



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Responsive Pricing of Fashion Products: The Effects of Demand Learning and Strategic Consumer Behavior

Yossi Aviv, Mike Mingcheng Wei, Fuqiang Zhang

To cite this article:

Yossi Aviv, Mike Mingcheng Wei, Fuqiang Zhang (2019) Responsive Pricing of Fashion Products: The Effects of Demand Learning and Strategic Consumer Behavior. Management Science 65(7):2982-3000. <https://doi.org/10.1287/mnsc.2018.3114>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2019, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Responsive Pricing of Fashion Products: The Effects of Demand Learning and Strategic Consumer Behavior

Yossi Aviv,^a Mike Mingcheng Wei,^b Fuqiang Zhang^c

^a Faculty of Management, Tel Aviv University, Tel Aviv 6997801, Israel; ^b School of Management, University at Buffalo, Buffalo, New York 14260; ^c Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130

Contact: yaviv@tau.ac.il,  <http://orcid.org/0000-0002-6479-2583> (YA); mcwei@buffalo.edu (MMW); fzhang22@wustl.edu,

 <http://orcid.org/0000-0003-2918-2613> (FZ)

Received: February 1, 2013

Revised: July 16, 2015; January 27, 2017; March 15, 2018

Accepted: April 25, 2018

Published Online in Articles in Advance: May 17, 2019

<https://doi.org/10.1287/mnsc.2018.3114>

Copyright: © 2019 INFORMS

Abstract. This paper studies the potential benefits of responsive pricing and demand learning to sellers of seasonal fashion goods. As typical in such markets, demand uncertainty is high at the beginning of a season, but there is a potential opportunity to learn about demand via early sales observations. Additionally, although the consumers have general preference for purchasing a fashion product earlier rather than later in the season, they may exhibit strategic behavior—contemplating the benefits of postponing their purchase in anticipation of end-of-season discounts. Our results demonstrate that the benefits of responsive pricing, in comparison with a benchmark case of a fixed-price policy, depend sharply on the nature of the consumers’ behavior. Interestingly, in stark contrast to markets of myopic consumers, when the consumers are all strategic, the benefits of responsive pricing tend to worsen when there is a higher potential for learning. We explain this counterintuitive outcome by pointing to two phenomena: the spread effect and information shaping. For example, sellers of fashion products that consider upgrading their pricing systems to incorporate “*accurate response*” strategies (i.e., integrating learning and responsive pricing) should be aware of the possibility that such action might lead them to a new and potentially worse equilibrium, particularly when there is a higher opportunity to learn. Despite the fact that price commitment completely eliminates the seller’s ability to learn, it appears to increasingly dominate responsive pricing as the portion of strategic consumers in the market increases. But, although performing better than responsive pricing, a price-commitment policy is typically limited in performing effective discrimination. Finally, we studied the potential benefits of quick response strategies—ones that embed both dynamic pricing and quick inventory replenishment during the sales season—and found that they are particularly significant under strategic consumer behavior. We explain this result by arguing that quick response provides the seller with a real option that serves as an effective implicit threat to the consumers: encouraging them to buy earlier at premium prices rather than wait for discounts at the end of the season.

History: Accepted by Martin Lariviere, operations management.

Funding: F. Zhang’s research was supported by the National Natural Science Foundation of China [Grants 71628203 and 71520107002].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2018.3114>.

Keywords: demand learning • dynamic pricing • strategic consumer behavior • revenue management • game theory

1. Introduction

Dynamic pricing of fashion products is a complex but potentially rewarding revenue management activity (see, e.g., Elmaghraby and Keskinocak 2003, Van Ryzin and Talluri 2005, and Aviv and Vulcano 2012 for comprehensive reviews). The complexity arises because of several factors, such as the need to take capacity levels into account, the relatively high degree of market uncertainty, and the nature of the consumers’ purchasing behavior. Research in this area has studied two important facets of this topic: *learning* and *strategic consumer behavior*.

Learning in dynamic pricing systems is of particular relevance to fashion products (see, e.g., Heun 2001). For

example, Fisher (2009) argues that, in retail settings of innovative products, a seller can learn about demand from early sales. Based on data gathered from a ski-wear supply chain, he observed that the demand forecast error could be reduced from 55% to a mere 8% after utilizing just two weeks of sales information. The implication of this observation on the potential value of dynamic pricing is significant. When selling a fashion product, a retailer can benefit from the ability to change the price later in the sales season in response to sales observations—in other words, the ability to act upon information; see research papers such as Balvers and Cosimano (1990), Aviv and Pazgal (2005), Besbes and

Zeevi (2009), Araman and Caldentey (2009), and Farias and van Roy (2010) as well as the case study in Ghemawat and Nueno (2003).

Unfortunately, in many retail environments, consumers are cognizant of the retailers' pricing strategies, or at least, they are aware of the possibility that price discounts may be offered later in the season. Consequently, the consumers may exhibit "strategic behavior"—that is, making a choice between purchasing the product at a premium price early in the season or postponing their purchasing decisions to a later time in the season when a price discount may be offered. Research on this topic in the economics literature goes back to the seminal paper of Coase (1972) and other key papers such as Stokey (1979) and Landsberger and Meilijson (1985). Over the past decade, the study of the effects of strategic consumer behavior on dynamic pricing of *finite capacities* has experienced significant growth; see, for example, Su (2007), Elmaghraby et al. (2008), Aviv and Pazgal (2008), Aviv et al. (2009), and Su and Zhang (2008, 2009). One of the key messages raised by the latter literature is that a retailer's ability to vary prices during a sales season can work to its detriment when the market consists of strategic consumers. Specifically, knowing that the retailer has the flexibility to change prices, consumers expect that the retailer will offer markdowns later in the season, which, in turn, encourages them to avoid paying the premium prices charged at the beginning of the sales season. In many cases, this phenomenon can turn a dynamic pricing policy to subperform even a fixed-price policy. For brevity of exposition, we refer the interested reader to Aviv and Vulcano (2012) for a detailed survey of the existing literature on strategic consumer behavior and to Netessine and Tang (2009), which contains several review chapters dedicated to this particular subject.

These contradicting messages, one in favor of *responsive pricing* (resulting from learning), and one against it (resulting from strategic consumer behavior), leave a serious question opened: are responsive pricing strategies effective or counterproductive when selling fashion-like products? We use the title "responsive pricing" to emphasize that the seller can adjust prices in reaction to early sales information. Similarly, we use the title "fashion-like products" to emphasize our focus on markets characterized by a high degree of uncertainty but with an opportunity to (partially) resolve it via early sales observations. To address our question, we propose an integrative theoretical model that reflects both learning and strategic consumer behavior. Our model is stylized in nature and considers a monopolist retailer that sells a finite quantity of a given fashion product over a two-period horizon. The seller selects a price for the first period based on limited demand information. Then, the seller adjusts the price in the second period based on the sales made in the first period. We utilize a Bayesian learning mechanism

to model the demand learning process, through which information about the average market size for the product evolves. As we discuss later, this learning process brings up an important connection between pricing decisions and the value of information extracted from sales data. Finally, our model considers a heterogeneous pool of strategic consumers whose valuations are drawn from a uniform distribution.

Our model enables us to examine four interesting subjects pertaining to the use of responsive pricing strategies. First, we analyze the influence of the degree of resolvable market uncertainty on the effectiveness of responsive pricing. Without strategic consumers in the market, the more one can learn about the market in the first period, the more one is interested in being able to act upon such information. However, this logic breaks down in the face of strategic consumer behavior. Interestingly, we find that strategic consumer behavior becomes increasingly detrimental as the market provides a higher potential for learning. We explain this counterintuitive outcome by pointing to two phenomena: the spread effect and the active learning effect. Furthermore, we argue that the latter effect is driven by information shaping—the conscious collective attempt of the consumers to influence the seller's interpretation of the market size. Second, we constructed a hypothetical model in which we eliminated the seller's ability to learn from sales, and surprisingly, we observe that the seller's expected revenue performance actually improves. This outcome is of potential interest to sellers of fashion products who consider upgrading their pricing systems to incorporate "accurate response" strategies (i.e., integrating learning and responsive pricing). Third, we explore the effectiveness of *price-commitment strategies*. Our results suggest that despite the fact that price commitment completely eliminates the seller's ability to learn, it appears to increasingly dominate responsive pricing as more strategic consumers are present in the market. Moreover, the most significant benefits of commitment were observed for the cases in which the degree of market uncertainty was high. However, although performing better than responsive pricing, when the market size is highly uncertain and when the seller's inventory is sufficiently large, price commitment is very limited in its ability to perform effective discrimination. Finally, we studied the potential benefits of *quick-response strategies*—ones that embed both dynamic pricing and quick inventory replenishment during the sales season. Our results suggest that the option to replenish inventory close to the end of the season serves as an effective implicit threat by the seller to the consumers in a way that discourages strategic waiting. Consistent with the earlier paper of Cachon and Swinney (2009), we also argue that the benefits of quick response strategies might be significantly

underestimated in the presence of markets with strategic consumers.

The rest of the paper is organized as follows. In Section 2, we introduce and analyze the main model. Section 3 presents our numerical study and our key discussions and insights. In Section 4, we discuss markets in which only a portion of the consumers behave strategically. Section 5 explores the effectiveness of price-commitment strategies, and Section 6 studies the potential benefits of quick-response strategies. The paper's conclusions are presented in Section 7. All proofs are given in the online appendix.

2. Main Model

We develop a game-theoretic model to study the research questions listed in the introduction. Consider a seller who sells Q units of a product to consumers in a relatively short season, divided into two periods: a regular selling period (we shall refer to it as the “first period”) and a final sales period (the “second period”). At the beginning of the first period, the seller sets a price, p_1 . Then, at the beginning of the second period, the seller can change the price to p_2 , after observing the initial sales.

To reflect the demand uncertainty as well as the seller's ability to learn from sales, we adopt a special type of statistical framework known as the *Gamma–Poisson model*. Specifically, we assume that the number of consumers in the market is drawn from a Poisson distribution with a mean value of Λ that is unknown. The value of Λ —the average market size—has a *prior* Gamma distribution with a mean $E[\Lambda]$ and a coefficient of variation denoted by $CV \doteq \sqrt{\text{Var}(\Lambda)}/E[\Lambda]$. Clearly, because sales realizations depend on the market size, the seller is able to resolve part of the uncertainty about Λ after observing the number of units sold in the first period. Note that the a priori (i.e., before sales observation) mean and standard deviation of the demand are equal to $E[\Lambda]$ and $\sqrt{E[\Lambda] + E^2[\Lambda] \cdot CV^2}$, respectively.

