Personalized Dynamic Pricing of Limited Inventories

Goker Aydin* Serhan Ziya**

*Department of Industrial and Operations Engineering,

University of Michigan,

Ann Arbor, MI, U.S.A.

ayding@umich.edu

**Department of Statistics and Operations Research

University of North Carolina

Chapel Hill, NC, U.S.A.

ziya@unc.edu

Abstract

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 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. Through a numerical study, we investigate the interactions between personalized pricing and dynamic pricing. In an extension model where signals from some of the customers are not available, we find that it might in fact be better for the firm if customers with high willingness-to-pay do not provide their signals, since it helps the firm better price-discriminate.

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 et al. (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] Ramasastry, a professor of law at the University of Washington, writes: "The reality is that Internet price customization does exist – and, contrary to popular opinion, is typically legal"(Ramasastry, 2005).

[1]Robinson-Patman Act of 1936 prohibits anti-competitive price discrimination, primarily in the context of wholesale prices charged to retailers and distributors.

It appears that the real challenge for the firms is to manage customers' perception of the practice since 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 where 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. (A simple Google search will yield newspaper articles that claim the Saturday stay over rule is on its way to extinction. Interestingly, however, such articles have been published as early as 2001 and as recently as 2006!)

One popular way of implementing personalized pricing is through discount coupons. In fact, a white paper by Blank et al. (2001) suggests that personalized pricing is more acceptable to consumers when it is framed as offering tailored discounts from a fixed price. Turow et al. (2005) also makes an interesting observation. Even though 90% of customers participating in their survey found personalized pricing unacceptable, only 64% of the same customers stated that they would be bothered if other customers got better discount offers than they did.

The use of personalized pricing requires that the seller acquire customer-specific information. There are many ways in which firms can collect such information. For example, a customer may simply identify herself as belonging to a certain demographic group (e.g., a student or a senior citizen), or reveal potentially useful information through a choice (e.g., by opting for a round-trip ticket with a Friday return as opposed to a Sunday return). In addition, 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. Of course, online retailers have a wealth of information about registered customers. In this paper, we use a model where each customer furnishes the retailer with a signal, which captures the information that the seller has about the individual.

Typically, the information about a customer is only a partial indicator of the customer's willingness-to-pay. For example, while students may tend to be more price sensitive, surely there are students who have much deeper pockets and higher willingness-to-pay than an average customer. In other words, even though a customer can be reliably identified as belonging to a certain group, this information may not be a perfect indicator of the customer's willingness-to-pay. In our model, the customer signal, which captures the information available about the customer, is assumed to be correlated with the reservation price of the customer. Hence, the signal provides 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*. However, implementing personalization to such a full extent may be impractical. Instead, 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 where a range of signals are bunched together and offered a single price.

One of the objectives of this paper is to provide insights as to what type of signal should be used for it to be a clear indicator of how the firm should set its 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 Section 4, the answer is "no". Positive correlation is not sufficient. A stronger condition is needed. In Section 4, we give a precise description of this condition, which leads to useful relationships between the signal and optimal pricing policies. For example, when this "strong" correlation condition is satisfied, under full personalization, 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 exist in the absence of the strong correlation condition.

A successful implementation of personalized pricing that does not create customer backlash would increase profits. In this paper, through a numerical study, we investigate conditions under which price personalization appears to be a particulary profitable strategy. First, we find that profit improvements are higher when initial inventory levels are higher. Second, personalization seems to have similar benefits regardless of whether it is used together with static pricing or dynamic pricing. And third, when compared with dynamic pricing, it is especially profitable when inventory levels are high since dynamic pricing does not bring much improvement in that case as it has been previously shown (see, e.g., Gallego and van Ryzin 1994).

The presence of personalization may give customers an incentive to act strategically. For example, a customer may wish to hide her identity if she feels that the seller will charge a lower price to an anonymous customer. A customer may 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, we do consider an extension where there are two types of customers in the population: those who provide signals and those who do not. An interesting finding here is that the firm does not necessarily prefer that all of its customers provide signals. The fact that a customer does not provide a signal has some informational value, in some cases, such information may be more useful to the firm than any other signal the customer could provide.

The rest of the paper is organized as follows: In Section 2, we review the relevant literature, and in Section 3, we describe our base model. Sections 4 and 5 present our results for the full personalization and partial personalization models, respectively. In Section 6, we report our numerical analysis on the benefits of personalized pricing and in Section 7, we extend our model to investigate the possibility that the signals from some of the customers may not be available. Section 8 gives our concluding remarks. All the proofs are given in the Appendix.

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. Among this work, Gallego and van Ryzin (1997), Zhang and Cooper (2005), and Maglaras and Meissner (2006) consider dynamic pricing of multiple products. Zhao and Zheng (2000) study a model that assumes time-dependent customer demand while Aviv and Pazgal (2005b) 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. Feng and Gallego (2000), and Feng and Xiao (2000) study pricing decisions when there is a finite menu of prices and there is possibly a restriction on the number of price changes. Smith and Achabal (1998) consider a model where demand depends on the inventory level in addition to time and price. Aviv and Pazgal (2005a), Liu and van Ryzin (2005), Elmaghraby, Gulcu, and Keskinocak (2006), and Su (2007) investigate dynamic pricing decisions when customers act strategically. Monahan, Petruzzi, and Zhao (2004) deal with the decision of determining the initial inventory for products that are under dynamic pricing. Popescu and Wu (2006) and Ahn, Gumus, and Kaminsky (2007) consider models where price in a given period influences the demand in other periods. For comprehensive reviews of the literature on dynamic pricing pre-2003, see Elmaghraby and Keskinocak (2003), Bitran and Caldentey (2003), and Talluri and van Ryzin (2004).

Our work mainly differs from the above 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, Ahn, and Aydin (2007) investigate negotiation while Netessine, Savin, and Xiao (2006) and Aydin and Ziya (2007) investigate upseling/cross-selling, in which case the price offered to a customer depends on what the customer already bought from the seller. All of these papers provide insights that are relevant within the specific form of personalized pricing they are considering. 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. (2007) the signal is the offer made by a bargainer. In Aydin and Ziya (2007), 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 (1988). The subject has received significant attention in the last few years. Among the earlier work, Narasimhan (1984) develops a theory of price discrimination through discount coupons. 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 since 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, Narasimhan, and Zhang (2001) find that when *customer targetability* is sufficiently low, personalized pricing is profitable for two competing firms, but for high targetability levels, firms may be worse off since 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 ketchup market, Besanko, Dubé, and Gupta (2003) also find the practice to be a profitable strategy.

Chen and Iyer (2002) investigate how consumer *addressability* (the ability to reach individuals and make personalized offers) affects market competition when customers can be identified perfectly and show that higher levels of customer addressability might lead to increased levels of competition between symmetric firms. Iyer, Soberman, and Villas-Boas (2005) find that targeted advertising is more valuable to competing firms than personalized pricing. Choudhary et al. (2005) investigate personalized pricing strategies in a duopoly where the firms compete both over the price and quality. Taylor (2004), Acquisti and Varian (2005), and Calzolari and Pavan (2006) are interested in consumer privacy issues and consumers' privacy preferences on pricing policies, firm's

[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.

profit, and consumer welfare. 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 lifecycles 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 (cdf's) Φ_1 and Φ_2 (with corresponding probability density function (pdf's) ϕ_1 and ϕ_2), if the failure rate of Φ_1 is less than that of Φ_2, i.e., if $\frac{\phi_1(x)}{1-\Phi_1(x)} \leq \frac{\phi_2(x)}{1-\Phi_2(x)}$, $\forall x$, then we say Φ_1 dominates Φ_2 in failure rate ordering, and we write $\Phi_1 \geq_{fr} \Phi_2$. Furthermore, if $\frac{\phi_1(x)}{\phi_2(x)} \geq \frac{\phi_1(y)}{\phi_2(y)}$ for any $x > y$, then we say Φ_1 dominates Φ_2 in likelihood ratio ordering, and we write $\Phi_1 \geq_{lr} \Phi_2$. If Φ_1 and Φ_2 are cdf's for discrete random variables, the pdf's in the definition of likelihood ratio ordering are replaced by corresponding probability mass functions (pmf's). See Appendix A 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 non-positive/negative.

Suppose 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 in the population. 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 $\overline{F}_i(\cdot) := 1 - F_i(\cdot)$. We make the following assumptions on the reservation price distributions.

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

(A2) $F_i(\cdot), i = 1, 2$, *have strictly increasing generalized failure rates, i.e.,* $\frac{x f_i(x)}{\overline{F}_i(x)}$ *is strictly increasing.*

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

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, Ayhan and Foley, 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 has 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 (2007) use a similar formulation within the context of patient triage, where patients are classified as critical or non-critical patients depending on an initial inspection by a triage officer.

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, may 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, since 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:

(A4) $G_i(\cdot), i = 1, 2$, *are strictly increasing functions, and they both have the same non-negative*

support.

Let μ_i denote the expected reservation price and s_i the expected signal of a segment-i customer. We make the following additional assumption on the signals from the two segments.

(A5) $s_1 > s_2$; i.e., the expected signal of segment-1 customers is strictly greater than that of segment-2.

The following proposition describes how the signals and reservation prices are related in our model.

Proposition 1 *Let S and R denote, respectively, the signal and the reservation price of an individual chosen at random from the consumer population. Suppose (A1) through (A5) hold. Then, S and R are positively correlated.*

Since the signal and reservation price are positively correlated, one might intuitively expect that a higher signal would prompt a higher price from the seller. We will investigate this in the following section.

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. As Jewel-Osco and Albertsons stores have been doing, one way of implementing such a practice would be to ask customers to identify themselves with their loyalty cards and offer them customized discounts. Such a practice is also technologically feasible for many online retailers, since the retailer can present different users with different prices after observing, for example, the IP address of the user, which can be used to identify the customer, or customers may simply choose to log in and identify themselves. Although this practice may be feasible, it may 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 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 $\widehat{q}_i(x)$ denote the probability that a customer with signal x belongs to segment i. Using Bayes' rule, $\widehat{q}_i(x)$ is given by

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

Suppose 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.

$$\begin{aligned} V_t(y) &= E_S\left[\max_p \left\{ \begin{array}{l} \lambda\left(\widehat{q}_1(S)\overline{F}_1(p)+\widehat{q}_2(S)\overline{F}_2(p)\right)(p+V_{t-1}(y-1)) \\ +\left[1-\lambda\left(\widehat{q}_1(S)\overline{F}_1(p)+\widehat{q}_2(S)\overline{F}_2(p)\right)\right]V_{t-1}(y) \end{array}\right\}\right], y>0, t=1,\ldots,T, \\ V_t(0) &= 0, t=1,\ldots,T, \text{ and } V_0(\cdot)=0. \end{aligned}$$

In the remainder of this section, let

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

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:

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

where

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

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.) Since the reservation price and the signal of a customer are positively correlated under our assumptions, one might expect that $p^*(x,y,t)$ is increasing in x. Interestingly, however, this is not necessarily true. Consider the following numerical example.

Example 1: Suppose that signals are discrete random variables. Suppose that the probability mass function for S_1 is given by $g_1(1)=0.1$, $g_1(2)=0.3$, $g_1(3)=0.2$, and $g_1(4)=0.4$. On the other hand, S_2 is uniform so that $g_2(1)=g_2(2)=g_2(3)=g_2(4)=0.25$. Also, 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 $F_1 \geq_{fr} F_2$.

In the example above, the signal of the first segment dominates the signal of the second segment in failure rate ordering. Therefore, $s_1 > s_2$, satisfying our assumption $(A5)$. Assumptions $(A1)$ through $(A4)$ are also satisfied in this example. Hence, the signal and reservation price of an individual are positively correlated (by Proposition 1). Nevertheless, the optimal price may be lower for a higher signal. 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 also 1 period. In fact, optimal prices satisfy the same relationship for many different pairs of inventory level and time. In this example, a segment-1 customer is more likely to provide a signal of "2" than a signal of "3" whereas a segment-2 customer is equally likely to provide all four signals. Hence, when a customer provides a signal of "2", that customer is more likely to be in segment-1 than in segment-2, which is why a signal of "2" prompts a larger price.

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 across the two signals – which, in our model, is stronger than positive correlation – may not be enough to guarantee an intuitive relationship between the signal and the optimal price. In the remainder of this paper, we will strengthen $(A5)$ as follows:

$(A5')$ $G_1(\cdot) \geq_{lr} G_2(\cdot)$.

Assumption $(A5')$ imposes a strong stochastic order between the signals from the two segments. (Likelihood ratio ordering implies failure rate ordering.[3]) As the next theorem shows, Assumption $(A5')$ leads to an intuitive relationship between the signal and the optimal price.

