OPERATIONS RESEARCH
Vol. 57, No. 6, November–December 2009, pp. 1523–1531
ISSN 0030-364X | EISSN 1526-5463 | 09 | 5706 | 1523

informs®

DOI 10.1287/opre.1090.0701

TECHNICAL NOTE

Personalized Dynamic Pricing of Limited Inventories

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,
Chapel Hill, North Carolina 27559, ziya@unc.edu

Prior work has investigated time- and inventory-level-dependent pricing of limited inventories with finite selling horizons. We consider a third dimension—in addition to time and inventory level—that the firms can use in setting their prices: the information that the firm has at the individual customer level. An arriving customer provides a *signal* to the firm, which is an imperfect indicator of the customer's willingness to pay, and the firm makes a *personalized* price offer depending on the signal, inventory level, and time. We consider two different models: *full personalization* and *partial personalization*. In the full personalization model, the firm charges any price it wishes given the customer signal, while in the partial personalization model, the firm can charge one of two prices. We find that a mere correlation between the signals and customers' willingness to pay is not sufficient to ensure intuitive relationships between the signal and the optimal prices. We determine a stronger condition, which leads to several structural properties, including the monotonicity of the optimal price with respect to the signal in the full personalization model. For the partial personalization model, we show that the optimal pricing policy is of threshold-type and that the threshold is monotonic in the inventory level and time.

Subject classifications: inventory/production: policies; marketing: dynamic pricing, personalized pricing.
Area of review: Revenue Management.
History: Received October 2007; revisions received April 2008, September 2008; accepted November 2008. Published online in *Articles in Advance* August 17, 2009.

1. Introduction

Firms that sell a limited inventory of a perishable or seasonal product (e.g., airlines, apparel retailers) have long been used to adjusting their prices over time based on inventory levels, a practice commonly referred to as *dynamic pricing*. In addition to adjustments based on time and inventory levels, a firm can also tailor the price it charges to its customers based on information available at the individual customer level. This practice is usually referred to as *personalized pricing*. In this paper, our goal is to explore the interactions between personalized pricing and dynamic pricing. In particular, we analyze how the availability of a product interacts with customer information to determine the price offered to a customer.

Many sellers have the ability to identify individual customers, collect information about them, and track their purchasing behavior (e.g., catalog retailers, online retailers, and stores that issue loyalty cards to their customers). Such sellers can and do use personalized pricing. In 1996, Denise Katzman found that a Victoria's Secret catalog sent to a male colleague offered a deeper discount than the nearly identical catalog she received and sued the company (Weiss and Mehrotra 2001). Victoria's Secret was not alone in its practice of charging different prices to different customers. In fact, Blank (2001) report that catalogers engage in geographical price discrimination, that is, they use different prices in catalogs mailed to different zip codes. Online retailers are also known to engage in personalized pricing. For example, an online retailer charged different prices for the same digital camera, depending on whether customers had previously visited a price-comparison site (Bridis 2005). According to a Forrester report, out of 30 online retailers interviewed in 2000, 57% had plans to try some form of personalized pricing (Johnson et al. 2000). Likewise, traditional grocery stores make personalized discount offers: More than 300 Jewel-Osco and Albertsons stores have installed kiosks where customers insert their membership cards to be presented with customized discount offers (Desjardins 2007).

A natural question is whether implementing such discriminatory pricing policies is legal. With some exceptions, it is widely considered to be so. It is telling that consumers sued Victoria's Secret for mail fraud as opposed to trying to make the case that the company's practice violated the

Robinson-Patman Act.[1] According to Ramasastry (2005), "The reality is that Internet price customization does exist—and, contrary to popular opinion, is typically legal."

It appears that the real challenge for the firms is to manage customers' perception of the practice because it can cause customer ill-will when not framed appropriately. For example, in the summer of 2000, Amazon was caught selling an X-Files DVD box at prices ranging from $80 to $100. There was an uproar from customers who felt that Amazon was tracking their purchase history to charge higher prices to more loyal customers, while Amazon denied the claim and said the prices were chosen at random as part of price testing (Adamy 2000). Furthermore, a survey by Turow et al. (2005) finds that 90% of customers disagree with the statement that "It's OK if a store charges me a price based on what it knows about me." Nonetheless, there are many cases in which customers do not appear to be bothered by personalized pricing. In fact, it appears that customers have little problem with different people paying different prices for the same product, as long as the pricing scheme is perceived to be fair. For example, hardly anyone objects to student or senior discounts, which are ultimately a form of personalized pricing based on demographic information. Two different customers buying two identical vehicles are likely to pay different prices, depending on their propensity to negotiate. Airlines charge higher prices to travelers who are not willing to stay the Saturday night, presumably because these are business travelers who are less price-sensitive.