Among the possible alternatives for statistical models of demand (such as the normal–normal or the beta–binomial models; see, e.g., Raiffa and Schlaifer 2000), the Gamma–Poisson model is appropriate as it allows the specification of large demand uncertainty scenarios that are often observed in fashion product sales. For example, the well-known article of Fisher (1997) reports that the demand forecasts for new products offered by a fashion skiwear designer (Sport Obermeyer) often err by as much as 200%. For the sake of illustration, suppose that $E[\Lambda] = 10$ and that the a priori standard deviation of the demand is equal to 100%–200% of that mean demand. This would translate to CV levels in between one and two.

The consumers in the market are heterogeneous in their valuations, which are independently drawn from a uniform distribution on $[0, 1]$. If a consumer chooses to wait for the sales period, then the consumer's

valuation for the product will be discounted by a factor $\delta \in (0, 1]$. Such a decline in valuation reflects a situation in which a product consumed later in the selling season provides a lower utility (e.g., purchasing a bathing suit at the end of a summer) or a situation in which consumers are “style conscious.” Unlike perishable products for which one would expect low δ values or durable goods for which one would expect relatively high δ values, we believe that fashion-like products are generally associated with midlevels of δ . In our main model, the consumers are *strategic* in the sense that they select the timing of their purchases by comparing the expected surpluses associated with “buy now” versus “wait” (to the second period) decisions. When appropriate, we also discuss and present results for the case in which the consumers in the market are all *myopic*; namely they do not consider future pricing in their purchasing decisions.

Our analyses are based on the study of the following sequential game. In the first stage, the seller posts the price p_1 . Thereafter, the consumers and the seller become engaged in a competitive situation (the *subgame*) in which the consumers decide on whether to buy in the first period or wait for a possible discount, and the seller determines its second-period price as a function of the quantity sold in the first period. We analyze the subgame by solving for a Nash equilibrium between the seller and the consumers. All problem parameters are common knowledge in this game. In addition, we assume that the seller and the consumers are risk neutral and aim at maximizing their expected payoff (revenue and surplus, respectively). Because the seller can set the second-period price in response to the initial sales and leftover inventory information, we refer to the dynamic pricing strategy as “responsive pricing.”

In the next two subsections, we first conduct the equilibrium analysis of the subgame (Section 2.1); then, we study the seller's optimal pricing decision in the first stage (Section 2.2).

2.1. Analysis of the Subgame

The analysis of the subgame is challenging because of two complicating factors. First, the consumers' purchasing decisions must be based on the seller's second-period pricing policy, which, in turn, depends on the consumers' purchasing strategy. Second, as in most models that study strategic consumer behavior, the consumers also compete against each other because their decisions affect the product availability in both periods. Any equilibrium in the subgame requires that each player adopts the optimal (payoff-maximizing) strategy given the rest of the players' strategies. The equilibrium analysis proceeds in two steps as follows. First, we analyze the consumers' and the seller's problems in the subgame separately; then, we characterize the equilibrium by tying the two parties' problems together. Because the initial level of inventory Q and the first-period

price p_1 are common knowledge at this stage, we generally avoid referring to them explicitly in the analysis of this section.

2.1.1. The Consumers' Purchasing Problem. Let $\{p_2(q) : q = 0, \dots, Q\}$ be any arbitrary second-period pricing strategy adopted by the seller, where q denotes the remaining inventory immediately after the first-period sales. The consumers need to decide whether to purchase in the first period. Specifically, a consumer with valuation v compares between the expected surplus gained from an immediate purchase and the expected surplus associated with waiting for the second period. Both assessments of the surpluses are complicated by the fact that the consumers must weigh in the likelihood that a unit will be available to them during these periods. Theorem 1 establishes that for any arbitrary pricing strategy of the seller and any arbitrary purchasing behavior of the other consumers in the market, it is optimal for an individual consumer to adopt a so-called *threshold policy*.

Theorem 1. *Consider a consumer with a given valuation v . For any initial inventory level Q , pricing scheme $\{p_1, p_2(\cdot)\}$, and any set of purchasing rules adopted by all other consumers in the market, it is optimal for that consumer to follow a threshold policy θ . Namely, the consumer will attempt to purchase a unit at the price p_1 if $v > \theta$; otherwise, the consumer will wait for the second period. Furthermore, the optimal threshold value is unique.*

The proof of Theorem 1 appears in the online appendix. Note that, given the specification of the information structure in our model, the probabilistic assessment of the purchasing rules adopted by all consumers in the market must be the same from each consumer's perspective. Therefore, Theorem 1 implies that, in equilibrium, all consumers must adopt the same unique threshold value. Thus, mixed-strategy equilibria are not predicted in this game.

Consider a setting in which all consumers adopt a threshold purchasing policy θ under a given pricing strategy $\{p_1, p_2(\cdot)\}$. Let $X(\theta)$ be the number of immediate purchases in the first period, which follows a Poisson distribution with a mean of $\Lambda \cdot (1 - \theta)$. In the following, we consider a particular consumer with valuation v and evaluate the expected surplus that is gained by each of the two possible actions. There are two common approaches for modeling the expected surplus values: the conditional¹ and the unconditional probability calculations. Under the unconditional approach (see, e.g., Yin et al. 2009), which we adopt in this paper, a consumer evaluates the current surplus by multiplying the potential surplus $(v - p_1)$ by the probability that a unit would be available in the current period if the consumer demands it. This is naturally the

case if a consumer buys the product online and there is no indication of product availability or if the consumer considers traveling to the store to purchase the product based on the advertised price.

To calculate the probability of a unit being available to an individual consumer during a given period, we consider the following model setup. Suppose that all consumers interested in purchasing the product arrive during a very small interval of time at the very beginning of the first period, according to a Poisson process with a rate $a = \Lambda(1 - \theta)$. Then, from any individual consumer's perspective, the size of the rest of the market (say Y) remains Poisson with the same parameter. Furthermore, suppose that if demand cannot be fully satisfied during a given period, all available units are randomly allocated among the consumers (who attempt to buy) with equal probabilities. Consequently, the likelihood that a unit is allocated to a given consumer can be written using the expression

$$H(Q, a) \doteq E \left[\frac{Q}{\max\{Q, Y + 1\}} \right] \\ = \sum_{y=0}^{Q-1} a^y e^{-a} / y! + \sum_{y=Q}^{\infty} \left(\frac{Q}{y+1} a^y e^{-a} / y! \right).$$

Hence, if the consumer makes a buy-now decision, then the consumer expects to gain a surplus of

$$s_1(v|\theta) \doteq (v - p_1) \cdot E_{\Lambda}[H(Q, \Lambda(1 - \theta))]. \quad (1)$$

If the consumer decides to wait for the second period, the consumer's expected surplus would be

$$s_2(v|\theta) \doteq E_{\Lambda} \left[E_{X(\theta)} \left[(\delta v - p_2(Q - X(\theta)))^+ \cdot H(Q - X(\theta), \Lambda(\theta - p_2(Q - X(\theta))/\delta)^+) \mid \Lambda \right] \right]. \quad (2)$$

(Note that both functions depend on $\{Q, p_1, p_2(\cdot)\}$, but we omit those parameters for brevity of exposition.) The expression within the inner expectation is once again a product of the surplus that could be gained from a purchase in the second period and the likelihood that a unit will be available in that period. Note that all (and only) consumers in the segment $(p_2/\delta, \theta]$ will be interested in buying the product in the second period if a price p_2 is set at that time.

As shown in the proof of Theorem 1, the consumer's best action is given by a unique threshold value, say $\theta^*(\theta)$, for which $s_1(v|\theta) > s_2(v|\theta)$ for all $v > \theta^*(\theta)$, and $s_1(v|\theta) \leq s_2(v|\theta)$ for all $v \leq \theta^*(\theta)$. The following proposition states that an equilibrium in the consumers' game always exists.

Proposition 1. *For any given Q and pricing scheme $\{p_1, p_2(\cdot)\}$, there exists a Nash equilibrium $\theta^* \in [p_1, 1]$ in the*

consumers’ game, in which all consumers adopt the same θ^* threshold purchasing policy.

We emphasize that the equilibrium threshold θ^* may not be unique in the consumers’ game for some arbitrary pricing strategies. In fact, we have found numerical examples in which the game possesses up to three equilibria. This issue need not be of concern. First, the technical analysis of the subgame between the seller and the consumers remains the same regardless of the predicted equilibria in the consumers’ game. Second, as we demonstrate later, no multiple equilibria appeared in our analysis of the game in which optimal prices (rather than arbitrary) are used.

2.1.2. The Seller’s Second-Period Pricing Problem. We now turn to the analysis of the seller’s best pricing policy $p_2(\cdot)$ in the second period in response to any given consumer threshold policy θ and in the presence of demand learning. In the second period, prices are determined on the basis of the first-period sales and the anticipated consumer purchasing behavior. Clearly, the first-period sales not only affect the remaining quantity, but they also influence the seller’s perception of the market size. Suppose that the seller updates its belief about the expected market size, Λ , using a Bayesian updating scheme, and recall that Λ has a Gamma prior distribution. In particular, the characteristics of the posterior distribution of Λ can be fully determined using two variables: the first-period sales (say x) and the consumers’ policy parameter (say θ). Note that although the first period price (p_1) and the initial inventory (Q) influence learning, we do not need to include them explicitly in the updating mechanism because the knowledge of x and θ is sufficient. By construction, the updated distribution of the market size is still Poisson, and we refer to that posterior random variable as $\tilde{\Lambda}(x, \theta)$.

Let q be the remaining quantity at the beginning of the second period. Then, the seller’s pricing problem at that time can be written in the following form:

$$\pi_2^*(q, x, \theta) = \max_{p_2 \in [0, \delta\theta]} \{p_2 \cdot E_{\tilde{\Lambda}(x, \theta)}[N(q, \tilde{\Lambda}(x, \theta)(\theta - p_2/\delta))]\}, \quad (3)$$

where $N(q, a)$ is defined as the expected number of units sold when there are q units available and the number of consumers interested in purchasing is Poisson-distributed with a mean of a . Theorem 2 establishes that the function within the maximization problem (3) is strictly concave and possesses a unique and interior solution in $[0, \delta\theta]$.

Theorem 2. *The objective function in (3) is strictly concave in p_2 . The optimal solution to (3), $p_2^*(q, x, \theta)$, lies in the interior of the range $[0, \delta\theta]$, and satisfies the first-order condition*

$$E_{\tilde{\Lambda}(x, \theta)} \left[N(q, \tilde{\Lambda}(x, \theta) \cdot (\theta - p_2/\delta)) - p_2 \cdot \frac{\tilde{\Lambda}(x, \theta)}{\delta} \sum_{y=0}^{q-1} P_y(\tilde{\Lambda}(x, \theta) \cdot (\theta - p_2/\delta)) \right] = 0, \quad (4)$$

where $P_y(\lambda)$ is the probability mass function at y of Poisson distribution with mean λ . Furthermore, the optimal solution $p_2^*(q, x, \theta)$ is increasing in θ and x and decreasing in q .