Theorem 1 *Suppose* $(A5')$ *holds in addition to* $(A1)$ *through* $(A4)$*. Then, 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.*

When signal distributions are ordered in the likelihood ratio, $\widehat{q}_1(x)$ increases with x, i.e., the higher the signal of a customer, the more likely the customer is to belong to segment 1, the segment with lower price sensitivity. As a result, optimal price is also higher for customers with higher signals. When the ordering between signal distributions is weaker, i.e. when there is only failure rate ordering, $\widehat{q}_1(x)$ is not necessarily increasing in x and thus the optimal price is also not necessarily larger. In the rest of the paper, we assume that Assumptions $(A1)$ through $(A4)$ and Assumption $(A5')$ hold.

Prior work in dynamic pricing literature has established the monotonicity of the optimal prices with respect to inventory level and time under various settings. The following proposition shows that similar monotonicity properties continue to hold when personalized pricing is introduced.

[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.

Proposition 2 *The (smallest) optimal price, $p^*(x,y,t)$, is:*

(a) increasing in remaining time, t, given signal x and inventory level y, and

(b) decreasing in inventory level, y, given signal x and time t.

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 it 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 will 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 $p_i^*(\Delta) := \arg\max_p (p-\Delta)\overline{F}_i(p),\ i = 1,2$. Note that, with this definition, $p_i^*(\Delta)$ 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 $p_1^*(\Delta) \geq p_2^*(\Delta)$, since 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 3 *Suppose a seller has y units of inventory with t periods to go. Then:*

(a) $p_1^*(\Delta_t(y)) \geq p^*(x,y,t) \geq p_2^*(\Delta_t(y))$ 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.

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, may be harder to justify and also implement than a pricing strategy where customers are grouped into a few classes, with each class being charged a different price. While customers may be uncomfortable with full personalization, it may be much easier to gain customer acceptance for group pricing. In this section, we consider such a group pricing model with 'partial personalization.' In this model, 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:

$$
\begin{aligned}
V_t^P(y) &= \max_{p_1,p_2} E_S\left[\max_{p\in\{p_1,p_2\}}\left\{\begin{array}{l}\lambda\left(\widehat{q}_1(S)\overline{F}_1(p)+\widehat{q}_2(S)\overline{F}_2(p)\right)\left(p+V_{t-1}^P(y-1)\right)\\ +\left[1-\lambda\left(\widehat{q}_1(S)\overline{F}_1(p)+\widehat{q}_2(S)\overline{F}_2(p)\right)\right]V_{t-1}^P(y)\end{array}\right\}\right],\\
&\quad y>0, t=1,\ldots,T,\\
V_t^P(0) &= 0, t=1,\ldots,T, \text{ and } V_0^P(\cdot)=0.
\end{aligned}
$$

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 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 $z \geq \bar{z}(y,t)$, 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 which of the two prices to offer to a given customer. The likelihood ratio condition $(A5')$ 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.

If the seller knew the segment of each arriving customer, then the seller would simply charge each customer the price that is optimal for the customer's segment. In the absence of such information, the seller is using the customer signals as a proxy to divide the customers into two classes: those whose signals exceed the threshold (*class-1 customers*) and those whose signals are below the threshold (*class-2 customers*). Knowing that a class-1 customer is more likely than a class-2 customer to be in segment-1 (the segment with lower price sensitivity), the seller charges class-1 customers a higher price. In doing so, the seller can make one of two types of classification errors: misclassify a segment-1 customer as a class-2 customer because the customer's signal happens to be low, which happens with probability $G_1(z)$ if the threshold is set at z, or misclassify a segment-2 customer as a class-1 customer because the customer's signal happens to be high, which happens with probability $\overline{G}_2(z)$. In the case of the former error, the seller is charging the customer the lower price when it could profitably charge a higher price. In the case of the latter error, the seller is charging the customer the higher price when it would be more profitable to charge a lower price. The seller's choice of the threshold signal clearly affects the likelihood of these two types of errors. For higher threshold values, the probability of misclassifying a segment-1 customer is also higher while the probability of misclassifying a segment 2 customer is lower.

In the rest of this section, we investigate how the remaining time and inventory affect the prices, the threshold level and the misclassification probabilities. In addition, we explore the interaction between optimal prices and the threshold level.

5.1 Fixed prices, dynamic threshold signal

In order to gain further insight into the choice of the threshold signal z, we now focus on the problem where the prices $p_1 \geq p_2$ 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 a seller who uses such a threshold setting strategy, given the seller has y units of inventory with t periods to go. The optimality equations for the seller's problem are as follows:

$$\begin{aligned} V_t^{FP}(y) &= \max_z \left\{ \begin{array}{l} \lambda\left(q_1\overline{G}_1(z)\overline{F}_1(p_1)+q_2\overline{G}_2(z)\overline{F}_2(p_1)\right)\left(p_1+V_{t-1}^{FP}(y-1)\right) \\ +\lambda\left(q_1G_1(z)\overline{F}_1(p_2)+q_2G_2(z)\overline{F}_2(p_2)\right)\left(p_2+V_{t-1}^{FP}(y-1)\right) \\ +\left[1-\lambda\left(q_1\overline{G}_1(z)\overline{F}_1(p_1)+q_2\overline{G}_2(z)\overline{F}_2(p_1)\right)\right. \\ \left.-\lambda\left(q_1G_1(z)\overline{F}_1(p_2)+q_2G_2(z)\overline{F}_2(p_2)\right)\right]V_{t-1}^{FP}(y) \end{array} \right\}, \qquad (4) \\ & \quad y>0, t=1,\ldots,T, \\ V_t^{FP}(0) &= 0, t=1,\ldots,T, \text{ and } V_0^{FP}(\cdot)=0. \end{aligned}$$

The first (second) term of the summation inside the curly brackets is the expected revenue-to-go if a class-1 (class-2) customer arrives and purchases the product at price p_1 (p_2). The third term is the expected revenue-to-go when no purchase is made in period t. As before, after some algebraic manipulation, one can rewrite $V_t^{FP}(y)$ as follows:

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

where

$$\Delta_t^{FP}(y)=V_{t-1}^{FP}(y)-V_{t-1}^{FP}(y-1), y>0, t=1,\ldots,T. \qquad (5)$$

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 describes the effect of y and t on the optimal threshold:

Proposition 4 *Suppose a seller has y units of inventory with t periods to go.*
(a) If y increases, $z^*(y,t)$ increases. Consequently, the probability of misclassifying a segment-1 customer as class-2 increases and the probability of misclassifying a segment-2 customer as class-1 decreases.
(b) If t increases, $z^*(y,t)$ decreases. Consequently, the probability of misclassifying a segment-1 customer as class-2 decreases and the probability of misclassifying a segment-2 customer as class-1 increases.

The 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 be classified as a class-2 customer and qualify for the lower price, p_2. This is rather intuitive: if the seller has a large inventory to be sold over a short span of time, then the seller is motivated to move inventory quickly, which can be done by offering the lower price to a larger set of consumers. At the same time, such an adjustment of the signal shows that the seller would rather make the error of misclassifying a segment-1 customer than misclassifying a segment-2 customer. In the case of the former error, the seller is charging a misclassified segment-1 customer a lower price than it could, but this error becomes less punishing as the pressure to move inventory increases.

5.2 Fixed signal threshold, dynamic prices

In the previous subsection, we assumed the prices are fixed exogenously at the beginning of the horizon. In order 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 groups of customers that will receive a discount regardless of time and inventory level; e.g., committing to offering discounts to students. 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 (4), but the maximization is over p_1 and p_2 instead of z. In order to obtain an equivalent, but more tractable representation of the optimality equations, we introduce $\widetilde{q}_{ij}(z)$, the probability that a customer is from segment i, given that she has been classified as a class-j customer, $i,j = 1,2$. For example, given threshold signal z, $\widetilde{q}_{12}(z)$ is the probability that a customer who provided a signal less than z (and, thus, were classified as a class-2 customer) is in fact from segment-1. Using Bayes' rule, $\widetilde{q}_{ij}(z)$ is given by

$$\begin{aligned}\widetilde{q}_{i1}(z) &= \frac{q_i\overline{G}_i(z)}{q_1\overline{G}_1(z)+q_2\overline{G}_2(z)},\ i=1,2,\\ \widetilde{q}_{i2}(z) &= \frac{q_iG_i(z)}{q_1G_1(z)+q_2G_2(z)},\ i=1,2.\end{aligned} \tag{6}$$

Furthermore, let

$$\beta_j(z,p) := \widetilde{q}_{1j}(z)\overline{F}_1(p)+\widetilde{q}_{2j}(z)\overline{F}_2(p).$$

Note that $\beta_j(z,p)$ is the probability that a customer will buy the product at price p, given that the customer has been classified as class-j when the threshold signal is z. Using this notation one can show that $V_t^{FT}(y,z)$ can be written as follows:

$$\begin{aligned} V_t^{FT}(y,z) &= V_{t-1}^{FT}(y,z) + \lambda\left(q_1\overline{G}_1(z) + q_2\overline{G}_2(z)\right)\max_{p_1}\left\{\beta_1(z,p_1)\left(p_1 - \Delta_t^{FT}(y,z)\right)\right\} \\ &+ \lambda\left(q_1 G_1(z) + q_2 G_2(z)\right)\max_{p_2}\left\{\beta_2(z,p_2)\left(p_2 - \Delta_t^{FT}(y,z)\right)\right\}, y > 0, t = 1,\ldots,T \end{aligned}$$

where

$$\Delta_t^{FT}(y,z) = V_{t-1}^{FT}(y,z) - V_{t-1}^{FT}(y-1,z), y > 0, t = 1,\ldots,T. \tag{7}$$

In the optimality equation above, the first maximization corresponds to choosing the price to be charged to a class-1 customer and the latter maximization to choosing the price for a class-2 customer. 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 until the end of the season (or the smallest optimal price when there are multiple optimizers). The following proposition describes how the optimal prices depend on the threshold signal.

Proposition 5 *Suppose 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., $p_1^*(z,y,t) \geq p_2^*(z,y,t)$.*
(*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 5(a) is not surprising given Theorem 2 and is similarly not necessarily true under some conditions that are weaker than the likelihood ratio ordering. Proposition 5(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 5(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.

Note that Proposition 5(b) does not answer the following question: Consider two firms each using a threshold signal that is fixed at the beginning of the horizon. Which firm will charge higher prices: the one with the higher threshold or the one with the lower threshold? The intuition behind Proposition 5(b) suggests that the firm with the higher threshold should charge higher prices as well. Although this insight appears to hold in many cases, there are instances where this is not the case, as the following example shows.

Example 2: Suppose that signal distributions are discrete with $g_1(1) = 0.05, g_1(2) = 0.1, g_1(3) = 0.15, g_1(4) = 0.25, g_1(5) = 0.45, g_2(1) = 0.45, g_2(2) = 0.25, g_2(3) = 0.15, g_2(4) = 0.1, g_2(5) = 0.05$, and reservation price distributions, $F_1(\cdot)$ are $F_2(\cdot)$ are Weibull with a common shape parameter 2 and scale parameters, 100 and 50, respectively. Suppose also that $\lambda = 0.5$ and $q_1 = 0.3$. For this example, it turns out that $p_1^*(4, 1, 24) \cong 123.36$ and $p_1^*(5, 1, 24) \cong 123.27$. That is, a firm using a threshold signal of "5" will charge a lower price than a firm with threshold "4", if both firms have one unit of inventory and 24 periods to go. In fact, this order between the two optimal prices appears to hold for any $t \geq 20$ and $y = 1$, while for all other values of inventory level and time, optimal prices are ordered in the other direction.

What drives the behavior illustrated in Example 2 is the marginal value of inventory. When both firms have one unit of inventory and 24 periods to go, the marginal value of inventory is larger for the firm using a threshold of "4" than the firm using a threshold of "5;" i.e., $\Delta_{24}^{FT}(1, 4) > \Delta_{24}^{FT}(1, 5)$. Although a higher threshold tends to drive the price up, the lower marginal value of inventory for the threshold of "5" more than compensates for such effects of the higher threshold, resulting in the lower price charged by the firm with a threshold of "5."

The next proposition shows that the monotonicity of prices in time and inventory level is preserved in this model.

Proposition 6 *Suppose the threshold signal is z and the seller has y units of inventory with t periods to go. Then:*

(a) If y increases, both $p_1^*(z, y, t)$ and $p_2^*(z, y, t)$ decrease.

(b) If t increases, both $p_1^*(z, y, t)$ and $p_2^*(z, y, t)$ increase.

