Chapter 12
Counteracting Strategic Consumer Behavior in Dynamic Pricing Systems

Yossi Aviv, Yuri Levin, and Mikhail Nediak

Abstract Dynamic pricing and revenue management practices are gaining increasing popularity in the retail industry, and have engendered a large body of academic research in recent decades. When applying dynamic pricing systems, retailers must account for the fact that, often, strategic customers may time their purchases in anticipation of future discounts. Such strategic consumer behavior might lead to severe consequences on the retailers' revenues and profitability. Researchers have explored several approaches for mitigating the adverse impact of this phenomenon, such as rationing capacity, making price and capacity commitments, using internal price-matching policies, and limiting inventory information. In this chapter, we present and discuss some relevant theoretical contributions in the management science literature that help us understand the potential value of the above mitigating strategies.

12.1 Introduction

In the 30 years since the successes of revenue management systems in airlines were first reported, applications have spread steadily into other business areas. Revenue management is now common in such service businesses as passenger railways, cruise lines, hotel and motel accommodation, and car rentals. Other applications have been proposed in such diverse areas as broadcast advertising, sports and entertainment event management, medical services, real estate, freight transportation,

Yossi Aviv
Olin Business School, Washington University, St. Louis, MO, USA
e-mail: aviv@wustl.edu

Yuri Levin
School of Business, Queen's University, Kingston, ON, Canada
e-mail: ylevin@business.queensu.ca

Mikhail Nediak
School of Business, Queen's University, Kingston, ON, Canada
e-mail: mnediak@business.queensu.ca

S. Netessine, C.S. Tang (eds.), *Consumer-Driven Demand and Operations Management Models*, International Series in Operations Research & Management Science 131,
DOI 10.1007/978-0-387-98026-3_12,

and manufacturing. Dynamic pricing and revenue management practices are also gaining increasing popularity in the retail industry, and have engendered a growing body of academic research in recent years. When applying dynamic pricing schemes, sellers need to account for key characteristics of the sales environment, including the scarcity of goods, demand uncertainty, and consumer behavior. In particular, sellers must account for the fact that, often, strategic customers may time their purchases in anticipation of future discounts. When supply is limited, strategic customers need to consider not only future prices, but also the likelihood of stockouts. And since the level of remaining inventory depends on the individual purchase decisions, each customer has to take into account the behavioral pattern of other customers. In recent years, a significant body of research on the topic of *strategic consumer behavior* has emerged in the management science literature.

Previous chapters in this book explain the notion of strategic waiting, and clearly articulate that it is present in many dynamic pricing environments, and that it has severe consequences on revenues and profitability. The purpose of this chapter, then, is to provide a framework for addressing the adverse impact of this phenomenon; namely, How could designers of dynamic pricing systems counteract strategic consumer behavior? To achieve this goal, we chose to focus on scientific models and emerging theoretical contributions in the management science literature. In fact, rather than providing a complete list of all relevant research, we chose to focus on a few papers with greater details. This allows us to present to the reader a meaningful description of modeling structures, different assumptions on system parameters, and ways in which equilibria settle in the dynamic pricing settings. Our research community has explored several approaches for dealing with strategic consumer behavior: (i) making credible price commitments (also known as "announced pricing schemes"); (ii) rationing capacity; (iii) making credible capacity commitments; (iv) using internal price-matching policies; and (v) limiting inventory information. All of these strategies can potentially be used to reduce the incentives of high-valuation customers to wait, and consequently mitigate the adverse impact of strategic waiting. Below, we provide a section-by-section description of some recent research papers that explore each of these strategies, and summarize and explain their findings.

12.2 The Effectiveness of Price Segmentation in Face of Strategic Customers

Research on price commitment goes back to the famous paper of Coase (1972), which considers a monopolist that sells a durable good to a large set of consumers with different valuations. Coase begins his paper with a qualitative discussion of how, ideally, the seller would want to price the product in a way that results in *perfect segmentation*. Namely, charge (initially) a high price from customers with high valuation, and then sequentially reduce the prices to capture customers with smaller valuations. Such strategy is called a *price-skimming strategy*, and if it works as planned, it results in extracting most or all of the consumer surplus. However, Coase argues that if high-valuation customers anticipate that prices will decline (in our terms – if customers behave strategically), they would wait for a price reduction rather than buy at premium costs. This, in equilibrium, will effectively lead the seller to offer the product at marginal cost.

Coase suggests a number of ways for the seller to avoid this result. For example, the seller can make a special contractual arrangement with the purchasers not to sell more than a given quantity of the product. In fact, this idea is close in spirit to the subject of *capacity rationing* we discuss in the next section. Alternatively, the seller could offer customers to buy back the product if it was offered in the future at a lower price. This idea is very similiar to the subject of *internal price guarantees* that we explore in a later section. The seller could also lease the product for relatively short periods of time and, e.g., announce that he would not change the rental price during the lease period.

Besanko and Winston (1990) introduce a game-theoretic model of a monopolistic seller facing a market of strategic consumers with heterogeneous valuations. The distribution of these valuations is uniform over the interval $[0, v^+]$. Customers know their individual valuations, but the monopolist is only privy to the statistical distribution characteristics. The population size is known and equal to N; in fact, all customers are present in the "store" from the beginning of the game. The monopolist has T periods of time to set the price, with p_t denoting the price in period t. It is assumed that the seller cannot make any price commitment (see extended discussion of this aspect later in this section), and thus the authors look for a subgame perfect equilibrium, in which the seller makes the optimal pricing decision in each given period. The seller's production capacity is unlimited, so any desired qauntity can be produced at any period, at a cost c per unit (where $c < v^+$). The monopolist and the seller are assumed to have the same discount factor δ.

The first observation made in the paper has to do with consumer behavior and the dynamics of information in the game. It is argued that if in period $t-1$ it was optimal for a customer with valuation v to buy the product, then it means that all customers with valuations of v and above have also purchased the product by that time. As a consequence, it suffices to consider a *threshold value* at the end of each period: For instance, v_{t-1} means that all customers with valuations larger or equal to v_{t-1} have purchased the product, whereas all other customers (with valuations below v_{t-1}) are still in the market. Therefore, let $p_t^*(v_{t-1})$ be the seller's equilibrium pricing strategy at time t, when faced with the state v_{t-1}. Taking the customers' perspective now, let $v_t^*(p_t, v_{t-1})$ be the threshold value that sets their buying policy in period t. Besanko and Winston derive the following dynamic program that is based initially on a "guess" that prices and thresholds in subsequent periods are set in a way that makes a customer with valuation v_t indifferent about immediately purchasing the product (and gaining a surplus of $v_t - p_t$) and waiting for the next period, and gaining a *discounted* surplus (of $\delta \cdot (v_t - p_{t+1}^*(v_t))$). With this structural restriction in mind, they show that the seller's equilibrium discounted profit over periods t through T is given by the recursive scheme:

$$H_t^*(v_{t-1}) = \max_{p_t, v_t} \left\{ (p_t - c) \cdot \frac{v_t - v_t}{v^+} \cdot N + \delta H_{t+1}^*(v_t) \right\}$$

$$\text{s.t.} \quad v_t \le v_{t-1}, \quad v_t - p_t = \delta \cdot \left(v_t - p^*_{t+1}(v_t)\right)$$

with $v_0 = v^+$. Indeed, they prove that this dynamic program formulation can serve as a basis for calculating a subgame perfect Nash equilibrium for the game.

Proposition 1 of Besanko and Winston (1990). *A subgame perfect Nash equilibrium exists and can be described as follows.*

$$H^*_t(v_{t-1}) = A_t \cdot (\max\{v_{t-1} - c, 0\})^2 \cdot \frac{N}{v^+}, \qquad t = 1,\ldots,T$$

$$p^*_t(v_{t-1}) = \min\{2A_t v_{t-1} + (1-2A_t)c, v_{t-1}\}, \qquad t = 1,\ldots,T$$

$$v^*_t(v_{t-1}) = \min\{\lambda_t v_{t-1} + (1-\lambda_t)c, v_{t-1}\} \qquad t = 1,\ldots,T$$

$$v^*_t(p_t, v_{t-1}) = \min\left\{\frac{p_t - \delta(1-2A_{t+1})c}{1-\delta+2A_{t+1}}, c\right\} \qquad t = 1,\ldots,T$$

where $\{A_t\}$ and $\{\lambda_t\}$ are sequences defined by the recursions specified in (8) and (9) in the paper.

The authors show that in the above type of equilibrium, prices monotonically decline over time; in other words, *price skimming* arises in equilibrium. Additionally, they develop a benchmark model in which all customers are myopic. Here, it is easy to see that the seller's optimal policy is provided by a solution to the dynamic program

$$H_t(v_{t-1}) = \max_{p_t : p_t \le v_{t-1}} \left\{(p_t - c) \cdot \frac{v_{t-1} - p_t}{v^+} \cdot N + \delta H_{t+1}(p_t)\right\}$$

Comparing the two models, Besanko and Winston show that with myopic consumers, the price is always higher in any *given* state than it is with strategic consumers. In other words, the *first-period price* with myopic consumers is higher than the first-period price with strategic consumers. It is noteworthy, however, that because the time paths of sales in the two cases differ, it is possible that the price with myopic consumers will fall below the price with strategic consumers in later periods. Two other interesting observations were made in the paper. First, the authors illustrate a situation in which a seller that commits to a price path that is based on the "myopic case," and use it when customers are actually strategic, might significantly hurt his expected profit. Second, the authors argue that for any v, $H^*_t(v)$ is *increasing* in t. In contrast, $\hat{H}_t(v)$ (the optimal expected profit in the myopic case) is *decreasing* in t. Consequently, starting in any state v, a monopolist prefers a *shorter* time horizon if faced with strategic consumers, but a *longer* time horizon if faced with myopic consumers. The intuition behind this is that the shorter is the time horizon, the smaller is the power of strategic consumers. In contrast, with myopic consumers, the monopolist prefers a longer time horizon because it gives him more flexibility in setting prices over time and hence extracting more revenues.

