Price Commitments with Strategic Consumers: Why It Can Be Optimal to Discount More Frequently … Than Optimal

Gérard P. Cachon, Pnina Feldman

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

Price Commitments with Strategic Consumers: Why It Can Be Optimal to Discount More Frequently...Than Optimal

Gérard P. Cachon
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, cachon@wharton.upenn.edu

Pnina Feldman
Haas School of Business, University of California, Berkeley, Berkeley, California 94720, feldman@haas.berkeley.edu

In many markets consumers incur search costs, and firms choose a long-run pricing strategy that determines how they respond to market conditions. A pricing strategy may involve commitments to take actions that do not optimize short-term revenue given the information the firm learns about demand. For example, as already suggested in the literature, the firm could commit to a single price no matter whether demand is strong or weak. We introduce a new strategy—charge a "high" price only if demand is indeed "high," otherwise offer a discount. This strategy discounts more frequently than would maximize revenue conditional on demand. Nevertheless, the frequent discounts attract consumers. We show that (i) the discount-frequently strategy is optimal (whether capacity is adjustable or not), (ii) discount-frequently is often much better than other pricing strategies, especially if no price commitment is made, and (iii) "overbuying" capacity (e.g., inventory) to attract consumers (by signaling availability and the likelihood of discounts) is a poor strategy. Contrary to some recommendations in the literature to limit markdowns and to purchase ample capacity, our results provide support for a strategy that embraces frequent discounts and moderate capacity.

1. Introduction

In many markets consumers incur costs to search/visit a firm, so they search only if it is worth the effort. In particular, consumers care about (1) what price they pay and (2) do they get a unit. A great price does no good for an item that is out of stock or a service that cannot be offered (e.g., no appointments or seats available). And availability is not useful if the price is too high. So in these environments the firm needs to attract consumers with a good deal (price and availability). The firm can do this with two levers: a pricing strategy and a capacity choice.

A pricing strategy communicates to consumers how the firm behaves. It often involves commitments to act in certain ways that incur some cost to the firm in the short term. For example, a firm often has better information about overall demand for a specific product or service than consumers do. A clothing retailer may know that a dress is more popular than usual, or a movie theater may learn that the latest movie is not pulling in audiences as was expected. The dress retailer may be tempted to raise its price (or not discount it), knowing that it can sell all dresses even at a higher price. The theater may be tempted to drop its price in an effort to attract bargain hunters. In either case, failing to make the price adjustment in response to the firm's updated information is costly in the short term because changing the price can increase revenue conditional on what the firm knows, as has been shown in many settings with nonstrategic consumers (i.e., consumers whose search decisions do not depend on the firm's pricing or capacity choices); e.g., Gallego and van Ryzin (1994), Talluri and van Ryzin (2004), Elmaghraby and Keskinocak (2003).

This paper studies a model with two time horizons. In the long term the firm, armed with an uncertain demand forecast, chooses a pricing strategy and possibly a capacity. In the short term the firm learns useful information about demand and then chooses a price that is consistent with the adopted pricing strategy. We consider three strategies. With the no-commitment strategy, the firm charges a price that maximizes revenue conditional on observed demand—if demand is "high," the firm charges a high list price, leaving consumers with no residual value, whereas if demand is "low," the firm charges a discounted price to ensure that all inventory is sold. With a static-pricing strategy, the firm commits to a single price no matter what it learns about demand. Although a static-pricing strategy does not allow a

firm to exploit new information, it has been found that committing to a fixed price (or set of prices) can help to reduce the propensity of consumers to strategically wait for discounts (e.g., Besanko and Winston 1990, Aviv and Pazgal 2008).

We introduce a third strategy, called discount frequently, in which the firm charges either the high (not discounted) list price or a discounted price, just as in the no-commitment strategy. The commitment with this policy is with the frequency the firm chooses to discount—in some cases, the firm discounts off the list price even if it is not in the interest of the firm to discount conditional on what it observes about demand. With the discount-frequently strategy, a consumer knows that sometimes she does not get a good deal (the price is high), but the firm nevertheless offers a discount often enough to justify the effort to visit the firm. For example, given demand, the firm may be indifferent between selling a portion of its capacity at the high price or discounting to sell all of its capacity. In those cases, the firm chooses to discount—it loses nothing but makes customers happier. It also chooses to discount in some cases in which it prefers the higher price. Although this is costly (again, conditional on demand), this commitment attracts more consumers to the firm, which is clearly beneficial.

For a fixed capacity, we find that (i) the static-pricing strategy can perform well, and in many cases much better than the no-commitment strategy, (ii) but the discount-frequently strategy is optimal among all possible strategies. Thus, the best strategy is to charge both high and low prices but also to avoid the temptation to maximize the short term (by not offering a discount often enough) at the expense of long-term profitability. Interestingly, given the same potential demand and the same capacity, a firm implementing the no-commitment strategy discounts with the same actual frequency as the firm that implements the discount-frequently strategy. The difference is that the discount-frequently firm attracts a higher fraction of potential demand, which makes the discount-frequently strategy more profitable. A naive manager may conclude that a discount-frequently firm is discounting too frequently given the demand it receives, not recognizing that the firm receives that demand precisely because it maintains the same frequency of discounting as the no-commitment firm. However, this can be a costly conclusion—we find that the short-term profit gain from avoiding discounts is generally considerably less than the loss in long-run profit if the firm loses its reputation for discounting frequently.

If the firm can choose its capacity, then it can also use capacity as a "carrot" to attract consumers. Dana and Petruzzi (2001) show that if consumers

care about availability—they do not like incurring the cost to search the firm only to learn that they cannot purchase the good or service—then the firm should invest in more capacity than it would if consumers were nonstrategic. (Gaur and Park 2007 find an analogous result in a model with consumer learning and multiple firms.) The additional capacity is meant to reassure consumers that the product will be available for them so they become more likely to search the firm—the firm cannot sell to a customer that does not visit. Similarly, adding capacity increases the likelihood that a clearance sale is justified, thereby exciting consumers with the greater prospect of a good deal. However, Dana and Petruzzi (2001) did not consider the possibility that the firm could choose its pricing strategy, in particular the frequency of clearance sales to help draw consumers. We find that it is ineffective to only use excess capacity to attract consumers. For example, the firm that is unable or unwilling to make price commitments to attract consumers purchases much more capacity and earns substantially lower profit than the firm that makes price commitments. As with a fixed capacity, the discount-frequently strategy is optimal. In fact, when capacity is expensive, discount frequently is the only strategy that makes a positive profit.