The use of personalized pricing requires that the seller acquire customer-specific information. In this paper, we use a model where each customer furnishes the seller with a signal, which captures the information that the seller has about the individual. This signal might take many forms and there might be many ways in which the seller can collect such signals. For example, the signal might take one of three possible values indicating whether the customer carrying the signal is a student, a senior citizen, or neither, and this signal can be obtained by simply asking students and senior citizens to identify themselves. In this case, the firm might be willing to bank on the fact that students and senior citizens typically have lower willingness to pay and offer them a discount. In many instances, this might be the best action for the firm, but offering a discount to a student (or senior citizen) is not always the best possible action because surely there are students (or senior citizens) who have much deeper pockets and higher willingness to pay than an average customer. That is, as useful as the signal might be to the firm, it is not necessarily a perfect indicator of the customers' willingness to pay. Firms can actually develop more sophisticated and more useful signals. For example, loyalty programs give firms not only demographic information about a customer, but also possibly a way to track the purchases of the customer over time. In cases where loyalty programs are used, customer signal could be the amount the customer spent with the seller in the last year or the time elapsed since the customer last made a purchase from the seller. Such signals are revealed to the seller every time the customer identifies herself, for example, by swiping a loyalty card. Similarly, online retailers have access to a wealth of information about registered customers. Another example of a signal is the zip code, which is easily available to and frequently used by catalog retailers. What is common in all these examples is that even though the signal from a customer is a useful bit of information, this information is not a perfect indicator of the customer's willingness to pay. Hence, in this paper, we develop a model where customer signals provide only limited information about the customer's reservation price.

If a seller were to use personalization to its full extent, then each unique signal would prompt a unique price from the seller. We consider such a model with *full personalization*. If personalization cannot be implemented to that extent, a seller may want to use a personalization strategy where there is an announced price and a single discount level that can possibly be offered to a customer, depending on the signal from the customer. To analyze such *partial personalization* of prices, we consider a model in which a range of signals are bunched together and offered a single price.

This paper characterizes the properties that a signal should have for it to be an intuitive determinant of the firm's prices. For example, if there is a positive correlation between the signal and customers' reservation prices, does that mean that the firm should charge higher prices to customers with higher signals? As we see in §4, the answer is "no." Positive correlation is not sufficient. A stronger condition is needed. In §4, we give a precise description of this condition, which leads to useful relationships between the signal and optimal prices. For example, when this "strong" correlation condition is satisfied, under full personalization, the optimal price is higher for customers with higher signals. Under partial personalization, the optimal pricing policy is of threshold type, meaning that customers above the threshold pay the higher price while those below the threshold pay the lower price. The stronger correlation condition also leads to several monotonic properties for the optimal policies. For example, the threshold level changes monotonically with respect to the inventory level and time when prices are kept at fixed levels. These are all properties that would make implementation of price personalization easier in practice. However, none of these structural properties exists in the absence of the strong correlation condition.

The presence of personalization may give customers an incentive to act strategically. For example, a customer might wish to hide her identity if she feels that the seller will charge a lower price to an anonymous customer. A customer might even provide false information to a seller if she believes that such information will qualify for a lower price. In this paper, we do not allow such strategic behavior on the part of customers. However, in the online supplement, we consider an extension where some customers in the population do not provide signals.

The rest of this paper is organized as follows. In §2, we review the relevant literature, and in §3, we describe our base model. Sections 4 and 5 present our results for the full personalization and partial personalization models, respectively. Section 6 gives our concluding remarks. All the proofs are given in the online supplement. An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

2. Literature Review

Following the pioneering work of Gallego and van Ryzin (1994) and Bitran and Mondschein (1997), there has been a growing interest in dynamic pricing of perishable products based on inventory levels and time. Both of these papers are concerned with the pricing of a single product over a finite horizon with no replenishment opportunities, but they use different formulations (one in continuous-time setting, the other in discrete-time setting) and come up with different insights. There has been a significant volume of research on dynamic pricing since then. For comprehensive reviews of this research pre-2003, we refer the reader to Elmaghraby and Keskinocak (2003), Bitran and Caldentey (2003), and Talluri and van Ryzin (2004). Among the more recent work, Maglaras and Meissner (2006) and Zhang and Cooper (2009) consider dynamic pricing of multiple products. Aviv and Pazgal (2005) consider a model where there is a high level of uncertainty about the demand (as a function of the price) but the firm *learns* more about the demand throughout the sales horizon by observing customer reaction to the prices. Aviv and Pazgal (2008), Liu and van Ryzin (2005), Elmaghraby et al. (2008), and Su (2007) investigate dynamic pricing decisions when customers act strategically. Monahan et al. (2004) deal with the decision of determining the initial inventory for products that are under dynamic pricing. Popescu and Wu (2004) and Ahn et al. (2007) consider models where price in a given period influences the demand in other periods.

Our work differs from the above mainly in that in our models, the firm has the flexibility of adjusting the price not only depending on the inventory level and time but also depending on the information that the firm has at the individual customer level. Some of the recent work investigated very specific forms of such use of personalized pricing. For example, Kuo et al. (2009) investigate negotiation, while Netessine et al. (2006) and Aydin and Ziya (2008) investigate upselling/cross-selling, in which case the price offered to a customer depends on what the customer already bought from the seller. All these papers provide insights that are relevant within the specific forms of personalized pricing they consider. Furthermore, from more of a technical point of view, as a consequence of this focused interest in specialized forms of personalized pricing, they consider signaling formulations that are fairly restrictive. For example, in Kuo et al. (2009) the signal is the offer made by a bargainer. In Aydin and Ziya (2008), the signal is whether the customer bought another product at an advertised price. In this paper, we are not restricting ourselves to any particular form of personalized pricing. We use a fairly general signaling formulation so as to make our results relevant to a large class of personalized pricing practices. Perhaps more importantly, unlike the other three papers, this paper is primarily focused on developing a better understanding about the signal itself. For example, we provide insights on the properties the signal should have in order for it to be a more intuitive determinant of customers' reservation price and thus be more useful in practice.