12.2.1 Models with Limited Inventories

Elmaghraby et al. (2007) consider a setting in which a seller uses a pre-announced markdown pricing mechanism, to sell a finite inventory of a product. Specifically, the seller's objective is to maximize expected revenues by optimally choosing the number of price steps over the season, and the price at each step. All potential buyers are present at the start of the selling period and remain until all the units have been sold or their demand has been satisfied. The buyers, who demand multiple units, may choose to wait and purchase at a lower price, but they must also consider a scarcity in supply. The authors study the potential benefits of segmentation; namely, the difference between the seller's profit under the optimal markdown mechanism and that under the optimal single price. They also provide a detailed discussion on the design of profitable markdown mechanisms.

Su (2007) presents a pricing control model in which consumers are infinitesimally small and arrive continuously according to a deterministic flow of constant rate. The customer population is heterogeneous along two dimensions: valuations and degree of patience (vis-à-vis waiting). The seller has to decide on pricing and a rationing policy which specifies the fraction of current market demand that is fulfilled. Given these retailer's choices, customers decide whether or not to purchase the product and whether to stay or leave the market. The paper shows how the seller can determine a revenue-maximizing selling policy in this game. Su demonstrates that the heterogeneity in valuation and degree of patience jointly influence the structure of optimal pricing policies. In particular, when high-valuation customers are proportionately less patient, markdown pricing policies are effective. On the other hand, when the high-valuation customers are more patient than the low-valuation customers, prices should increase over time in order to discourage waiting.

Aviv and Pazgal (2008) consider a seller that has Q units of an item available for sale during a sales horizon of length H. The sales season $[0,H]$ is split into two parts, $[0,T]$ and $[T,H]$, for a given fixed value T. During the first part of the season, a "premium" price p_1 applies, and during the second phase of the season a "discount" price p_2 is offered (where $p_2 \leq p_1$). The seller's objective is to set the premium and discount prices in order to maximize the expected total revenues collected during the sales horizon. An important feature of their model is that it includes two types of *demand uncertainty*: the total market size and the time of arrivals. Specifically, it is assumed that customers arrive to the store following a Poisson process with a rate of λ. Customers' valuations of the product vary across the population, and decline over the course of the season according to

$$V_j(t) = V_j \cdot e^{-\alpha t}$$

for every customer j. Specifically, customer j's *base valuation* V_j is drawn from a given continuous distribution form F. Then, depending on the particular time of purchase (t), the realized valuation is discounted appropriately by a known exponential decline factor $\alpha \geq 0$, fixed across the population. Customers that arrive prior to time T behave according to the following strategy: A given customer j, arriving at time t, will purchase immediately upon arrival (if there is inventory) if two conditions are satisfied about his current surplus $V_j e^{-\alpha t} - p_1$: (i) it is non-negative; and (ii) it is larger or equal to the *expected* surplus he can gain from a purchase at time T (when the price is changed to p_2). Of course, the latter expected surplus depends on the customer's belief about p_2 as well as the likelihood that a unit will be available to the customer. If the customer purchases a unit, he leaves the store immediately. Otherwise, the customer stays until time T. At time T, all existing customers take a look at the new price p_2 and if they can gain a non-negative surplus, they request a unit of the remaining items (if any). In case there are fewer units than the number of customers who wish to buy, the allocation is made randomly. After time T, new customers buy according to whether or not they can immediately gain a non-negative surplus. Clearly, it does not make sense for a customer to wait in the store after time T, since prices will not drop. The seller observes the purchases, or equivalently, his level of inventory. Hence, the discounted price p_2 depends on the remaining inventory at time T.

Aviv and Pazgal identify a subgame perfect Nash equilibrium for the game between the customers and the seller. Note that the seller's strategy is characterized by the initial premium price p_1 and the discounted price menu $\{p_2(q)\}_{q=1}^{Q}$. The customers' strategy is one that prescribes purchasing decisions for every possible pair of individual arrival time t and base valuation V.

They first study the best response of the customers to a given seller's *pure* strategy of the form $p_1, p_2(1), \ldots, p_2(Q)$. The response strategy is based on a competitive situation that exists among consumers, which arises due to the fact that an individual consumer's decision impacts the product availability for others. It is shown that a *time-dependent* threshold emerges, as follows.

Theorem 1 of Aviv and Pazgal (2008). *For any given pricing scheme* $\{p_1, p_2(1), \ldots, p_2(Q)\}$*, it is optimal for the customers to base their purchasing decisions on a threshold function* $\theta(t)$*. Specifically, a customer arriving at time* $t \in [0, H]$ *will purchase an available unit immediately upon arrival if* $V(t) \geq \theta(t)$*. Otherwise, if* $V(t) < \theta(t)$ *and* $t < T$*, the customer will revisit the store at time* T*, and purchase an available unit if* $V(T) \geq \theta(T)$*. The threshold function* $\theta(t)$ *is given by*

$$\theta(t) = \begin{cases} \psi(t) & 0 \leq t < T \\ \\ p_2 & T \leq t \leq H \end{cases} \tag{12.1}$$

where $\psi(t)$ *is the unique solution to the implicit equation*

$$\psi - p_1 = \mathrm{E}_{Q_T}\left[\max\{\psi e^{-\alpha(T-t)} - p_2(Q_T), 0\} \cdot \mathbf{1}\{\mathscr{A} | Q_T\}\right] \tag{12.2}$$

The random variable Q_T *represents the remaining inventory at time* T*, and the event* A *represents the allocation of a unit to the customer upon request.*

The function ψ hence defines the customers' purchasing strategy. The left-hand side of (12.2) represents the current surplus the customer can gain by purchasing a unit, whereas the right-hand side of the equation represents the expected surplus that will be gained by the customer if he postpones his purchase to time T. The latter expected value takes into account two conditions. The first condition is that the discounted price needs to leave the customer with a non-negative surplus. This is simply given by the condition $\psi e^{-\alpha(T-t)} - p_2(Q_T) \geq 0$, where Q_T is a random variable. The second condition is that in order to provide a surplus, a unit needs to be available and be allocated to the customer. Given a specific realization of Q_T, the allocation probability depends on the statistical distribution of the number of other customers that postpone their purchases to time T. This is taken into account by the indicator expression $\mathbf{1}\{\mathscr{A}|Q_T\}$. The strength of this theorem is that it demonstrates the optimality of a threshold-type policy for each individual customer, under *any arbitrary* purchasing strategies of the others. Using this result, Aviv and Pazgal argue that for any given pricing scheme $\{p_1, p_2(1), \ldots, p_2(Q)\}$, the equilibrium in the game between the customers is unique, and consists of symmetric strategies; i.e., ψ is the same for all customers. Additionally, they show that the threshold function $\psi(t) : [0,T) \rightarrow [p_1, \infty)$ is increasing in t.

Next, the seller's strategy is studied; namely, the best contingent pricing $\{p_2(1), \ldots, p_2(Q)\}$ in response to a given purchasing strategy ψ and a given initial premium price p_1. As a basis for the analysis, it is useful to distinguish between five types of customers, as illustrated in Figure 12.1.

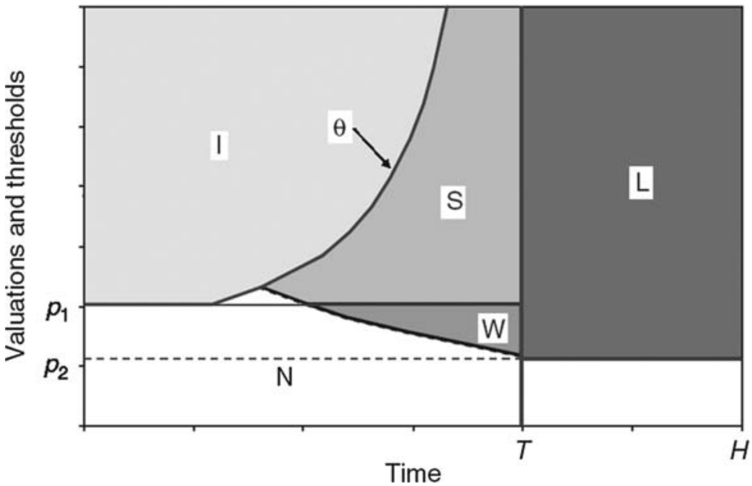

Fig. 12.1 For a given realized price path (p_1, p_2) and a customer threshold purchasing policy $\psi(t)$, the customer's space (arrival times and valuations) is split into five areas: (i) "I" = immediate buy at premium price; (ii) "S" = strategic wait and buy at discounted price; (iii) "W" = non-strategic wait and buy at discounted price; (iv) "L" = immediate purchasers at discounted price; and (v) 'N' = no buyers. In this context "buy" means a desire to buy.

The first group is of customers that arrive during $[0,T]$ and purchase[1] immediately at price p_1 (denoted by ‘I’). It is easy to see that the expected number of customers of this type is

$$\Lambda_I(\psi) \doteq \lambda \int_{t=0}^{T} \bar{F}\left(\psi(t)\, e^{\alpha t}\right) dt$$

We define the remaining four types of customers with respect to a *realized* value of p_2 at time T. The second group of customers (denoted by “S”) consists of those who could get a non-negative surplus upon their arrival during $[0,T]$ but decided to wait “strategically” (i.e., in anticipation of a better expected surplus at time T), and indeed want to purchase a unit at the realized price p_2 (note that it may happen that the actual surplus realized is lower than the original surplus a customer could gain if purchased a unit immediately upon arrival). The number of customers that falls in this category has a Poisson distribution with a mean

$$\begin{aligned}&\Lambda_S(\psi, p_1, p_2)\\ &\quad\doteq \lambda \int_{t=0}^{T} \left[\bar{F}\left(\min\left\{\max\left\{p_1 e^{\alpha t}, p_2 e^{\alpha T}\right\}, \psi(t)\, e^{\alpha t}\right\}\right) - \bar{F}\left(\psi(t)\, e^{\alpha t}\right)\right] dt\end{aligned}$$

The third group (denoted by “W”) includes those customers who waited for time T because their valuations upon arrival were below the premium price p_1, and want to purchase a unit at the realized discount price p_2. The number of such customers is Poisson distributed with a mean

$$\Lambda_W(p_1, p_2) \doteq \lambda \int_{t=0}^{T} \left[\bar{F}\left(\min\left\{p_1 e^{\alpha t}, p_2 e^{\alpha T}\right\}\right) - \bar{F}\left(p_1 e^{\alpha t}\right)\right] dt$$

The fourth group of customers (denoted by “L” for “late”) includes those who arrive at or after time T with a valuation higher than the posted discounted price p_2. Clearly, the number of customers in this group is Poisson with a mean of

$$\Lambda_L(p_2) \doteq \lambda \int_{t=T}^{H} \bar{F}\left(p_2 e^{\alpha t}\right) dt$$

Finally, a fifth group (denoted by “N” in Figure 12.1) includes those customers who do not wish to purchase a unit of the product at any point of time. Note that in this group, we include customers who decided (strategically) to wait for time T in anticipation of a better surplus, but then find out that the realized discounted price p_2 is higher than their individual valuations.