Theorem 2 demonstrates that the second-period optimal prices are monotonically increasing in the consumer threshold policy θ . This result can be viewed as a combination of two effects, which we qualitatively explain as follows. The first effect is direct: a larger θ for given values of q and x means that the seller is selling the same quantity but to a larger segment of customers (with equal or higher valuations); hence, the optimal second-period price is expected to increase. The second effect is indirect, driven by the learning mechanism. For any given x , the seller needs to infer the value of Λ based on purchases that are generated via customers in the valuation segment $[\theta, 1]$. Thus, a larger θ leads the seller to statistically estimate the value of Λ to be higher. Consequently, the seller is expecting a larger number of customers to be interested in purchasing the product in the second period; in turn, this leads to a similar effect on the optimal second-period price. The increase in x can be explained using precisely the latter argument. The monotonicity of the optimal price in q —higher inventory leads to lower second-period prices for given x and θ —is also intuitive, given that the learning effect is kept constant.

By combining these analyses of the consumers’ and the seller’s problems, we are ready to prove the existence of subgame equilibrium for any arbitrary first-period price p_1 . Again, in any equilibrium $\{\theta^*, p_2^*(\cdot)\}$, each player must adopt its optimal strategy in response to all other players’ strategies.

Proposition 2. *For any arbitrary first-period pricing p_1 , there exists a Nash equilibrium in the subgame between the consumers and the seller.*

Unfortunately, because of the complexity of the players’ payoff functions, we were unable to analytically prove that the subgame possesses a unique equilibrium. However, we fully address this standing concern in Proposition 3.

2.2. Optimal Pricing in the First Stage

The seller’s first-period pricing decision is an optimization problem that involves a search for the best price p_1^* given the initial level of inventory Q . This problem is quite challenging even without the significant complexity introduced by the presence of strategic consumer behavior. Note that the seller must optimally settle the following trade-off. On one hand, the seller may wish to charge a relatively high initial price so as to collect high

revenues in the first period; on the other hand, a high price limits the volume of early sales, which constrains the seller’s learning ability and, hence, imposes a negative effect on future revenues. This dilemma is referred to as the *exploitation-versus-exploration* trade-off in the literature; see, for example, Araman and Caldentey (2009). In our model, when setting the initial price p_1 , the seller must take into account the subgame discussed in the previous section. In particular, for any given choice of p_1 and any predicted consumer behavior θ , the seller’s revenue performance can be calculated using the conditional expectation form

$$\pi^R(Q, p_1, \theta^*) \doteq p_1 \cdot E_\Lambda[N(Q, \Lambda \cdot (1 - \theta^*))] + E_\Lambda \left[\sum_{x=0}^{Q-1} \pi_2^*(Q - x, x, \theta^*) \cdot P_x(\Lambda \cdot (1 - \theta^*)) \right].$$

(We use the superscript R to denote responsive pricing.) Thus, the seller’s first price–decision problem can be stated as follows:

$$\pi^R(Q) \doteq \max_{p_1 \in [0,1]} \{ \pi^R(Q, p_1, \theta^*(p_1)) \}, \quad (5)$$

where $\theta^*(p_1)$ is the predicted consumer behavior in the subgame (again, the parameter Q is omitted for convenience). To solve this problem, we could have treated $\pi^R(Q)$ as a constrained nonlinear optimization on p_1 . However, we took an alternative approach in which we present the seller’s pricing problem as a *segmentation problem*. That is, instead of choosing p_1 , the seller chooses the θ value to separate the consumers who will buy or wait in the first period. This way, the seller’s decision problem is presented as an optimization over the consumers’ threshold value θ . The significant advantage of this approach is that for any given θ , the determination of the second-period price becomes independent of p_1 and is done via an optimization procedure using (4). Once the second-period prices are identified for a given θ value, we simply use the indifference equation

$$s_1(\theta|\theta) = s_2(\theta|\theta) \quad (6)$$

(see (1) and (2)), to set the corresponding first-period price $p_1(\theta)$. We then evaluate the expected revenue using the equation

$$\begin{aligned} \pi^R(Q|\theta) \doteq & p_1(\theta) \cdot E_\Lambda[N(Q, \Lambda \cdot (1 - \theta))] \\ & + E_\Lambda \left[\sum_{x=0}^{Q-1} p_2(Q - x, x, \theta) \right. \\ & \cdot N \left(Q - x, \Lambda \cdot \left(\theta - \frac{p_2(Q - x, x, \theta)}{\delta} \right) \right) \\ & \left. \cdot P_x(\Lambda \cdot (1 - \theta)) \right]. \end{aligned} \quad (7)$$

Finally, we sequentially calculated the expected revenue $\pi^R(Q|\theta)$ over the range $\theta \in [0, 1]$ to identify the best segmentation (say θ^*) that leads to the optimal expected revenue for the seller; that is, $\theta^* = \arg \max_{\theta} \pi^R(Q|\theta)$. Interestingly, we have noticed throughout our numerical studies (see Section 3), without exception, that the first-period price that makes the marginal consumer with valuation θ indifferent is *strictly* increasing in θ . The implication of this numerical observation is significant in view of the next proposition.

Proposition 3. *If $p_1(\theta)$, the implicit function defined by (6), is strictly increasing in θ , then there is a unique Nash equilibrium in the subgame analyzed in Section 2.1.*

We can now conclude that for each and every case studied in our extensive numerical study, the subgame between the seller and consumers (see Section 2.1) possesses a *unique* equilibrium.

3. The Benefits of Responsive Pricing: A Numerical Study

To measure the potential benefits of responsive pricing, we follow a standard approach. As a baseline, we consider a setting in which the seller employs a single price across the season. Clearly, in such situations, strategic consumer behavior is irrelevant. Moreover, the pricing problem practically reduces to a single-period price optimization as follows:

$$\pi^F(Q) = \max_{p \in [0,1]} \{ E_\Lambda[p \cdot N(Q, \Lambda \cdot (1 - p))] \}. \quad (8)$$

Consequently, one can gauge the percentage benefit of responsive pricing by the metric $\pi^R(Q)/\pi^F(Q) - 1$.

Because of the complexity of our model setting, closed-form expressions for the seller’s expected revenue functions are practically impossible to obtain. Therefore, we have conducted a numerical study that spans across 3,000 instances, covering all possible parameter combinations of (i) 15 levels of the degree of uncertainty about the average market size, measured by the coefficient of variation $CV \doteq \sqrt{\text{Var}(\Lambda)}/E[\Lambda] \in \{0.2, 0.4, 0.6, \dots, 3\}$; (ii) 10 levels of the consumer discount factor $\delta \in \{0.1, 0.2, \dots, 1\}$; and (iii) 20 levels of the initial inventory $Q \in \{1, 2, 3, \dots, 20\}$. For all instances and without any significant loss of scope, we held the average value of the market size fixed at $E[\Lambda] = 10$.²

3.1. Preliminary Results

Imagine for a moment a market consisting of myopic consumers only. In this case, responsive pricing provides the seller with a valuable option. Not only that sales realized in the first period result in up-to-date inventory information (inventory status update), but the seller is also able to resolve part of the uncertainty

regarding the market size (demand learning) prior to setting the second-period price. Therefore, it is obvious that responsive pricing dominates a fixed-price policy. For example, across all 3,000 scenarios in our study, the benefit of responsive pricing in the case of myopic consumers fell in the range (+0.8%, +33.3%) with an average value of +13.2%. In contrast, the corresponding values for the case of strategic consumers were (−13.0%, +9.6%), and −5.6%, respectively. In 86% of the instances (2,584 out of 3,000) the benefits of responsive pricing were negative.

Consistent with the extant literature (e.g., Aviv and Pazgal 2008 and Lai et al. 2010), we find that the benefits of responsive pricing are highly dependent on the consumers’ discount rate δ ; see Figure 1. The pattern observed in the figure for the myopic consumers case is intuitive. As δ increases, the valuations of the consumers in the second period become proportionally higher, and hence, the seller can extract higher revenues under responsive pricing. As expected, the benefits of responsive pricing are significantly affected by strategic consumer behavior; see the right-hand side of Figure 1. But here, the value of δ has two opposing effects. The first, positive effect is similar to that discussed in the case of myopic consumers. The second, negative effect stems from the fact that a higher δ value corresponds to a higher willingness to wait among strategic consumers (i.e., more patient consumers). Together, these effects lead to a nonmonotonic pattern of the benefits with respect to δ with the worst impact predicted for midvalues of δ (say 0.5–0.8). As explained in the introduction, midvalues of δ are reasonably expected in markets of fashion goods.

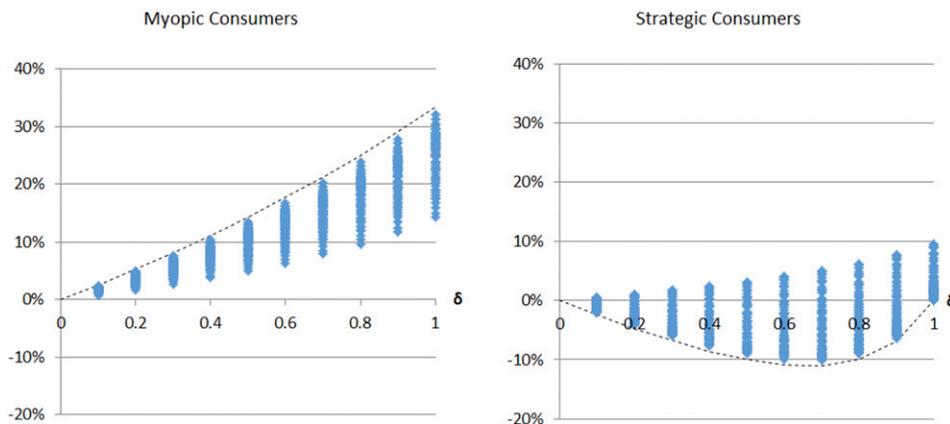
3.2. The Adverse Impact of Responsive Pricing in Fashion Markets Consisting of Strategic Consumers

Confining our attention to midvalues of δ and our core model of strategic consumers, we have taken a closer look at the benefits of responsive pricing as a function of the

CV value and the level of inventory. Hereafter, we generally refer to fashion markets as ones characterized by mid- δ values, low-to-medium levels of inventory (see further discussion in Section 3.8), and relatively large CV values. For example, Figure 2 shows the benefits of responsive pricing for the case of $\delta = 0.7$. Interestingly, we found that responsive pricing tends to become increasingly *counterproductive* as the value of CV increases. In other words, it is in those settings, in which the seller has a larger opportunity to learn from early sales realizations, that responsive pricing might significantly hurt revenue performance rather than contribute. This observation has important ramifications for sellers of fashion goods that consider adopting dynamic pricing–decision technologies. When the market consists of myopic consumers only, the benefits (excluding the system deployment costs) are expected to be high, but on the other extreme—when the consumers are all strategic—the impact is expected to be substantially negative. This obviously emphasizes the critical importance of understanding the extent to which strategic consumer behavior persists in the specific market; see, for example, Li et al. (2014) as well as Section 4 in this paper.