6 Benefits from Personalized Pricing

In this section, we examine how the benefits from personalized pricing depend on time and inventory level with a numerical study. In addition, we check if personalized pricing is more beneficial in the presence or absence of dynamic pricing. To this end, we consider six different pricing strategies, each with a differing degree of dynamic pricing and personalization:

Dynamic pricing with full personalization: This strategy corresponds to the model discussed in Section 4. At the beginning of each period, the seller picks the optimal announced price based on time and inventory. After the customer arrives and provides a signal, the seller offers the optimal discount given the customer's signal.

Dynamic pricing with partial personalization: At the beginning of each period, the seller picks a price and a discount level based on time and inventory, and offers the discount to customers whose signals are below a threshold signal. The price, discount level and threshold signal are chosen optimally at the beginning of each period.

Dynamic pricing with no personalization: At the beginning of each period, the seller picks the optimal announced price based on time and inventory without any subsequent personalization.

Static pricing with full personalization: At the beginning of the horizon, the seller picks an announced price and a discount level for each signal. During the horizon, every time a customer arrives, the seller offers the discount corresponding to the customer's signal. The static price and discount levels are chosen optimally at the beginning of the horizon.

Static pricing with partial personalization: At the beginning of the horizon, the seller picks an announced price, a discount level, and a threshold signal. During the horizon, every time a customer arrives, if the customer's signal is below the threshold, then the seller offers the discount to this customer. The static price, discount level and threshold signal are chosen optimally at the beginning of the horizon.

Static pricing with no personalization: At the beginning of the horizon, the seller picks an announced price that will be in effect throughout the horizon. The announced price is chosen optimally.

We considered several scenarios which differed in the signal and reservation price distributions and segment identity probabilities q_1 and q_2. We observed that our conclusions were consistent across all the scenarios. For brevity, here we report our findings over a single example. In this example, signals are discrete random variables. Suppose that the probability mass function for S_1

is given by $g_1(1) = 0.2$, $g_1(2) = 0.3$, $g_1(3) = 0.5$ while the probability mass function for S_2 is given by $g_2(1) = 0.5$, $g_2(2) = 0.3$, and $g_2(3) = 0.2$. Also, 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.

Figure 1 depicts the effect of the inventory level on the benefits from full and partial personalization, both under dynamic pricing. (These benefits correspond to revenue improvement over the base case of "dynamic pricing with no personalization.") As the figure indicates, the benefits are larger when inventory level is larger. This is due to two reasons. First, when inventory level is larger, the firm has the opportunity to serve a larger number of customers. Therefore, the firm can engage in personalization more frequently, which increases the benefits from personalization. Secondly, when the product is scarce, the seller is reluctant to sell the product to any customer but those that are willing to pay very high prices. Therefore, the seller announces a very high price and does not offer much of a discount even after seeing the signal of a customer. In contrast, as inventory level gets larger, the seller is more inclined to let the product go even at lower prices. In such a setting, in the presence of personalization, the seller can start with a high announced price without significantly increasing its risk of excess inventory, because the seller has the option to offer a generous discount after seeing the customer's signal. Therefore, personalized pricing makes it easier for the seller to manage the trade-off between revenue per unit sold and the risk of excess inventory, resulting in larger benefits in this region. Kuo et al. (2007) observe similar behavior regarding the benefits of negotiation, which is not surprising, since negotiation is ultimately a form of personalized pricing.

Figure 2 compares the benefits from full personalization under dynamic and static pricing. The figure suggests that benefit of personalization in dynamic pricing is comparable to the benefit of personalization in static pricing across all initial inventory levels. For high levels of initial inventory this is not surprising, because as the inventory level grows optimal dynamic prices begin to approximate the optimal static price and the benefits from personalization under the two scenarios converge. When inventory level is low, one might expect that the benefits from personalization under static pricing will be larger, because the introduction of personalization to a static pricing regime makes price discrimination possible, while some price discrimination is already taking place under dynamic pricing even in the absence of personalization. However, we observe from Figure 2 that the benefits are similar even when the inventory level is low. This happens because, under static pricing, the seller must commit to a price for each signal at the beginning of the horizon,

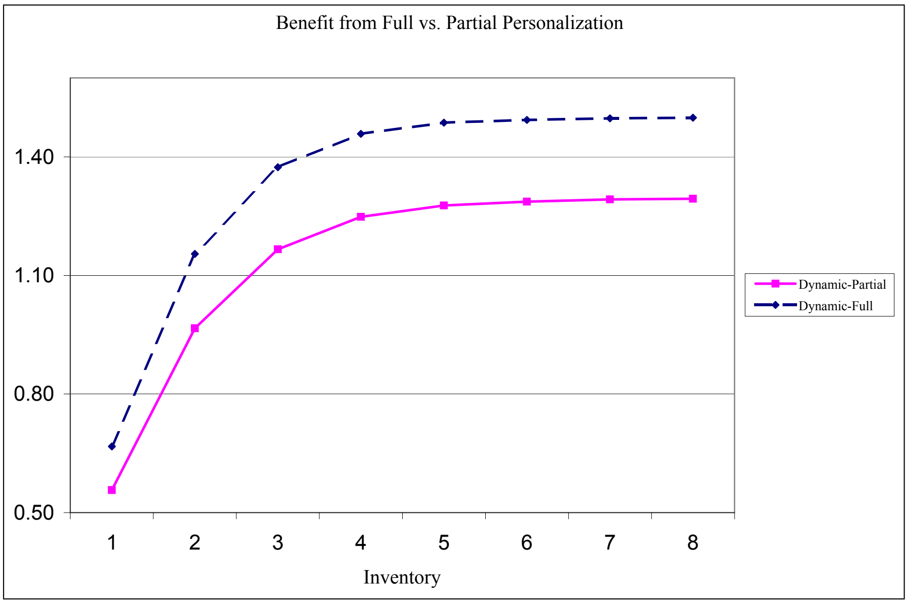

Figure 1: Benefits from full and partial personalization under dynamic pricing. The benefits are increasing in inventory level.

which puts a damper on the benefits from personalization. In contrast, under dynamic pricing, the seller can customize the price for each signal depending on time and inventory level, which gives personalization under dynamic pricing an additional edge.

Consider a seller who is using neither personalization nor dynamic pricing in the status quo. What would be more beneficial for such a seller: adding dynamic pricing or personalization? Figure 3 depicts the benefits from personalization and benefits from dynamic pricing for a seller who is not using either in the status quo. As one would expect, the benefit from dynamic pricing is larger when inventory level is low and gets smaller as inventory level gets larger. The benefit from personalized pricing is small at lower inventory levels, but grows as inventory level gets larger and eventually exceeds the benefits from dynamic pricing. Therefore, if a seller needs to pick one strategy over the other, then an overstock seller should opt for personalization, while a seller with a scarce product should opt for dynamic pricing. Kuo et al. (2007) make a similar observation when comparing the benefits of negotiation and dynamic pricing.

7 Extension: Non-signaling Customers

In this section, we extend the model of Section 4 to consider the possibility that the firm may not be able to observe signals from all the customers. In the case where signals are volunteered by the

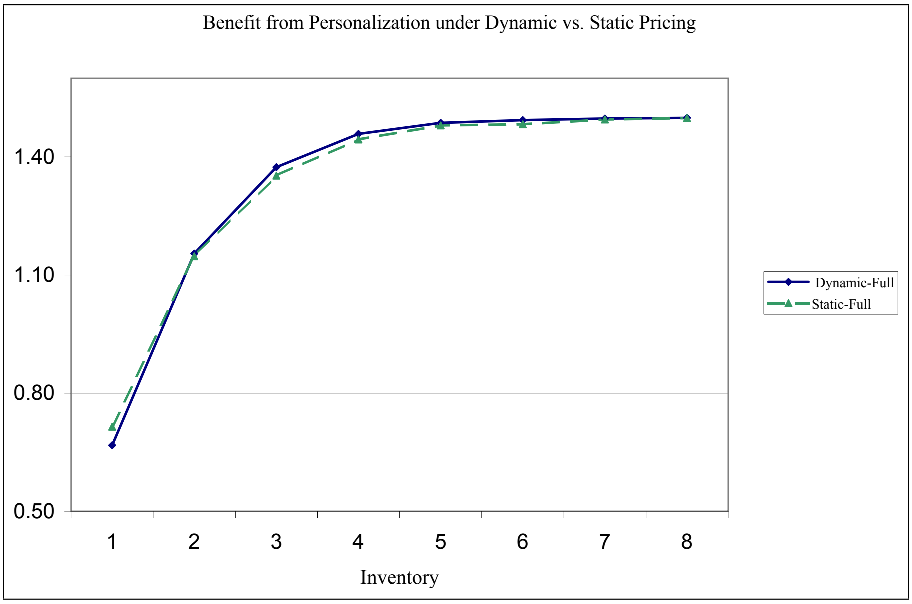

Figure 2: Benefits from full personalization under dynamic versus static pricing. The benefits are more or less the same under both pricing regimes.

customers themselves, some of the customers may simply choose to hide their signals for privacy concerns. For example, customers can choose to hide their identities online by simply deleting the cookies that track them (see Acquisti and Varian 2005). In the case where signals are typically directly observed by the firm without any need for customer volunteering, the signal may simply be unavailable for a group of customers. For instance, if the signal uses past purchasing data, the firm may not have sufficient information on relatively new customers.

We assume that the signal of a segment i customer is available with probability r_i.[4] At the beginning of each period, the firm announces a price that will be charged to customers who do not provide a signal. On the other hand, customers who provide signals may be offered a discount from the announced price. The amount of the discount, if offered, depends on the customer's signal. It is unlikely that a firm would announce a price and then charge a premium to customers who do provide a signal. Therefore, we assume that the firm never charges a signaling customer a price that is higher than the announced price, the price asked from the non-signaling customers.

[4]An interesting extension would be to consider a model where customers act strategically and choose to volunteer their signal or not depending on the value of their signal. In such a model, it would be reasonable to assume that customers with lower signals would be more willing to provide their signals.

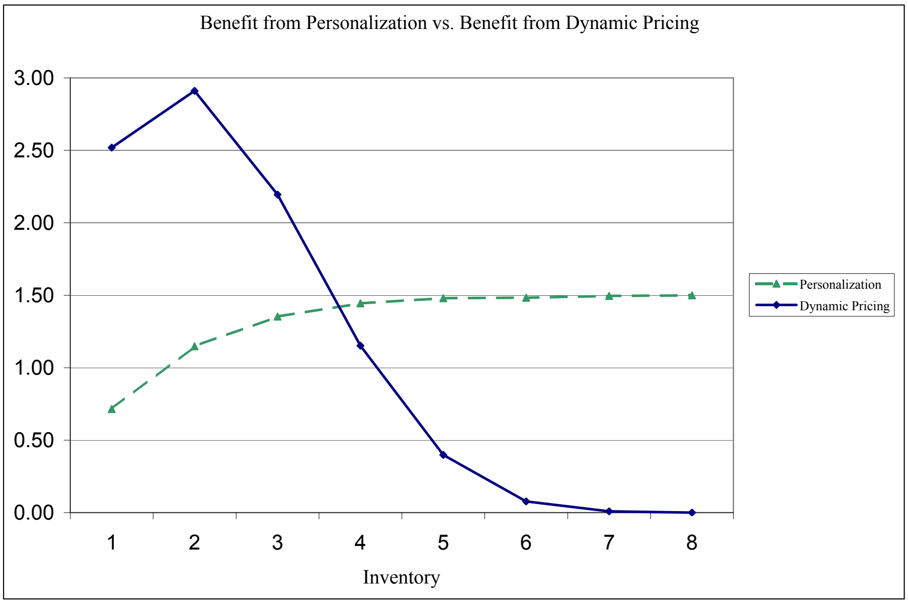

Figure 3: Benefits from dynamic pricing versus benefits from personalization. Dynamic pricing is preferable at lower inventory levels and personalization at higher inventory levels.

Using the notation of Section 4, we can write the optimality equation as follows:

$$V_t^{NS}(y) = \max_p \left\{ \begin{array}{l} \lambda(q_1(1-r_1)\overline{F}_1(p) + q_2(1-r_2)\overline{F}_2(p))(p + V_{t-1}^{NS}(y-1)) \\ +\lambda(q_1(1-r_1)F_1(p) + q_2(1-r_2)F_2(p))V_{t-1}^{NS}(y-1) \\ +\lambda(q_1r_1+q_2r_2)E_S\left[\max_{p_d \le p} \left\{ \begin{array}{l} (\widehat{q}_1(S)\overline{F}_1(p_d) + \widehat{q}_2(S)\overline{F}_2(p_d))(p_d + V_{t-1}^{NS}(y-1)) \\ +(\widehat{q}_1(S)F_1(p_d) + \widehat{q}_2(S)F_2(p_d))V_{t-1}^{NS}(y-1) \end{array} \right\}\right] \\ +(1-\lambda)V_{t-1}^{NS}(y) \end{array} \right\},$$

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

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

In each period, the firm determines a price p that will be charged to all the customers who do not provide a signal. This price is also an upper bound on the price charged to a signaling customer, p_d, which is chosen optimally based on the signal. The first two terms in the summation above capture the revenue-to-go if a non-signaling customer arrives and the third term captures the revenue-to-go if a signaling customer arrives. The last term in the summation is the revenue-to-go if no customer arrives in the current period.