2.   Related Literature

We are not the first to study price commitments even though we introduce a novel type of price commitment. A number of authors consider models in which consumers decide whether to purchase at the current price or to wait to purchase at a future price (e.g., Besanko and Winston 1990; Feng and Gallego 1995; Aviv and Pazgal 2008; Liu and van Ryzin 2008; Su and Zhang 2009; Cachon and Swinney 2009, 2011; Swinney 2011, 2012; Ovchinnikov and Milner 2012; Özer and Zheng 2015; Whang 2015). To combat this strategic behavior, it has been suggested that a firm commit to restrict discounts. This commitment is costly because it limits the firm's ability to react to updated demand information. However, a price commitment can detract consumers from strategically waiting for a discount. In our model consumers are offered a single price, so they do not consider whether to "buy now or wait for the discount." Hence, our motivation for price commitments is exclusively to attract consumers rather than to prevent strategic waiting.

A critical feature of our model is that consumers incur search costs—consumers choose to visit the firm only if they anticipate that the reward for doing so (i.e., purchasing a product at a good price) justifies the effort. Others have incorporated similar search costs, generally in a single-firm setting: Baye and Morgan (2001), Dana and Petruzzi (2001), Çil and Lariviere

(2013), Alexandrov and Lariviere (2012), and Su and Zhang (2009). Dana and Petruzzi (2001) have fixed prices and focus instead on how search costs influence the firm's capacity choice. Baye and Morgan (2001) study a marketplace for price information for which customers may pay to subscribe or, alternatively, incur search costs to learn the price set by their local firm only. Çil and Lariviere (2013) study the allocation of capacity across two market segments, and Alexandrov and Lariviere (2012) study why firms may offer reservations. Prices are exogenously fixed in both of those papers. Su and Zhang (2009) focus on capacity commitments and availability guarantees.

Pricing and availability are considered in a number of papers that model competition across two or more firms, e.g., Deneckere and Peck (1995), Bernstein and Fedegruen (2004), Cachon and Harker (2002), Gaur and Park (2007), and Allon and Fedegruen (2007). These models assume firms choose a price without the benefit of updated demand information. Several papers empirically document that consumers do value higher availability, e.g., Masta (2011) and Cachon et al. (2013).

Other papers that compare different pricing schemes when consumers are strategic include single versus priority pricing (Harris and Raviv 1981), subscription versus per-use pricing (Barro and Romer 1987, Cachon and Feldman 2011), and markdown regimes with and without reservations (Elmaghraby et al. 2009). In all of these papers the firm selects its pricing strategy before learning some updated demand information, whereas in our study we allow the firm to choose a price after potential demand is observed. Thus, consumers in our model are not sure what is the firm's price or the product's availability before they choose whether to search the firm.

3. Model Description

A single firm with k units of capacity sells to two types of consumers, high types and low types, all of whom require one unit of capacity to be served and are indistinguishable to the firm. There is a potential number of X high-value consumers, where X is a nonnegative random variable that is drawn from a cumulative distribution function $F(\cdot)$, probability density function $f(\cdot)$, and mean $\mu = \mathbb{E}[X]$. Let $\bar{F}(\cdot) = 1 - F(\cdot)$. The high-value consumers have zero "mass" as each is unable to have any influence on the market individually. They have value v_h for the firm's service. As in Dana and Petruzzi (2001) (and other papers), high-value consumers must incur a positive cost, $c < v_h$, to search the firm to purchase the good or service. These search costs include the time and effort to physically travel to the firm and the mental effort associated with a purchasing decision. They

Table 1 Summary of Consumer Types

Segment	Number	Value	Search cost
High type	$X \sim F(\cdot)$	v_h	c
Low type	∞	v_l	0

receive zero value if they choose not to search. They can implement mixed strategies: let $\gamma \in [0, 1]$ be the probability that a high-type consumer searches the firm. Mixed strategies have also been used to describe consumer behavior in the context of joining a service, modeled as a queue (e.g., Edelson and Hildebrand 1975, Lariviere and van Mieghem 2004), paying to have access to a list of prices from multiple firms (e.g., Baye and Morgan 2001), and whether to visit a restaurant (e.g., Çil and Lariviere 2013).

There is an ample number of low-value consumers, and each has v_l value for the firm's service. These consumers have zero (or low) search costs. We assume that $v_l < v_h - c$: a high-type consumer who visits the firm generates more value than a low type (net of search cost). If $v_l > v_h - c$, then the firm prefers selling exclusively to low-type consumers, which is not interesting. Table 1 summarizes the consumer types.

The firm seeks to maximize revenue and consumers seek to maximize their net value, the value of the service minus search costs and the price paid to the firm.

Events can be divided into two periods. In the first, or "long-term," period the firm chooses a pricing strategy. The strategy determines how the firm behaves in the second, or "short term" period. All consumers observe the firm's pricing strategy. At the start of the second period the number of high-type consumers, X, is realized. The firm observes X but consumers do not. Next, the firm chooses a price, p, and at the same time the high-type consumers choose to search the firm or not (i.e., they select γ). At the end of the short-term period consumers who visit the firm purchase an item, if available. If there are more than k high-type consumers who want to purchase, the k units are randomly rationed among them (our allocation rule). While we model the "short-term" period as a single period, results are equivalent to a model that has one "long-term" period followed by multiple but independent "short-term" periods (hence the names).

We consider three pricing strategies. With the first, called "no commitment," the firm chooses a price in the short-term period that maximizes revenue conditional on observed demand, X, capacity, k, and the firm's expectation of consumer behavior. There is no commitment with this strategy because the firm is maximizing its short-term revenue conditional on all the information it knows. The second pricing strategy is called "static pricing" because in the long-term

period the firm commits to charge a single price, p_s, in the short-term period. This requires a commitment because p_s may not be the price that maximizes the firm's revenue in the short-term period. However, by committing to $p_s < v_h$, the firm potentially increases the return a high type receives from search, thereby increasing the number of high-type consumers who search. Finally, we consider a "discount-frequently" strategy in which the firm chooses the same prices as in the no-commitment strategy (v_l or v_h), but chooses the discount price with a higher likelihood (i.e., more frequently) than with the no-commitment strategy. Hence, the firm sometimes does not maximize revenue in the short-term period conditional on its information, but this commitment also encourages high types to search the firm because they expect to receive the discounted price with a higher probability.

There are additional pricing strategies beyond the three we consider. For example, instead of committing to the frequency of the two focal prices (v_l or v_h), the firm could commit to always offer a discount by either choosing to maximize short-term revenue with an intermediate price v_m, $v_l < v_m < v_h$, or the deep discount, v_l. (If $v_m = v_h$, then this is the no-commitment strategy. If $v_m = p_s$, then this strategy dominates static pricing.) Using the set of scenarios described in our numerical study (§§5 and 6), we find this "discount always" strategy is reasonably effective when capacity is fixed (it achieves on average 96.9% of the revenue earned with discount frequently), but less effective when capacity can be chosen (it earns on average only 79% of the profit earned by discount frequently). (A complete analysis of the discount-always policy is available from the authors.)