It is instructive to note that the values of Λ_S, Λ_W, and Λ_L depend on the value of p_2, which is generally unknown prior to time T. Therefore, such uncertainty needs to be taken into account by customers who contemplate between an immediate purchase and a strategic wait. For the seller, the above values play a key role in

[1] We use the term “purchase” to reflect a desire to purchase a unit. Clearly, when inventory is depleted, we assume that the store closes and no further purchases are feasible.

setting the contingent pricing menu. Let $Q_T \in \{0,\dots,Q\}$ be the current inventory at time T. Clearly, if $Q_T = 0$, then pricing is irrelevant. When Q_T is positive, the optimal discounted price $p_2(Q_T)$ is chosen so as to maximize the revenues collected from customers of types "S", "W", and "L". Specifically, for $Q_T \in \{1,\dots,Q\}$, let

$$p_2(Q_T) \in \arg\max_{z \le p_1} \{z \cdot N(Q_T, \Lambda_S(\psi, p_1, z) + \Lambda_W(p_1, z) + \Lambda_L(z))\} \tag{12.3}$$

where $N(q,\Lambda) \doteq \sum_{x=0}^{\infty} \min(x,q) \cdot P(x|\Lambda)$ is the expected value of a Poisson random variable (with mean Λ) truncated at q.

Now, consider the subgame that begins after the premium price p_1 is set. A contingent pricing scheme $\{p_2(q)\}$ for the seller and a purchase strategy ψ for the consumers form a Nash equilibrium in the subgame, if the following conditions are satisfied. First, each price $p_2(q)$ needs to satisfy (12.3); i.e., it is a best response of the seller if all customers follow the equilibrium strategy ψ. Second, the strategy ψ needs to satisfy the conditions of theorem 1 stated above; i.e., it needs to be an optimal response to the contingent pricing scheme $\{p_2(q)\}$. For brevity of exposition, we refer the reader to Section 4.3 of the paper for a discussion of an iterative algorithm developed and employed by Aviv and Pazgal to find an equilibrium to the *subgame*. Finally, in search of maximizing the expected total revenue over the sales horizon, the seller needs to pick the best premium price p_1. This is done with the anticipation that a choice of p_1 will be followed by the subgame we described above. Given a Nash equilibrium $\psi^*(p_1)$ and $\{p_2^*(q|p_1)\}_{q=1}^{Q}$, the seller is faced with an optimization problem of the following type. Let $J(q|p_1)$ be the expected revenues collected during the second part of the season ($[T,H]$) given the subgame Nash equilibrium strategies:

$$J(q|p_1) = p_2^*(q|p_1)$$
$$\cdot N(q, \Lambda_S(\psi^*(p_1), p_1, p_2^*(q|p_1)) + \Lambda_W(p_1, p_2^*(q|p_1)) + \Lambda_L(p_2^*(q|p_1)))$$

Then, the seller's task is to maximize the expression

$$\pi_{C/S}^* \doteq \max_{p_1} \left\{ p_1 \cdot Q \cdot \left(1 - \sum_{x=0}^{Q-1} P(x \mid \Lambda_I(\psi^*(p_1)))\right) \right. \tag{12.4}$$
$$\left. + \sum_{x=0}^{Q-1} (p_1 \cdot x + J(Q-x|p_1)) \cdot P(x|\Lambda_I(\psi^*(p_1))) \right\}$$

The subscript "C/S" denotes the case of contingent pricing policies with strategic customers.

In order to measure the effectiveness of price segmentation, Aviv and Pazgal consider the expected increase in revenues obtained by moving from an optimal *static-pricing* strategy (π_F^*) to a two-price strategy. In other words, they propose the metric $\left(\pi_{\text{C/S}}^* - \pi_F^*\right)/\pi_F^*$, where

$$\pi_F^* \doteq \max_p \left\{ p \cdot N \left(Q, \lambda \int_0^H \bar{F} \left(p \cdot e^{\alpha t} \right) dt \right) \right\}$$

Like Besanko and Winston (1990), they also consider benchmark models in which all customers are myopic. Their findings are summarized below.

In general, the benefits of using contingent pricing schemes appear to be very significant under the case of myopic customers. Unlike the case of myopic customers, they show that strategic customer behavior clearly interferes with the benefits of price segmentation to the retailer. Strategic consumer behavior suppresses the benefits of segmentation, under medium-to-high values of heterogeneity and modest rates of decline in valuations. An underlying reason for this is that when the rate of decline in valuation is small, customers are patient in waiting for the discount time. Not only that, but unlike the case of fixed-price strategies, customers rationally expect discounts to take place. However, when the level of consumer heterogeneity is small, the rate of decline is medium to high, and the time of discount can be optimally chosen (in advance) by the seller, segmentation can be used quite effectively even with strategic consumers. The rationale for the result is that under these conditions, there is typically a little difference in price valuations at the time of discount. Hence, the discount price is generally set in a way that does not offer a substantial surplus to consumers.

The seller cannot effectively avoid the adverse impact of strategic behavior even under low levels of initial inventory. When the initial inventory level is low, the seller expects customers to be more concerned about product availability at discount time, and thus act in a similar way to myopic customers. However, for myopic consumers the seller should benefit from *high-price experimentation*. In other words, the seller would set a high price for the initial period of time in order to bet on collecting large revenues, speculating that small number of unsold units can be easily sold at the end of the season. But high-price experimentations in the case of strategic consumers could drive customers to wait even if the availability probability is relatively low. Thus, high-price betting cannot be sustained in equilibrium and for more moderate prices customers are more inclined to purchase immediately than wait and take the risk of a stockout.

When the seller incorrectly assumes that strategic customers are myopic in their purchasing decisions, it can be quite costly, reaching up to 20% loss of potential revenues. When the level of heterogeneity is large, misclassification results in offering high discounts. Now, if valuations do not decline significantly during the horizon, strategic customers would most likely wait. The dependency of the sub-optimality gap on rate of decline in valuations works in two different ways. On one hand, when it is small, customers are typically more inclined to wait, but their valuations do not decline significantly. When the rate of decline is high, fewer customers will decide to wait, but the rapid decline in these customers' valuations would hurt the seller's ability to extract high revenues at the time of discount. Finally, when customers' valuations are homogeneous and the decline rate is large, strategic customers do not have a substantial incentive to wait, and so misclassification is not expected to lead to significant loss.

12.3 Price Commitments

In addition to the models described above, Aviv and Pazgal (2008) also consider a two-period pricing problem in which the seller commits *upfront* to a *fixed*-price path (p_1, p_2). Specifically, under this policy class, named "announced discount strategies," the discounted price p_2 is not contingent upon the remaining inventory at time T. Interestingly, sellers such as Filene's Basement (Bell and Starr 1994), Land's End, and Syms use this method for pricing some of their products. Consider for a moment the case of myopic customers. Clearly, in such settings, a fixed discount is generally not optimal. The seller could obviously increase his expected revenues by waiting until time T and *then* determine the best-price discount according to the remaining amount of inventory on hand. The same logic does not straightforwardly apply in face of strategic customers. In fact, under strategic consumer behavior there could be cases in which announced pricing schemes (denoted by "A/S") may perform better than contingent pricing policies (denoted by "C/S"). For the sake of illustration, let us consider the case of unlimited inventory ($Q = \infty$). In this case, it is easy to show that "A/S" is at least as good as "C/S". This is because the discounted price under "C/S" can be determined accurately in advance. Thus, that same discounted price could be used under "A/S", yielding the same expected revenues as in the best strategy under "C/S". But, of course, the seller could perhaps do even better by announcing (and committing to) a smaller discount in order to discourage strategic waiting. Indeed, game theory provides us with variety of examples in which limiting one's choices of future actions (burning the bridges) may put one in a better position at equilibrium. The following result is proven.

Theorem 2 of Aviv and Pazgal (2008). *Consider any given (and credible) announced pricing path $\{p_1, p_2\}$. Then, the threshold-type policy defined by the function θ below constitutes a Nash equilibrium in the customers' purchasing strategies. Let*

$$\theta(t) \doteq \begin{cases} \psi_A(t) \doteq \max\left\{p_1, \dfrac{p_1 - wp_2}{1 - we^{-\alpha(T-t)}}\right\} & 0 \le t < T \\ p_2 & T \le t \le H \end{cases} \tag{12.5}$$

where w is a solution to the equation:

$$w \doteq \sum_{x=0}^{Q-1} P(x|\Lambda_I(\psi_A)) \cdot A(Q - x|\Lambda_S(\psi_A, p_1, p_2) + \Lambda_W(p_1, p_2)) \tag{12.6}$$

(Note that ψ_A is dependent on w, as given in (12.5).) Specifically, a customer arriving at time $t \in [0, H]$ will purchase an available unit immediately upon arrival if $V(t) \ge \theta(t)$. Otherwise, if $V(t) < \theta(t)$ and $t < T$, the customer will revisit the store at time T, and purchase an available unit if $V(T) > \theta(T) = p_2$.

As in the case of contingent pricing policies, in equilibrium, every customer needs to take into account the behavior of the other customers. This is reflected by the parameter w which represents the likelihood of receiving a unit of the product at time T. Clearly, if Q is relatively large, the interaction between customers is negligible, and their optimal purchasing policy is given in (12.5), with $w \approx 1$.