The optimality equation for $y > 0, t = 1, \ldots, T$ can more simply be written as

$$V_t^{NS}(y) = V_{t-1}^{NS}(y) + \lambda \left[\max_p \left\{ \begin{array}{l} (q_1(1-r_1)\overline{F}_1(p) + q_2(1-r_2)\overline{F}_2(p))(p - \Delta_t^{NS}(y)) \\ +(q_1 r_1 + q_2 r_2) E_S \left[\max_{p_d \le p} \left\{ (\widehat{q}_1(S)\overline{F}_1(p_d) + \widehat{q}_2(S)\overline{F}_2(p_d))(p_d - \Delta_t^{NS}(y)) \right\} \right] \end{array} \right\} \right] \tag{8}$$

where

$$\Delta_t^{NS}(y) = V_{t-1}^{NS}(y) - V_{t-1}^{NS}(y-1)$$

The trade-off involved in determining p can be more easily observed from the optimality equation above. The firm would like to charge the price that will maximize the expected revenue from the customers who do not give signals but at the same time would like to make this price as high as possible so that it will have more flexibility in determining the personalized prices for the signaling customers. We should also mention that in fact non-signaling customers do provide a signal to the firm by simply not signaling. Since customers from different segments have different probabilities of providing a signal, the fact that the customer did not signal tells something about the customer's likelihood of belonging to one segment versus the other.

7.1 Pricing decisions

Our objective here is to investigate how changes in the probabilities r_1 and r_2 will reflect on the optimal announced price and the discounts that the signaling customers will be offered. Let the optimal announced price for a given inventory level y and time t be denoted by p^*. Suppose that, at time t, there is a change in the signaling probabilities r_1 and r_2. The following proposition describes how the optimal price p^* is affected by such a change.

Proposition 7 *The optimal announced price increases if:*
(a) the signaling probability for segment 2, r_2 increases, while the signaling probability for segment 1, r_1 remains the same, or
(b) the signaling probabilities are the same across both segments ($r_1 = r_2$) and they both increase by the same amount.

Proposition 7 identifies scenarios in which the direction of change in the optimal price is known. Under the first scenario (Proposition 7(a)), after the changes in signaling probabilities, a higher percentage of non-signaling customers are of segment 1, the segment with lower price sensitivity.

Thus, the firm would like to charge the non-signaling customers a higher price. Since charging a higher announced price also increases the profits from signaling customers (as the firm will be more flexible in making personalized offers), the firm chooses to increase the announced price.

In the other scenario (Proposition 7(b)), the signaling probabilities are the same for both segments before and after the increase in signaling probabilities. Here, the firm chooses to increase the announced price, because higher announced price is always good for the profit that will be obtained from signaling customers and a larger portion of the population now chooses to signal.

Proposition 7 does not say anything about what happens when the percentage of signaling segment 1 customers increases while the percentage of signaling segment 2 customers remains the same. In such a case, the ratio of segment 1 customers among the non-signaling group decreases, which is an incentive for the firm to decrease the price. However, now, more customers provide a signal and, thus, there is also the incentive to increase the price so as to better price discriminate among signaling customers. Thus, it is not clear how the optimal price would change.

To summarize our findings in this section, it appears that changes in the number of customers who provide signals have different effects on the optimal prices depending on who these new signaling customers are. Are they segment 1 customers with high willingness-to-pay or segment 2 customers with low willingness-to-pay? Non-signaling customers might feel that as signals from more customers become available, they will be worse-off if they insist on not signaling. That will probably be the case, if the new signaling customers mostly consist of segment 2 customers. From the signaling customers' point of view, things are a little more complicated. Those who already get a discount continue to pay the same discounted price if the price charged to non-signaling customers increases (in response to changes in the signaling probabilities), but those who are not given a discount might end up paying more. On the other hand, if the announced price drops, some signaling customers who normally would get a discount may end up paying the announced price.

7.2 The firm's profits

In this section, our objective is to investigate the firm's profits under different scenarios each differing only in their signaling probabilities r_1 and r_2. We are mainly interested in identifying cases where the firm's profits are particularly large. With that objective, we have carried out a numerical study observing the effects of changes in the probabilities on the expected profits. Although we have considered several different examples (each assuming different distributions for the signals and different parameters for the reservation price distributions), the general characteristics of the relationship between the expected profits and the signaling probabilities did not appear to change from example to example. For brevity we report our observations on a single example.

The example we consider is the same example considered in Section 6. Here, we compute the optimal expected profit for the initial inventory level of $y = 8$ and a sales horizon of $t = 24$ for different levels of signaling probabilities r_1 and r_2. Figure 7.2 gives the plots of the optimal expected profits for each scenario. The data for the figure is provided in Table 1.

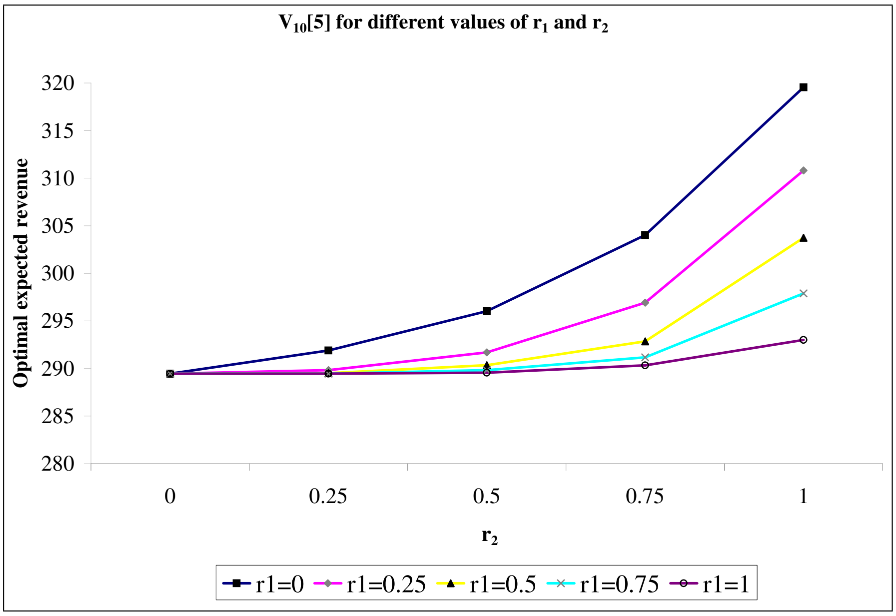

Figure 4: Plot of optimal expected revenues with an initial inventory level of 8 and a sales horizon of 24 under different signaling probabilities. The x-axis is for r_2. Each curve corresponds to the optimal expected revenue for different values of r_1, indicated by different types of markers.

	$\mathbf{r_2 = 0}$	$\mathbf{r_2 = 0.25}$	$\mathbf{r_2 = 0.5}$	$\mathbf{r_2 = 0.75}$	$\mathbf{r_2 = 1}$
$\mathbf{r_1 = 0}$	289.462	291.893	296.021	304.021	319.540
$\mathbf{r_1 = 0.25}$	289.462	289.832	291.697	296.918	310.818
$\mathbf{r_1 = 0.5}$	289.462	289.504	290.343	292.860	303.730
$\mathbf{r_1 = 0.75}$	289.462	289.462	289.826	291.177	297.893
$\mathbf{r_1 = 1.0}$	289.462	289.462	289.554	290.340	293.001

Table 1: Optimal expected revenues with an initial inventory level of 8 and a sales horizon of 24 under different signaling probabilities.

What would the firm want its customers to do? Intuition might suggest that since signals help the firm make more informed decisions, the firm would want all the customers to signal (i.e.,

$r_1 = r_2 = 1$). However, the flaw with this argument is that it ignores the fact that the customers actually give another signal by simply not signaling. In fact, depending on the circumstances this other signal might be much more valuable to the firm than the signal they are hiding. Consider the case where $r_1 = 0$ and $r_2 = 1$. In other words, all segment 2 customers signal while all segment 1 customers hide their signal. This is the perfect scenario for the firm since it knows the segment identity of every customer and therefore can price discriminate across the two segments. In this case, the signals that the customers voluntarily give are in fact useless since just the fact that they signal reveals all the information that the firm can possibly know anyway. Now, suppose that more customers signal. That means all segment 2 customers plus some of the segment 1 customers are signaling. In this case, if a customer does not signal the firm still knows that she is of segment 1, but how about a signaling customer? Now, the message that the firm is getting from the signaling customers is blurred as the firm no longer knows with certainty what segment a signaling customer belongs to. Therefore, there is not a perfect price discrimination opportunity here and the expected profits are lower. This can be observed from Figure 4.

Now, suppose that $r_1 = 1$ and $r_2 = 0$ so that all segment 1 customers signal but none of the segment 2 customers does so. Then, again, the firm can perfectly identify the segment identities of the customers. However, in this case, the optimal expected profit is not as large (in fact it is the scenario with the lowest profits along with other scenarios for which $r_2 = 0$) since the firm cannot price discriminate due to the fact that signaling customers cannot be charged prices that are higher than the price charged to non-signaling customers. Hence, the firm has perfect information but it cannot make any use of it.

The firm would then ideally want all segment 2 customers to signal while all segment 1 customers not to signal. As we can see from Figure 4, the farther away we are from this ideal scenario, the lower are the optimal expected profits. More specifically, profits decrease with r_1 for any fixed value of r_2 and increase with r_2 for any fixed value of r_1.

Our numerical results also indicate that if the signaling probabilities do not differ across the two customer segments, the firm's optimal expected profit increases as more customers volunteer their signals. This can more easily be seen from Table 1. Since changes in signaling probabilities have conflicting effects depending on the customer segment, when signaling probabilities for both segments increase at equal amounts, the increase in the optimal expected profit is very small.

Finally, from Table 1, we observe that the optimal expected profits for several scenarios are the same. (These are all five scenarios for $r_2 = 0$ and the two scenarios with $r_1 = 0.75, r_2 = 0.25$ and

$r_1 = 1, r_2 = 0.25$.) In all of these cases, since most (in some cases all) of the signaling customers are of segment 1, the firm chooses not to offer any discounts and all the customers end up being offered the same price. Consequently, expected optimal profits also turn out to be the same.

8 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.[5] 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 towards 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 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?

[5]Chapter 12 of Phillips (2005) provides an excellent discussion of fairness and customer acceptance issues in relation to pricing practices.

In the 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 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 does not charge more than 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.

Our numerical analysis provides insights on several questions mostly related to the interactions between dynamic and personalized pricing. It appears that the benefit of adding personalization on top of dynamic pricing is higher for higher levels of inventory. Also, personalization seems to be almost equally effective regardless of whether it is used along with static pricing or dynamic pricing. When inventory levels are high relative to the selling horizon, it is known that dynamic pricing does not bring much additional profit over static prices. In such a case, however, our analysis shows that personalizing prices would have significant benefits. On the other hand, for low inventory levels, personalization has marginal benefits, which suggests that the benefit from personalization might not be worth the effort. Finally, we observed that when signals from some of the customers are not available, the unavailability of the signal provides another signal that the firm can utilize in setting prices. In fact, the firm might even prefer some of the customers (those with high willingness-to-pay) not to signal so that it can better price-discriminate among its customers.

Paralleling the companies' increasing interest in price personalization, the recent marketing and economics literature has shown growing interest in the practice. However, work that investigates operational questions related to personalized pricing seems to be lagging behind. For example, more research is needed to better understand the implications of such pricing policies on inventory or production decisions. The signal/reservation price formulation that we propose in this paper might be used in other models that aim to provide insights into some of these operational decisions.

Appendix A - Definitions of Stochastic Orders Used in the Paper

The following definitions are based on Müller and Stoyan (2002) and Shaked and Shanthikumar (2007). Here, we make slight changes such as restricting the definitions to non-negative random variables having the same support.

Suppose that X and Y are two non-negative random variables having the same support with corresponding cumulative distribution functions $F_X(\cdot)$ and $F_Y(\cdot)$, respectively.