There are several important features of our model, which we discuss next.

The firm is able to commit to a pricing strategy. Like Aviv and Pazgal (2008), Liu and van Ryzin (2008), Elmaghraby et al. (2008), Yin et al. (2009), Liu and Shum (2013), and Whang (2015), (i) we allow the firm to commit to a pricing strategy that is not always sub-game perfect, and (ii) we presume these commitments are credible. Credibility is generally achieved through repeated interaction (e.g., Fudenberg and Levine 1989). Although we consider a model with only one short-term period, we have in mind a situation in which the firm interacts with consumers over multiple short-term periods (e.g., multiple months or quarters). Consequently, the firm is able to establish a long-run reputation for how it conducts business. For example, in January 2012, JC Penney, a large U.S. department store chain, announced a new, simplified pricing strategy that involved far fewer discounts. However, by the spring of 2013, the company decided that the strategy did not work, and they returned to a more aggressive use of price promotions

(Clifford and Rampell 2013). Thus, a firm commits to a pricing strategy through advertising and subsequent behavior—consumers who visited JC Penney after the announcement indeed noticed that they were not promoting as frequently.

Consumers must incur a search cost before observing availability and price. Search costs associated with product availability are likely to be inconsequential only if a consumer knows exactly which item they want (at the brand/model level) and they have the ability to find availability information quickly and accurately. Those conditions are likely to apply only in specialized situations. It is more common that consumers might not know exactly the item they want, or shop at retailers that are unable to easily provide accurate availability information. A department store, such as JC Penney, fits this description—a consumer might know that they want to purchase a blender from their housewares department, but they do not know the exact brand and model, and even if they did, the website (assuming they have easy Internet access) may not provide timely and accurate inventory information for their local store. (See DeHoratius and Raman 2008 for evidence that firms struggle to maintain accurate inventory records, even for their own internal use.) It is also likely that consumers incur significant search costs for price. Again, for the search cost to be inconsequential, the consumer must know the precise item they intend to purchase, which does not always apply. But even if that is known, search costs can remain. For example, a search on a retailer's website may provide the price of an item at one store near the consumer, but retailers do not always charge the same price across all stores. To find the full list of prices in the nearby stores may require calling the individual stores, which is clearly time consuming (i.e., a costly activity). Finally, Hann and Terwiesch (2003) discover that consumers act as if price search is costly even when intuition suggests it should not be, possibly because of cognitive effort or limitations (e.g., Miller 1956, Roberts and Lattin 1997, Kuksov and Villas-Boas 2010). In sum, it is likely that in many markets consumers behave as if searching for price and availability is consequential (i.e., costly).

Low types have no search costs. In contrast to the high types that are limited in number and incur search costs, low types are ample and search is inconsequential to them. One interpretation is that the low-type consumers are bargain hunters who visit the store without the intention to purchase in the category of interest but nevertheless are willing to make a purchase if they notice a very good deal (e.g., the price is no more than v_l). Alternatively, even if low-type consumers have a search cost, the qualitative results of the model remain. To explain, as long as their search costs are sufficiently low, a certain fraction of them

will be willing to visit the store (like the high types, they would make a trade-off between the cost of visiting and the potential value of visiting). Thus, the firm would continue to make the trade-off between pricing "high" to sell only to the high types and pricing "low" to sell to both types. The firm might not always be able to clear remaining inventory with the "low" price, but the firm could still generate a considerable sales boost by discounting its price.

The firm observes demand before selecting its price. The firm may use early season sales to quickly determine if the product has excess demand or not (e.g., Raman and Fisher 1996, Iyer and Bergen 1997, Caro and Gallien 2012). The firm uses this information when it chooses its price, constrained by its pricing strategy (and capacity) commitments. Although we assume the firm observes a perfect signal of demand, we suspect our qualitative results continue to hold (though the analysis becomes more cumbersome) if the firm is imperfectly informed but remains better informed than consumers—when the firm has more information than consumers, consumers know that the firm may be tempted to use that information to increase its revenue, and, thus, price commitments can still be used to attract more demand.

High types receive priority in the allocation rule. This allocation rule is most favorable to the firm because it encourages high-type consumers to visit the firm (they know that they have priority). (Su and Zhang 2008 and Tereyagoglu and Veeraraghavan 2012 also adopt this allocation rule.) Alternatively, as in Cachon and Swinney (2009), high-type and bargain hunting consumers could form a queue in which every $1/\theta$ customer is a high type until there are no more high types, where $\theta \in [0, 1]$. If $\theta = 1$, high types are given full priority, which is the allocation rule we consider. As θ decreases, high types are more likely to be rationed. Nevertheless, all our results apply even if the high-type consumers are not given full priority. (Details are available from the authors.)

4. Analysis
This section analyzes the three pricing strategies already discussed: no commitment, static pricing, and discount frequently.

4.1. No-Commitment Pricing
Under the no-commitment strategy, the firm chooses either v_h or v_l. Given $\{v_l, v_h\}$, the firm can price at $p = v_l$ and earn revenue $v_l k$. Alternatively, it can price at $p = v_h$ and earn revenue $v_h \min\{\gamma x, k\}$. Consequently, the firm chooses $p = v_l$ when

$$x \leq \frac{v_l}{v_h} \frac{k}{\gamma}, \tag{1}$$

which occurs with probability $F(v_l k/(v_h \gamma))$ and chooses $p = v_h$ otherwise.

The high-type consumers only earn positive utility if the price is v_l and they are able to obtain the unit. In all other cases, consumers get zero surplus. Thus, to find the high-type consumer surplus from visiting the firm, let ψ be the high-type consumer's expectation for the probability that the firm charges v_l and he is able to get a unit. A high-type consumer is indifferent between searching the firm or not if

$$\psi(v_h - v_l) = c. \tag{2}$$

In equilibrium, the belief about the probability ψ must be consistent with the actual probability. Given the rationing rule, because v_l is charged only when $\gamma x \leq (v_l/v_h)k < k$ (from (1)), high-type consumers are *guaranteed* to get the unit when the price is v_l—the firm discounts the product only when demand is sufficiently low, which means that a unit is available for everyone. They may not be able to get a unit if the price is v_h, but in this case, their surplus is zero either way. Thus, high-type consumers do not face a rationing risk if their utility from getting a unit is positive.