It is easily seen that the game between the seller and the customers has a Stackelberg form, with the seller being the leader in his announcement of the strategy (p_1, p_2) and the customers following by selecting their strategy w. Given the seller's knowledge of the customers' response to any particular pair of prices, his task is to maximize the expected revenue[2]:

$$\pi_{A/S}(p_1,p_2) \doteq p_1 \cdot Q \cdot \left(1 - \sum_{x=0}^{Q-1} P(x|\Lambda_I)\right) + \sum_{x=0}^{Q-1} P(x|\Lambda_I) \cdot [p_1 \cdot x + p_2 \cdot N(Q-x, \Lambda_S + \Lambda_W + \Lambda_L)]$$

In other words, find the solution to $\pi^*_{A/S} = \max_{p_1,p_2:p_2 \le p_1} \{\pi_{A/S}(p_1,p_2)\}$. Aviv and Pazgal find that announced fixed-discount strategies perform essentially the same as contingent pricing policies in the case of myopic consumers. However, they suggest caution in interpreting this result. First, they consider the *optimal* announced discount. If a seller picks an arbitrary discount level, the sub-performance with respect to contingent pricing can be very high. Second, announced discounts prevent the seller from acting upon *learning* about demand; see, e.g., Aviv and Pazgal (2005) and Mantrala and Rao (2001). Under strategic consumer behavior, they found that announced pricing policies can bring an advantage to the seller (up to 8.32% increase in expected revenues), compared to contingent pricing schemes. Particularly, they observed that announced pricing schemes are advantageous compared to contingent pricing schemes under the following conditions: (i) the number of units are sufficiently high; (ii) the level of heterogeneity in base valuations is high; (iii) the discounts are offered at a late part of the season; and (iv) the rate of decline in valuations is at a medium level. The underlying reason for the better performance of announced discount strategies is that a *credible* pre-commitment to a fixed-discount level removes the rational expectation of customers that at discount time the seller will optimally offer large discounts. Interestingly, they found that in those cases that announced discount strategies offer a significant advantage compared to contingent pricing policies, they appear to offer only a minimal advantage in comparison to fixed pricing policies.

12.4 Capacity Rationing

In order to mitigate the adverse impact of strategic consumer behavior, firms may consider the use of *capacity rationing strategies*. Under such strategies, a seller deliberately understocks a product, hence creating a shortage risk for customers.

[2] The functions Λ_I, Λ_S, Λ_W, and Λ_L stand for $\Lambda_I(\psi_A)$, $\Lambda_S(\psi_A, p_1, p_2)$, $\Lambda_W(p_1, p_2)$, and $\Lambda_L(p_2)$, respectively.

As a consequence, this motivates high-valuation customers to purchase early in the season at premium prices.

Liu and van Ryzin (2008) propose a model for examining the potential value of rationing strategies. In their model, a (monopoly) firm pre-announces a single markdown pricing policy over two periods: A premium price p_1 for the first period, and a *discounted* price, p_2, for the second period. Similarly to some of the models surveyed in the previous section, this situation represents a setting in which the seller is able to make a price commitment. The market size, denoted by N, is deterministic, and consists of strategic consumers that are present at the "store" from the beginning of the horizon. Consumers have heterogeneous valuations that are independently drawn from a known distribution $F(v)$, and they enjoy a utility of $u(v-p)$ when they have valuation v and purchase a unit at price p. The model assumes that the market size is large enough so that strategic interaction between customers can be ignored. The firm seeks to maximize profits by choosing its stocking quantity (capacity) at the beginning of the sales season.

When making their purchasing decisions, customers weigh the immediate utility $u(v-p_1)$ (reflecting a "buy now" choice) against the expected utility that they can gain in the second period. But to calculate the latter value, the customers need to multiply the utility $u(v-p_2)$ by the likelihood that a unit will be available. This probability, denoted by q, is assumed to be estimated exactly at the same value by all customers. It follows that, for any given q, the customers optimally follow a threshold policy with a parameter $v(q)$. This threshold value is implicitly defined by the equation:

$$u(v-p_1) = q \cdot u(v-p_2)$$

Consequently, customers with valuations larger than $v(q)$ buy in period 1; the other customers wait for period 2.

The firm needs to determine its optimal capacity C, by taking into account the per-unit procurement cost of α (assumed to be lower than p_2). But before we get into showing the formulation of the firm's decision problem, it is important to consider the connection between the capacity choice and the customers' behavior. Liu and van Ryzin show that under rational expectations[3]:

$$q = \frac{C - N \cdot \bar{F}(v(q))}{N \cdot (F(v(q) - F(p_2)))}$$

the probability q is referred to as the *fill rate*. Note that $N \cdot \bar{F}(v(q))$ represents the number of customers that purchases the product (or more precisely, attempt to) in period 1. The term $N \cdot (F(v(q) - F(p_2)))$ is equal to the number of customers that

[3] In fact, they use the equation

$$q = \min\left\{\max\left(\frac{C - N \cdot \bar{F}(v(q))}{N \cdot (F(v(q) - F(p_2)))}, 0\right), 1\right\}$$

but argue that under all reasonable policies, the condition $\bar{F}(p_1) \le C/N \le \bar{F}(p_2)$ needs to be satisfied, and so the "min" and "max" operands are redundant.

is still in the market at the beginning of period 2, and has valuations that are at least as high as the discounted price p_2. The firm's profit maximization is given by

$$\begin{aligned}\max\ & \{N\cdot(p_1-p_2)\cdot\bar{F}(v)+(p_2-\alpha)\cdot C\}\\ \text{s.t.}\ & u(v-p_1)=\frac{C-N\cdot\bar{F}(v(q))}{N\cdot(F(v(q)-F(p_2)))}\cdot u(v-p_2)\\ & p_1\le v\le\bar{U}\end{aligned}$$

where $\bar{U}$ is an upper bound on the customers' valuations. Let us now follow the special case studied in that paper, where F is uniform over $[0,\bar{U}]$, and the utility function takes the form $u(x)=x^{\gamma}$ $(0<\gamma<1)$. The parameter γ corresponds with the degree of *risk aversion* (lower values of γ correspond to more risk aversion). Under these assumptions, the authors show that the potential value of rationing depends on the number of high-valuation customers in the market (reflected by the parameter $\bar{U}$). Specifically, with a sufficiently large number of high-valuation customers in the market, it makes sense to adopt a rationing policy; otherwise, the firm should serve the entire market at the low price. The intuition behind this observation is simple. Note that by increasing the degree of rationing (i.e., less capacity brought to market), the firm induces higher demand in period 1 at the expense of missing the opportunity to serve some demand in period 2. Therefore, if the number of high-valuation customers in the market is small, the benefits gained in period 1 cannot justify the loss incurred in period 2. The authors find that when the firm can optimally select the prices, rationing is *always* an optimal strategy, for *any* value of $\bar{U}$. Liu and van Ryzin also explore the way in which the level of risk aversion (γ) influences the value of capacity rationing. They argue that when γ approaches 1 (i.e., customers becoming risk-neutral), the rationing risk that is needed in order to induce segmentation is very high. In other words, the planned leftover capacity level for period 2 should practically be set to 0 as γ approaches 1. Consequently, when the market consists of a sufficiently large number of high-valuation customers, it is optimal for the firm to serve the market only at the high price in period 1; otherwise, the firm serves the entire market at the low price only.

Liu and van Ryzin also study a model of oligopolistic competition. In sake of brevity, we refer the interested reader to their paper for technical details (see Section 4.4 there). They show that competition makes it more difficult to support segmentation using rationing, and explain this as follows. When competing against a large number of other firms, a focal firm is very limited in his ability to create a sense shortage risk. Thus, by reducing its capacity, a focal firm severely influences his own ability to serve demand, rather than drive high-valuation customers to buy at high prices. Particularly, the authors prove that there exists a critical number of firms beyond which creating rationing risk is never a sustainable equilibrium. Thus, rationing is more likely to be used in cases where a firm has some reasonable degree of market power.

Cachon and Swinney (2008) explore a two-period model that allows for dynamic planning of pricing and inventory. Recall, from our discussion in the previous section, that whether or not a retailer commits to a price path can significantly influence consumer behavior and retailer's performance. The key features of the model setup are as follows. A retailer sells a product over two periods. The price for the first period, p, is *exogenously* given. In the second period, the product is sold at a given markdown price s. The retailer's objective is to maximize his expected profit by selecting the optimal sale price and the initial level of inventory. The unit cost to the retailer is c.

The total number of customers that may purchase during the first period is random and denoted by D, having a distribution F. Customers are heterogeneous in terms of their behavior. A portion of the customers ($\alpha \cdot D$) are strategic, and the remainder ($(1-\alpha) \cdot D$) are myopic consumers that exist in the market in the first period *only*. All customers have the same (known) valuation in the first period, equal to v_M ($v_M \geq p$). The valuations of strategic customers in the second period change in a random manner, assumed to be uniformly distributed between $[\underline{v}, \bar{v}]$ (where $\bar{v} \leq p$). It is assumed that strategic consumers know their individual second-period valuations in advance, from the beginning of period 1. In addition to the above mix of customers, the model introduces a third group of *bargain hunters*. These customers arrive only in period 2, and in large numbers; they are all assumed to share the same common valuation, v_B. The value of v_B is assumed to be lower than the cost c, meaning that targeting this set of customers could be used as a mechanism for salvaging unsold inventory.