Definition 1 ***Usual Stochastic Ordering:*** *Suppose that $F_X(x) \leq F_Y(x)$ for all $x \in (0, \infty)$. Then F_X is said to be greater than F_Y in the usual stochastic order (denoted by $F_X \geq_{st} F_Y$).*

Definition 2 ***Failure Rate Ordering:*** *Suppose that $F_X(\cdot)$ and $F_Y(\cdot)$ are absolutely continuous with failure rate functions $r_X(\cdot)$ and $r_Y(\cdot)$, respectively. If $r_X(x) \leq r_Y(x)$ (or equivalently $\frac{1-F_X(x)}{1-F_Y(x)}$ is increasing) over the common support of X and Y, then we say that F_X is greater than F_Y in failure rate ordering (denoted by $F_X \geq_{fr} F_Y$).*

Definition 3 ***Likelihood Ratio Ordering:*** *Suppose that the following condition holds:*

$$P\{X \in A\}P\{Y \in B\} \leq P\{X \in B\}P\{Y \in A\} \tag{9}$$

for all measurable sets A and B in $\mathbb{R}^+$ such that $A \leq B$, where $A \leq B$ means that $x \in A$ and $y \in B$ implies that $x \leq y$. Then, F_X is said to be greater than F_Y in the likelihood ratio ordering (denoted by $F_X \geq_{lr} F_Y$).

If X and Y are continuous random variables, and $f_X(\cdot)$ and $f_Y(\cdot)$ are the corresponding probability density functions, Condition (9) is equivalent to $\frac{f_X(x)}{f_Y(x)}$ being increasing over the common support of X and Y.

If X and Y are discrete random variables, and $f_X(\cdot)$ and $f_Y(\cdot)$ are the corresponding probability mass functions, Condition (9) is equivalent to $\frac{f_X(x)}{f_Y(x)}$ being increasing over the common support of X and Y.

Appendix B - Proofs of the Results

Proof of Proposition 1: S and R are positively correlated if $Cov(S, R) > 0$.

$$\begin{aligned} Cov(S, R) &= E[SR] - E[S]E[R] \\ &= (q_1 s_1 \mu_1 + q_2 s_2 \mu_2) - (q_1 s_1 + q_2 s_2)(q_1 \mu_1 + q_2 \mu_2) \end{aligned}$$

After some algebra, one can show that

$$Cov(S,R) \;=\; q_1 q_2(\mu_1-\mu_2)(s_1-s_2)$$

The result follows from the equality above. □

Proof of Theorem 1 and Propositions 2 and 3: For the purposes of this proof, define $\widehat{\Pi}(x,p,\Delta) := \alpha(x,p)(p-\Delta)$. Let $\widehat{p}(x,\Delta)$ denote the smallest optimizer of $\widehat{\Pi}(x,p,\Delta)$. Using this definition, note that $p^*(x,y,t)$ is given by $\widehat{p}(x,\Delta_t(y))$. In addition, define $\overline{\Pi}_i(p,\Delta) := (p-\Delta)\overline{F}_i(p)$, $i=1,2$. Let $\overline{p}_i(\Delta)$ denote the optimizer of $\overline{\Pi}_i(p,\Delta)$.

We now apply Lemma 1 to prove Theorem 1. By Lemma 1(a), $\widehat{p}(x,\Delta_t(y)) \in [\overline{p}_2(\Delta_t(y)), \overline{p}_1(\Delta_t(y))]$. By Lemma 1(b), $\Pi(x,p,\Delta_t(y))$ is supermodular in p and x for $p \in [\overline{p}_2(\Delta_t(y)), \overline{p}_1(\Delta_t(y))]$ and $x \geq 0$. Therefore, $\widehat{p}(x,\Delta_t(y))$ is increasing in x, which concludes the proof of Theorem 1.

Likewise, by Lemma 1(c), $\Pi(x,p,\Delta_t(y))$ is supermodular in p and $\Delta_t(y)$ for $p \geq 0$ and $\Delta_t(y) \geq 0$. Therefore, $\widehat{p}(x,\Delta_t(y))$ is increasing in $\Delta_t(y)$. Proposition 2(a),(b) now follow from the fact that $\Delta_t(y)$ is decreasing in y and increasing in t (by Lemma 5(a)).

Proposition 3(a) follows from Lemma 1(a). The rest of the proposition follows directly from part (a) of the proposition. □

Proof of Theorem 2: Given time t, inventory level y, and the customer signal x, it is optimal to charge p_1 if

$$\alpha(x,p_1)\left(p_1-\Delta_t(y)\right) \geq \alpha(x,p_2)\left(p_2-\Delta_t(y)\right)$$

where $\Delta_t(y)$ is the marginal value of the inventory. This condition can be equivalently written as

$$\Delta_t(y) \geq \frac{p_2\alpha(x,p_2)-p_1\alpha(x,p_1)}{\alpha(x,p_2)-\alpha(x,p_1)}.$$

Now let

$$A(x) = \frac{p_2\alpha(x,p_2)-p_1\alpha(x,p_1)}{\alpha(x,p_2)-\alpha(x,p_1)}.$$

Then, it is sufficient to show that $A(x_1) \leq A(x_2)$ for $x_1 > x_2$. After some algebra, one can show that $A(x_1)-A(x_2)$ is equivalent in sign to

$$(p_1-p_2)\left(\alpha(x_1,p_2)\alpha(x_2,p_1)-\alpha(x_1,p_1)\alpha(x_2,p_2)\right).$$

Since $p_1 > p_2$, the sign of $B(x) := \alpha(x_1, p_2)\alpha(x_2, p_1) - \alpha(x_1, p_1)\alpha(x_2, p_2)$ determines the sign of $A(x_1) - A(x_2)$. Substituting for $\alpha(x, p)$ from (2), we find:

$$B(x) = -\left(\overline{F}_1(p_1)\overline{F}_2(p_2) - \overline{F}_1(p_2)\overline{F}_2(p_1)\right)\left(\widehat{q}_1(x_1)\widehat{q}_2(x_2) - \widehat{q}_1(x_2)\widehat{q}_2(x_1)\right).$$

Now, $\overline{F}_1(p_1)\overline{F}_2(p_2) - \overline{F}_1(p_2)\overline{F}_2(p_1) > 0$ follows from the facts that $F_1 >_{fr} F_2$ and the definition of failure rate ordering given in Appendix A. Therefore, we will conclude the proof if we can show that $\widehat{q}_1(x_1)\widehat{q}_2(x_2) - \widehat{q}_1(x_2)\widehat{q}_2(x_1)$ is positive. Substituting for $\widehat{q}_i(x)$ from (1) and after some algebra, one can verify that $\widehat{q}_1(x_1)\widehat{q}_2(x_2) - \widehat{q}_1(x_2)\widehat{q}_2(x_1)$ is equivalent in sign to

$$q_1 q_2 \left(g_1(x_1)g_2(x_2) - g_1(x_2)g_2(x_1)\right)$$

The term in the parenthesis above is positive, due to the assumption that $G_1 \geq_{lr} G_2$ and the definition of likelihood ratio ordering (see Appendix A). This concludes the proof. $\square$

Proof of Proposition 4: Note that $z^*(y,t)$ is given by the (smallest) value of z that maximizes the function $\Pi(z, p_1, p_2, \Delta_t^{FP}(y))$ defined in the statement of Lemma 2. By Lemma 2, $\Pi(z, p_1, p_2, \Delta_t^{FP}(y))$ is submodular in z and $\Delta_t^{FP}(y)$. Therefore, $z^*(y,t)$ is decreasing in $\Delta_t^{FP}(y)$. The result now follows from the fact that $\Delta_t^{FP}(y)$ is decreasing in y and increasing in t (by Lemma 5(c)). $\square$

Proof of Propositions 5 and 6: For the purposes of this proof, define $\widetilde{\Pi}_j(z, p, \Delta) := \beta_j(z, p)(p - \Delta)$. Let $\widetilde{p}_j(z, \Delta)$ denote the smallest optimizer of $\widetilde{\Pi}_j(z, p, \Delta)$. Note that, with this definition, $p_j^*(z, y, t)$ is given by $\widetilde{p}_j(z, \Delta_t^{FT}(y, z))$. In addition, define $\overline{\Pi}_i(p, \Delta) := (p - \Delta)\overline{F}_i(p)$, $i = 1, 2$. Let $\overline{p}_i(\Delta)$ denote the optimizer of $\overline{\Pi}_i(p, \Delta)$. We can now utilize Lemma 3 to prove the propositions.

The proof of Proposition 5(a) is by contradiction. Suppose $\widetilde{p}_1(z, \Delta_t^{FT}(y, z)) < \widetilde{p}_2(z, \Delta_t^{FT}(y, z))$. For the sake of exposition, we suppress the arguments of $\widetilde{p}_1$ and $\widetilde{p}_2$ in what follows. By definition of $\widetilde{p}_j$, we have

$$\widetilde{\Pi}_2(z, \widetilde{p}_2, \Delta_t^{FT}(y, z)) > \widetilde{\Pi}_2(z, \widetilde{p}_1, \Delta_t^{FT}(y, z)).$$

In addition, due to the assumption that $\widetilde{p}_1 < \widetilde{p}_2$, it follows from Lemma 3(a),(b) that

$$\widetilde{\Pi}_1(z, \widetilde{p}_2, \Delta_t^{FT}(y, z)) - \widetilde{\Pi}_2(z, \widetilde{p}_2, \Delta_t^{FT}(y, z)) \geq \widetilde{\Pi}_1(z, \widetilde{p}_1, \Delta_t^{FT}(y, z)) - \widetilde{\Pi}_2(z, \widetilde{p}_1, \Delta_t^{FT}(y, z)).$$

The last two inequalities together imply that

$$\widetilde{\Pi}_1(z, \widetilde{p}_2, \Delta_t^{FT}(y,z)) > \widetilde{\Pi}_1(z, \widetilde{p}_1, \Delta_t^{FT}(y,z)),$$

which yields a contradiction to the optimality of $\widetilde{p}_1$ for $\widetilde{\Pi}_1$, thus concluding the proof.

We next prove Proposition 5(b). By Lemma 3(a), $\widetilde{p}_j(z, \Delta_t^{FT}(y,z)) \in [\overline{p}_2(\Delta_t^{FT}(y,z)), \overline{p}_1(\Delta_t^{FT}(y,z))]$. By Lemma 3(c), for a given $\Delta > 0$, $\widetilde{\Pi}_j(z,p,\Delta)$ is supermodular in p and z for $p \in [\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$ and $z \geq 0$. Note that when z changes in period t only, $\Delta_t^{FT}(y,z)$ is not affected (since $\Delta_t^{FT}(y,z) = V_{t-1}^{FT}(y,z) - V_{t-1}^{FT}(y-1,z)$). Therefore, we can apply Lemma 3(c) to conclude that $\widetilde{p}_j(z, \Delta_t^{FT}(y,z))$ increases if z increases in period t only.

Likewise, by Lemma 3(d), $\Pi(z, p, \Delta_t^{FT}(y))$ is supermodular in p and $\Delta_t^{FT}(y)$ for $p \geq 0$ and $\Delta_t^{FT}(y) \geq 0$. Therefore, $\widetilde{p}(z, \Delta_t^{FT}(y))$ is increasing in $\Delta_t^{FT}(y)$. Proposition 6(a), (b) now follow from the fact that $\Delta_t^{FT}(y)$ is decreasing in y and increasing in t (by Lemma 5(b)). □

Proof of Proposition 7:

Proof of (a): To simplify the presentation, we first define the following two functions:

$$\Gamma_1(p, r_2) = \left(q_1(1-r_1)\overline{F}_1(p) + q_2(1-r_2)\overline{F}_2(p)\right)(p - \Delta)$$

and

$$\Gamma_2(p, r_2) = (q_1 r_1 + q_2 r_2)\Theta(p)$$

where

$$\Theta(p) = E_S\left[\max_{p_d \leq p}\left\{(\widehat{q}_1(S)\overline{F}_1(p_d) + \widehat{q}_2(S)\overline{F}_2(p_d))(p_d - \Delta)\right\}\right].$$

and $\Delta = \Delta_t^{NS}(y)$. Now, for a given r_2, notice that the optimization problem in (8) can be written as

$$\max_p\{\Gamma_1(p, r_2) + \Gamma_2(p, r_2)\}.$$

For any value of r_2, using arguments similar to those in Lemma 1(a), it can be shown that the optimal value of p resides in the interval $[\bar{p}_2(\Delta), \bar{p}_1(\Delta)]$. We will next show that both $\Gamma_1(\cdot)$ and $\Gamma_2(\cdot)$ are supermodular in p and r_2 for $p \in [\bar{p}_2(\Delta), \bar{p}_1(\Delta)]$. That will be sufficient to conclude that the optimal value of p is non-decreasing in r_2, since the summation of two supermodular functions is also supermodular.