Personalized pricing is closely related to price discrimination, which has been studied extensively in economics and marketing literature. (This body of work is fundamentally different from ours in that inventory considerations typically play no role in pricing decisions made by the firms.) For a survey of this literature and detailed bibliography, see Varian (1989). More recent work has concentrated on the profitability of personalized pricing. Some researchers found that in a competition environment personalized pricing may actually hurt the firms because it intensifies competition. For example, see Thisse and Vives (1988), Shaffer and Zhang (1995), Fudenberg and Tirole (2000), and Chen and Iyer (2002). Villas-Boas (2004) and Liu and Zhang (2006) show that the practice may not be profitable even for a monopolist (due to strategic behaviors of the customers in the case of Villas-Boas 2004, and misalignment of incentives within the supply chain in the case of Liu and Zhang 2006). On the other hand, Chen et al. (2001) find that when *customer targetability* is sufficiently low, personalized pricing is profitable for two competing firms, but for high targetability levels, firms might be worse off because a "prisoner's dilemma" occurs.[2] Shaffer and Zhang (2002) show that when two competing firms are asymmetric, one of the two competing firms can profit from personalized pricing. Based on these conflicting findings it appears that, perhaps not surprisingly, whether or not personalized pricing is profitable is context specific. For supermarkets, even when there is competition, Kumar and Rao (2006) find that using past purchase information for personalized pricing is profitable. Using data from the ketchup market, Besanko et al. (2003) also find the practice to be a profitable strategy. For an extensive survey of marketing literature on behavior-based price personalization, see Fudenberg and Villas-Boas (2006). For a broader survey of research on personalization in marketing, see Murthi and Sarkar (2003).

3. Model Description

Consider a firm with I units of inventory that will be salvaged at the end of a finite selling season. Following the standard approach in dynamic pricing models, we assume that the firm does not have an option to replenish inventory during the selling season. This is certainly true for an airline that has a fixed number of seats on a flight or an apparel retailer that faces long procurement lead times coupled with short life cycles for its seasonal products. Following the approach first used by Bitran and Mondschein (1997), we assume that the selling season is divided into T periods, where each period is short enough that at most one customer arrives in a period. Let λ denote the probability that a customer arrives in a given period.

As we will discuss shortly, our model uses two different forms of stochastic ordering: *failure rate ordering* and *likelihood ratio ordering*. For two cumulative distribution functions (cdfs) Φ_1 and Φ_2 (with corresponding probability density functions (pdfs) ϕ_1 and ϕ_2), if the failure rate of Φ_1 is less than that of Φ_2, i.e., if $\phi_1(x)/(1-\Phi_1(x)) \leqslant \phi_2(x)/(1-\Phi_2(x))\ \forall x$, then we say Φ_1 dominates Φ_2 in failure rate ordering, and we write $\Phi_1 \geqslant_{fr} \Phi_2$. Furthermore, if $\phi_1(x)/\phi_2(x) \geqslant \phi_1(y)/\phi_2(y)$ for any $x > y$, then we say Φ_1 dominates Φ_2 in likelihood ratio ordering, and we write $\Phi_1 \geqslant_{lr} \Phi_2$. If Φ_1 and Φ_2 are cdfs for discrete random variables, the pdfs in the definition of likelihood ratio ordering are replaced by corresponding probability mass functions (pmfs). See Online Appendix B for more detailed definitions of the forms of stochastic ordering used in this paper. We refer the reader to Müller and Stoyan (2002) or Shaked and Shanthikumar (2007) for more on stochastic orderings, but here it is useful to mention that likelihood ratio ordering is a stronger condition implying failure rate ordering, which implies first-order stochastic dominance, which in turn implies ordering of the means. Note that throughout the paper we use increasing/decreasing and positive/negative in the weak sense unless specifically qualified as strictly increasing/decreasing or nonpositive/negative.

Suppose that the consumer population is divided into two segments, one with higher willingness to pay than the other. Let q_i, $i = 1, 2$ denote the fraction of segment i customers. We assume that all customers in a given segment have independent and identically distributed (iid) reservation prices. Let F_i denote the cdf of the reservation prices of customers in segment i, and f_i their pdf. Let $\bar{F}_i(\cdot) := 1 - F_i(\cdot)$. We make the following assumptions on the reservation price distributions.

Assumption (A1). $F_i(\cdot)$, $i = 1, 2$, *are twice-continuously-differentiable, strictly increasing functions, and they both have the same nonnegative support.*

Assumption (A2). $F_i(\cdot)$, $i = 1, 2$, *have strictly increasing generalized failure rates, i.e.,* $xf_i(x)/\bar{F}_i(x)$ *is strictly increasing.*

Assumption (A3). $F_1(\cdot) \geqslant_{fr} F_2(\cdot)$.

Assumption (A2) is satisfied by a large family of distributions, including all Weibull distributions and the positive part of the normal distribution. (For a comparison of various assumptions on reservation prices used in revenue management problems, see Ziya et al. 2004.) The ordering stated in (A3) implies that the absolute price elasticity of demand is smaller for customers in segment 1; i.e., segment 1 is less price-sensitive. A further implication of (A3) is that the reservation price of a customer in segment 1 stochastically dominates that of a customer in segment 2.

We assume that the customer's segment is not directly observable by the seller. In other words, the seller cannot say with certainty if an individual is coming from the population with high or low price elasticity of demand. However, each arriving customer provides the seller with a signal, which embodies the information available about the customer. (Argon and Ziya 2009 use a similar formulation within a queueing context, where customers' signals determine their priority levels.)