The retailer's and customers' actions are driven by a rational expectation equilibrium in the game we will describe below. But let us dwell for a moment on the types of beliefs that drive actions. In contemplating a "buy now versus buy later" decision, a customer needs to assess the likelihood of getting a unit in period 2. Of course, this is influenced by the retailer's level of inventory, which is not directly observed. Hence, $\hat{q}$ is defined as the customers' (common) belief about the initial stocking level of the retailer. Similarly, in order for the retailer to set the optimal inventory quantity, he needs to be able to anticipate the way in which customers make their purchasing decisions. Let $\hat{v}$ represent this belief.[4] The authors identify a *subgame perfect Nash equilibrium* with *rational expectations* to the game. Such equilibrium is denoted by (q^*, v^*) and needs to satisfy a set of three conditions. First, the retailer is assumed to act optimally under his belief about the consumer behavior. In other words,

$$q^* \in \arg\max_{q \geq 0} \pi(q, \hat{v})$$

where $\pi(q, \hat{v})$ is equal to the retailer's expected profit under a given choice of initial inventory q, and if all customers behaved according to the purchasing threshold policy $\hat{v}$. Second, the consumers' purchasing policy should be optimal if the seller indeed sets the initial quantity to $\hat{q}$. Let $v^*(q)$ denote an optimal threshold policy for a known inventory q. Then this condition is reflected by

$$v^* \in v^*(\hat{q})$$

[4] The authors show that in equilibrium, there exists some $v^* \in [\underline{v}, \bar{v}]$ such that all strategic consumers with second-period valuation less than v^* purchase in period 1, and all consumers with valuation greater than v^* wait for period 2.

Third, the beliefs need to be consistent. In other words,

$$\hat{q} = q^*, \quad \hat{v} = v^*$$

Note that the retailer's expected profit under a belief $\hat{v}$ can be presented in a stochastic dynamic programming format, where the first action (at the beginning of period 1) is the initial level of inventory q, and the second action (at the beginning of period 2) is the markdown price s. The retailer knows that if indeed the customers' policy is $\hat{v}$, then the portion of strategic consumers that will purchase in period 1 is given by $(\hat{v} - \underline{v}) / (\bar{v} - \underline{v})$. Together with the myopic customers who purchase in the first period, the demand in the first period sums up to

$$D_1(\hat{v}, D) \doteq \alpha \frac{\hat{v} - \underline{v}}{\bar{v} - \underline{v}} \cdot D + (1 - \alpha) \cdot D = \left(1 - \alpha \cdot \frac{\bar{v} - \hat{v}}{\bar{v} - \underline{v}}\right) \cdot D$$

The sales in period 1 are hence given by $S_1(\hat{v}, D) \doteq \min(q, D_1(\hat{v}, D))$, and the remaining level of inventory by the end of period 1 is

$$I(\hat{v}, D) \doteq \max(q - D_1(\hat{v}, D), 0)$$

With this, we get

$$\pi(q, \hat{v}) = -c \cdot q + E_D \left[p \cdot S_1(\hat{v}, D) + \max_s \{R(s, I(\hat{v}, D))\}\right] \tag{12.7}$$

where the function $R(s, I)$ represents the expected revenues collected during the second period if the retailer has a residual capacity of I units, and sets the price to s; see paper for details.

Note that at the beginning of period 2, the retailer knows the actual value of D. Specifically, for a given $I > 0$, the retailers infer that $D_1 = q - I$, and so

$$D = (q - I) \cdot \left(1 - \alpha \cdot \frac{\bar{v} - \hat{v}}{\bar{v} - \underline{v}}\right)^{-1}$$

Cachon and Swinney now show that if demand D is sufficiently large (which also means that the residual inventory level I is relatively small), the retailer should set the highest price that clears the Inventory I. If the demand is in some medium level, the retailer needs to optimally pick a price that maximizes the revenues collected from the remaining *strategic* customers. Typically, this will result in partial sales of the inventory I. Finally, if the demand is at a relatively low level, the retailer then prices at an inventory clearance price of v_B (making the product attractive to the large set of bargain hunters). To find the optimal level of inventory (namely, the solution to $\max_{q \geq 0} \pi(q, \hat{v})$), the optimal sales price in period 2 is substituted into (12.7). The authors prove that the function $\pi(q, \hat{v})$ is quasi-concave in q, and that the optimal order quantity is determined by the unique solution to the first-order condition $d\pi(q, \hat{v}) / dq = 0$.

Let us now examine how Cachon and Swinney analyze the customers' purchasing policy. As in most papers we previously discussed, strategic customers need to compare the current surplus (namely, $v_M - p_1$) with the expected surplus to be gained if they wait for period 2. Let us then consider a focal strategic customer with valuation equal to v^* (i.e., a valuation at the indifference point). A key technical observation in the analysis is that for such focal customer, the only scenario under which he can gain a positive surplus in period 2 is if the retailer sets the clearance price v_B. Recall, this event will happen if the demand D turns to be sufficiently small (say, below a given value D_l). Thus, the expected surplus for the focal consumer is

$$\left(v^*\left(\hat{q}\right) - v_B\right) \cdot \Pr\left\{D < D_l \text{ and the consumer receives a unit}\right\}$$

In sake of brevity, we refer the reader to Section 5 in the paper, which presents a rationing mechanism for allocating the inventory in case the demand in period 2 exceeds the residual quantity I.

Cachon and Swinney prove the following theorem.

Theorem 1 of Cachon and Swinney (2008). *There exists a subgame perfect Nash equilibrium with rational expectations* (q^*, v^*) *to the game between the retailer and strategic consumers, and any equilibrium satisfies* $q^* \leq q^m$ *and* $\pi^* \leq \pi^m$.

The superscript "m" refers to the *benchmark* case of myopic customers, reflected by the exact same model above, with α reset to 0. For example, we can learn from this theorem that the retailer orders less with strategic consumers compared to the case where he faces myopic customers only. In other words, this behavior can be viewed as an act of capacity rationing. It is noteworthy to consider the authors' statement vis-à-vis this result: *Other [researchers] have also found that the presence of strategic consumers causes a firm to lower its order quantity (e.g., Su and Zhang [2008a] and Liu and van Ryzin [2008]). However, the mechanism by which this result is obtained is different: they depend on rationing risk, whereas in our model the result is due to price risk – strategic consumers expect they will receive a unit in the markdown.* This observation is quite important in appreciating the value of the second model of Cachon and Swinney that we next present.

Consider now a situation in which the retailer can replenish his inventory at the beginning of period 2, *after* observing the demand D. Such type of replenishment option is generally feasible when the supply side (in terms of procurement, production, and delivery) is sufficiently responsive. Thus, the authors use the title *Quick Response* to reflect this situation. The analysis of the quick response setting appears to be very similar to that of the previous model. But, more importantly, the authors show that for a *given* value of $\hat{q}$, the customers' behavior does *not* change in comparison to the situation with no replenishment opportunity. The driver of this result is that in both cases (with or without quick response), strategic customers can gain positive surplus in period 2 only if the demand happens to be sufficiently low. If that happens, then the retailer will not exercise the quick response option anyway. But the reader should not be confused! As can be expected, the retailer will not order the same number of units (q) under both cases; we anticipate the initial order quantity to be lower with quick response, since the retailer can always order more later. In other words, the customers' behavior in equilibrium will be different (in general), depending on the feasibility of a quick response delivery. It is shown that

Theorem 2 of Cachon and Swinney (2008). *There exists a subgame perfect Nash equilibrium with rational expectations (q_r^*, v_r^*) to the game between the retailer with quick response and strategic consumers. It yields equilibrium expected profit π_r^* and satisfies $q_r^* \le q^*$ and $\pi_r^* \ge \pi^*$. Furthermore, if*

$$\frac{v_M - p}{\bar{v} - v_B} \ge \frac{c_2 - c_1}{c_2 - v_B},$$

then in equilibrium all strategic consumers purchase in the first period (c_2 is the per-unit procurement cost in period 2).

The result $q_r^* \le q^*$, together with the discussion above, means that quick response enables the retailer to increase the sense of rationing risk among strategic consumers, by ordering less. By driving strategic customers to purchase at the premium price p, the retailer gains an increased profit. Not only that, but with quick response, the retailer has an option to better match the supply with demand after observing D. Based on a numerical analysis, Cachon and Swinney note that the profit increase due to quick response can be dramatically higher when a retailer faces strategic consumers, than under settings with myopic customers only. This observation is key to the understanding of the potential benefits of quick response systems, and the ways in which it depends on consumer behavior.

Su and Zhang (2008a) propose an extension of the traditional "*newsvendor*" inventory model to incorporate the impact of strategic consumer behavior. In their setting, a single seller makes a choice of the capacity Q to bring to the market, as well as the price (p) to charge during the main season. Per-unit cost to the seller is c, and at the end of the season the seller must set the price to s (not a decision variable). Demand, denoted by X, is a random variable and is interpreted as the total mass of infinitesimal consumers in the market. The random variable X follows a distribution F. Consumers' valuations for the product are fixed at v. It is assumed that $s < c < v$. Customers' purchasing behavior is assumed to be governed by a threshold policy with a *reservation price* r. Specifically, if $r \ge p$, they all attempt to buy immediately at price p; otherwise, they all wait for the salvage price.

It is noteworthy that the level of inventory, picked by the seller, is not observable to the customers. Similarly, the reservation price r is known to the customers, but not directly observed by the seller. Thus, Su and Zhang propose to study a rational expectation equilibrium in which estimates of these values are formed by the seller and customers. Specifically, the seller forms a belief ξ_r about the customers' reservation price, while the customers form a belief ξ_{prob} about the probability of availability on the salvage market (which obviously depends on the seller's choice of Q). Based on these expectations, customers need to compare between the surplus gained by a "buy now" decision (i.e., $v - p$) and the expected surplus to be gained by waiting (i.e., $(v - s)\,\xi_{prob}$). Since r has an interpretation of an indifference point (between a "buy now" and "wait"), it is easy to see that

$$r\left(\xi_{\text{prob}}\right) = v - (v-s)\,\xi_{\text{prob}}$$

(where r is stated as a function of the customers' belief ξ_{prob}). From the seller's perspective, it is obviously optimal to set the price to be equal to the customers' reservation price. However, since this value is not observed, we can write, for the moment, $p = \xi_r$. The optimal quantity (for any given price p) is given by

$$Q(p) = \arg\max_{Q} \left\{\Pi(Q,p) = (p-s)\cdot E\left[\min\{X,Q\}\right] - (c-s)\,Q\right\}$$

The authors show that a rational expectation equilibrium is given by the solution to the equation $p = v-(v-s)\,F\left(Q(p)\right)$. They further show that

Proposition 1 of Su and Zhang (2008a). *In the rational expectation equilibrium, all customers buy immediately, and the seller's price and quantity are characterized by* $p_c = s + \sqrt{(v-s)(c-s)}$ and $F(Q_c) = 1 - \sqrt{(c-s)/(v-s)}$.