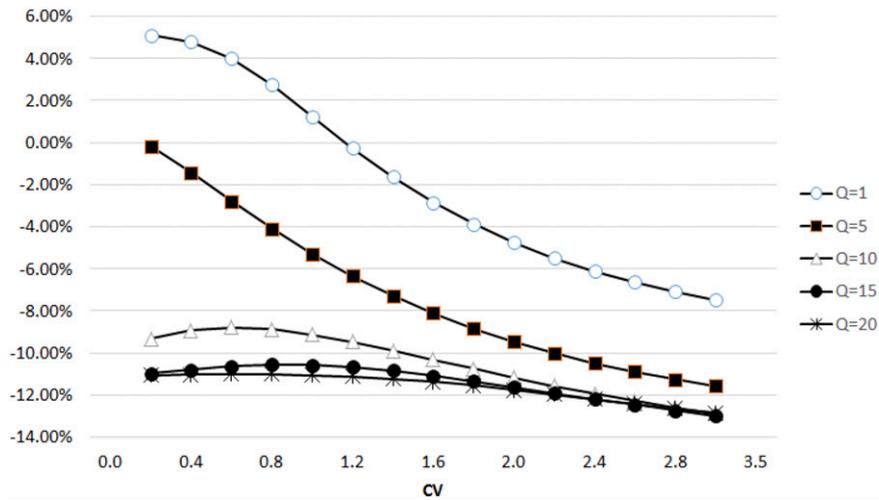
To better understand the drivers behind the negative benefits observed in our study, we propose the following methodology. Note that the value of responsive pricing is driven by two key factors: (i) the mere ability to offer two prices, to which we refer as the *base value of segmentation*; and (ii) the ability to observe information, to which we refer as the *value of information*. To measure the base value of segmentation, we consider a *hypothetical* setting in which the seller adopts an *open-loop* pricing–decision rule. As commonly interpreted in the literature (e.g., Puterman 2005), an open-loop policy does not require the observation of the system’s state (sales data in our paper) except for the very beginning state (i.e., initial inventory position). In this case, the seller utilizes a planning process that specifies the prices

Figure 1. (Color online) The Percentage Benefits of Responsive Pricing Under (Left) Myopic and (Right) Strategic Consumer Behavior for All 3,000 Instances Studied in Our Numerical Analysis



Note. Data are sliced according to the δ values.

Figure 2. (Color online) The Benefits of Responsive Pricing Under Strategic Consumer Behavior $(\pi^R(Q)/\pi^F(Q) - 1)$ for the Case $\delta = 0.7$



for both periods up front but without a commitment power. To avoid potential confusion, we emphasize that the second-period price is still determined by identifying a Nash equilibrium in the subgame. In particular, the seller’s best response to any given consumers’ purchasing strategy θ can be calculated as follows:

$$\max_{p_2 \in [0, \delta\theta]} \{p_2 \cdot E_{\Lambda}[N(Q, \Lambda \cdot (1 - p_2/\delta)) - N(Q, \Lambda \cdot (1 - \theta))]\}.$$

The complete equilibrium analysis of this setting is identical to that of the main model stated in Section 2, and we refer to this model as the case of dynamic pricing with *no sales information*, whose performance is denoted by $\pi^{NI}(Q)$. By comparing the latter value with the revenue performance of a fixed-price strategy, we obtain the base value of segmentation: $\pi^{NI}(Q)/\pi^F(Q) - 1$. In other words, it is the benefit the seller can gain by merely utilizing two prices but without using sales information. Similarly, we define the value $(\pi^R(Q) - \pi^{NI}(Q))/\pi^F(Q)$ as the value of information, and the division by $\pi^F(Q)$ is used to obtain a standardized reference basis:

$$\begin{aligned} \text{Benefits of Responsive Pricing} &= \left(\frac{\pi^{NI}(Q) - \pi^F(Q)}{\pi^F(Q)} \right) \\ &\quad + \left(\frac{\pi^R(Q) - \pi^{NI}(Q)}{\pi^F(Q)} \right) \\ &= \text{Base Value of Segmentation} + \text{Value of Information.} \end{aligned}$$

3.3. The Base Value of Segmentation and the Spread Effect

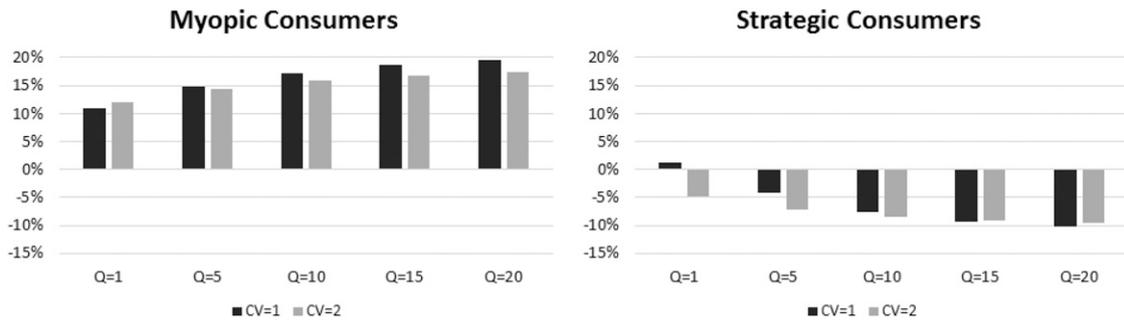
When the consumers are all *myopic*, the base value of segmentation is always positive. Across all 3,000

instances in our numerical study, this value ranged between 0.9% and 33.3% with an average of 12.8%. In stark contrast but as expected, strategic consumer behavior can dramatically affect the base value of segmentation. Under such a market setting, this value was mostly negative (about 86% of the instances) and ranged between -11.1% and 9.6% with an average of -4.8%.

In a market consisting of myopic consumers, the base value of segmentation is driven by the fact that a *noncontingent (on sales)* but rationally expected two-price scheme allows the seller to target two types of consumers: charge a relatively high price in the first period to exploit scenarios in which high-valuation consumers arrive to the store. Then, charge a reduced price in the second period to gain revenues from lower-valuation consumers. Let us look at Figure 3 as a representative illustration of our results. Restricting our attention to the case of myopic consumers presented on the left-hand side of the figure, we can see that the base value of segmentation increases as a function of Q in a concave fashion. However, the CV value has two opposing effects on the base value of segmentation. When Q is small (say $Q = 1$) and CV gets larger, the seller increasingly benefits from an aggressive price *betting strategy*, under which a high first-period price is posted and followed by a significant markdown in the second period. This is merely because high market-size scenarios are more likely to happen. As Q increases, the betting strategy becomes less appropriate. Here, a larger CV may negatively affect the base value of segmentation by the virtue of uncertainty in the decision-making process. In other words, with increased uncertainty regarding the market size, the seller’s ability to accurately segment the market becomes impaired.

Let us turn our attention to markets consisting of strategic consumers, presented on the right-hand side

Figure 3. Base Value of Segmentation ($\pi^{NI}(Q)/\pi^F(Q) - 1$) Under (Left) Myopic and (Right) Strategic Consumer Behavior for Selected Values of Q and CV and $\delta = 0.7$



of Figure 3. First, consider the smallest value of Q (i.e., $Q = 1$). When the spread in the possible market-size scenarios increases (i.e., higher CV values), the “extreme” scenarios of very small market-size and very large market-size realizations are more likely to happen. This has an interesting effect on the expected surplus values associated with buy-now and wait decisions. If the market size is very large, both surplus values become negligible because of the low likelihood of obtaining the unit. But if the market size is small, a consumer is likely to find a unit available in the second period at the reduced price. In fact, with a larger CV value and lack of a sale in the first period, the seller is rationally expected to offer a drastic discount. In sum, a higher market-size uncertainty makes waiting preferable. As we increase the value of Q , the consumers face a lower shortage risk in the second period, and they are more inclined to wait (see, e.g., Liu and van Ryzin 2008). In fact, the base value of segmentation becomes *independent* of CV as the shortage risk declines to zero (see, e.g., the cases of $Q = 15$ and $Q = 20$ in the figure).

Let us define the *degree of strategic behavior* in equilibrium by the ratio $(\theta - p_1)/(1 - p_1)$, namely the percentage of consumers who attempt to wait for sales among those who can afford the price p_1 . Figure 4 shows the value of this measure as a function of CV for different levels of inventory. As can be seen, the degree of strategic behavior is generally increasing in CV value, and the difference is particularly large under the smallest value of Q as we argued earlier. We consequently define the following phenomenon.

Definition 1 (The Spread Effect). In markets consisting of strategic consumers, the degree of strategic behavior in equilibrium tends to increase as a function of the market-size uncertainty (CV). The increase is particularly significant under small levels of inventory (Q).

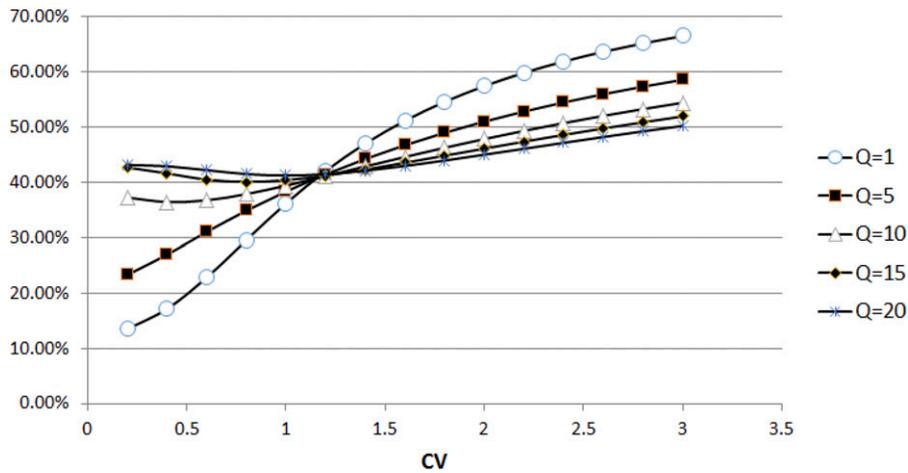
(Note that the graphs are nonmonotone for larger Q values. As we explained before, this is merely because the influence of Q and CV on the shortage risk is intertwined.)