To show the supermodularity of $\Gamma_1(p, r_2)$, we note that

$$\frac{\partial^2 \Gamma_1(p, r_2)}{\partial p \partial r_2} = f_2(p)(p - \Delta) - (1 - F_2(p)) \geq 0$$

where the inequality follows from the fact that $p \geq \bar{p}_2(\Delta)$. Thus, $\Gamma_1(p, r_2)$ is supermodular for $p \in [\bar{p}_2(\Delta), \bar{p}_1(\Delta)]$.

Now, in order to establish the supermodularity of $\Gamma_2(p, r_2)$, it is sufficient to show that $\Gamma_2(p, r_2)$ has increasing differences. Let $\bar{p}_1(\Delta) \geq p_1 \geq p_2 \geq \bar{p}_2(\Delta)$ and $1 \geq r_2^1 \geq r_2^2 \geq 0$. Then, we need to show that

$$\Gamma_2(p_1, r_2^1) - \Gamma_2(p_1, r_2^2) \geq \Gamma_2(p_2, r_2^1) - \Gamma_2(p_2, r_2^2).$$

After some algebra, one can show that the above inequality is equivalent to $(r_2^2 - r_2^1)(\Theta(p_2) - \Theta(p_1)) \geq 0$, which holds because $\Theta(p_1) \geq \Theta(p_2)$ and $r_2^1 \geq r_2^2$. Hence, $\Gamma_2(p, r_2)$ is supermodular, concluding the proof.

Proof of (b): The optimization problem in (8) can be written as

$$V_t(y) = V_{t-1}(y) + \lambda \max_p \{(1 - r)\Lambda_1(p) + r\Lambda_2(p)\}$$

where

$$\Lambda_1(p) = (q_1\overline{F}_1(p) + q_2\overline{F}_2(p))(p - \Delta)$$

and

$$\Lambda_2(p) = E_S \left[\max_{p_d \leq p} \left\{ (\widehat{q}_1(S)\overline{F}_1(p_d) + \widehat{q}_2(S)\overline{F}_2(p_d))(p_d - \Delta) \right\} \right].$$

Consider two scenarios: One in which $r_1 = r_2 = r$ with the corresponding optimal price p^* (or the smallest optimizer if there is more than one) and the other in which $r_1 = r_2 = \widehat{r} > r$ with the corresponding optimal price $\widehat{p}^*$ (again the smallest optimizer if there is more than one). Suppose for contradiction that $\widehat{r} > r$ but $p^* > \widehat{p}^*$. Then,

$$\Lambda_2(p^*) \geq \Lambda_2(\widehat{p}^*). \tag{10}$$

By optimality of p^* for $r_1 = r_2 = r$, we have

$$(1 - r)\Lambda_1(p^*) + r\Lambda_2(p^*) > (1 - r)\Lambda_1(\widehat{p}^*) + r\Lambda_2(\widehat{p}^*),$$

which can also be written as

$$\Lambda_1(p^*) - \Lambda_1(\widehat{p}^*) > \frac{r}{1 - r}(\Lambda_2(\widehat{p}^*) - \Lambda_2(p^*)). \tag{11}$$

In addition, by optimality of $\widehat{p}^*$ for $r_1 = r_2 = \widehat{r}$, we have

$$(1 - \widehat{r})\Lambda_1(p^*) + \widehat{r}\Lambda_2(p^*) \leq (1 - \widehat{r})\Lambda_1(\widehat{p}^*) + \widehat{r}\Lambda_2(\widehat{p}^*),$$

which can also be written as

$$\Lambda_1(p^*) - \Lambda_1(\widehat{p}^*) \leq \frac{\widehat{r}}{1-\widehat{r}}(\Lambda_2(\widehat{p}^*) - \Lambda_2(p^*)). \tag{12}$$

But, (11) together with (12) is a contradiction to (10). Thus, we must have $p^* \leq \widehat{p}^*$. □

Lemma 1 *Define* $\widehat{\Pi}(x,p,\Delta) := \alpha(x,p)(p-\Delta)$. *Let* $\widehat{p}(x,\Delta)$ *denote the smallest optimizer of* $\widehat{\Pi}(x,p,\Delta)$. *Define* $\overline{\Pi}_i(p,\Delta) := (p-\Delta)\overline{F}_i(p)$, $i=1,2$. *Let* $\overline{p}_i(\Delta)$ *denote the optimizer of* $\overline{\Pi}_i(p,\Delta)$. *Then:*

(a) (Aydin and Ziya, 2007) $\overline{p}_1(\Delta) \geq \widehat{p}(x,\Delta) \geq \overline{p}_2(\Delta)$.

(b) At a fixed $\Delta \geq 0$, $\widehat{\Pi}(x,p,\Delta)$ *is supermodular in* x *and* p *for* $x \geq 0$ *and* $p \in [\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$.

(c) At a fixed $x \geq 0$, $\widehat{\Pi}(x,p,\Delta)$ *is supermodular in* Δ *and* p *for* $\Delta \geq 0$ *and* $p \geq 0$.

Proof of Lemma 1:

Proof of (a): The proof of this part is similar to Lemma 1 in the online supplement to Aydin and Ziya (2007). We include the proof here for the sake of completeness. It is not difficult to show that $\overline{\Pi}_i(p,\Delta)$ is strictly unimodal in p due to assumption $(A3)$. (See, for example, Lariviere and Porteus, 2001.) Since $\overline{\Pi}_i(p,\Delta)$ is strictly unimodal in p, $\widehat{p}_i(\Delta)$ must satisfy the following first-order condition (FOC):

$$\overline{F}_i(p) - (p-\Delta)f_i(p) = \overline{F}_i(p)\left(1 - (p-\Delta)\frac{f_i(p)}{\overline{F}_i(p)}\right) = 0, i = 1,2. \tag{13}$$

Now the FOCs in (13) along with assumption $(A4)$ imply that $\overline{p}_1(\Delta) \geq \overline{p}_2(\Delta)$. To prove that $\widehat{p}(x,\Delta) \in [\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$, we first observe that $\widehat{\Pi}(x,p,\Delta)$ can be written as

$$\widehat{\Pi}(x,p,\Delta) = \widehat{q}_1(x)\overline{\Pi}_1(p,\Delta) + \widehat{q}_2(x)\overline{\Pi}_2(p,\Delta)$$

Now, since $\overline{\Pi}_i(p,\Delta)$, $i=1,2$ are unimodal in p and $\overline{p}_1(\Delta) \geq \overline{p}_2(\Delta)$, it follows that $\widehat{\Pi}(x,p,\Delta)$ is increasing in p for $p \leq \overline{p}_2(\Delta)$ and decreasing in p for $p \geq \overline{p}_1(\Delta)$. Therefore, the optimizer of $\widehat{\Pi}(x,p,\Delta)$, denoted by $\widehat{p}(x,\Delta)$, must be in $[\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$.

Proof of (b): First, note that $\widehat{\Pi}(x,p,\Delta)$ being supermodular in x and p is equivalent to $\widehat{\Pi}(x,p,\Delta)$ having increasing differences in x and p. Therefore, we will prove that, for $x_1 > x_2$, the difference $\widehat{\Pi}(x_1,p,\Delta) - \widehat{\Pi}(x_2,p,\Delta)$ is increasing in p.

$$\frac{\partial\left(\widehat{\Pi}(x_1,p,\Delta) - \widehat{\Pi}(x_2,p,\Delta)\right)}{\partial p} = (p-\Delta)\left(\frac{\partial\alpha(x_1,p)}{\partial p} - \frac{\partial\alpha(x_2,p)}{\partial p}\right) + \alpha(x_1,p) - \alpha(x_2,p) \tag{14}$$

where

$$\begin{aligned}
\alpha(x,p) &= \widehat{q}_1(x)\overline{F}_1(p) + \widehat{q}_2(x)\overline{F}_2(p) && (15)\\
\frac{\partial\alpha(x,p)}{\partial p} &= -\widehat{q}_1(x)f_1(p) - \widehat{q}_2(x)f_2(p) && (16)
\end{aligned}$$

Substituting from (15) and (16) in (14) and rearranging the terms, we obtain:

$$\begin{aligned}
\frac{\partial\left(\widehat{\Pi}(x_1,p,\Delta)-\widehat{\Pi}(x_2,p,\Delta)\right)}{\partial p} &= \left[-(p-\Delta)f_1(p)+\overline{F}_1(p)\right]\left(\widehat{q}_1(x_1)-\widehat{q}_1(x_2)\right)\\
&+ \left[-(p-\Delta)f_2(p)+\overline{F}_2(p)\right]\left(\widehat{q}_2(x_1)-\widehat{q}_2(x_2)\right)
\end{aligned}$$

Because $\widehat{q}_1(x)+\widehat{q}_2(x)=1$ for any x, we have $\widehat{q}_1(x_1)-\widehat{q}_1(x_2)=-\widehat{q}_2(x_1)+\widehat{q}_2(x_2)$. Using this observation, we can write the above equality as:

$$\begin{aligned}
\frac{\partial\left(\widehat{\Pi}(x_1,p,\Delta)-\widehat{\Pi}(x_2,p,\Delta)\right)}{\partial p} &= \left(\widehat{q}_1(x_1)-\widehat{q}_1(x_2)\right)\\
&\times \left\{\left[-(p-\Delta)f_1(p)+\overline{F}_1(p)\right]-\left[-(p-\Delta)f_2(p)+\overline{F}_2(p)\right]\right\} && (17)
\end{aligned}$$

Observe that

$$\overline{F}_1(p)-(p-\Delta)f_1(p)=\frac{\partial\overline{\Pi}_1(p,\Delta)}{\partial p}\geq 0 \text{ for } p\leq\overline{p}_1(\Delta),$$

and

$$\overline{F}_2(p)-(p-\Delta)f_2(p)=\frac{\partial\overline{\Pi}_2(p,\Delta)}{\partial p}\leq 0 \text{ for } p\geq\overline{p}_2(\Delta).$$

Therefore, the term in curly brackets in (17) is positive, and we will conclude the proof if we can show that $\widehat{q}_1(x_1)-\widehat{q}_1(x_2)\geq 0$. Now:

$$\widehat{q}_1(x_1)-\widehat{q}_1(x_2)=\frac{q_1q_2\left(g_1(x_1)g_2(x_2)-g_1(x_2)g_2(x_1)\right)}{\left(q_1g_1(x_1)+q_2g_2(x_1)\right)\left(q_1g_1(x_2)+q_2g_2(x_2)\right)}\geq 0$$

where the inequality follows from the fact that $\frac{g_1(x_1)}{g_2(x_1)}\geq\frac{g_1(x_2)}{g_2(x_2)}$, because $G_1(\cdot)\geq_{lr}G_2(\cdot)$ by assumption $(A5')$.

Proof of (c): First, observe that

$$\frac{\partial^2\widehat{\Pi}(x,p,\Delta)}{\partial\Delta\partial p}=-\frac{\partial\alpha(x,p)}{\partial p} \tag{18}$$

where

$$\frac{\partial\alpha(x,p)}{\partial p} = -\widehat{q}_1(x)f_1(p)-\widehat{q}_2(x)f_2(p)<0. \tag{19}$$

The result follows directly from (18) and (19). □

Lemma 2 *Given $p_1 > p_2$, define*

$$\begin{aligned}\Pi(z,p_1,p_2,\Delta) \ &:= \ \left(q_1\overline{G}_1(z)\overline{F}_1(p_1)+q_2\overline{G}_2(z)\overline{F}_2(p_1)\right)(p_1-\Delta)\\ &+ \ \left(q_1G_1(z)\overline{F}_1(p_2)+q_2G_2(z)\overline{F}_2(p_2)\right)(p_2-\Delta).\end{aligned}$$

Then, $\Pi(z,p_1,p_2,\Delta)$ is submodular in Δ and z for $\Delta \geq 0$ and $z \geq 0$.