Depending on the information a seller collects and uses, the signal could be demographic information about the customer (e.g., age, gender, zip code) or information regarding transaction history (e.g., the last time the customer made a purchase from the seller, the amount the customer spent with the seller in the last year). This signal, while not enough in itself to determine the segment of the customer, might still be valuable for the seller in updating its belief about the customer's segment. The signals from customers in segment i will be distributed over a range because a segment consists of similar yet heterogenous customers. We assume that the signals of the customers in segment i are iid random variables, denoted by the generic random variable S_i. We allow S_i to be either discrete or continuous, but not a mixture of the two. Let G_i, $i = 1, 2$ denote the cdf of S_i, and g_i denote the pdf of S_i if S_i is continuous and the pmf if S_i is discrete. We impose the following technical assumption on the signal distributions.

Assumption (A4). $G_i(\cdot)$, $i = 1, 2$, *are strictly increasing functions on the support set, and they both have the same nonnegative support.*

In addition, we assume the following ordering between the signals from the two segments:

Assumption (A5). $G_1(\cdot) \geqslant_{lr} G_2(\cdot)$.

Assumption (A5) imposes a strong stochastic order between the signals from the two segments. (Likelihood ratio ordering implies failure rate ordering, which in turn implies first-order stochastic dominance.) In the next section, we investigate the signal's effect on the optimal price and illustrate the significance of Assumption (A5).

4. Dynamic Pricing with Personalization

We first consider a seller who does not have to commit to a price prior to the arrival of the customer. Upon arrival, the customer furnishes the seller with a signal. The seller then quotes a price to the customer based on the signal. One way to implement such a practice is to ask customers to identify themselves with their loyalty cards and offer them customized discounts, as done by Jewel-Osco and Albertsons stores. Such a practice is also technologically feasible for many online retailers because the retailer can present different users with different prices after observing, for example, the IP address of the user, or customers might simply choose to log in and identify themselves. Although this practice might be feasible, it might not be easily implementable, as evidenced by the customer complaints that Amazon received during its random price testing.

Let S denote the signal furnished by a customer, a random variable prior to the arrival of the customer. Suppose that a customer arrives with signal $S = x$. After observing this signal, the seller can update its belief about the segment this particular customer belongs to. Let $\hat{q}_i(x)$ denote the probability that a customer with signal x belongs to segment i. Using Bayes' rule, $\hat{q}_i(x)$ is given by

$$\hat{q}_i(x) = \frac{q_i g_i(x)}{q_1 g_1(x) + q_2 g_2(x)}, \quad i = 1, 2. \tag{1}$$

Suppose that the seller has y units of inventory with t periods to go until the end of the season. Let $V_t(y)$ denote the seller's optimal expected revenue to go. The following optimality equations characterize the dynamic program to be solved by the seller:

$$V_t(y) = E_S\Big[\max_p\big\{\lambda\big(\hat{q}_1(S)\bar{F}_1(p) + \hat{q}_2(S)\bar{F}_2(p)\big)\big(p + V_{t-1}(y-1)\big) + \big[1 - \lambda(\hat{q}_1(S)\bar{F}_1(p) + \hat{q}_2(S)\bar{F}_2(p))\big]V_{t-1}(y)\big\}\Big], \quad y > 0,\ t = 1, \ldots, T,$$

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

In the remainder of this section, let

$$\alpha(x, p) := \hat{q}_1(x)\bar{F}_1(p) + \hat{q}_2(x)\bar{F}_2(p). \tag{2}$$

Hence, $\alpha(x, p)$ is the probability that a customer with a signal x is willing to purchase the product at price p. After some algebraic manipulation, one can rewrite $V_t(y)$ as follows:

$$V_t(y) = V_{t-1}(y) + \lambda E_S\Big[\max_p\{\alpha(S, p)(p - \Delta_t(y))\}\Big], \quad y > 0,\ t = 1, \ldots, T,$$

where

$$\Delta_t(y) = V_{t-1}(y) - V_{t-1}(y-1), \quad y > 0,\ t = 1, \ldots, T. \tag{3}$$

Here, $\Delta_t(y)$ can be interpreted as the marginal value of inventory. Let $p^*(x, y, t)$ denote the optimal price quoted to a customer who furnishes signal x when the seller has y units in inventory with t periods to go until the end of the season. (In case there are multiple optimal prices, we define $p^*(x, y, t)$ to be the smallest optimizer.)

We next discuss how the optimal price $p^*(x, y, t)$ depends on the signal x. Suppose that a customer provides signal x and consider $\hat{q}_1(x)$, the probability that this customer belongs to segment 1. This probability, given by (1), can be rewritten as

$$\hat{q}_1(x) = \frac{q_1 g_1(x)/g_2(x)}{q_1 g_1(x)/g_2(x) + q_2}. \tag{4}$$

Notice from (4) that $\hat{q}_1(x)$ is increasing in the signal x due to our assumption that the signal from segment 1 dominates the signal from segment 2 in likelihood ratio ordering. In other words, the larger the signal from a customer, the more likely the customer is to belong to segment 1, which is the segment with higher willingness to pay. Hence, one would expect that the optimal price $p^*(x, y, t)$ increases with the signal x, which is formalized in the following theorem.

Theorem 1. *Given time t and inventory level y, the (smallest) optimal price, $p^*(x, y, t)$, is increasing in the signal x provided by the customer.*

Theorem 1 is not surprising, given that a higher signal implies a higher probability that the customer belongs to the segment with higher willingness to pay. What is perhaps surprising is that Theorem 1 does not hold under weaker versions of Assumption (A5). Consider the following example.