From (1), the high-type consumer knows that the firm charges a low price when demand is sufficiently low. From the perspective of a consumer, the probability density function of x is $\tilde{f}(x) = xf(x)/\mu$. (See Deneckere and Peck 1995 for a detailed derivation of the demand density conditional on a consumer's presence in the market, $\tilde{f}(x)$.) Therefore, this consumer anticipates that the price is v_l with probability

$$\psi = \int_0^{(v_l/v_h)(k/\gamma)} \tilde{f}(x) \, dx = \frac{\int_0^{(v_l/v_h)(k/\gamma)} xf(x) \, dx}{\mu}. \tag{3}$$

Given $v_l < v_h - c$, there exists some γ that satisfies (3). A symmetric equilibrium strategy for high-type consumers is a $\gamma \in [0, 1]$ such that γ is optimal for each customer given that all other consumers choose γ as their strategy.

Let γ_0 be the fraction of consumers who visit the firm in equilibrium under the no-commitment policy. The following lemma characterizes γ_0. (Proofs are provided in the online supplement, available as supplemental material at http://dx.doi.org/10.1287/msom.2015.0527.)

Lemma 1. *With the no-commitment pricing strategy, the fraction of high-type consumers who visit the firm in equilibrium, γ_0, is unique. Furthermore, $\gamma_0 = 1$ if*

$$\int_0^{(v_l/v_h)k} xf(x) \, dx \geq \frac{\mu c}{v_h - v_l};$$

otherwise, γ_0 is implicitly defined by

$$(v_h - v_l) \int_0^{(v_l/v_h)(k/\gamma_0)} xf(x) \, dx = \mu c. \tag{4}$$

The firm's revenue under no-commitment, $R_0(\gamma_0)$, is

$$R_0(\gamma_0) = F\left(\frac{v_l}{v_h}\frac{k}{\gamma_o}\right)v_l k + v_h \gamma_0 \int_{(v_l/v_h)(k/\gamma_o)}^{k/\gamma_o} x f(x)\, dx$$
$$+ \bar{F}\left(\frac{k}{\gamma_o}\right)v_h k$$
$$= v_l k + v_h \gamma_0 \left(S\left(\frac{k}{\gamma_o}\right) - S\left(\frac{v_l}{v_h}\frac{k}{\gamma_o}\right)\right). \quad (5)$$

4.2. Static Pricing

With a static-pricing strategy, the firm commits to a single price, p, before observing demand, so consumers know that the price will indeed be p before deciding whether or not to search the firm. All high-value consumers who search the firm receive a net value equal to $v_h - p - c$ if they obtain a unit, and if they do not obtain a unit, their net value is $-c$. A customer searches the firm if net utility is not negative, i.e., if

$$\phi(v_h - p) \geq c, \quad (6)$$

where ϕ is the customer's expectation for the probability of getting a unit conditional on searching the firm. The underlying potential demand distribution, X, the high-value customers' strategy, γ, and the rationing rule used to allocate scarce capacity determine ϕ. All else being equal, as γ increases, more high-type customers search the firm, thereby reducing the chance that any one of them gets a unit. Consequently, the probability she gets a unit is

$$\phi = \int_0^\infty \frac{\min\{\gamma x, k\}}{\gamma x}\tilde{f}(x)\, dx = \int_0^\infty \frac{\min\{\gamma x, k\}}{\gamma x}\frac{x f(x)}{\mu}\, dx$$
$$= \frac{S_{\gamma x}(k)}{\gamma \mu} = \int_0^\infty \frac{\min\{x, k/\gamma\}}{\mu}f(x)\, dx = \frac{S(k/\gamma)}{\mu}, \quad (7)$$

where $S_D(q) = \mathbb{E}_D[\min\{D, q\}]$ is the sales function given demand D and $S(\cdot)$ is shorthand for $S_X(\cdot)$. (Note, $S_{\gamma X}(k) = \gamma S(k/\gamma)$.) Given a static price, p, a symmetric equilibrium strategy for high-type consumers is a $\gamma(p) \in [0, 1]$ such that $\gamma(p)$ is optimal for each consumer given that all other consumers choose $\gamma(p)$ as their strategy. If p is low enough, there is an equilibrium in which all high-type consumers visit the firm, i.e., $\gamma(p) = 1$. From (6) and (7), that occurs if

$$\frac{S(k)}{\mu}(v_h - p) \geq c$$

or

$$p \leq v_h - \frac{\mu c}{S(k)} \equiv \bar{p}.$$

If $p > \bar{p}$, the unique symmetric equilibrium has $\gamma(p) < 1$, where $\gamma(p)$ is the unique solution to

$$S\left(\frac{k}{\gamma(p)}\right) = \frac{\mu c}{v_h - p}. \quad (8)$$

Using (8), the firm's revenue function can be written as a function of γ alone. Define $R_s^h(\gamma)$ as the firm's revenue function from only high-type consumers:

$$R_s^h(\gamma) = S_{\gamma X}(k)\left(v_h - \frac{\mu c}{S(k/\gamma)}\right)$$
$$= \gamma S\left(\frac{k}{\gamma}\right)v_h - \gamma \mu c. \quad (9)$$

The next lemma finds the equilibrium fraction of high-type consumers who search the firm under static-pricing, γ_s.

LEMMA 2. *Define $\tilde{k}$ implicitly as*

$$v_h \int_0^{\tilde{k}} x f(x)\, dx = \mu c.$$

With static pricing, the firm's revenue function from high-type consumers, $R_s^h(\gamma(p))$, is concave. Let $\gamma_s = \arg\max R_s^h(\gamma)$. The price charged to the high types is p_s^h. If $k \geq \tilde{k}$, then $\gamma_s = 1$ and $p_s^h = \bar{p}$. Otherwise, $\gamma_s = k/\tilde{k}$ and

$$p_s^h = v_h - \frac{\mu c}{S(\tilde{k})}.$$

Instead of choosing p_s^h and selling only to high-type consumers, the firm also has the option to choose $p_s = v_l$, in which case the firm sells all its capacity and its revenue is $p_s k$. Finally, the firm chooses the static price, $p_s \in \{p_s^h, v_l\}$ to maximize revenues. When $\gamma_s < 1$, $R_s^h(\gamma_s) = k v_h \bar{F}(\tilde{k})$. Thus, under static pricing, the optimal price is $p_s = p_s^h$ when $v_h \bar{F}(k/\gamma_s) \geq v_l$, otherwise $p_s = v_l$.

4.3. Discount Frequently

Static pricing commits to charge some price that is suboptimal once demand is observed. Although it is costly, it may be done to increase demand from high-type consumers—they do not visit if the expectation of what they can receive is too low. However, there is another way to make searching the firm attractive to consumers. Instead of always providing an intermediate discount (as in static pricing), the firm could commit to provide the deep discount (to v_l) more frequently than would be optimal given the realization of demand. In particular, with the discount-frequently strategy the firm chooses to price at either v_l or v_h (the two optimal prices ex post) but commits to markdown to v_l whenever potential demand is $\delta k/\gamma_f$ or lower and charge v_h otherwise, where $\delta \in [v_l/v_h, 1]$. This implies that the firm charges v_l with probability $F(\delta k/\gamma_f)$. Under this policy the firm does not commit to limit its prices—the firm charges the same prices as in the no-commitment strategy. Rather, it commits to discount often and deeply, thereby encouraging high-type consumers to search. The following lemma characterizes the fraction of high-type consumers who search the firm under the discount-frequently strategy.