(The subscript "c" is used to refer to a centralized setting with a single seller.) One of the lessons taken of this proposition is that the seller lowers his stocking quantity under strategic consumer behavior. We will get back to this paper in the next section, in which we discuss the topic of capacity commitment.

Su and Zhang (2008b) propose a similar "newsvendor"-type model as in the previous setting, but with a slightly different feature. Here, instead of strategic customers contemplating between a "buy now" versus a "wait" decision, they are contemplating whether to "go to the store," or not. A "go to the store" action involves a search cost of h, which could be justified if the surplus from a purchase exceeds this value. However, customers face uncertainty about product availability, and so they risk losing the search cost if they came to the store and did not find an item available. The authors explore the outcome of the rational expectation game. The setting of the game is very similar to the aforementioned framework: The seller announces a price and selects a capacity (not directly observed by customers), and customers form a belief about the likelihood of finding a product available (this likelihood is not directly observed by the seller).

Ovchinnikov and Milner (2005) consider a firm that offers last-minute discounts over a series of periods. Their model incorporates both stochastic demand and stochastic customer waiting behavior. Two waiting behaviors are considered in their paper. In the first, called "the smoothing case," customers interpolate between their previous waiting likelihood and their observation of the firm's policy. In the second, named "the self-regulating case," customers anticipate other customers' behavior and the likelihood that they will receive a unit on sale. They show that, under the self-regulating case, it is generally optimal for the firm to set some units on sale in each period and allow the customer behavior to limit the number of customers that enjoy the benefit of the reduced price. In contrast, in the smoothing case, the firm can increase its revenues by following a sales policy that regulates the number of customers waiting. The authors conduct numerical simulations to illustrate the value of making decisions optimally, as compared to a set of reasonable benchmark heuristics. They find that the revenues can increase by about 5–15%. The paper also discusses how these benefits are affected if overbooking is allowed, and show that the impact of overbooking is greatly dependent on the proportion of high-valuation customers in the market.

12.4.1 Capacity Commitments

In the last part of their paper, Cachon and Swinney (2008) report on some results related to the question of whether or not a price commitment can perform better than subgame perfect dynamic pricing. It is noteworthy that unlike the model of Aviv and Pazgal (2008) (who studied a similar type of question), the current model assumes that the initial price is set exogenously, and so commitment here is on the markdown price s *only*. Another important difference is that in the first period, all strategic and myopic customers share the *same valuation*, v_M. It is for these reasons that Cachon and Swinney could argue that if the retailer had to commit for a discount, it would be optimal to not markdown at all; in other words, the retailer should use a *static-pricing* policy, equal to the exogenously set price p. In settings with no quick response, making a "no markdown" commitment is beneficial *only* when the ratio between the margins in the first period $(p-c)$ and the cost of leftover inventory $(p-v_B)$ is sufficiently high. This is because the gain from inducing purchase during period 1 outweighs the loss due to the inability of the retailer to salvage inventory in case of a low-demand realization. When the retailer has quick response capability, static pricing tends to be more beneficial. The reason behind this is that quick response reduces the likelihood of having significant leftovers, and thus the loss due to price commitment (or, alternatively viewed – the likelihood of a need for inventory clearance) is not high.

We now get back to Su and Zhang (2008a). This paper considers two types of commitments that can be made by the seller: a capacity commitment and a price commitment. The first commitment represents a situation in which the seller can order Q units and convince customers that this is indeed the quantity level. In this case, there is no need for the customers to form a rational expectation about ξ_{prob}, as they can simply calculate the availability probability via $F(Q)$. The seller will then price

$$p_q(Q) = v - (v-s)F(Q)$$

Consequently, the seller's optimal quantity decision is given by Q_q^*, the maximizer of the expected profit

$$\begin{aligned}\Pi_q(Q) &= (p(Q)-s)E[\min\{X,Q\}] - (c-s)Q \\ &= (v-s)\bar{F}(Q)E[\min\{X,Q\}] - (c-s)Q\end{aligned}$$

Let us now look at their model of *price commitment*. Here, it is easy to verify that the seller should commit not to reduce the price; i.e., use a static price of v. Given this commitment, all customers attempt to buy the product at price v, and hence the expected profit is given by $\Pi_p(Q) = vE[\min\{X,Q\}] - cQ$. Su and Zhang show that price commitment may increase the seller's profit when the production cost c is relatively low and when the valuation v is relatively high. However, there exist situations in which price commitments is not desirable. This happens when the valuation v is relatively small. The intuition behind this is similar in spirit to that argued by Cachon and Swinney (2008); see above.

Su and Zhang (2008b) also discuss two commitment strategies that the seller can use to improve expected profits. When capacity commitment can be made, it can be very valuable in that it encourages customers to spend the search cost and visit the store. This effect increases expected profit margins and leads the seller to set a higher capacity. As a consequence, with higher inventory, the customers indeed experience a higher level of product availability. A second type of mechanism studied in this paper is an *availability guarantee*. Here, the seller promises to compensate consumers, ex post, if the product is out of stock. They find that the seller has an incentive to over-compensate consumers during stockouts, relative to the benchmark case under which social welfare is maximized. Finally, the authors argue that first-best outcomes (i.e., those achieved under the benchmark case) do not arise in equilibrium, but can be achieved when the seller uses some combination of commitment and availability guarantees.

12.5 Internal Price-Matching Policies

In the retail industry, many companies offer some form of price guarantee to encourage customers not to delay their purchases; see Arbatskaya et al. (2004). One such offering, called an *internal*[5] price-matching guarantee, reflects a situation in which a retailer ensures that a customer will be reimbursed the difference between the current purchase price and any lower price the retailer might offer within a fixed *future time* period. The practice of internal price matching is very common in the retail industry and has been adopted (either formally or informally) by companies such as Amazon.com, Circuit City, and Gap. For example, if a product you purchased at Amazon.com within the last 30 days has dropped in price, they will typically credit you back the difference. In fact, a price-matching policy is effectively practiced by retailers who offer "free-return" policies. This is because customers who witness a price drop can simply return their purchased item for a full refund, and then immediately repurchase the same item at a lower price.

Obviously, in order to extract the maximum benefit from price matching, customers need to monitor the price constantly. From a retailer's perspective, a worst case refund scenario is one in which all customers are willing and able to monitor prices, and then exercise their right for a refund when applicable. In fact, with the evolution of technology and web-based services, this concern is becoming increasingly realistic. For example, Refundplease.com is a company that offers a web

[5] An alternative offering, called an *external* price-matching guarantee, is also popular in practice. Here, the retailer offers to match the price advertised by any *other* retailer at the time of purchase. Nonetheless, our focus in this section is on *internal* price-matching mechanism only.

service that simplifies the process for customers. Customers no longer need to keep track of the latest prices in order to get a credit. After a customer makes a purchase, he visits their web site and enters the purchase information. Refundplease.com will then check the prices everyday and send the customer a message if the price has dropped, and provide a link right to the place where the customer can claim a credit.

There are many possible variations of internal price-matching guarantees. One way is to offer a refund equal to the difference between what the customer paid and the marked down price (if such action took place within a specified time window); see, e.g., Debo et al. (2008). Another way is to compute a compensation (to customers) by taking the difference between the price of the product at some future point in time and the *strike price* offered at the time of purchase. There are several possibilities for the selection of the time when the price is compared to the strike price. One of them is to allow a customer to select it. Another is to let it be the time when the price is the lowest. Another set of the instrument's details is related to the notion of "similar product." In the case of airlines, for example, this could be "a ticket for the exact same flight," "a ticket for the same route by the same airline on the same date," or "a ticket on the same date by any airline with flights to this destination." Each definition of similarity will result in its own price for the service. For example, Levin et al. (2007) analyze the case of the same items and the comparison of the price with the strike price made at the time when the price is the lowest.

A limited amount of research on price matching has been done by economists and marketing scientists; see Hess and Gerstner (1991), Moorthy and Winter (2005), and Srivastava and Lurie (2001). These studies focus on the economics of price matching and do not involve revenue management practices. For example, Butz (1990) studies the impact of posterior price-matching policies in a setting with a durable goods sold over an infinite horizon. The seller is also a producer and can meet any demand quantity; i.e., unlike the situation in typical revenue management systems, there is no constraint on capacity (inventory). In his model, the seller offers a price-matching guarantee to buyers with a prespecified time window. Finally, this model assumes that demand is exogenous and free of strategic consumer behavior.

Levin et al. (2007) present a dynamic pricing model that includes an internal price guarantee instrument. Their paper considers a situation with *limited capacity*, reflective of revenue management settings. The internal price guarantee they consider provides a customer with compensation if the price of the product drops below a specified level (called the strike price). Customers have the choice of either buying or declining the guarantee; if they buy, then need to pay a fee. A price guarantee can increase the probability that customers will purchase at or near the time they first inquire about a product because it reduces their risk of future opportunity loss. For the company, an increased number of early purchases can reduce the uncertainty of late-purchasing "rushes" and last-minute price reductions, facilitate forecasting and capacity planning, and improve customer satisfaction and retention. Furthermore, the price guarantee itself constitutes a service provided for a fee in addition to the regular product price. Because this fee can be set by the seller so that it exceeds potential average losses from paying compensations, the collected fees provide an additional revenue stream. We elaborate below on some of the model details.