Example. Suppose that signals are discrete random variables with probability mass functions shown in Table 1. In this example, the signal from the first segment does not dominate the signal from the second segment in likelihood ratio ordering, but in failure rate ordering (which is weaker).[3] Due to the failure rate ordering between the two signals, the mean signal from segment 1 is larger than the mean signal from segment 2. It is not difficult to check that, as a consequence of this ordering between the means, the signal and the reservation price of an individual are positively correlated. Therefore, intuition may suggest that the larger the signal, the higher the price must be. This, however, is not true. For example, here $p^*(1, 1, 1) \cong 38.56$, $p^*(2, 1, 1) \cong 44.16$, $p^*(3, 1, 1) \cong 41.46$, and $p^*(4, 1, 1) \cong 46.66$. Hence, the optimal price for a customer with a signal of 3 is smaller than the optimal price for a customer with a signal of 2 when the inventory level is 1 and the time remaining is 1 period. In fact, optimal prices satisfy the same relationship for many different pairs of inventory

Table 1. The probability mass functions $g_1(\cdot)$ and $g_2(\cdot)$ for the signals from segments 1 and 2, respectively.

x	1	2	3	4
$g_1(x)$	0.1	0.3	0.2	0.4
$g_2(x)$	0.25	0.25	0.25	0.25

Notes. In addition, suppose that $q_1 = 0.3$, $q_2 = 0.7$, $\lambda = 0.5$, F_1 is Weibull with shape and scale parameters, 2 and 100, respectively, and F_2 is Weibull with shape and scale parameters, 2 and 50, respectively, so that $F_1 \geqslant_{fr} F_2$.

level and time. This happens because a customer with a signal of 2 is more likely to belong to segment 1 than a customer with a signal of 3.

The above example shows that positive correlation between signals and reservation prices does not imply that customers with higher signals should be charged higher prices. In fact, the example says more: Even failure rate ordering between the two signals might not be enough to guarantee an intuitive relationship between the signal and the optimal price. For the optimal price to be increasing in the signal, we need the likelihood ratio ordering imposed in Assumption (A5).

Remark 1. Prior work in dynamic pricing literature has established the monotonicity of the optimal prices with respect to inventory level and time under various settings. One can show that similar monotonicity results hold in our problem as well: The optimal price $p^*(x, y, t)$ is increasing in the remaining time t and decreasing in the inventory level y.

In many cases, it is unlikely that a seller will delay the pricing decision until after the revelation of the signal because consumers are likely to find the resulting price discrimination unacceptable. However, for many retailers, be they online or traditional, it is quite possible to offer personalized discounts from an announced price. In effect, such a discounting strategy is similar to a retailer providing customers with special discount offers. While such a practice boils down to charging different prices to different customers, it somehow appears to be more palatable to consumers, either because the discrimination is less transparent or because the consumers find it "fair" that different individuals qualify for different discounts. Therefore, one way to implement the personalized pricing strategy discussed in this section is to announce a price at the beginning of each period and then provide each arriving customer with a personalized discount after seeing the customer's signal. Provided that the announced price is high enough, such a discounting strategy, which we refer to as *personalized discounting*, will achieve the same results as not committing to a price until after seeing the customer's signal. As we discuss next, there is a natural ceiling on how high the announced price needs to be, which makes personalized discounting all the more attractive.

Define $p_i^*(\Delta) := \arg\max_p (p-\Delta)\bar{F}_i(p),\ i = 1, 2$. Note that $p_i^*(\Delta)$ is the price that a seller would charge to a customer who is known to be from segment i, given that the seller's marginal value of one unit of inventory is Δ. Furthermore, one can show that $p_1^*(\Delta) \geqslant p_2^*(\Delta)$ because segment 1 customers are less price-sensitive. In our model, a seller does not know with certainty what segment an arriving customer belongs to, even after observing the signal from the customer. Therefore, regardless of what the customer's signal is, a seller using personalized pricing will quote a price somewhere between $p_1^*(\Delta)$ and $p_2^*(\Delta)$. This observation suggests that the following discounting strategy performs as well as not committing to a price until after seeing the customer's signal: Set the announced price equal to $p_1^*(\Delta)$, and offer the optimal personalized discount after seeing the customer's signal. The following proposition formalizes this result.

Proposition 1. *Suppose that a seller has y units of inventory with t periods to go. Then*:

(a) $p_1^*(\Delta_t(y)) \geqslant p^*(x, y, t) \geqslant p_2^*(\Delta_t(y))$ *for any signal x.*

(b) *Suppose that a seller uses the following personalized discounting strategy: At the beginning of period t, announce the price to be $p_1^*(\Delta_t(y))$ and then offer a discount of $p_1^*(\Delta_t(y)) - p^*(x, y, t)$ after observing the customer's signal x, thereby charging an effective price of $p^*(x, y, t)$ to this particular customer. The seller's expected revenue under such a strategy is the same as the optimal expected revenue of a seller who quotes prices only after seeing the signal of the customer.*

One caveat is that the higher announced price could result in some potential customers giving up on the product prematurely, even though some of those customers would qualify for lower prices had they made an attempt to purchase and revealed their signals in the process. In such a case, the potential loss of revenue must be weighed against the benefits from personalization.

5. Dynamic Pricing with Partial Personalization

In the previous section, we have considered a scenario where the seller uses not only dynamic pricing (adjusting the price over time based on inventory levels) but also personalized pricing (adjusting the price in response to the signal from the customer). In that model, at any given time and inventory level, each unique signal prompts a unique price from the seller. Such a pricing strategy, which uses personalization to its full extent, might be harder to justify and implement than a pricing strategy in which customers are grouped into a few classes, with each class being charged a different price. While customers might be uncomfortable with full personalization, it might be much easier to gain customer acceptance for group pricing. In this section, we consider such a group pricing model with "partial personalization." Here, the seller continues to use personalized dynamic pricing but picks only two different prices at the beginning of each period, before observing the customer's signal. After observing the customer's signal, the firm decides which of the two prices to offer.