Proof of Lemma 2: First, note that $\Pi(z,p_1,p_2,\Delta)$ being submodular in z and Δ is equivalent to $\Pi(z,p_1,p_2,\Delta)$ having decreasing differences in z and Δ. Therefore, we will prove that, for $z_1 > z_2$, the difference $\Pi(z_1,p_1,p_2,\Delta) - \Pi(z_2,p_1,p_2,\Delta)$ is increasing in Δ.

$$\begin{aligned}&\frac{\partial\left(\Pi(z_1,p_1,p_2,\Delta)-\Pi(z_2,p_1,p_2,\Delta)\right)}{\partial\Delta} =\\ - \ &\left(q_1\overline{G}_1(z_1)\overline{F}_1(p_1)+q_2\overline{G}_2(z_1)\overline{F}_2(p_1)\right)-\left(q_1G_1(z_1)\overline{F}_1(p_2)+q_2G_2(z_1)\overline{F}_2(p_2)\right)\\ + \ &\left(q_1\overline{G}_1(z_2)\overline{F}_1(p_1)+q_2\overline{G}_2(z_2)\overline{F}_2(p_1)\right)+\left(q_1G_1(z_2)\overline{F}_1(p_2)+q_2G_2(z_2)\overline{F}_2(p_2)\right)\end{aligned}$$

Noticing that $\overline{G}_i(z_1) - \overline{G}_i(z_2) = G_i(z_2) - G_i(z_1)$ and rearranging the terms:

$$\begin{aligned}&\frac{\partial\left(\Pi(z_1,p_1,p_2,\Delta)-\Pi(z_2,p_1,p_2,\Delta)\right)}{\partial\Delta} =\\ &q_1\left(G_1(z_2)-G_1(z_1)\right)\left(-\overline{F}_1(p_1)+\overline{F}_1(p_2)\right)+q_2\left(G_2(z_2)-G_2(z_1)\right)\left(-\overline{F}_2(p_1)+\overline{F}_2(p_2)\right)\end{aligned} \tag{20}$$

Observe that the right-hand side of the above equality is negative since $z_1 > z_2$ and $p_1 > p_2$, which concludes the proof. □

Lemma 3 *Define $\widetilde{\Pi}_j(z,p,\Delta) := \beta_j(z,p)(p-\Delta)$, $j = 1,2$. Let $\widetilde{p}_j(z,\Delta)$ denote the smallest optimizer of $\widetilde{\Pi}_j(z,p,\Delta)$. Define $\overline{\Pi}_i(p,\Delta) := (p-\Delta)\overline{F}_i(p)$, $i = 1,2$. Let $\overline{p}_i(\Delta)$ denote the optimizer of $\overline{\Pi}_i(p)$. Then:*

(a) $\overline{p}_1(\Delta) \geq \widetilde{p}_j(z,\Delta) \geq \overline{p}_2(\Delta)$.

(b) At a fixed $\Delta \geq 0$ and $z \geq 0$, $\widetilde{\Pi}_1(z,p,\Delta) - \widetilde{\Pi}_2(z,p,\Delta)$ is increasing in p for $p \in [\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$.

(c) At a fixed $\Delta \geq 0$, $\widetilde{\Pi}_j(z,p,\Delta)$ is supermodular in z and p for $z \geq 0$ and $p \in [\overline{p}_2(\Delta), \overline{p}_1(\Delta)]$.

(d) At a fixed $z \geq 0$, $\widetilde{\Pi}_j(z,p,\Delta)$ is supermodular in Δ and p for $\Delta \geq 0$ and $p \geq 0$.

Proof of Lemma 3:

Proof of (a): We can write

$$\widetilde{\Pi}_j(z,p,\Delta) = \widetilde{q}_{1j}(z)\overline{\Pi}_1(p,\Delta) + \widetilde{q}_{2j}(z)\overline{\Pi}_2(p,\Delta), j = 1,2$$

We already showed in Lemma 1(a) that $\overline{\Pi}_i(p,\Delta)$, $i=1,2$ are unimodal in p and $\overline{p}_1(\Delta)\geq\overline{p}_2(\Delta)$. Therefore, $\widetilde{\Pi}_j(z,p,\Delta)$ is increasing in p for $p\leq\overline{p}_2(\Delta)$ and decreasing in p for $p\geq\overline{p}_1(\Delta)$. Hence, the optimizer of $\widetilde{\Pi}_j(z,p,\Delta)$, denoted by $\widetilde{p}_j(z,\Delta)$, must be in $[\overline{p}_2(\Delta),\overline{p}_1(\Delta)]$.

Proof of (b): Note that

$$\frac{\partial(\widetilde{\Pi}_1(z,p,\Delta)-\widetilde{\Pi}_2(z,p,\Delta))}{\partial p}=(\beta_1(z,p)-\beta_2(z,p))+\left(\frac{\partial\beta_1(z,p)}{\partial p}-\frac{\partial\beta_2(z,p)}{\partial p}\right)(p-\Delta)$$

where

$$\begin{aligned}\beta_j(z,p) &= \widetilde{q}_{1j}(z)\overline{F}_1(p)+\widetilde{q}_{2j}(z)\overline{F}_2(p) && (21)\\ \frac{\partial\beta_j(z,p)}{\partial p} &= -\widetilde{q}_{1j}(z)f_1(p)-\widetilde{q}_{2j}(z)f_2(p) && (22)\end{aligned}$$

Therefore:

$$\begin{aligned}\frac{\partial(\widetilde{\Pi}_1(z,p,\Delta)-\widetilde{\Pi}_2(z,p,\Delta))}{\partial p} &= (\widetilde{q}_{11}(z)-\widetilde{q}_{12}(z))\left[\overline{F}_1(p)-(p-\Delta)f_1(p)\right]\\ &+ (\widetilde{q}_{21}(z)-\widetilde{q}_{22}(z))\left[\overline{F}_2(p)-(p-\Delta)f_2(p)\right]\\ &= (\widetilde{q}_{11}(z)-\widetilde{q}_{12}(z))\\ &\times \left\{\left[\overline{F}_1(p)-(p-\Delta)f_1(p)\right]-\left[\overline{F}_2(p)-(p-\Delta)f_2(p)\right]\right\}\end{aligned}$$

where the last equality follows from $\widetilde{q}_{1j}(z)+\widetilde{q}_{2j}(z)=1$. Observe that

$$\overline{F}_1(p)-(p-\Delta)f_1(p)=\frac{\partial\overline{\Pi}_1(p,\Delta)}{\partial p}\geq 0 \text{ for } p\leq\overline{p}_1(\Delta),$$

and

$$\overline{F}_2(p)-(p-\Delta)f_2(p)=\frac{\partial\overline{\Pi}_2(p,\Delta)}{\partial p}\leq 0 \text{ for } p\geq\overline{p}_2(\Delta).$$

Therefore, we will conclude the proof if we can show that $\widetilde{q}_{11}-\widetilde{q}_{12}\geq 0$. Now:

$$\widetilde{q}_{11}(z)-\widetilde{q}_{12}(z)=\frac{q_1q_2(G_2(z)-G_1(z))}{(q_1\overline{G}_1(z)+q_2\overline{G}_2(z))(q_1G_1(z)+q_2G_2(z))}\geq 0$$

where the inequality follows from the fact that $G_1(\cdot)<G_2(\cdot)$, because $G_1\geq_{lr}G_2$ by assumption $(A5')$.

Proof of (c): First, note that $\widetilde{\Pi}_j(z,p,\Delta)$ being supermodular in z and p is equivalent to $\widetilde{\Pi}_j(z,p,\Delta)$ having increasing differences in z and p. Therefore, we will prove that, for $z_1>z_2$, the difference $\widetilde{\Pi}_j(z_1,p,\Delta)-\widetilde{\Pi}_j(z_2,p,\Delta)$ is increasing in p. Note that

$$\frac{\partial\left(\widetilde{\Pi}_j(z_1,p,\Delta)-\widetilde{\Pi}_j(z_2,p,\Delta)\right)}{\partial p}=(p-\Delta)\left(\frac{\partial\beta_j(z_1,p)}{\partial p}-\frac{\partial\beta_j(z_2,p)}{\partial p}\right)+\beta_j(z_1,p)-\beta_j(z_2,p) \quad (23)$$

Substituting from (21) and (22) in (23) and rearranging the terms, we obtain:

$$\frac{\partial\left(\widetilde{\Pi}_j(z_1,p,\Delta)-\widetilde{\Pi}_j(z_2,p,\Delta)\right)}{\partial p} = \left[-(p-\Delta)f_1(p)+\overline{F}_1(p)\right](\widetilde{q}_{1j}(z_1,p)-\widetilde{q}_{1j}(z_2,p)) + \left[-(p-\Delta)f_2(p)+\overline{F}_2(p)\right](\widetilde{q}_{2j}(z_1,p)-\widetilde{q}_{2j}(z_2,p))$$

As before, observe that $\overline{F}_1(p)-(p-\Delta)f_1(p)\geq 0$ for $p\leq \overline{p}_1(\Delta)$, and $\overline{F}_2(p)-(p-\Delta)f_2(p)\leq 0$ for $p\geq \overline{p}_2(\Delta)$. Furthermore, note that $\widetilde{q}_{1j}(z_1,p)\geq \widetilde{q}_{1j}(z_2,p)$ for $j=1,2$ and $\widetilde{q}_{2j}(z_1,p)\leq \widetilde{q}_{2j}(z_2,p)$ for $j=1,2$ (by Lemma 4). Therefore,the right-hand side of the above equality is non-negative, which concludes the proof.

Proof of (d): First, observe that

$$\frac{\partial^2\widetilde{\Pi}_j(z,p,\Delta)}{\partial\Delta\partial p} = -\frac{\partial\beta_j(z,p)}{\partial p} \tag{24}$$

where

$$\frac{\partial\beta_j(z,p)}{\partial p} = -\widetilde{q}_{1j}(z)f_1(p)-\widetilde{q}_{2j}(z)f_2(p)<0. \tag{25}$$

The result follows directly from (24) and (25). □

Lemma 4 *For $z_1>z_2$, we have $\widetilde{q}_{1j}(z_1)\geq\widetilde{q}_{1j}(z_2)$ and $\widetilde{q}_{2j}(z_1)\leq\widetilde{q}_{2j}(z_2)$, where $\widetilde{q}_{ij}(z)$ is as defined by (6).*

Proof of Lemma 4: Note from (6) that

$$\widetilde{q}_{1j}(z_1)-\widetilde{q}_{1j}(z_2)=\frac{q_1q_2\left(\overline{G}_1(z_1)\overline{G}_2(z_2)-\overline{G}_2(z_1)\overline{G}_1(z_2)\right)}{\left(q_1\overline{G}_1(z_1)+q_2\overline{G}_2(z_1)\right)\left(q_1\overline{G}_1(z_2)+q_2\overline{G}_2(z_2)\right)}$$

Therefore, we will conclude the proof if we can show that $\overline{G}_1(z_1)\overline{G}_2(z_2)-\overline{G}_2(z_1)\overline{G}_1(z_2)\geq 0$. Now:

$$\overline{G}_1(z_1)\overline{G}_2(z_2)-\overline{G}_2(z_1)\overline{G}_1(z_2)=\overline{G}_1(z_1)(\overline{G}_2(z_2)-\overline{G}_2(z_1))-\overline{G}_2(z_1)(\overline{G}_1(z_2)-\overline{G}_1(z_1))$$

To conclude the proof, we note that S_1 with cdf G_1 dominates S_2 with cdf G_2 in likelihood ratio ordering. We then apply the definition of likelihood ratio ordering given by (9) in Appendix A, by setting $A=\{z:z_2\leq z\leq z_1\}$ and $B=\{z:z\geq z_1\}$. The proof of $\widetilde{q}_{2j}(z_1)\leq\widetilde{q}_{2j}(z_2)$ follows from symmetric arguments. □

Lemma 5 *For $\Delta_t(y)$, $\Delta_t^{FT}(y,z)$ and $\Delta_t^{FP}(y)$ as defined by, respectively, (3), (7) and (5), we have:*
(a) $\Delta_t(y) \geq \Delta_t(y+1)$ *and* $\Delta_{t+1}(y) \geq \Delta_t(y)$ *for* $t \geq 1$ *and* $y \geq 1$.
(b) $\Delta_t^{FT}(y,z) \geq \Delta_t^{FT}(y+1,z)$ *and* $\Delta_{t+1}^{FT}(y,z) \geq \Delta_t^{FT}(y,z)$ *for* $t \geq 1$ *and* $y \geq 1$.
(c) $\Delta_t^{FP}(y) \geq \Delta_t^{FP}(y+1)$ *and* $\Delta_{t+1}^{FP}(y) \geq \Delta_t^{FP}(y)$ *for* $t \geq 1$ *and* $y \geq 1$.