LEMMA 3. *With the discount-frequently strategy, the fraction of high-type consumers who visit the firm in equilibrium, γ_f, is unique. Furthermore, $\gamma_f = 1$ if*

$$(v_h - v_l) \int_0^{\delta k} x f(x)\, dx > \mu c; \tag{10}$$

otherwise, γ_f is implicitly defined by

$$(v_h - v_l) \int_0^{\delta k/\gamma_f} x f(x)\, dx = \mu c. \tag{11}$$

The revenue function is

$$R_f(\delta, \gamma_f) = v_l k F(\delta k/\gamma_f) + \gamma_f v_h \int_{\delta k/\gamma_f}^{k/\gamma_f} x f(x)\, dx + v_h k \bar{F}(k/\gamma_f).$$

If the firm chooses $\delta = v_l/v_h$, then it replicates the no-commitment strategy. (Hence, discount frequently cannot be worse than no commitment.) However, if $\delta > v_l/v_h$, then the firm discounts more frequently than would be optimal to maximize revenue conditional on realized demand (i.e., holding γ_f fixed), which may entice enough high-type consumers to visit to justify the cost of discounting even though demand is high enough to keep the price high. The next theorem characterizes the optimal discount frequency, δ^*, and the corresponding fraction of high-type consumers who visit given that discount frequency, $\gamma_f^* = \gamma_f(\delta^*)$.

THEOREM 1. *Define $\hat{k}$ implicitly as*

$$(v_h - v_l) \int_0^{\hat{k}} x f(x)\, dx = \mu c. \tag{12}$$

With the discount-frequently strategy, let δ^* and γ_f^* be the optimal discount frequency and the resulting high-type consumer search probability, respectively:

$$\delta^* = \begin{cases} 1 & k \le \hat{k}, \\ \hat{k}/k & \hat{k} < k \le \hat{k} v_h/v_l, \\ v_l/v_h & \hat{k} v_h/v_l < k; \end{cases} \quad \gamma_f^* = \begin{cases} k/\hat{k} & k \le \hat{k}, \\ 1 & \hat{k} < k \le \hat{k} v_h/v_l, \\ 1 & \hat{k} v_h/v_l < k. \end{cases}$$

According to Theorem 1, the actual probability of a discount is $F(\hat{k})$ when $k \le \hat{k} v_h/v_l$ (and $F(v_l k/v_h)$ otherwise), which is exactly the same probability of a discount with the no-commitment strategy. The difference is that with discount frequently, the firm enjoys higher demand (more high types choose to search), precisely because the firm maintains frequent discounts even with the higher demand. For example, when $k < \hat{k}$, the firm discounts whenever realized high-type demand is less than capacity—the firm charges the higher price, v_h, only when it is able to sell its entire capacity to high-type consumers. Consequently, unlike with static pricing, the firm always

clears its entire inventory with the best discount-frequently strategy. Hence, it could be named the "everything must go" strategy.

The firm's revenue function with the optimal discount-frequently strategy is

$$R_f^* = \begin{cases} k(v_l F(\hat{k}) + v_h \bar{F}(\hat{k})) & k \le \hat{k}, \\ \\ v_l k F(\hat{k}) + v_h \int_{\hat{k}}^{k} x f(x)\, dx + v_h k \bar{F}(k) \\ \qquad\qquad \hat{k} < k < \hat{k} v_h/v_l, \\ \\ v_l k F\left(\dfrac{v_l}{v_h} k\right) + v_h \int_{(v_l/v_h)k}^{k} x f(x)\, dx + v_h k \bar{F}(k) \\ \qquad\qquad \hat{k} v_h/v_l \le k. \end{cases}$$

Our main result, reported in Theorem 2, is that discount frequently is optimal for the firm considering all possible policies. This is a strong motivation for the use of the discount-frequently strategy. By offering discounts sufficiently often, and in some cases when the firm's short-term demand does not justify the discount, the firm provides the motivation needed for consumers to visit, thereby increasing the firm's demand, which in turn leads to the maximum revenue.

THEOREM 2. *The discount-frequently strategy maximizes the firm's revenue.*

Why is discount frequently optimal? Commitments are made to make searching the firm more attractive to consumers, but they are also costly to the firm. Thus, the best commitment substantially improves the attractiveness of the firm relative to its cost. Suppose the firm observes that realized demand, x, equals $v_l k/(v_h \gamma)$. The firm is indifferent between charging v_h and selling γx units or charging v_l and selling k units. As the revenue function is *flat* at that point, a slight increase in the demand level that triggers a discount, say to $v_l k/(v_h \gamma) + \epsilon$, causes essentially no loss in revenue, yet it generates an immediate benefit in that it attracts more high-type customers to visit. Thus, this slight distortion from the short-term optimal causes little cost initially in terms of revenue conditional on γ but provides an immediate benefit in that it strictly increases γ, the fraction of high types that search. As the distortion increases, the cost increases, but this distortion in the short-term decision provides the greatest benefit relative to its cost. Cachon and Lariviere (2001) observe a similar finding in the context of signaling capacity in a supply chain—some signals are cheaper to make than others, and the best signal distorts a decision near the "flat area" of a maximum.

**Table 2   Parameter Values Used in the Numerical Study**

Parameter	Values
Demand distribution	Gamma
μ	1
σ	$\{0.25\,\mu, 0.5\,\mu, \mu, 1.5\mu, 2\,\mu\}$
v_h	1
v_l	$\{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.4, 0.5, 0.6\}$
c	$\{0.01, 0.02, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$
$F(k)$	$\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99\}$

5. Comparison of Strategies with Fixed Capacity

Table 2 lists the parameters we use in a numerical study to compare the pricing strategies discussed in the previous section. To compare revenue, it is sufficient to manipulate k, σ, $\mu c/v_h$, and v_l/v_h. Hence, we fix the values of v_h and μ to 1. To represent different levels of demand uncertainty, with the gamma distribution, the coefficient of variation ranges from a low level of uncertainty (e.g., $\sigma/\mu = 0.25$) to a high level of uncertainty (e.g., $\sigma/\mu = 2$). The v_l and c parameters span the range $[0, v_h]$ while also satisfying $v_l \leq v_h - c$. There are 360 combinations of these parameters. We choose capacity to correspond to 10 different fractiles of the potential demand distribution, $k = \{F^{-1}(0.1), \ldots, F^{-1}(0.99)\}$. This yields a total of 3,600 scenarios. However, if capacity is greater than $v_h \hat{k}/v_l$, then the no-commitment and discount-frequently strategies are both optimal. Thus, we remove scenarios that have capacity above this threshold, leaving 3,146 scenarios.

