a bundle's transaction value. *J. Marketing Res.* **30** 350–358.
- Zhang, D., W. L. Cooper. 2005. Revenue management for parallel flights with customer-choice behavior. *Oper. Res.* **53** 415–431.
- Ziya, S., H. Ayhan, R. D. Foley. 2004. Relationship among three assumptions in revenue management. *Oper. Res.* **52** 804–809.