3.4. The Value of Information

As expected, we found that when all consumers are *myopic*, the value of information was always non-negative. However, the values observed were quite small, ranging between 0.0% and +2.1% with an average of +0.5% across all 3,000 instances. Obviously, for $Q = 1$, information has no value because the only situation that leads to pricing in the second period is if there were no sales in the first period. Similarly, when Q approaches infinity, the value of information declines to zero as the sales in the first period have no practical effect on the second-period price. The value of information under the myopic-consumer setting is relatively small in comparison with the base value of segmentation. This means that the benefits of dynamic pricing are driven almost solely by the ability to charge two prices with the “responsiveness” feature typically yielding no further significant advantage. For example, Gallego and van Ryzin (1994)—in a model with myopic consumers and no learning element—report that a fixed-price policy performs almost the same as an optimal responsive-pricing policy based on real-time sales information; see section 3.3 there. To reconcile our results with the research papers on dynamic pricing with learning (see references in the introduction), we note two important characteristics of our setting: first, we study a two-period setting, having an opportunity to incorporate learning only at the very end of the sales season. Second, pricing in the second period is focused on the segment of consumers with valuations below p_1 , further limiting the seller’s ability to take advantage of learning. We come back to the latter point shortly for the case of strategic consumers.

When all consumers are *strategic*, the value of information was mainly *negative* (82% of the 3,000 instances), ranging between -3.7% and +3.0% with an average of -0.8%. Among all 3,000 instances, the value of information exceeded 1% only under the special case in which the consumer’s δ was at the highest level ($\delta = 1$), the CV value was large ($CV > 1$), and the inventory level was in a medium value range (relative to

Figure 4. (Color online) Degree of Strategic Behavior in Equilibrium ($(\theta - p_1)/(1 - p_1)$) as a Function of CV for Different Levels of Inventory (Q)



Note. Results are presented for the “no information” reference model and $\delta = 0.7$.

the spread of the demand distribution). To explain this outcome, note that, under the aforementioned conditions, the consumers are highly patient and are expected to be strongly strategic (i.e., follow a high θ threshold value). The seller, hence, reaches the second period with a market of remaining with residual valuations in the range $[0, \delta\theta]$. Thus, market-size information can be very useful for setting the second-period price optimally. Interestingly, because the residual market valuation in the second period is lower in the myopic case than in the strategic case (i.e., $[0, \delta p_1]$ versus $[0, \delta\theta]$), the value of information was higher in the all-strategic consumers case than in the all-myopic consumers case.

As we shift our attention back to our focal range of medium δ levels, we note that the value of information under strategic consumer behavior is negative for all $Q > 1$ and sufficiently large degrees of market uncertainty (e.g., $CV > 0.5$); see Figure 5 as a representative illustration of our results.

In other words, the ability to act upon information might have a *negative* influence on the seller’s expected revenue performance when faced with strategic consumers. Generally, this is not necessarily a surprising result. Existing research indeed demonstrates that responsive pricing policies that utilize inventory status information can underperform fixed-price policies in the face of strategic consumers. However, in contrast the extant research (e.g., Aviv and Pazgal 2008 and Cachon and Swinney 2009), our modeling framework allows us to clearly separate the negative effect of strategic consumer behavior into two parts: the base value of segmentation and the value of information. Interestingly, when the market-size uncertainty is higher, which brings an attractive opportunity to learn via sales observations, the value of information might become *counterproductive*. This outcome stands in contrast

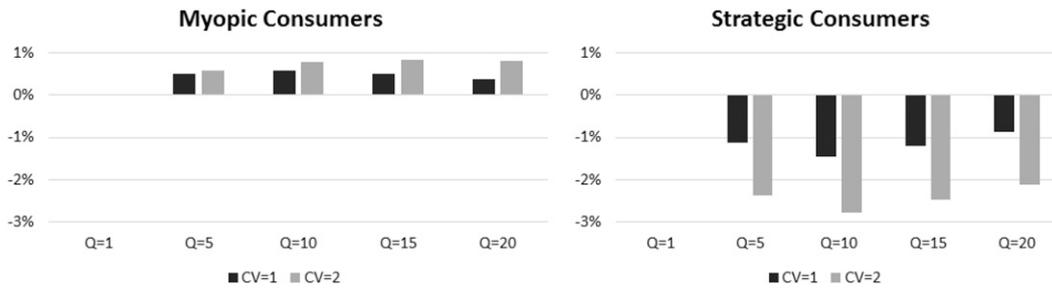
to the situation in the case of all-myopic consumers, in which the value of information tends to increase in CV; see the left- and right-hand sides of Figure 5. This leads us to the next part of our study.

3.5. Can Learning Be Bad in the Sales of Fashion Goods?

To better understand why the value of information can be negative in the sales of fashion goods (again, restricting our attention to the regime of medium levels of δ and high CV values), we have conducted an additional study. Note that the value of information is driven by two *interdependent* types of information: the up-to-date inventory status at the beginning of the second period (*value of inventory information*) and the revised forecast of the average market size at that same time (*value of demand learning*). To achieve some separation between these two interrelated information pieces and, in particular, to isolate the impact of demand learning, we constructed an auxiliary model that hypothetically assumes the seller is committed to *not* (or *unable to*) updating demand information using the first-period sales information. Instead, the seller will continue to rely on the prior distribution of Λ rather than the updated distribution $\tilde{\Lambda}(x, \theta)$ when setting the second-period prices, according to (4). The consumers are fully aware of this particular limitation. We refer to this auxiliary model by the title “no learning” and denote its expected revenue performance by $\pi^{NL}(Q)$.

Figure 6 shows the values of $(\pi^R(Q) - \pi^{NL}(Q))/\pi^F(Q)$ for the same subset of scenarios considered in the figures presented in the previous sections. As expected, the value of demand learning is nonnegative when the consumers are myopic. However, for the set of scenarios exhibited in the table, all values are negative for all-strategic consumers, meaning that the seller’s performance would be better if we forced it to use

Figure 5. Value of Information $((\pi^R(Q) - \pi^{NI}(Q))/\pi^F(Q))$ Under (Left) Myopic and (Right) Strategic Consumer Behavior for Selected Values of Q and CV and $\delta = 0.7$



the original demand-forecasting scenarios (and likelihoods) without updating them on the basis of early sales readings. As can be seen in the figure, for most scenarios, and particularly those with a high degree of resolvable market uncertainty, the seller might be able to perform better if it could credibly commit not to update its demand forecast. Looking at this from a slightly different angle, a seller that develops a sales-based forecasting process to support responsive pricing strategies may find itself settling in an equilibrium that yields a lower revenue performance.

3.6. The Adverse Impact of Active Demand Learning: Exploration (Seller) and Information Shaping (Consumers)

To appreciate the combined impact of demand learning and strategic consumer behavior on the value of demand learning, recall the two contradicting effects mentioned in the introduction. On one hand, learning enables the seller to react to market information. On the other hand, strategic consumers may respond to that capability in a way that negatively affects the seller’s performance.

Let us first consider the seller’s perspective. It is easy to see that the seller can impact its demand learning process via a single lever only—the first-period price. In our model, the seller optimally takes this consideration into account. More specifically, in its attempt to generate maximal revenue performance, the seller considers the subtle way in which the first-period price

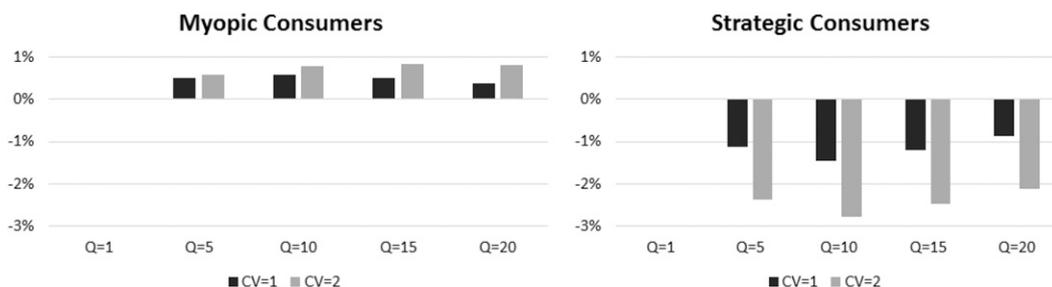
affects the consumers’ purchasing behavior, which, in turn, affects its learning process and effectiveness in charging the right price in the second period. The trade-off associated with the ways in which the ability to collect immediate revenue gains versus the ability to learn affect pricing decisions is named exploitation versus exploration in the economic and revenue management literature.

Turning to the consumers, note that, in determining their purchasing strategy, they also must consider the effect of their decisions on the seller’s ability to learn about the market size. For example, a large θ implies that a (statistically) larger number of consumers will deprive the seller of sales information and, consequently, limit the ability of the seller to resolve market uncertainty. Obviously, this affects the seller’s pricing in the second period, and for the consumers, it impacts the expected surplus associated with a “wait” decision. Before we proceed, let us define the following phenomenon.

Definition 2 (Information Shaping). In markets consisting of strategic consumers, information shaping refers to the conscious collective attempt of the consumers to influence the seller’s interpretation of the market size.

In using the term “collective attempt,” we obviously refer to the consumers’ selection of the value of θ in equilibrium. Here, it is also useful to recall our observation in Section 2.2 that the seller’s first-period

Figure 6. Benefits Gained by Demand Learning: The Seller’s Ability to Utilize Sales Realizations in Its Demand Forecasting Process $((\pi^R(Q) - \pi^{NL}(Q))/\pi^F(Q))$ Under (Left) Myopic and (Right) Strategic Consumer Behavior for Selected Values of Q and CV and $\delta = 0.7$



pricing problem can be translated into a segmentation problem, namely an optimal selection of θ . Thus, in effect, learning is solely controlled via the model parameter θ , but the ways in which the seller and the consumers affect learning are nontrivially intertwined. In this section, we study the way in which the *conscious* considerations of how θ affects learning—exploration (by seller) and information shaping (by consumers)—plays a role in the overall value of demand learning. Subsequently, we propose an auxiliary model that neutralizes the conscious consideration of the impact that θ has on learning. Specifically, we look at a model in which learning takes place and affects the seller's and consumers' decisions in equilibrium. However, in this model, neither the seller nor the consumers believe that they can affect the seller's demand learning process.

Consider a hypothetical case in which the seller cannot use sales information in updating its demand forecast. Instead, the seller is guaranteed to observe the (aggregate) number of consumers in the market with valuations in the range $[\bar{\theta}, 1]$ but just prior to the end of the season (i.e., prior to setting p_2). The seller can use that information to update the demand forecast. For instance, if $\bar{\theta} = 0$, the model reflects a situation in which the seller perfectly observes the total number of consumers (but not their individual valuations) in the market prior to setting p_2 . In contrast, when $\bar{\theta} = 1$, the model becomes equivalent to the no-learning model presented in the previous section.