The optimality equations can be written as follows:

$$V_t^P(y) = \max_{p_1, p_2} E_S\Big[\max_{p\in\{p_1,p_2\}}\big\{\lambda\big(\hat{q}_1(S)\bar{F}_1(p)+\hat{q}_2(S)\bar{F}_2(p)\big)\big(p+V_{t-1}^P(y-1)\big) + \big[1-\lambda\big(\hat{q}_1(S)\bar{F}_1(p)+\hat{q}_2(S)\bar{F}_2(p)\big)\big]V_{t-1}^P(y)\big\}\Big],$$
$$y > 0,\ t = 1, \ldots, T,$$

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

The outer maximization corresponds to the problem of picking the two prices at the beginning of period t, and the inner maximization corresponds to the problem of deciding what price to offer for a given signal, S. Although this is a relatively complicated optimization problem, it is possible to prove a certain structure for the optimal policy, stated in the following proposition.

Theorem 2. *Suppose that a seller has y units of inventory with t periods to go and charges the customers one of the two prices p_1 and p_2 with $p_1 > p_2$. Then, there exists a threshold $\bar{z}(y,t)$ such that if an arriving customer has a signal $z \geqslant \bar{z}(y,t)$, it is optimal to offer that customer price p_1; otherwise, it is optimal to offer the customer p_2.*

Theorem 2 shows that, for any given prices p_1 and p_2, a threshold-type policy will be used to determine which of the two prices to offer a given customer. The likelihood ratio condition imposed in (A5) is crucial for Theorem 2. The threshold structure does not necessarily exist when the ordering is not as strong, e.g., when the ordering is in failure rate sense.

5.1. Fixed Prices, Dynamic Threshold Signal

To gain further insight into the choice of the threshold signal z, we now focus on the problem where the prices $p_1 \geqslant p_2$ are fixed exogenously at the beginning of the horizon, but the seller picks the optimal threshold at the beginning of each period. (We know from Theorem 2 that a threshold-type policy is optimal when deciding which of the two prices to offer.) Let $V_t^{FP}(y)$ denote the optimal expected revenue of the seller, given that the seller has y units of inventory with t periods to go. The optimality equations for the seller's problem are as follows:

$$\begin{aligned} V_t^{FP}(y) &= \max_z \big\{\lambda(q_1\bar{G}_1(z)\bar{F}_1(p_1)+q_2\bar{G}_2(z)\bar{F}_2(p_1))(p_1+V_{t-1}^{FP}(y-1)) \\ &\quad + \lambda(q_1G_1(z)\bar{F}_1(p_2)+q_2G_2(z)\bar{F}_2(p_2))(p_2+V_{t-1}^{FP}(y-1)) \\ &\quad + [1-\lambda(q_1\bar{G}_1(z)\bar{F}_1(p_1)+q_2\bar{G}_2(z)\bar{F}_2(p_1)) \\ &\qquad - \lambda(q_1G_1(z)\bar{F}_1(p_2)+q_2G_2(z)\bar{F}_2(p_2))]V_{t-1}^{FP}(y)\big\}, \\ &\qquad\qquad y>0,\ t=1,\ldots,T, \end{aligned} \tag{5}$$

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

As before, after some algebraic manipulation, one can rewrite $V_t^{FP}(y)$ as follows:

$$\begin{aligned} V_t^{FP}(y) &= V_{t-1}^{FP}(y)+\lambda\max_z\big\{(q_1\bar{G}_1(z)\bar{F}_1(p_1)+q_2\bar{G}_2(z)\bar{F}_2(p_1)) \\ &\quad \cdot(p_1-\Delta_t^{FP}(y))+(q_1G_1(z)\bar{F}_1(p_2)+q_2G_2(z)\bar{F}_2(p_2)) \\ &\quad \cdot(p_2-\Delta_t^{FP}(y))\big\}, \quad y>0, t=1,\ldots,T, \end{aligned}$$

where

$$\Delta_t^{FP}(y)=V_{t-1}^{FP}(y)-V_{t-1}^{FP}(y-1), \quad y>0, t=1,\ldots,T. \tag{6}$$

Let $z^*(y,t)$ denote the optimal threshold signal (or the smallest optimal threshold when there is more than one optimal value). The following proposition indicates that if the seller has more inventory or less time until the end of the selling season, then the seller increases the threshold signal, thus making it more likely that an arriving customer will qualify for the lower price, p_2.

Proposition 2. *Suppose that a seller has y units of inventory with t periods to go.*
(a) *If y increases, $z^*(y,t)$ increases.*
(b) *If t increases, $z^*(y,t)$ decreases.*

5.2. Fixed Signal Threshold, Dynamic Prices

To explore the pricing problem further, we now analyze the problem where the firm chooses the two prices p_1 and p_2 dynamically, depending on the time and inventory level, but fixes the threshold signal at the beginning of the horizon. In essence, fixing the threshold signal ahead of time implies that the firm is designating classes of customers that will receive a discount regardless of time and inventory level; e.g., committing to offering discounts to students. We term the customers whose signals exceed the threshold as *class*-1 *customers* and those whose signals are below the threshold as *class*-2 *customers*. Let z denote the threshold signal chosen by the firm. Given the threshold z, let $V_t^{FT}(y,z)$ denote the firm's optimal expected revenue under partial personalization when starting period t with y units of inventory. Then, $V_t^{FT}(y,z)$ is given by the same optimality equations as in (5), but the maximization is over p_1 and p_2 instead of z.[4]

Let $p_j^*(z,y,t)$ denote the optimal price quoted to a class-j customer when the seller has y units in inventory with t periods to go (or the smallest optimal price when there are multiple optimizers). The following proposition describes how the optimal prices depend on the threshold signal.