Figure 1 illustrates the revenue achieved with each policy for a representative scenario and a wide range of capacities. (The 99th percentile of capacity is $k = 4.6$.) Although no commitment performs better than static pricing for very large capacities, for most reasonable capacities, static pricing performs best and

Figure 1   Comparison of Expected Revenue with Three Pricing Strategies for One of the Considered Scenarios
$(v_h = 1, v_l = 0.1, c = 0.25, X \sim \text{Gamma}(1, 1))$
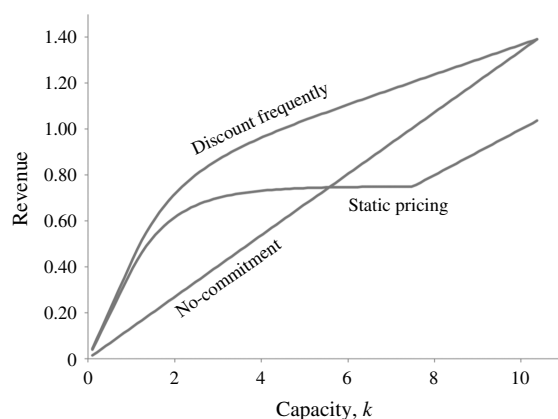
Figure 2   Revenue Ratios
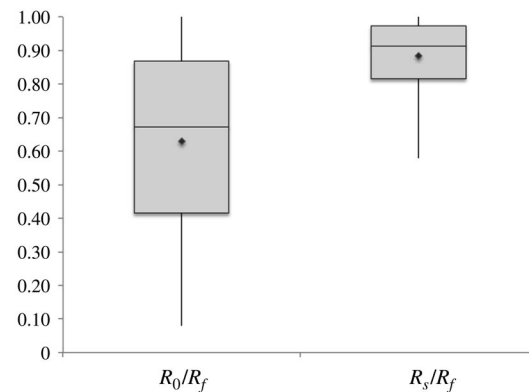
Note. Diamond = average; the bottom, middle, and upper horizontal lines correspond to the first quartile, median and third quartile, respectively; the bottom and top points of the vertical lines are the minimum and the maximum, respectively.

potentially much better. Overall, the figure suggests that the firm sacrifices a considerable amount of revenue if it does not make a price commitment (i.e., no commitment performs poorly). Furthermore, some types of commitments perform much better than others (i.e., discount frequently is best). These specific findings extend to our larger sample, as we discuss next.

Figure 2 shows a box-plot of the revenue ratios of the pricing strategies relative to the revenue obtained from discount frequently. As observed in Figure 1, the discount-frequently strategy performs substantially better than no commitment or static pricing. Though both schemes can approach the discount-frequently strategy in some cases (with maximum ratios of $R_0^*/R_f^* = 1$ and $R_s^*/R_f^* = 99.99\%$), in most cases they perform poorly in comparison: the average R_0^*/R_f^* ratio is 63.0% and the average R_s^*/R_f^* ratio is 88.3%.

The results in Figure 2 emphasize two points. First, failing to make a price commitment can substantially reduce a firm's revenue. Second, it is important to make the right price commitment. In particular, a commitment to a static price is often better than no commitment but not effective relative to discount frequently. Hence, our results do not support a static-pricing strategy. This runs counter to the results in the literature that suggest a firm should, in some cases, try to commit to either not discount, or not discount deeply, or both (e.g., Besanko and Winston 1990, Aviv and Pazgal 2008, and others). In those papers the motivation for a static-pricing strategy is to mitigate the negative consequences of consumers strategically waiting for a price discount. As already mentioned, that motivation is not present in our model, so our results in no way contradict those findings. Instead, our results provide a counterargument for the adoption of static pricing. Specifically, we find that static

pricing is not the best strategy for attracting consumers to the firm when search is costly. Whether it is best to adopt an aggressive discounting strategy (as our model recommends) or to commit to not offer discounts probably depends on the importance of attracting consumers to search a firm relative to the desire to prevent consumers from strategically timing their purchases. In the example of JC Penney, it appears that JC Penney's decision to limit promotions prevented consumers from even visiting their stores, and the negative effect of this loss in demand was larger than the benefit of preventing consumers from waiting for sales.

Although discount frequently maximizes expected profit over all strategies assuming the firm adheres to its commitment, in the short-term period the firm may have an incentive (depending on the demand realization) to deviate from its commitment by charging the higher price rather than a discount. To be specific, based on (5), let $R_0(\gamma_f)$, be the firm's maximum expected revenue conditional that consumers choose γ_f as their visit strategy. Hence, deviating from the discount-frequently strategy can increase the firm's expected revenue by $R_0(\gamma_f) - R_f^* \geq 0$. If the firm and consumers indeed interacted over only a single period, then the discount-frequently commitment would not be credible. However, in the settings we consider (e.g., a department store), the firm is likely to interact with consumers over multiple periods, each of which is like our short-term period. In such case, the commitment to follow the discount-frequently strategy could be credible if consumers implement a trigger strategy that penalizes the firm for deviating. With a trigger strategy, consumers follow the discount-frequently equilibrium until they detect a deviation, at which point they switch to the no-commitment equilibrium in all subsequent periods. If the trigger is implemented, the firm's revenue loss in each period is $R_f^* - R_0^* \geq 0$. Thus, the discount-frequently strategy is credible if the profit from deviation, $R_0(\gamma_f) - R_f^*$, is less than the discounted future loss of revenue. To determine if this is likely, we evaluate the ratio of the short-term gain from deviation to the single-period loss from detection:

$$\frac{R_0(\gamma_f) - R_f^*}{R_f^* - R_0^*}.$$

On average, in our sample, this ratio is 0.51 and in 90% of the scenarios the ratio is 1.0 or lower—on average the gain in a single-period deviation from discount frequently is only 51% of the loss that is incurred in each period after the trigger is activated. Unless the firm heavily discounts future revenue, the short-term gain is unlikely to be justified by the loss of future revenue, suggesting that discount frequently can be credible.

The credibility of any pricing strategy also hinges on the ability of consumers to detect a deviation. This cannot be achieved with a single period, but it can be achieved with several periods. As the required length of the detection period increases, the benefit to the firm from a deviation increases—$R_0(\gamma_f) - R_f^*$ is earned over multiple periods. Nevertheless, given that the single-period benefit is generally less than even a single period of loss, credibility can be achieved as long as the firm cares enough about future revenue. See Radnor (1986) and Abreu et al. (1990) for detailed studies of repeated games with imperfect monitoring. (This issue is also relevant in papers, such as Liu and van Ryzin 2008, that assume a firm can credibly commit to a fill rate in future periods.)