The auxiliary model is considerably more complex than our original model, and for the sake of clarity of exposition, we refer the interested reader to the online appendix for a brief overview of the technical details of our analysis. Most importantly, for any *exogenously* given value $\bar{\theta}$, we predict the market equilibrium given by the pair $\{p_1^*(\bar{\theta}), \theta^*(\bar{\theta})\}$ (and the resulting p_2 menu, which is less crucial for our discussion at the moment). We then apply a search procedure over $\bar{\theta} \in [0, 1]$ to identify *fixed points* of the type $\theta^*(\bar{\theta}) = \bar{\theta}$. Assuming that a fixed point indeed exists (which happened at a unique value of $\bar{\theta}$ for each one of the instances we studied), the model would predict a rational expectation equilibrium in which the consumers behave according to a policy $\bar{\theta}$, and the seller would observe the number of consumers in the range $[\bar{\theta}, 1]$, which is, in fact, equal to the sales quantity. Note that, in the equilibrium predicted for the hypothetical model, we have neutralized the seller's and consumers' ability to actively influence learning. In other words, learning takes place but in a *passive* way only. Let us use the notation $\pi^{PL}(Q)$ as the predicted expected revenue in equilibrium for this auxiliary model (*PL* stands for "passive learning").

Based on our analysis of this model (see, e.g., Table 4 in the online appendix), we conclude that the values of active learning (exploration and information shaping),

defined by the measure $(\pi^R(Q) - \pi^{PL}(Q))/\pi^F(Q)$, do not exhibit any particularly strong pattern with respect to the model parameters. The results suggest that the active learning interplay between the seller and the consumers leads to a loss of 1% to 2% in expected revenue performance. Because exploration is controlled by the seller, this leads us to the conjecture that information shaping—the conscious attempt of the consumers to influence demand learning via their purchasing decisions—brings a loss that is at least as strong as that just mentioned.

3.7. The Expected Market Size

Recall that our numerical study has focused on a unique value of the expected market size $E[\Lambda] = 10$. In many fashion-goods settings, this is a reasonably realistic order of magnitude for the market size; in other situations, the expected market size could vary. By fixing the expected market size at a given level, we were able to explore the impact of the inventory-to-mean market size ratio, $Q/E[\Lambda]$. However, if the market size was larger in orders of magnitude, one might wonder if the same phenomena discussed earlier would continue to exist. For example, could it be that the negative effect of demand learning takes place only in markets characterized by a small number of consumers as only then each consumer's decision is sizeable enough to affect the seller's perception of the market? To address this issue, we conducted a numerical study, including two different values of the mean market size $E[\Lambda] \in \{10, 1,000\}$, covering all possible combinations of 10 inventory-to-mean market size ratios $Q/E[\Lambda] \in \{0.1, 0.2, \dots, 1\}$, two levels of degree of uncertainty $CV \in \{1, 2\}$, and four levels of consumers' $\delta \in \{0.5, 0.6, 0.7, 0.8\}$; see illustrative results in Table 5 in the online appendix. Overall, our study indicates that changing the average market size does not qualitatively influence our managerial results and insights. Moreover, the increase in the average market size could even amplify the negative effects of demand learning. For example, the negative impact of demand learning worsens from a loss of 4.48% when $E[\Lambda] = 10$ to a loss of 7.13% when $E[\Lambda] = 1,000$ for $Q/E[\Lambda] = 0.2$. To summarize, the presence of a large number of consumers in the market (large $E[\Lambda]$) does not mitigate the adverse consequences of the demand-learning effect, which can be viewed as an outcome of the *collective* action of all consumers in the market.

3.8. Optimal Capacity Choice: Illustrations and Discussion

As the initial inventory plays a significant role in explaining the benefits of responsive pricing, it may be of interest to examine the case in which the seller selects not only the optimal prices but also the best level of

inventory Q . For brevity, the nature of this section is primarily illustrative, and our approach is standard. We consider a given per-unit cost, c , and identify the inventory level and best pricing strategy that optimize the seller’s expected profit. For example, in our main model, we have searched for $\Pi^R \doteq \max_Q \{\pi^R(Q) - cQ\}$. Similarly, we calculated the optimal expected profit for the fixed-price case (denoted by Π^F) and for a benchmark case in which the seller utilizes responsive pricing in a market consisting of myopic consumers (denoted by Π^M). For example, we have considered 44 scenarios, spanned by the combinations of $\delta = 0.7$; two cost parameter values, $c = 0.15$ and 0.25 ; and two CV values, one and two. Our results are summarized in Table 1. First, this table demonstrates that the influence of strategic consumer behavior on optimal profits can be quite severe. Moreover, as the value of CV increases, the difference between the benefits in the myopic-consumers and strategic-consumers settings tends to increase. In particular, under relatively high per-unit costs, higher CV values result in small-to-medium optimal inventory levels. This is exactly where the effect of strategic consumer behavior tends to be the most detrimental.

4. The Degree of Strategic Consumer Behavior in the Market

As shown earlier (see Figure 1), the benefits of responsive pricing depend heavily on the extent to which strategic consumer behavior is prevalent in the market. Indeed, our primary purpose in this paper is to juxtapose the two extreme cases—of all-myopic and all-strategic consumers—to derive interesting insights about the interplay between demand learning and strategic consumer behavior. Nonetheless, the consideration of markets characterized by a mixture of strategically and myopically behaving consumers has taken place in several key papers in the management science literature; see, for example, Su (2007), Cachon and Swinney (2009), Levin et al. (2009), Lai et al. (2010), and Mersereau and Zhang (2012).

To study markets in which not all consumers are strategic, we propose a variant of our model in which a portion (say, α) of the consumers are strategic and the rest are myopic. In particular, consider a case in which each individual consumer arriving to the store behaves strategically with a probability of α and suppose that the consumers are statistically independent in this regard. Such a setting lends itself to a relatively easy modification of our analysis, technically stemming from the fact that the market size can be separated into two parts that are Poisson distributed and statistically independent: one with a mean of $\alpha\lambda$ (strategic consumers) and one with a mean of $(1 - \alpha)\lambda$ (myopic consumers). For the details of the analysis of this model, we refer the reader to the online appendix.

Using this modified model, we conducted a numerical study that spans across 1,760 scenarios, covering all possible combinations of 20 levels of the initial inventory $Q \in \{1, 2, \dots, 20\}$, two levels of degree of uncertainty $CV \in \{1, 2\}$, four levels of the consumer discount factor $\delta \in \{0.5, 0.6, 0.7, 0.8\}$, and 11 values of the expected portion of strategic consumers in the market $\alpha \in \{0, 0.1, \dots, 1\}$. To be consistent with our main model, the average value of the market size for all instances was fixed at $E[\Lambda] = 10$. The benefits of responsive pricing over the fixed pricing (i.e., $\pi^R/\pi^F - 1$) ranged between -11.28% and 22.15% with an average value of 1.26% . For example, Table 2 shows a subset of the results for the case in which $\delta = 0.7$ and $CV = 2$. As discussed before, the benefits of responsive pricing decrease as consumers are more likely to behave strategically (i.e., α increases). Most impressive, in our opinion, is the magnitude of dependence of the benefits of responsive pricing with respect to α —across the range $[0, 1]$ —emphasizing the necessity of understanding the consumer mixture and developing empirical methodologies for measuring the extent to which strategic behavior is likely to occur in markets of fashion goods. The rest of our results exhibit a similar pattern

Table 1. The Benefits of Responsive Pricing Under Myopic and Strategic Consumer Behavior When Inventory Levels Are Chosen Optimally for the Case $\delta = 0.7$

CV	c = 0.15						c = 0.25					
	Q*	p ₁	Markup	Profit (Π)	θ	Benefits	Q*	p ₁	Markup	Profit (Π)	θ	Benefits
1												
Fixed	5	0.644	329%	0.878			3	0.692	177%	0.464		
Myopic	7	0.703	369%	1.162		32.38%	4	0.750	200%	0.655		41.10%
Strategic	5	0.595	297%	0.792	0.763	-9.81%	3	0.661	164%	0.434	0.796	-6.52%
2												
Fixed	4	0.706	371%	0.338			1	0.776	210%	0.089		
Myopic	5	0.777	418%	0.487		44.14%	2	0.821	228%	0.159		78.31%
Strategic	3	0.665	344%	0.266	0.859	-21.27%	1	0.737	195%	0.073	0.888	-18.06%

with respect to all other δ and CV value combinations and are, therefore, excluded from the table.

Consistent with our discussion in the previous section, we also observed that both the spread effect and the demand-learning effect increase with the expected portion of strategic consumers (α). Intuitively, it is expected that the information-shaping effect would diminish as more myopic consumers are present in the market because these consumers already convey market-size information to the seller; thus, strategic waiting becomes less effective in influencing the seller's perception of the market size.

5. The Effectiveness of Price Commitment

In this section, we explore whether price commitment can be used as an effective tool for counteracting strategic consumer behavior. The promise behind price commitment is that it ties the hands of the seller in the second period; consequently, consumers realize that their purchasing decisions will not affect the second-period price, and therefore, their incentive to wait can be controlled and reduced. But commitment in settings in which market-size uncertainties can be resolved via sales observations might be radically counterproductive.

Consider a setting in which the first-period and second-period prices, p_1 and p_2 , can be *credibly* announced in advance. It is easy to verify that it is always optimal for the seller to set $p_2 < \delta p_1$; otherwise, the seller effectively utilizes a fixed-price policy. In fact, because fixed-price policies can be viewed as a special case of price commitment, this means that the optimal price-commitment policy always dominates the optimal fixed-price policy. In light of our earlier results, this observation immediately establishes that price-commitment strategies are generally of potential value.

Clearly, for any given pair of announced prices, it follows directly from Theorem 1 that only threshold policies can hold in equilibrium. The following theorem characterizes the consumers' purchasing behavior

under a price-commitment policy. Let θ^* be the unique solution to

$$p_1 = \theta + \delta(1 - \theta) - (\delta - p_2) \frac{E_\Lambda[H(Q, \Lambda(1 - p_2/\delta))]}{E_\Lambda[H(Q, \Lambda(1 - \theta))]} \tag{9}$$

Theorem 3. For any pair of credibly announced prices (p_1, p_2) such that $p_1 > p_2/\delta$, there exists a unique threshold policy, $\theta \geq p_1$, in the consumer game. Furthermore, θ is given by

$$\theta = \begin{cases} \theta^* \text{ given by (9)} & \text{if } p_1 \leq 1 - \frac{(\delta - p_2) \cdot E_\Lambda[H(Q, \Lambda(1 - p_2/\delta))]}{E_\Lambda[H(Q, 0)]} \\ 1 & \text{otherwise.} \end{cases} \tag{10}$$

The significance of the condition in the first line of (10) is that it allows us to establish a bound on the first-period price. If p_1 is too high, all consumers wait, and the seller is not expected to generate any revenue in the first period. As a result, the seller can drop the price exactly to that bound without loss of revenue. This argument allows us to focus on Equation (9) for calculating p_1 in our optimization procedure as stated subsequently. For any value of $\theta \in [0, 1]$, we conduct a search for a value $p_2 \in (0, \theta\delta)$ such that p_1 is set according to Equation (10), and the following expected revenue is maximized:

$$\begin{aligned} \pi^C(Q) \doteq & \max_{\substack{\theta \in [0,1] \\ p_2 \in (0, \theta\delta)}} \{p_2 \cdot E_\Lambda[N(Q, \Lambda \cdot (1 - p_2/\delta))] \\ & + (p_1 - p_2) \cdot E_\Lambda[N(Q, \Lambda \cdot (1 - \theta))]\} \\ & \text{s.t. (9).} \end{aligned}$$

(We use the superscript C for commitment.) To understand the objective function stated herein, note that we separate the revenue into two parts: p_2 collected from the *total* sales made in both periods and a premium of $(p_1 - p_2)$ collected from the sales made in the first period only.