Proposition 3. *Suppose that the threshold signal is z and the seller has y units of inventory with t periods to go. Then*:
(a) *Class-1 customers are charged a higher price than class-2 customers; i.e., $p_1^*(z,y,t) \geqslant p_2^*(z,y,t)$.*
(b) *If the firm increases the threshold z to be used in period t while keeping the threshold signal unchanged in other periods, both $p_1^*(z,y,t)$ and $p_2^*(z,y,t)$ increase.*

Proposition 3(a) is not necessarily true under some conditions that are weaker than the likelihood ratio ordering. Proposition 3(a) suggests a way in which this pricing strategy can be implemented. Given a threshold signal z, the seller could announce the price to be $p_1^*(z,y,t)$ at the beginning of period t, and then give class-2 customers a discount in the amount of $p_1^*(z,y,t)-p_2^*(z,y,t)$.

Proposition 3(b) indicates that if the firm raises the threshold signal for a single period only, optimal prices for that period also increase. When the threshold is higher, a customer must exhibit a higher signal to be classified as a class-1 customer, which indicates that the customer is likely to have a higher reservation price as well. As a result, the price charged to a class-1 customer increases. As for class-2 customers, as the threshold signal increases, this class grows to include customers with larger signals who were earlier classified as class-1, thus growing to include customers who are likely to have higher reservation prices. Therefore, the higher the threshold is, the higher the price charged to a class-2 customer as well.

6. Conclusion

Arguably, the biggest challenge for firms in personalizing prices is to effectively manage their customers' perceptions of the practice. Amazon's experience clearly shows that a not well-thought-out implementation might seriously upset the customers. As Phillips (2005) argues, however, framing can make the whole difference. For example, a customer will be more satisfied by receiving a discount from a higher price as opposed to paying a premium on top of a lower price, even if both practices resulted in the same effective price. Depending on the industry, customers might also have different attitudes toward price personalization. For example, it is reasonable to expect that in industries where customers are already accustomed to paying different prices (such as travel, hotel, or apparel industries), personalizing prices might be more acceptable. For such industries, time- and inventory-level-dependent pricing of limited inventories has been well researched. In this paper, we investigated the optimal pricing policies in the presence of a third dimension (in addition to time and inventory level) that can be used in setting prices: information the firm has at the individual customer level.

When making pricing decisions based on customer-specific information, firms in some way need to relate the information they have about a particular customer with the customer's willingness to pay. Typically, firms use information that they believe is correlated with the willingness to pay so that they can set their prices accordingly. We have shown, however, that a mere correlation between the information used (which we call the signal) and the willingness to pay does not necessarily imply that customers with higher signals should be charged higher prices. A stronger technical condition than correlation might be needed (the likelihood ratio ordering condition in our model) to ensure that the signal is a more intuitive determinant of the price that needs to be charged. This suggests that firms need to carefully determine the bit of information they use when setting their prices, and it thus points to an interesting avenue for future research: How should the signal be determined, or more precisely, how should the signal be designed so that higher values of the signal would indeed imply higher prices to be charged?

In this paper, we mainly considered two different models for two different price personalization policies. In the first model, we assumed that there are no limitations on the number of different prices that the firm can charge, and we showed that under the likelihood ratio ordering condition that ensures a "strong" positive correlation between the signal and willingness to pay, optimal prices increase with the signal. For the second model, considering the possibility that firms would typically limit the number of different prices they would charge, we assumed that the firm at any one point can only charge one of two different prices. Here, we showed that regardless of whether prices are dynamically set or they are set at the beginning of the horizon, the optimal policy is of threshold-type, i.e., customers with signals that are above a certain level are charged the higher price while those with lower signals are charged the smaller price. We have also established several monotonicity properties. For example, we showed that when prices do not change dynamically, the optimal threshold level changes monotonically with the inventory level and time.

In the online supplement we include an extension to consider the case where the firm cannot observe the signals of some of the customers. The analysis of this extension shows that when signals from some of the customers are not available, the lack of a signal becomes a signal in itself, which the firm can utilize in setting prices. In fact, the firm prefers that the customers with high willingness to pay do not signal. When that is the case, the firm charges a high price to the nonsignaling, high-willingness-to-pay customers and tailors the price for the signaling, low-willingness-to-pay customers.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/.

Endnotes

1. The Robinson-Patman Act of 1936 prohibits anti-competitive price discrimination, primarily in the context of wholesale prices charged to retailers and distributors.
2. *Targetability* is defined as "the ability to predict the preferences and purchase behaviors of individual consumers for the purpose of customizing price or product offers." The authors consider a model where the firm can make errors when classifying customers, a possibility that we also allow in this paper.
3. The signal distributions assumed in Example 1 are taken from Shaked and Shanthikumar (2007). Shaked and Shanthikumar give these distributions as an example to demonstrate that it is possible that two distributions are not ordered in the likelihood ratio while they are ordered in failure rate and reversed failure rate.
4. Online Appendix C provides an equivalent yet more tractable formulation for the optimality equations.