6. Comparison of Strategies with Capacity Investment

In this section we determine the optimal capacity decisions under each of the pricing strategies and compare these capacity investments to a "myopic" benchmark that assumes all high types always search, i.e., $\gamma = 1$ for sure. The next theorem characterizes the optimal capacity levels under the different policies.

THEOREM 3. *Let c_k be the marginal cost of capacity. Then we have the following*:

1. *Myopic benchmark*: The optimal capacity, k_m, is given implicitly by $v_h \bar{F}(k) + v_l F(v_l k / v_h) = c_k$.

2. *No commitment*: If $c_k < c_k^0 = v_l + v_l (S(v_h \hat{k} / v_l) - S(\hat{k})) / \hat{k}$, then the optimal capacity is $k_0 = \max\{v_h \hat{k} / v_l, k_m\}$. Otherwise, $k_0 = 0$.

3. *Static pricing*: If $c_k < c_k^s = v_h \bar{F}(\tilde{k})$, then the optimal capacity, k_s, is given implicitly by $\bar{F}(k) = c_k / v_h$. Otherwise, $k_s = 0$.

4. *Discount frequently*: If $c_k < c_k^f = v_l + (v_h - v_l) \bar{F}(\hat{k})$, then the optimal capacity is $k_f = \max\{k'', k_m\}$, where k'' is given implicitly by $v_h \bar{F}(k'') + v_l F(\hat{k}) = c_k$. Otherwise, $k_f = 0$.

THEOREM 4. *Assume $c_k < c_k^f$. The following statements hold*:

1. $k_s < k_m \leq k_f$.

2. *Let $\underline{c}_k = v_l F(\hat{k}) + v_h \bar{F}(v_h \hat{k} / v_l)$. If $c_k \leq \underline{c}_k$, then $k_m = k_o = k_f$. If $\underline{c}_k < c_k \leq c_k^0$, then $k_f < k_0$.*

To compare the performance of the pricing strategies when capacity can be selected, we again considered the 360 parameter combinations detailed in Table 2 (excluding k). Since $c_k \in (v_l, c_k^f)$ is the interesting range for capacity costs (capacity levels greater than c_k^f result in zero profit for any policy), for each parameter combination we select five evenly distributed values from this range, i.e., $c_k^{(i)} = v_l + i(c_k^f - v_l)/6$, $i = 1, \ldots, 5$. This yields 1,800 scenarios. Table 3

Table 3   Optimal Capacity Levels Relative to the Myopic Consumers Optimal Capacity

	k_0/k_m	k_s/k_m	k_f/k_m
No. of observations with $k > 0$	662	1,254	1,800
Average	1.90	0.89	1.24
SD	0.77	0.17	0.30
Minimum	1.000	0.08	1.00
First quartile	1.22	0.84	1.03
Median	1.79	0.98	1.13
Third quartile	2.43	1.00	1.33
Maximum	4.25	1.00	2.94

Note. All statistics apply only in the scenarios in which it is profitable to invest in capacity, i.e., $k > 0$.

summarizes the ratios of the optimal capacity relative to the optimal capacity with myopic consumers.

Table 3 demonstrates that the capacity investment under no commitment can be large (e.g., on average $k_0/k_m = 1.90$). As in Dana and Petruzzi (2001), when the firm does not make a price commitment, then it can attract strategic consumers only by over investing in capacity. However, this can be inefficient, and in only a minority of scenarios (662 out of 1,800 scenarios) is the no-commitment strategy able to earn a positive profit. Static pricing performs better and of course discount frequently performs the best—because it is optimal, it yields a positive profit in all 1,800 scenarios. Relative to the myopic benchmark capacity, discount frequently invests in more capacity, 23.8% more on average. This "overinvestment" helps to attract consumers, because it provides additional availability and reduces the distortion needed in the frequency of discounts to attract the high-type consumers. However, discount frequently invests in far less capacity than the no-commitment strategy (assuming it is profitable to invest in capacity with no-commitment).

Figure 3 reports on the profits earned by each strategy relative to the discount-frequently strategy. We denote the optimal profit under each scheme by $\Pi_i^* = R_i - c_k k_i$ $i = \{f, s, 0\}$ and consider all 1,800 scenarios. Figure 3 confirms that when the firm has the flexibility to choose both its pricing strategy and its capacity, discount frequently outperforms the others, often by a large margin. In particular, overinvesting in capacity without a price commitment is not an effective strategy—it earns on average only 25% of what can be earned with discount frequently. Static-pricing also struggles—it earns on average only 48% of the optimal profit. Overall, the option to choose capacity accentuates the importance of making a price commitment and in particular, the correct price commitment.

As with a fixed capacity, we investigate the firm's short-term incentive to deviate from discount frequently. Let $\hat{\Pi}_f$ be the single-period profit the firm earns if it deviates from discount frequently even

Figure 3   Optimal Profit Ratios Considering All 360 Scenarios
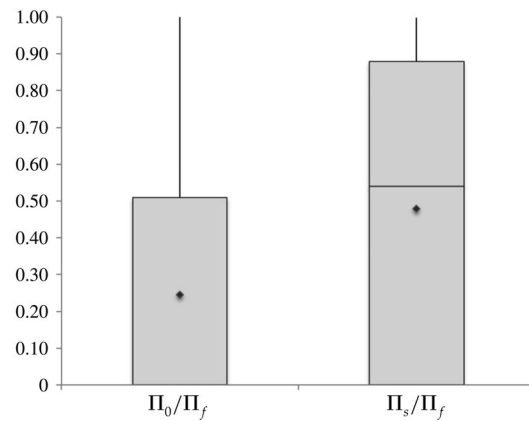
though consumers play the discount frequently equilibrium. On average the ratio of the short-term gain relative to the potential single-period loss of profit,

$$\frac{\hat{\Pi}_f - \Pi_f^*}{\Pi_f^* - \Pi_0^*},$$

is 0.76. This is higher than in the fixed capacity sample (0.51) but still suggests that the gain from deviating is potentially considerably lower than the loss in future discounted profit. Table 4 reveals that the consumer's search cost and the cost of capacity strongly influence this average. When search is not costly (first three rows), the short-term gain from deviating is small compared to the potential long-run loss in profit, and in some cases very small (upper left corner). Given the low search cost in these scenarios, one might assume that discount frequently does not provide much of an advantage relative to no commitment. But Table 5 reveals that even with low search costs, failing to make a price commitment reduces profit by a substantial amount (between 18%–100%). Table 4 also indicates that there are conditions in which credibility might be harder to achieve—when search is