Proof of Lemma 5: The results follows by slight modifications to the proof of Lemma 2 in Bitran and Mondschein (1993). We prove only part (a) here. The proofs of the other two parts are similar. Following Bitran and Mondschein (1993), we use induction on $t+y$ to prove the following inequalities hold:

$$I1(t,y) : V_{t+1}(y+1) - V_{t+1}(y) \geq V_t(y+1) - V_t(y), y \geq 0, t = 1, \ldots, T.$$
$$I2(t,y) : V_{t+1}(y) - V_t(y) \geq V_{t+2}(y) - V_{t+1}(y), y \geq 0, t = 1, \ldots, T.$$
$$I3(t,y) : V_t(y+1) - V_t(y) \geq V_t(y+2) - V_t(y+1), y \geq 0, t = 1, \ldots, T.$$

Note that $I1(t,y)$ and $I3(t,y)$ correspond to the properties stated in part (a) of the lemma. The three inequalities hold when $t+y=0$. Suppose they hold for $t+y=m-1$. We prove they hold when $t+y=m$ to complete the induction.
(i) $I1(t,y): \; V_{t+1}(y+1) - V_{t+1}(y) \geq V_t(y+1) - V_t(y)$.
It is easy to show that the result holds at $y=0$ for all $t=1,\ldots,T$. Suppose $y>0$. Recall that $p^*(S,y,t)$ is the optimal price that a seller with y units at time t will charge to a customer with signal S. Then:

$$V_{t+1}(y) = E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right)(p^*(S,y,t) + V_t(y-1)) \\ + \left[1 - \lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right)\right] V_t(y) \end{array} \right]$$

Subtracting $V_t(y)$ from both sides, we get

$$V_{t+1}(y) - V_t(y) = E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right) p^*(S,y,t) \\ + \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right)(V_t(y-1) - V_t(y)) \end{array} \right] \quad (26)$$

Now, we note that a seller with $y+1$ units in inventory at period $t+1$ could set its price for signal S to the optimal price of a seller with y units in inventory at period $t+1$, which is suboptimal. Therefore:

$$V_{t+1}(y+1) \geq E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right)(p^*(S,y,t) + V_t(y)) \\ + \left[1 - \lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t))\right)\right] V_t(y+1) \end{array} \right]$$

Subtracting $V_t(y+1)$ from both sides of the above inequality, we obtain

$$V_{t+1}(y+1) - V_t(y+1) \geq E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t)) \right) p^*(S,y,t) \\ +\lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t)) \right) (V_t(y) - V_t(y+1)) \end{array} \right] \tag{27}$$

From induction hypothesis $I3(t, y-1)$, we have

$$V_t(y) - V_t(y+1) \geq V_t(y-1) - V_t(y). \tag{28}$$

Thus, from equations (26)– (28), we have

$$I1(t,y) : V_{t+1}(y+1) - V_t(y+1) \geq V_{t+1}(y) - V_t(y).$$

(ii) $I2(t,y) : V_{t+1}(y) - V_t(y) \geq V_{t+2}(y) - V_{t+1}(y)$.

The case of $y = 0$ is trivial. Suppose $y > 0$. Note that

$$V_{t+2}(y) = E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) (p^*(S,y,t+2) + V_{t+1}(y-1)) \\ + \left[1 - \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) \right] V_{t+1}(y) \end{array} \right]$$

Subtracting $V_{t+1}(y)$ from both sides, we get

$$\begin{array}{l} V_{t+2}(y) - V_{t+1}(y) = \\ E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) p^*(S,y,t+2) \\ +\lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) (V_{t+1}(y-1) - V_{t+1}(y)) \end{array} \right] \end{array} \tag{29}$$

Now, we note that a seller with y units in inventory at period $t+1$ could set its price for signal S to the optimal price of a seller with y units in inventory at period $t+2$, which is suboptimal. Therefore:

$$V_{t+1}(y) \geq E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) (p^*(S,y,t+2) + V_t(y-1)) \\ + \left[1 - \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) \right] V_t(y) \end{array} \right]$$

Subtracting $V_t(y)$ from both sides, we get

$$V_{t+1}(y) - V_t(y) \geq E_S \left[\begin{array}{l} \lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) p^*(S,y,t+2) \\ +\lambda \left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y,t+2)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y,t+2)) \right) (V_t(y-1) - V_t(y)) \end{array} \right] \tag{30}$$

From induction hypothesis $I1(t, y-1)$, we have

$$V_t(y-1) - V_t(y) \geq V_{t+1}(y-1) - V_{t+1}(y). \tag{31}$$

Thus, from equations (29)– (31), we have

$$I2(t,y) : \; V_{t+1}(y) - V_t(y) \geq V_{t+2}(y) - V_{t+1}(y).$$

(iii) $I3(t,y): V_t(y+1) - V_t(y) \geq V_t(y+2) - V_t(y+1)$.

Using arguments similar to those in parts (i) and (ii), we obtain

$$V_t(y+2) - V_{t-1}(y+1) = \\ E_S \left[\begin{array}{l} \lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y+2,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y+2,t))\right) p^*(S,y+2,t) \\ +\lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y+2,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y+2,t))\right)(V_{t-1}(y+2) - V_{t-1}(y+1)) \end{array} \right] \tag{32}$$

and

$$V_{t+1}(y+1) - V_t(y) \geq E_S \left[\begin{array}{l} \lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y+2,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y+2,t))\right) p^*(S,y+2,t) \\ +\lambda\left(\widehat{q}_1(S)\overline{F}_1(p^*(S,y+2,t)) + \widehat{q}_2(S)\overline{F}_2(p^*(S,y+2,t))\right)(V_t(y+1) - V_t(y)) \end{array} \right] \tag{33}$$

From $I1(t-1,y)$ and $I3(t-1,y)$, we have

$$V_t(y+1) - V_t(y) \geq V_{t-1}(y+1) - V_{t-1}(y) \geq V_{t-1}(y+2) - V_{t-1}(y+1). \tag{34}$$

Therefore, from equations (32) – (34), we obtain

$$V_{t+1}(y+1) - V_t(y) \geq V_t(y+2) - V_{t-1}(y+1). \tag{35}$$

From $I2(t-1,y+1)$, we have

$$V_t(y+1) - V_{t-1}(y+1) \geq V_{t+1}(y+1) - V_t(y+1). \tag{36}$$

Finally, adding (35) and (36), we obtain

$$I3(t,y): V_t(y+1) - V_t(y) \geq V_t(y+2) - V_t(y+1).$$

$\square$

References

Acquisti, A. & Varian, H. R. (2005), 'Conditioning prices on purchase history', *Marketing Science* **24**, 367–381.

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

Ahn, H.-S., Gumus, M. & Kaminsky, P. (2006), 'Pricing and manufacturing decisions when demand is a function of prices in multiple periods', *Operations Research* . Forthcoming.

Argon, N. T. & Ziya, S. (2007), Identifying the critical patient: Finding the balance between undertriage and overtriage. Working Paper, University of North Carolina, Department of Statistics and Operations Research.

Aviv, Y. & Pazgal, A. (2005*a*), Optimal pricing of seasonal products in the presence of forward-looking consumers. Working Paper, Washington University, Olin School of Business.

Aviv, Y. & Pazgal, A. (2005*b*), 'A partially observed markov decision process for dynamic pricing', *Management Science* **51**, 1400–16.

Aydin, G. & Ziya, S. (2007), 'Pricing promotional products under upselling', *Manufacturing and Service Operations Management* . Forthcoming.

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

Bitran, G. R. & Caldentey, R. (2003), 'Commissioned paper: An overview of pricing models for revenue management', *Manufacturing and Service Operations Management* **5**, 203–29.

Bitran, G. R. & Mondschein, S. V. (1993), Pricing perishable products: An application to the retail industry. Working Paper, MIT, Sloan School of Management.

Bitran, G. R. & Mondschein, S. V. (1997), 'Periodic pricing of seasonal products in retail', *Management Science* **43**, 64–79.

Blank, J. (2001), 'Pricing gets personal', Jupiter Research Report.

Bridis, T. (June 2, 2005), 'Online retailers charge shoppers different prices', *Chicago Sun-Times* .

Calzolari, G. & Pavan, A. (2006), 'On the optimality of privacy in sequential contracting', *Journal of Economic Theory* **130**, 168–204.

Chen, Y. & Iyer, G. (2002), 'Consumer addressability and customized pricing', *Marketing Science* **21**, 197–208.

Chen, Y., Narasimhan, C. & Zhang, Z. J. (2001), 'Individual marketing with imperfect targetability', *Marketing Science* **20**, 23–41.

Choudhary, V., Ghose, A., Mukhopadhyay, T. & Rajan, U. (2005), 'Personalized pricing and quality differentiation', *Management Science* **51**, 1120–1130.

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

Elmaghraby, W., Gulcu, A. & Keskinocak, P. (2004), 'Optimal markdown mechanisms in the presence of rational customers with multi-unit demands', *Manufacturing and Service Operations Management* . Forthcoming.

Elmaghraby, W. & Keskinocak, P. (2003), 'Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions', *Management Science* **49**, 1287–309.

Feng, Y. & Gallego, G. (2000), 'Perishable asset revenue management with markovian time dependent demand intensities', *Management Science* **46**, 941–956.

Feng, Y. & Xiao, B. (2000), 'A continuous-time yield management model with multiple prices and reversible price changes', *Management Science* **46**, 644–657.

Fudenberg, D. & Tirole, J. (2000), 'Customer poaching and brand switching', *RAND Journal of Economics* **31**, 634–657.

Fudenberg, D. & Villas-Boas, J. M. (2006), Behavior-based price discrimination and customer recognition, *in* T. J. Hendershott, ed., 'Handbook on Economics and Information Systems', Elsevier, Amsterdam, The Netherlands.

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

Gallego, G. & van Ryzin, G. (1997), 'A multiproduct dynamic pricing problem and its applications to network yield management', *Operations Research* **45**, 24–41.

Iyer, G., Soberman, D. & Villas-Boas, J. M. (2005), 'The targeting of advertising', *Marketing Science* **24**, 461–476.

Johnson, C. A., Allen, L., Young, D. & Tong, J. (2000), 'Personal value pricing', The Forrester Report.

Kumar, N. & Rao, R. (2006), 'Using basket composition data for intelligent supermarket pricing', *Marketing Science* pp. 188–199.

Kuo, C., Ahn, H. & Aydin, G. (2007), Dynamic pricing of limited inventories when customers negotiate. Working Paper, University of Michigan, Department of Industrial and Operations Engineering.

Liu, Q. & van Ryzin, G. (2005), Strategic capacity rationing to induce early purchases. Working Paper, Columbia University Business School.

Liu, Y. & Zhang, Z. J. (2006), 'The benefits of personalized pricing in a channel', *Marketing Science* **25**, 97–105.

Maglaras, C. & Meissner, J. (2006), 'Dynamic pricing strategies for multi-product revenue management problems', *Manufacturing and Service Operations Management* **8**, 136–48.

Monahan, G. E., Petruzzi, N. C. & Zhao, W. (2004), 'The dynamic pricing problem from a newsvendor's perspective', *Manufacturing and Service Operations Management* **6**, 73–91.

Muller, A. & Stoyan, D. (2002), *Comparison Methods for Stochastic Models and Risks*, Wiley, New York.

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

Narasimhan, C. (1984), 'A price discrimination theory of coupons', *Marketing Science* **3**, 128–47.

Netessine, S., Savin, S. & Xiao, W. (2006), 'Revenue management through dynamic cross-selling in e-commerce retailing', *Operations Research* **54**, 893–915.

Phillips, R. L. (2005), *Pricing and Revenue Optimization*, Stanford University Press, Stanford, California.

Popescu, I. & Wu, Y. (2006), 'Dynamic pricing strategies with reference effects', *Operations Research* . Forthcoming.

Ramasastry, A. (2005), 'Web sites change prices based on customers' habits', Retrieved from http://www.cnn.com/2005/LAW/06/24/ramasastry.website.prices/index.html on 8/2/2007.

Shaffer, G. & Zhang, Z. J. (1995), 'Competitive coupon targeting', *Marketing Science* **14**, 395–416.

Shaffer, G. & Zhang, Z. J. (2002), 'Competitive one-to-one promotions', *Management Science* **48**, 1143–1160.

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

Smith, S. A. & Achabal, D. D. (1998), 'Clearance pricing and inventory policies for retail chains', *Management Science* **44**, 285–300.

Su, X. (2007), 'Intertemporal pricing with strategic customer behavior', *Management Science* **53**, 726–741.

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

Taylor, C. R. (2004), 'Consumer privacy and the market for customer information', *RAND Journal of Economics* **35**, 631–650.

Thisse, J.-F. & Vives, X. (1988), 'On thestrategic choice of spatial price policy', *American Economic Review* **78**, 122–137.

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

Varian, H. R. (1989), Price discrimination, *in* R. Schmalensee & R. D. Willig, eds, 'Handbook of Industrial Organization', North Holland, Amsterdam, The Netherlands.

Villas-Boas, J. M. (2004), 'Price cycles in markets with customer recognition', *RAND Journal of Economics* **35**, 486–501.

Zhang, D. & Cooper, W. L. (2005), ‘Revenue management for parallel flights with customer-choice behavior’, *Operations Research* **53**, 415–31.

Zhao, W. & Zheng, Y.-S. (2000), ‘Optimal dynamic pricing for perishable assets with nonhomogeneous demand’, *Management Science* **46**, 375–388.

Ziya, S., Ayhan, H. & Foley, R. D. (2004), ‘Relationship among three assumptions in revenue management’, *Operations Research* **52**, 804–9.