Table 2. The Benefits of Responsive Pricing as a Function of the Expected Portion of Strategic Consumers in the Market (α) and Inventory (Q) Presented for the Case of $CV = 2$, $\delta = 0.7$, and $E[\Lambda] = 10$

Initial inventory level Q	Benefits of responsive pricing $(\pi^R - \pi^F)/\pi^F$			Base value of segmentation $(\pi^{NI} - \pi^F)/\pi^F$			Value of demand learning $(\pi^R - \pi^{NL})/\pi^F$		
	$\alpha = 0.2$	$\alpha = 0.5$	$\alpha = 0.8$	$\alpha = 0.2$	$\alpha = 0.5$	$\alpha = 0.8$	$\alpha = 0.2$	$\alpha = 0.5$	$\alpha = 0.8$
1	7.93%	0.58%	-2.96%	7.93%	0.59%	-2.93%	1.14%	-2.23%	-3.24%
5	8.99%	-1.56%	-7.21%	8.20%	-0.95%	-5.08%	1.11%	-2.67%	-4.86%
10	9.25%	-2.88%	-8.78%	8.13%	-1.94%	-6.29%	0.95%	-2.43%	-4.35%
15	9.26%	-3.82%	-9.23%	8.05%	-2.56%	-7.04%	0.83%	-2.24%	-3.35%
20	9.20%	-4.30%	-9.40%	8.00%	-2.95%	-7.50%	0.75%	-1.87%	-2.46%

As discussed earlier, price commitment enables the seller to regain part of the loss of market power caused by the dynamic competitive environment of the responsive pricing model, but it cannot eliminate strategic consumer behavior when the seller employs price discrimination (i.e., when $p_2 < \delta p_1$). In other words, the seller's revenue performance could still be dramatically lower than that which can be achieved in a market composed of myopic consumers. Among all 3,000 instances, the percentage benefit measure ($\pi^C(Q)/\pi^R(Q) - 1$) ranges between -2.3% and $+15.0\%$ with an average of $+6.6\%$. As discussed in Section 3.4, the extreme results (both negative and positive) were observed for the high CV values with the nature of the impact highly sensitive to δ . For $\delta = 1$, the impact of price commitment was negative whereas for midvalues of $\delta \in [0.5, 0.8]$ the impact of price commitment was highly positive. In summary, price commitment is most helpful in the cases in which strategic consumer behavior has its most adverse impact on revenue performance. But although beneficial, it does not fully eliminate the negative impact of such consumer behavior.

5.1. Commitment vs. Contingency: Further Discussion and Analysis

The price-commitment strategy discussed earlier maintains two important characteristics: commitment and lack of contingency option. As we observed, the commitment provides a seller with a possible advantage. However, the lack of contingency is potentially counterproductive to the seller. To understand the degree to which the two opposing forces affect the benefits of price-commitment strategies, we studied an *auxiliary* model in which a seller deploys a *sales-dependent* price-commitment strategy. In this model, the seller commits in advance to a price function $p_2(Q - x)$ rather than to a single price p_2 , where x is the number of units sold in the first period. Note that, unlike the simple price-commitment model, this hypothetical model is *responsive* to early sales. Let us denote its optimal expected revenue performance by $\pi^{RC}(Q)$ (with the superscript standing for both "responsive" and "commitment"). With a commitment to the second-period price function, the subgame in Section 2.1 maintains a Stackelberg game form, in which the seller enjoys the benefits of being the leader.

We have developed an optimization procedure to identify the best responsive pricing for our auxiliary model. Technically, some pricing schemes $\{p_1, p_2(\cdot)\}$ could induce multiple equilibria (θ values) in the consumers' game; see also the discussion in Section 2.1.1. Thus, when searching for the optimal announced sales-dependent pricing policy, we evaluated the seller's expected revenue for all possible equilibria points (θ values). Nevertheless, we found that, for all the numerical instances, the

seller's *optimal* announced sales-dependent prices $\{p_1^*, p_2^*(\cdot)\}$ always induced a unique equilibrium in the consumers' game.

As expected, a sales-dependent price-commitment strategy dominates both the simple price-commitment strategy and the responsive-price strategy (i.e., our main model, in which no form of commitment exists). Among all 3,000 instances, the percentage benefit measure ($\pi^{RC}(Q)/\pi^R(Q) - 1$) ranges between 0.0% and $+25.6\%$ with an average of $+9.9\%$. The *marginal* benefits of the contingency feature, that is, $(\pi^{RC}(Q) - \pi^C(Q))/\pi^F(Q)$ —taken again in reference to the baseline value π^F —ranges between 0.0% and $+15.6\%$ with an average of $+3.0\%$. A deeper examination of our results suggests that to increase revenue performance via discrimination, some sort of contingency capability—with commitment—is needed. Specifically, when the initial inventory is relatively small and the level of CV is at a medium level, the simple-price commitment $\{p_1, p_2\}$ performs similarly to the hypothetical sales-dependent price-commitment policy. But as the levels of inventory and CV increase, discrimination via a larger set of possible second-period prices becomes valuable, particularly for large δ values.

5.2. Optimal Choice of Inventory Revisited

Similarly to Section 3.8, we have also considered settings in which the seller makes an optimal selection of the initial inventory level Q . Table 3 displays the results pertaining to the same illustrative scenarios that we examined earlier (see Table 1). This numerical study echoes the results we have found earlier: price commitment can be substantially beneficial for modest-to-high Q and CV values, and *practical* price-commitment policies (i.e., those with no sales-dependent price menu) can be valuable to a seller. Nevertheless, when CV is high—that is, when the seller can resolve a high degree of market-size uncertainty by reading early sales—price commitment considerably suffers from its lack of ability to respond to sales.

6. Quick Response: The Effect of an Inventory Replenishment Option

Our analysis thus far has focused on fashion retail settings in which the seller can utilize responsive pricing for a fixed inventory of a product, brought to the store prior to the beginning of the season. Nonetheless, some fashion retail supply chains may possess the operational capacity to procure additional inventory *during* the season in response to the sales and demand signals for the product. For a rich background and a review of the relevant research, we refer the reader to Cachon and Swinney (2009). In that paper, the authors present a model of a market consisting of a mixture of myopic and strategic consumers and study the benefits of *quick*

Table 3. The Value of Price Commitment and the Value of Contingency Under Commitment, Under Strategic Consumer Behavior, When Inventory Levels Are Chosen Optimally

CV	c	Responsive pricing		Price commitment with single price		Price commitment with sales-dependent prices		Value of commitment ($\Pi^C/\Pi^R - 1$)	Marginal value of contingency ($\Pi^{RC}/\Pi^C - 1$)
		Q*	Profit (Π^R)	Q*	Profit (Π^C)	Q*	Profit (Π^{RC})		
1	0.15	5	0.792	5	0.894	6	0.920	12.94%	2.86%
	0.25	3	0.434	3	0.485	3	0.491	11.71%	1.33%
2	0.15	3	0.266	4	0.342	4	0.378	28.34%	10.56%
	0.25	1	0.073	1	0.091	2	0.096	25.11%	5.01%

Note. All four scenarios share the same value of $\delta = 0.7$.

response (QR) strategies in which the seller can obtain perfect demand information prior to the beginning of the season and procure additional inventory (instantaneously) on the basis of such information but at a higher cost per unit. One of their main findings is that QR may become significantly more valuable to the seller as the proportion of strategic consumers in the market increases. They explain this by the fact that the seller’s ability to observe demand and respond to it—something valuable already in the all-myopic consumers setting—becomes even more valuable in the presence of strategic consumers because of a rationing effect (e.g., if the seller learns that the demand is low, then the seller would source fewer units, thus mitigating the adverse consequence of strategic consumer behavior).

In this section, we report on results stemming from an alternative model variant of QR in which the inventory replenishment takes place after the interaction between the seller and the consumers has begun and the consumers *endogenously* affect the seller’s perception of the market. Specifically, we constructed a model in which the seller has an option to replenish its inventory—in a *quick fashion*—at the end of the first period. After observing the first-period demand and prior to setting the price at the beginning of the second period, the seller can order additional units of the product. Such units would be expeditiously procured and delivered at a higher per-unit cost rate than that available prior to the beginning of the season. For the details of the analysis of this model, we refer the reader to the online appendix.

Obviously, the seller is able to utilize the demand information gathered during the first period to refine its understanding of the market potential and accordingly determine its replenishment quantity and second-period price. When the consumers are strategic, we find that the loss in expected profit (because of strategic behavior) is *mitigated* under QR in comparison with the original (no-QR) setting (see the results in the online appendix in and around Table 6). We argue that QR provides a *strategic option* to the seller in the following sense. In the no-QR case, the seller must make

an up-front commitment to inventory; thereafter, it becomes engaged with the consumers in the pricing-and-purchasing game. The negative outcome of such an interaction is known by now. With QR, the seller does not make a full commitment to inventory. This gives rise to a market equilibrium that is settled in the presence of an *implicit threat* by the seller to the consumers. When fewer consumers purchase the product in the first period and particularly if they are successful in their information shaping attempt—to create a perception of a poor market—the seller would respond by not ordering additional units. Of course, the consumers who are now faced with increased shortage risk and no hope for significant price markdowns at the end of the season would be leaning more toward immediate purchasing at premium prices. In summary, we argue that, as more strategic consumers are present in the market, responsive inventory becomes more effective than responsive pricing.

We also found that the benefits of QR are higher in the case of all-strategic consumers than in the case of all-myopic consumers across all scenarios (see, again, some representative numerical results in the online appendix). This result is consistent with the finding of Cachon and Swinney (2009), who argue that the benefits of QR might be significantly underestimated in the presence of markets with strategic consumers if such behavior is not taken into account. Furthermore, we conjecture that, by using a QR strategy, the seller becomes less susceptible to an error regarding the estimation of the portion of strategic consumers in the market. The seller can order first a small level of inventory, observe the market, and then react by ordering a proper amount of additional inventory. If the market seems to be poor, responding by ordering a small number of units (or none) would be proper regardless of the nature of the consumers.