Assume that customers arrive according to a discrete-time counting process $N(t)$ with at most one arrival per time period and the probability of arrival is λ in each period. (The actual sales process is non-homogeneous in time in a manner described below.) The company has a total of Y items in inventory available for sale during T time periods. The company's policy is specified by a triple of values, representing the price, strike price (for the guarantee option), and the option purchase fee, given by the "quote": $\Pi(t) = (p(t), k(t), f(t))$, respectively. Thus, the quote is a three dimensional (stochastic) process. The guarantee is assumed to have a duration equal to D time periods, and the price guarantee *payments* are always made at the end of the planning period (i.e., at time T) regardless of the value of D. Specifically, if a customer bought an item with the price guarantee at time t, then the price guarantee payment to this customer will be equal to

$$\max\{\max(k(t) - p(\tau), 0) : t \leq \tau \leq \min\{t + D, T\}\}$$

Customers, who are assumed to be *myopic*, choose between not making a purchase, making a purchase without the price guarantee at price $p(t)$, or paying $p(t) + f(t)$ for the purchase with the guarantee. The consumer choice model is described by two probability functions. First, it is assumed that the probability that a customer makes either type of purchase upon arrival at time t is given by a function $u(p,k,f,t)$ .Thus, the effective probability of sale at time t is $\lambda(\Pi(t),t) = u(\Pi(t),t) \cdot \lambda$. Second, the conditional probability that a customer also purchases the price guarantee (given that a purchase is made) is given by $\upsilon(\Pi(t),t)$. The sales process is a two-dimensional counting process $(N_1(t), N_2(t))$, where $N_1(t)$ specifies the number of sales without price guarantees and $N_2(t)$ specifies the number of sales with price guarantees.

The problem of identifying an optimal policy for the retailer is not simple! At any point in time, the state of information consists of the values $(N_1(t), N_2(t))$ and the history of the process Π. It is easy to verify that the retailer's dynamic planning problem does not possess a Markovian property with respect to $(N_1(t), N_2(t))$, and hence an optimal policy will need to get involved with a solution of a dynamic program with a prohibitively large state space. To tackle this technical problem, Levin et al. confine themselves to a policy class that prescribes the values $\Pi(t)$ on the basis of the values $(N_1(t), N_2(t), t)$. They propose a nonlinear programming (NLP) approach to identify the best policy within this class. This approach can be implemented in relatively small problems. For large-scale problems, the paper offers a computationally tractable heuristic.

12.5.1 Internal Price Guarantees Under Strategic Consumer Behavior

Png (1991) considers a monopolist that sells a limited capacity (of size k) to customers whose valuations are either low (v_l) or high ($v_h; v_h > v_l$). A random portion of the customers $x \in [0,1]$ belongs to the high-valuation set; x follows a given statistical distribution Φ. There are two periods in the interaction between the seller and customers. If at any time there is excess demand, the available units are allocated at random. All customers are present at the "store" from the beginning of period 1. To study the equilibrium in the game, it is useful to begin from the second period, where it is clear that the price $p_2 = v_l$ should be set. Going back to the first period, we now consider the customers' choice. Suppose that a high-valuation customer decides to purchase in period 1. Then, the likelihood of getting a "unit" of the product is given by

$$\sigma_1(k) = \Phi(k) + \int_k^1 \frac{k}{x} d\Phi(x)$$

and therefore, the expected utility for buying in advance is $(v_h - p_1) \cdot \sigma_1(k)$. The latter value is based on an implicit assumption that all other high-valuation customers will act the same way as that of the "focal" customer. Now, let us consider the situation in which just the "focal" customer *deviates* from the buy now action, and decides to wait for period 2. In this case, he will receive the product in the second period with the probability

$$\sigma_2(k) = \int_0^k \frac{k-x}{1-x} d\Phi(x)$$

and hence his expected utility will be $(v_h - v_l) \cdot \sigma_2(k)$. Png argues that in order to maximize profit, the seller should set the first-period price p_1 so that each high-valuation customers will be indifferent between buying immediately and waiting. In other words,

$$p_1(k) = v_h - (v_h - v_l) \frac{\sigma_2(k)}{\sigma_1(k)}$$

The seller's expected revenues are given by

$$R(k) = v_l \cdot E[\max(k-x, 0)] + p_1(k) \cdot E[\min(x, k)]$$

In his second model, Png (1991) considers a *most favorable customer* (MFC) protection plan. Under this policy, the seller promises customers who buy early that they will receive a refund to cover for any subsequent price cut. Let $\hat{p}_1$ and $\hat{p}_2$ denote the prices in the first and second periods, respectively. It is clear that for high-valuation customers, the best strategy is to purchase in the first period. This is because when customers buy early, they increase their likelihood of product availability and they have nothing to lose on price. Therefore, it is optimal for the seller to set $\hat{p}_1 = v_h$. Given this choice, the low-valuation customers will wait to period 2. Consequently, there are two cases that need to be analyzed: $x \geq k$ and $x < k$. The first case is simple, since all units are purchased by the high-valuation customers, and the seller's profit is given by $v_h \cdot x$. In the second case, the seller faces the following trade-off in setting $\hat{p}_2$. If $\hat{p}_2$ is set to v_l, the seller can extract a revenue of $v_l \cdot (k - x)$, but will need to refund the high-valuation customers. Or, more simply, the seller will effectively charge all customer the price v_l, and gain a revenue of $v_l \cdot k$. If $\hat{p}_2$ is set to v_h, then the seller's revenue is $v_h \cdot \min(k, x) = v_h \cdot x$ (in view of $x < k$). The seller's revenue under the second case is therefore $\max(v_l \cdot k, v_h \cdot x)$, and therefore the total expected revenue with an MFC provision is equal to

$$\hat{R}(k) = E[\max(v_l \cdot k, v_h \cdot \min(x,k))]$$

A key result in Png's paper is that the seller always (weakly) prefers to guarantee MFC treatment to first-period buyers and sell over two periods than to sell in one period only (see Proposition 1 in his paper). He explains the intuition behind this finding by arguing that selling over two periods enables the seller to collect and use information about the customer demand. Formally, this type of advantage is described in the following inequality:

$$\hat{R}(k) = E[\max(v_l \cdot k, v_h \cdot \min(x,k))] \geq \max(v_l \cdot k, E[v_h \cdot \min(x,k)])$$

with the right-hand side representing the maximal expected revenues that can be gained under a fixed-price policy. Another important observation made by Png has to do with the comparison of MFC to no-MFC policies. He finds that MFC protection is the favorable choice for the seller when the capacity is large. The logic behind this result is that when capacity is high, customers have a large confidence that waiting will not significantly harm their likelihood of receiving the product at v_l (under the no-MFC policy). Therefore, customers will wait for the second period, resulting in minimal expected revenues. Png also finds that when customers are more uncertain about the degree of excess demand in the first period, they tend to buy early at the high price. As a consequence, such customer base is a good candidate for price discrimination, and no-MFC is the right choice for the seller.

Xu (2008) studies the optimal choice of internal price-matching policies, to which she refer by the term *best-price policies* (BP). Unlike the previous paper, Xu characterizes the *best choice of policy parameters*; namely, the time window during which the BP applies and the portion of the price difference that is refunded (*refund scale*). The models of this paper are set on the basis of the following set of assumptions. A seller and customers interact over an infinite horizon, with the seller essentially able to continuously change prices.[6] Customers are either high valuation or low valuation, and they all have the same discount rate (the seller, too, share the same discount factor). At some point (and only one point) of time the valuations of all customers jump simultaneously into a lower state, according to some probabilistic mechanism. This *random-shock* phenomenon is embedded in the model in order to represent situations in practice where a product is going out of fashion, or become obsolete. The main model in the paper considers the case where the seller offers a BP policy, and cannot commit to prices. Two benchmark models are also analyzed: A case in which no BP is offered, but the seller can *commit to prices* (prices are contingent on the information history), and a case in which no BP is offered and the seller cannot make a price commitment. Xu finds that a *finite* and positive BP policy can be optimal for the seller, when the likelihood for a sharp drop in evalua-

---

[6] The paper assumes that the price-change points are confined to the time epochs $t = 0, \Delta, 2\Delta, \ldots$, but it then focuses on the limiting case $\Delta \to 0$.

tion can take place. In other words, BP policies can be effective for retailers whose products may go out of fashion or become obsolete. The seller's equilibrium profit under the optimal BP policy falls between the profits of the two benchmark models. In general, the optimal BP policy cannot achieve the profit that can be gained with full commitment because of the uncertainty in the time of demand drop, and the fact that a BP policy cannot be contingent on the event of demand drop (since such event is unverifiable and non-contractible).

Debo et al. (2008) consider a *posterior price-matching policy*, a marketing policy offered by a seller to match the lower prices if the seller marks down within a specified time. In their model, the market consists of high-end (valuation = V_H) and low-end (valuation = V_L; $V_L < V_H$) consumers. They assume that the number of low-end consumers in the market is infinite, and that there is a large volume of potential high-end consumers, such that an individual consumer has a negligible effect on demand. The total volume of high-end consumers, denoted by λ, is unknown in advance, but can be characterized by a known distribution F, with a mean μ and a standard deviation σ. Information about the actual value of λ is gained via sales observations during the first period. In the second period, the high-end customer's valuations decline from V_H to V_h (where $V_L \leq V_h \leq V_H$). In contrast, the low-end consumer's valuations remain constant over the two periods. Among the high-end consumers, a fraction ϕ is strategic, whereas the rest are myopic. Strategic customers always request a refund, whereas only a fraction (γ) of the myopic customers do so. The seller determines whether or not to offer a posterior price matching (denoted by a binary variable υ), sets the first-period price (p_1), and invests in inventory (Q). These decisions are made *before* the market volume is realized. The unit acquisition cost, c, satisfies the condition $V_L < c < V_H$.