References

Adamy, J. 2000. E-tailer price tailoring may be wave of future. *The Chicago Tribune* (September 25) 4.

Ahn, H.-S., M. Gumus, P. Kaminsky. 2007. Pricing and manufacturing decisions when demand is a function of prices in multiple periods. *Oper. Res.* **55**(6) 1039–1057.

Argon, N. T., S. Ziya. 2009. Priority assignment under imperfect information on customer type identities. *Manufacturing Service Oper. Management*, ePub ahead of print February 10, http://msom.journal.informs.org/cgi/content/abstract/msom.1080.0246v1.

Aviv, Y., A. Pazgal. 2005. A partially observed Markov decision process for dynamic pricing. *Management Sci.* **51** 1400–1416.

Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* **10**(3) 339–359.

Aydin, G., S. Ziya. 2008. Pricing promotional products under upselling. *Manufacturing Service Oper. Management* **10** 360–376.

Besanko, D., J.-P. Dubé, S. Gupta. 2003. Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Sci.* **49** 1121–1138.

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. 1997. Periodic pricing of seasonal products in retail. *Management Sci.* **43** 64–79.

Blank, J. 2001. Personal value pricing. Jupiter Research Report, New York.

Bridis, T. 2005. Study: Most Web shoppers unaware of "price customization"; Online stores are able to charge people different amounts based on data. *Telegraph Herald*, Dubuque, IA, B11.

Chen, Y., G. Iyer. 2002. Consumer addressability and customized pricing. *Marketing Sci.* **21** 197–208.

Chen, Y., C. Narasimhan, Z. J. Zhang. 2001. Individual marketing with imperfect targetability. *Marketing Sci.* **20** 23–41.

Desjardins, D. 2007. Kiosks: Supervalue loyalty on avenue to success? *Retailing Today* (January 8) 27.

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. Gulcu, P. Keskinocak. 2008. Designing optimal preannounced markdown mechanisms in the presence of rational customers with multiunit demands. *Manufacturing Service Oper. Management* **10**(1) 126–148.

Fudenberg, D., J. Tirole. 2000. Customer poaching and brand switching. *RAND J. Econom.* **31** 634–657.

Fudenberg, D., J. M. Villas-Boas. 2006. Behavior-based price discrimination and customer recognition. T. J. Hendershott, ed. *Handbook on Economics and Information Systems*. Elsevier, Amsterdam, 377–436.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizions. *Management Sci.* **40** 999–1020.

Johnson, C. A., L. Allen, D. Young, J. Tong. 2000. Pricing gets personal. The Forrester Report, Cambridge, MA.

Kumar, N., R. Rao. 2006. Using basket composition data for intelligent supermarket pricing. *Marketing Sci.* **25** 188–199.

Kuo, C., H. Ahn, G. Aydin. 2009. Dynamic pricing of limited inventories when customers negotiate. Working paper, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor.

Liu, Q., G. van Ryzin. 2005. Strategic capacity rationing to induce early purchases. Working paper, Business School, Columbia University, New York.

Liu, Y., Z. J. Zhang. 2006. The benefits of personalized pricing in a channel. *Marketing Sci.* **25** 97–105.

Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management* **8** 136–148.

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.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. Wiley, New York.

Murthi, B. P. S., S. Sarkar. 2003. The role of the management sciences in research on personalization. *Management Sci.* **49** 1344–1362.

Netessine, S., S. Savin, W. Xiao. 2006. Revenue management through dynamic cross-selling in e-commerce retailing. *Oper. Res.* **54** 893–915.

Phillips, R. L. 2005. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA.

Popescu, I., Y. Wu. 2007. Dynamic pricing strategies with reference effects. *Oper. Res.* **55**(3) 413–429.

Ramasastry, A. 2005. Web sites change prices based on customers' habits. Accessed August 2, 2007, http://www.cnn.com/2005/LAW/06/24/ramasastry.website.prices/index.html.

Shaffer, G., Z. J. Zhang. 1995. Competitive coupon targeting. *Marketing Sci.* **14** 395–416.

Shaffer, G., Z. J. Zhang. 2002. Competitive one-to-one promotions. *Management Sci.* **48** 1143–1160.

Shaked, M., J. G. Shanthikumar. 2007. *Stochastic Orders*. Springer, New York.

Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Sci.* **53** 726–741.

Talluri, K. T., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Springer, New York.

Thisse, J.-F., X. Vives. 1988. On the strategic choice of spatial price policy. *Amer. Econom. Rev.* **78** 122–137.

Turow, J., L. Feldman, K. Meltzer. 2005. Open to exploitation: American shoppers online and offline. Annenberg Public Policy Center Report, Philadelphia.

Varian, H. R. 1989. Price discrimination. R. Schmalensee, R. D. Willig, eds. *Handbook of Industrial Organization*. North-Holland, Amsterdam.

Villas-Boas, J. M. 2004. Price cycles in markets with customer recognition. *RAND J. Econom.* **35** 486–501.

Weiss, R. M., A. K. Mehrotra. 2001. Online dynamic pricing: Efficiency, equity, and the future of e-commerce. *Virginia J. Law and Tech.* **6**.

Zhang, D., W. L. Cooper. 2009. Pricing substitutable flights in airline revenue management. *Eur. J. Oper. Res.* **197** 848–861.

Ziya, S., H. Ayhan, R. D. Foley. 2004. Relationship among three assumptions in revenue management. *Oper. Res.* **52** 804–809.