7. Conclusion and Discussion

Over the past two decades, the subject of responsive pricing has received an increasing level of attention from both practitioners and academics. Two major

factors that affect the value of price segmentation via responsive pricing are demand learning and strategic consumer behavior. Existing research on strategic consumer behavior demonstrates that the ability to dynamically change prices over the course of a season can be counterproductive to a seller's revenue performance. In contrast, research on demand learning emphasizes the critical (positive) role of responsive pricing in increasing revenue performance via actionable information; namely recording early sales, updating the forecast of the market condition, and appropriately changing the price in reaction to the newly available information. Although the effects of these two factors have been studied extensively, there has been little research on optimal pricing under the presence of both. This paper aims to fill this void in the literature.

We study the benefits of responsive pricing with a primary focus on fashion-like products, characterized by relatively large demand uncertainty—an opportunity to learn about the demand from sales observations, medium levels of the consumers' discount factor, and modest levels of inventory deployment. Our results overwhelmingly demonstrate that the benefits of responsive pricing, in comparison with a benchmark case of a fixed-price policy, depend sharply on the nature of the consumers' behavior. In stark contrast to markets of myopic consumers, when the consumers are all strategic, the benefits tend to worsen (in fact, they become mostly negative) when there is a higher potential for learning. We explain this counterintuitive outcome by pointing to two phenomena: the spread effect and the active-learning effect. The spread effect is the phenomenon in which a higher degree of market-size uncertainty makes the expected surplus associated with a wait-for-markdown decision increasingly more attractive than that associated with a purchase-now decision because of the higher spread in market-size scenarios. The active learning effect refers to the way in which market equilibrium is influenced by the seller's demand exploration (via pricing) and by the consumers' act of information shaping. We define information shaping as the conscious collective attempt of the consumers to influence the seller's interpretation of the market size via their purchasing decisions. For example, by waiting for markdowns, the consumers withhold sales information from the seller, possibly leading the seller to interpret the market as weak and, hence, reduce prices later.

The information-shaping effect has the following ramification. Sellers of fashion products that consider upgrading their pricing systems to incorporate "accurate response" strategies (i.e., integrating learning and responsive pricing) should be aware of the possibility that this might lead them to a new and potentially worse equilibrium, particularly when there is a higher opportunity to learn. We demonstrated this

insight using a hypothetical model in which we eliminated the seller's ability to learn from sales (Section 3.5). However, we do not wish to make a case against learning. In contrast, we argue that sellers in fashion markets should extensively seek market signals other than mere sales figures on which to base their forecasting process (Fisher and Rajaram 2000). For example, our model presented in Section 3.6 is readily available to study the performance of responsive pricing in a market in which the number of consumers is fully visible to the seller (use $\bar{\theta} = 0$, *ibid*). The reliance on indirect demand data may considerably mitigate the degree of the information-shaping effect. For example, Zara, a well-known "fast-fashion" Spanish retailer relies on store managers and supporting information technology to gather and analyze customer-preference data. In particular, store managers provide Zara's design teams with information that is not captured in the data gathered by their automated point-of-sale system (Ghemawat and Nueno 2003).

Despite the fact that price commitment completely eliminates the seller's ability to learn, it appears to increasingly dominate responsive pricing as more strategic consumers are present in the market. Moreover, the most significant benefits of commitment were observed for the cases in which the degree of market uncertainty is high. However, it is noteworthy that, in cases in which the market size is highly uncertain and when the seller's inventory is large, price commitment may be highly limited in its ability to perform effective discrimination because of its lack of ability to prescribe sales-dependent price changes.

Cachon and Swinney (2009, 2011) study the value of quick response in the presence of strategic consumers. For example, in their earlier paper, they find that the ability to quickly replenish inventory at the very beginning of the selling period is dramatically more beneficial in a market of strategic consumers than in a market of myopic consumers. A key driver behind this result is that, by better matching supply with demand, the intensity of strategic consumer behavior can be considerably mitigated. Building upon this interesting insight, we studied the potential benefits of *quick response* strategies—ones that embed both dynamic pricing and quick inventory replenishment *during* the sales season. We also find that the benefits of quick response strategies might be significantly underestimated in the presence of markets with strategic consumers. We explain that the option to replenish inventory close to the end of the season serves as an effective implicit threat by the seller to the consumers in a way that discourages strategic waiting. If fewer consumers purchase the product at the premium price and, particularly if they are successful in their information-shaping attempt, the seller would likely respond by not ordering additional units. But

the consumers, now faced with increased shortage risk and no hope for significant price markdowns at the end of the season, would be leaning more toward immediate purchasing at premium prices. Similarly to Cachon and Swinney (2009), we conjecture that a quick-response seller may be less susceptible to an error regarding the estimation of the portion of strategic consumers in the market. The seller can order first a small level of inventory, observe the market, and then react by ordering a proper amount of additional inventory. If the market seems to be poor, responding by ordering a small number of units (or none) would be proper regardless of the nature of the consumers.

The results reported herein highlight opportunities for further research. Our model considers a case in which a seller can resolve uncertainty regarding the mean market size. Although technically challenging, this framework may be expanded to consider settings in which uncertainty about price elasticity could also be learned via early sales; see, for example, Aviv and Pazgal (2005) for such a treatment in the case of myopic consumers. We speculate that when strategic consumers are present in the market, these phenomena—the spread effect and information shaping—would still play a major role and perhaps even intensify. On one hand, sellers can derive important information from sales data to effectively price in the second period. However, in equilibrium, the sellers' ability to learn and deploy responsive pricing may put them at significant disadvantage because of the aforementioned effects. Other obvious directions of research include empirical studies of strategic consumer behavior in fashion markets in which demand-learning opportunities exist. For example, one type of study may involve an estimation of the extent to which strategic behavior is prevalent (again, see Li et al. 2014). Or one can also propose and validate novel models for consumer behavior. Recall that in our model consumers are surplus maximizers, fully rational, and privy to all model information. Another theme of research is the study of competitive settings involving more than a single seller (see, e.g., Levin et al. 2009 and Liu and van Ryzin 2008). It is not up-front obvious if the spread and information-shaping effects would be diminished or intensified when demand is split among several retailers.

Endnotes

¹ Under the conditional approach, we assume that the consumer can verify the availability of a unit when making the purchasing decision. If a unit is available, then the immediate surplus would be $(v - p_1)$. But, if the customer decides to wait, the expected surplus needs to be calculated based on the information that a unit is currently available. In other words, the expression in (2) must be replaced by conditional expectation. We refer the reader to Osadchiy and Vulcano (2010), in which such approach is adopted.

² We have also tested a set of an additional 12,000 instances with $E[\Lambda] = 50$. For this set, we extended the values of Q , ranging from one

to 80. The results based on this augmented study were fully aligned with our reported observations, and hence, we omit them.

References

- Araman VF, Caldenteu R (2009) Dynamic pricing for nonperishable products with demand learning. *Oper. Res.* 57(5):1169–1188.
- Aviv Y, Pazgal A (2005) A partially observed Markov decision process for dynamic pricing. *Management Sci.* 51(9):1400–1416.
- Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.
- Aviv Y, Vulcano G (2012) Dynamic list pricing. Ozer O, Phillips R, eds. *Handbook of Pricing Management* (Oxford University Press, Oxford, UK), 522–584.
- Aviv Y, Levin Y, Nediak M (2009) Counteracting strategic consumer behavior in dynamic pricing systems. Netessine S, Tang CS, eds. *Consumer-Driven Demand and Operations Management Models* (Springer, New York), 323–352.
- Balvers RJ, Cosimano TF (1990) Actively learning about demand and the dynamics of price adjustments. *Econom. J.* 100(402):882–898.
- Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.* 57(6):1407–1420.
- Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3):497–511.
- Cachon GP, Swinney R (2011) The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management Sci.* 57(4):778–795.
- Coase RH (1972) Durability and monopoly. *J. Law Econom.* 15(1): 143–149.
- Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* 49(1):1287–1309.
- Elmaghraby W, Gulcu A, Keskinocak P (2008) Designing optimal preannounced markdowns in the presence of rational customers with multi-unit demands. *Manufacturing Service Oper. Management* 10(1):126–148.
- Farias VF, van Roy B (2010) Dynamic pricing with a prior on market response. *Oper. Res.* 58(1):16–29.
- Fisher M (1997) What is the right supply chain for your product? *Harvard Bus. Rev.* 75(2):105–117.
- Fisher M (2009) Rocket science retailing: The 2006 Philip McCord Morse lecture. *Oper. Res.* 57(3):527–540.
- Fisher M, Rajaram K (2000) Accurate retail testing of fashion merchandise: Methodology and application. *Marketing Sci.* 19(3):266–278.
- Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.
- Ghemawat P, Nueno JL (2003) ZARA: Fast fashion. Harvard Business Case 9-703-416, Harvard Business School, Boston.
- Heun CT (2001) Dynamic pricing boosts bottom line. *InformationWeek* (October 29), <https://www.informationweek.com/dynamic-pricing-boosts-bottom-line/d/d-id/1012413>.
- Lai G, Debo LG, Sycara K (2010) Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *Manufacturing Service Oper. Management* 12(1):33–55.
- Landsberger M, Meilijson I (1985) Intertemporal price discrimination and sales strategy under incomplete information. *RAND J. Econom.* 16(3):424–430.
- Levin Y, McGill J, Nediak M (2009) Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management Sci.* 55(1):32–46.
- Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9):2114–2137.

- Liu Q, van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.
- Mersereau AJ, Zhang D (2012) Markdown pricing with unknown fraction of strategic customers. *Manufacturing Service Oper. Management* 14(3):355–370.
- Netessine S, Tang CS (2009) *Consumer-Driven Demand and Operations Management Models: A Systematic Study of Information-Technology-Enabled Sales Mechanisms* (Springer, New York).
- Osadchiy N, Vulcano G (2010) Selling with binding reservations. *Management Sci.* 56(12):2173–2190.
- Puterman ML (2005) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley, Hoboken, NJ).
- Raiffa H, Schlaifer R (2000) *Applied Statistical Decision Theory* (Wiley, Hoboken, NJ).
- Stokey NL (1979) Inter-temporal price discrimination. *Quart. J. Econom.* 93(3):355–371.
- Su X (2007) Inter-temporal pricing with strategic customer behavior. *Management Sci.* 53(5):726–741.
- Su X, Zhang F (2008) Strategic customer behavior, commitment and supply chain performance. *Management Sci.* 54(10):1759–1773.
- Su X, Zhang F (2009) On the value of commitment and availability guarantees when selling to strategic consumers. *Management Sci.* 55(5):713–726.
- Van Ryzin GJ, Talluri KT (2005) *The Theory and Practice of Revenue Management* (Springer, New York).
- Yin R, Aviv Y, Pazgal A, Tang CS (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic consumers. *Management Sci.* 55(8):1391–1408.