Debo et al. use the following dynamic procedure to evaluate the outcome of the two-period interaction between the seller and customers. Considering the second period, the seller is equipped with the information $(\upsilon, p_1, Q, q_s, s)$; the first three values are the seller's actions in the first period, whereas the latter values are observed during period 1. Specifically, q_s is the seller's belief about the customers' purchasing probabilities q, and s is the realized amount of sales during the first period. The seller's best choice for the second period's price is given by

$$p_2^o \in \arg\max_{p_2} R_2\left(p_2; \upsilon, p_1, Q, q_s, s\right) \tag{12.8}$$

where R_2 is the seller's second-period profit. Going back to the first period, in which customers need to determine their purchasing parameter q, they do so by weighing the expected utilities gained from a purchase in each period (the functions u_1 and u_2), as follows:

$$q^o \in \arg\max_{q} \left\{(1-q) \cdot u_1\left(\upsilon, p_1, Q_c, q_c, p_2^o\right) + q \cdot u_1\left(\upsilon, p_1, Q_c, q_c, p_2^o\right)\right\}$$

where Q_c is the seller's level of inventory that the customers rationally expect. Getting one step back to the seller's initial decision, the authors consider a two-level optimization process. First, for any given choice of p_1, the seller determines the optimal quantity of inventory through

$$Q^o(\upsilon, p_1) \in \arg\max_Q \{\pi(Q; \upsilon, p_1, q_s, p_2^o)\}$$

Note that the q_s represents the seller's belief about the customers' purchasing policy parameter q. The function π represents the seller's two-period profit with p_2^o satisfying (12.8). Finally, the seller's first-period price is determined via

$$p_1^o(\upsilon) \in \arg\max_{p_1} \{\pi(Q^o(\upsilon, p_1); \upsilon, p_1, q^o, p_2^o)\}$$

Obviously, to establish a rational expectation equilibrium, we need to find a solution to the above system, in which $q_s = q_c = q^o$ and $Q_c = Q^o$.

We present below a part of the solution to the models, focusing on the second-period pricing and the customers' choices. This will be sufficient for the reader to develop a sense of the fundamental difference between the cases with and without posterior price matching. Of course, the interested reader is referred to the paper for complete details.

Consider first the model without posterior price matching. Here, it can be shown that the price p_2 can be restricted to the two values V_h and V_L (either sell exclusively to strategic customers, or mark down further and clear the inventory). To compute the function R_2, consider two cases: (i) if $p_2 = V_L$, all remaining $Q - s$ units will be purchased; (ii) if $p_2 = V_h$, then only strategic customers will purchase. In this case, the seller needs to infer the size of the market from the sales volume s (note that only the case $s < Q$ is of interest, since if there is no inventory left, the second price has no meaning). In this case, the authors argue that the size of the high-end market is given by $s/(1-\phi q_s)$, and hence the number of high-end strategic customers that will purchase in period 2 is estimated by $q_s \phi s/(1-\phi q_s)$. To summarize, we get

$$R_2 = \begin{cases} (Q-s) \cdot V_L & \text{if } p_2 = V_L \\ \min\{Q-s, q_s\phi s/(1-\phi q_s)\} \cdot V_h & \text{if } p_2 = V_h \end{cases}$$

As a consequence (see Proposition 1 in the paper), the seller's second-period price can be shown to satisfy

$$p_2^o = \begin{cases} V_L & \text{if } s/(1-\phi q_s) < \alpha(q_s) Q \\ V_h & \text{if } \alpha(q_s) Q \le s/(1-\phi q_s) \le Q/(1-\phi q_s) \\ \text{n/a} & \text{if } s = Q \end{cases}$$

where $\alpha(q) \doteq V_L/(V_L + q\phi(V_h - V_L))$. Let us now consider the strategic customers' decisions. If the customer purchases immediately (i.e., in period 1), he gains a surplus of $(V_H - p_1)$. However, a unit needs to be available. Clearly, if the total demand is given by a quantity $\lambda \le Q_c/(1-\phi q_c)$, a focal customer will be allocated a unit upon purchase. If, however, the latter inequality is reversed, the Q_c units will need to be allocated among the $(1-\phi q_c)\lambda$ customers who decide to purchase immediately. Thus,

$$u_1 = \left[F\left(\frac{Q_c}{1-\phi q_c}\right) + \int_{Q_c/(1-\phi q_c)}^{\infty} \frac{Q_c}{(1-\phi q_c)\lambda} f(\lambda)\,d\lambda\right](V_H - p_1)$$

If a focal customer decides to postpone the purchase to the second period, then he expects the price p_2 to be either V_h (with probability $1-F(\alpha(q_c)Q_c)$), or V_L with the complementary probability. Thus, we get

$$u_2 = F(\alpha(q_c)Q_c)\cdot(V_h - V_L)$$

Let us consider now the case in which a posterior price-matching policy is offered. Here, it is also useful to consider the values V_L or V_h for p_2. However, the seller may want to consider charging $p_2 > \max\{p_1, V_h\}$. The rational behind the last price possibility is as follows: if a price $p_2 > V_h$ is expected, then the no strategic customer is expected to wait for period 2. Then, setting the price below p_1 is ineffective since it will result in myopic high-end customers asking for refund. When the seller perception about the delay probability (i.e., q_s) is sufficiently small, the best pricing in period 2 is to either clear leftover inventory ($p_2 = V_L$) or price above $\max\{p_1, V_L\}$. It is only when q_s passes a certain threshold that it is worthwhile to consider the price $p_2 = V_h$; see Proposition 5 in the paper. Interestingly, when it comes to purchasing behavior of customers, the authors state (Lemma 4): The unique purchasing equilibrium is to buy immediately; i.e., $q^o = 0$. It is easy to see the reasoning behind this result. For a focal strategic customer, the posterior price-matching policy enables to obtain refund if prices decline. Furthermore, the likelihood of obtaining a unit can only decline from period 1 to period 2. As a consequence of this observation, it is easy to verify that the optimal first-period price is $p_1^o = V_H$. As can be seen, the seller's first-period price and the customers' waiting strategy are both independent of ϕ – the fraction of strategic customers. The fraction ϕ influences the salvage value of the leftover inventory in a monotonic way. The largest is ϕ, the larger are the refunds in case of inventory clearance. Consequently, the authors show that the seller's equilibrium inventory level Q^o and his optimal expected profit are both monotonically decreasing in ϕ; see Proposition 8 in the paper.

Based on the above analytical findings and further numerical analyses, the authors conclude that price-matching policies eliminate strategic consumers' waiting incentive and thus allows the seller to increase the price in the regular selling season. When the market consists of a sufficiently large fraction of strategic consumers with declining valuations (over time), the matching policy can be very effective. In contrast, price-matching policies can be detrimental when there are only a few strategic consumers in the market, or if the strategic consumers' valuations do not decline much during the sales horizon. Finally, they find that the ability to credibly commit to a fixed-price path is not very valuable when the seller can implement price matching.

12.6 Limiting Inventory Information

Recently, Yin et al. (2008) have proposed a game-theoretic model of a retailer who sells a limited inventory of a product over a finite selling season, using one of the two inventory display formats: Display All (DA) and Display One (DO). Under the DA format, the retailer displays all available units so that each arriving customer has perfect information about the actual inventory level. Under the DO format, the retailer displays only one unit at a time so that each customer knows about product availability but not the actual inventory level. Clearly, display formats can be used as a tool to influence customers' perceptions about the risk of stockouts if they decide to wait. Therefore, by optimally selecting the display format, a retailer could discourage high-valuation customers from waiting to the clearance sales. Focusing on price-commitment strategies, Yin et al. address the following questions: When considering the influence of the display formats on the level of inventory information conveyed to customers, which one of the two formats is better for the retailer? Furthermore, can a move from one display format to another be effective in mitigating the adverse impact of strategic consumer behavior? They find support to the hypothesis that the DO format could potentially create an increased perception of scarcity among customers, and hence it is better than the DA format. However, while potentially beneficial, the move from a DA to a DO format is very far from eliminating the adverse impact of strategic consumer behavior. Since this paper is surveyed in great detail on a separate chapter in this book, we omit the technical details.

References

Arbatskaya M, Hviid M, Shaffer G (2004) On the incidence and variety of low price guarantees. Journal of Law and Economics XLVII:307–332

Aviv Y, Pazgal A (2005) A partially observed Markov decision process for dynamic pricing. Management Science 51(9):1400–1416

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

Bell DE, Starr D (1994) Filene's basement. Harvard Business School Case 9-594-018

Besanko D, Winston WL (1990) Optimal price skimming by a monopolist facing rational consumers. Management Science 36:555–567

Butz D (1990) Durable goods monopoly and best-price provisions. Amer Econ Rev 80(5):659–683

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. Management Science 55(3): 497–511

Coase RH (1972) Durability and monopoly. J Law Econom 15:143–149

Debo L, Lai G, Sycara K (2009) Buy now and match later: The impact of posterior price matching on profit with strategic consumers. Manufacturing & Service Operations Management, Published online in *Articles in Advance*, February 10, 2009 DOI: 10.1287/msom.1080.0248

Elmaghraby W, Gulcu A, Keskinocak P (2008) Designing optimal pre-announced markdowns in the presence of rational customers with multi-unit demands. Manufacturing and Service Operations Management 10(1): 126–148

Hess J, Gerstner E (1991) Price matching policies: An empirical case. Managerial Decision Economics 12:305–315

Levin Y, McGill J, Nediak M (2007) Price guarantees in dynamic pricing and revenue management. Operations Research 55(1):75–97

Liu Q, van Ryzin G (2008) Strategic capacity rationing to induce early purchases. Management Science 54(6):1115–1131

Mantrala MK, Rao S (2001) A decision-support system that helps retailers decide order quantities and markdowns for fashion goods. Interfaces 31(3 part 2):S146–S165

Moorthy S, Winter R (2006) Price-matching guarantees. RAND Journal of Economics 37(2): 449–465.

Ovchinnikov A, Milner JM (2005) Strategic response to wait-or-buy: Revenue management through last minute deals in the presence of customer leaning. University of Toronto. Working paper

Png IP (1991) Most-favored-customer protection vs price discrimination over time. Journal of Political Economy 99:1011–1028

Srivastava J, Lurie N (2001) A consumer perspective on price-matching refund policies: Effect on price perceptions and search behaviour. J Consumer Res 28(2):296–307

Su X (2007) Inter-temporal pricing with strategic customer behavior. Management Science 53(5):726–741

Su X, Zhang F (2008) Strategic customer behavior, commitment and supply chain performance. Management Science 54(10): 1759–1773

Su X, Zhang F (2009) On the value of commitment and availability guarantees when selling to strategic consumers. Management Science 55(5): 713–726

Xu Z (2008) Optimal best-price policy. Working paper

Yin R, Aviv Y, Pazgal A, Tang CS (2008) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. Management Science.