Table 4   Average Ratio of Short-Term Profit Gain, $\hat{\Pi}_f - \Pi_f^*$, to Long-Run Per Period Profit Loss, $\Pi_f^* - \Pi_0$, as a Function of the Search Cost, c, and the Cost of Capacity, $c_k^{(i)}$

c	$c_k^{(i)} = v_l + i(c_k^f - v_l)/6$				
	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$
0.01	0.01	0.02	0.04	0.14	0.23
0.02	0.03	0.05	0.12	0.18	0.35
0.05	0.11	0.20	0.22	0.30	0.70
0.10	0.28	0.28	0.34	0.53	1.28
0.15	0.35	0.37	0.48	0.9	1.89
0.20	0.44	0.48	0.66	1.10	2.56
0.25	0.54	0.63	0.87	1.48	3.35
0.30	0.68	0.81	1.16	1.95	4.33

Table 5   Average Ratio of No-Commitment Profit, Π_0^*, to Discount-Frequently Profit, Π_f^*

	$c_k^{(i)} = v_l + i(c_k^f - v_l)/6$				
c	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$
0.01	0.82	0.73	0.61	0.43	0.14
0.02	0.77	0.65	0.51	0.30	0.06
0.05	0.67	0.51	0.32	0.13	0.00
0.10	0.54	0.35	0.07	0.04	0.00
0.15	0.44	0.24	0.10	0.01	0.00
0.20	0.35	0.17	0.05	0.00	0.00
0.25	0.28	0.11	0.02	0.00	0.00
0.30	0.22	0.07	0.00	0.00	0.00

costly and capacity is expensive (lower right corner), the gain from a deviation is a reasonable amount of the potential loss (up to 4.33 periods). In these cases, credibility requires the firm to be sufficiently farsighted (which remains possible, just harder). Nevertheless, in these cases the firm's bigger problem is that the product is barely profitable (and generally not profitable with the no-commitment strategy). Overall, even when the firm can choose its capacity, the firm loses a substantial amount of profit if it either chooses not to make a commitment or if it loses the credibility of its commitment.

7. Conclusion

Firms that do not make a price commitment can optimally select a price to respond to available demand information to maximize revenue. This is the best pricing strategy for the firm if consumers are not strategic. However, with strategic consumers, even though price commitments are costly to the firm in the short term, they are useful for attracting demand.

With fixed capacity, we show that a firm can do much better by committing to static pricing relative to no price commitment, despite the fact that the commitment reduces the firm's ability to match supply with demand. However, static pricing is not the best price commitment strategy. Discount frequently is the optimal strategy, and it performs substantially better than the other strategies we study. Thus, while it has been suggested that a static-pricing strategy can mitigate the negative effects of consumers strategically waiting for end-of-season discounts, we do not recommend that strategy when it is important to attract consumers to the firm (due to search costs).

When the firm can choose capacity, we show that a firm without a price commitment overinvests in capacity to attract consumers and earns substantially lower profit, if it can even earn a profit. As with fixed capacity, the discount-frequently strategy is optimal and much better than the other policies.

According to our model, adopting a "simplified pricing policy" or curtailing discounts, as done by

JC Penney in 2012 (Clifford and Rampell 2013), can backfire considerably. Thus, we conclude that in the presence of strategic consumers and search costs, (i) price commitments are generally necessary; (ii) the right commitment is discount frequently—the firm should give consumers a deep discount even if doing so lowers revenue conditional on demand because the higher frequency of discounts attracts consumers to the firm; and (iii) the firm should not exclusively use excess capacity as a tool to attract consumers.

Our results theoretically justify the implementation of a discount-frequently strategy. Additional work could seek to determine if firms currently utilize some form of discount frequently and if discount frequently is better in practice than alternatives. A firm could be said to be using discount frequently if they choose deeper discounts than would be optimal given existing inventory and current unbiased estimates of demand. Although not definitive, further evidence could be provided via surveys and interviews with managers—managers may intuitively be using a strategy like discount frequently if they say they are willing to provide discounts at the expense of short-term profits to avoid losing their long-term reputation for offering "good deals" or "competitive prices." Finally, JC Penney's experience provides some support for the effectiveness of discount frequently, but additional, controlled experiments would offer more direct evidence.

References

Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5):1041–1063.

Alexandrov A, Lariviere MA (2012) Are reservations recommended? *Manufacturing Service Oper. Management* 14(2): 218–230.

Allon G, Fedegruen A (2007) Competition in service industries. *Oper. Res.* 55(1):37–55.

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.

Barro RJ, Romer PM (1987) Ski-lift pricing, with applications to labor and other markets. *Amer. Econom. Rev.* 77(5):875–890.

Baye M, Morgan J (2001) Information gatekeepers on the Internet and the competitiveness of homogeneous product markets. *Amer. Econom. Rev.* 91(3):454–474.

Bernstein F, Fedegruen A (2004) Dynamic inventory and pricing models for competing retailers. *Naval Res. Logist.* 51(2):258–274.

Besanko D, Winston WL (1990) Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* 36(5):555–567.

Cachon GP, Feldman P (2011) Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manufacturing Service Oper. Management* 13(2):244–260.

Cachon GP, Harker P (2002) Competition and outsourcing with scale economies. *Management Sci.* 10(10):1314–1333.

Cachon GP, Lariviere MA (2001) Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* 47(5):629–646.

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.

Cachon GP, Gallino S, Olivares M (2013) Does adding inventory increase sales? Evidence of a scarcity effect in U.S. automobile dealerships. Working paper, Columbia Business School, New York.

Caro F, Gallien J (2012) Clearance pricing optimization for a fast-fashion retailer. *Oper. Res.* 60(6):1404–1422.

Çil EB, Lariviere MA (2013) Saving seats for strategic customers. *Oper. Res.* 61(6):1321–1332.

Clifford S, Rampell C (2013) Sometimes, we want prices to fool us. *New York Times* (April 13), http://nyti.ms/183VhFb.

Dana JD, Petruzzi NC (2001) Note: The newsvendor model with endogenous demand. *Management Sci.* 47(11):1488–1497.

DeHoratius N, Raman A (2008) Inventory record inaccuracy: An empirical analysis. *Management Sci.* 54(4):627–641.

Deneckere R, Peck J (1995) Competition over price and service rate when demand is stochastic: A strategic analysis. *RAND J. Econom.* 26(1):148–162.

Edelson N, Hildebrand K (1975) Congestion tolls for poisson queuing processes. *Econometrica* 43(1):81–92.

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations. *Management Sci.* 49(10):1287–1309.

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

Elmaghraby W, Lippman S, Tang C, Yin R (2009) Pre-announced pricing strategies with reservations. *Production Oper. Management* 18(4):381–401.

Feng Y, Gallego G (1995) Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Sci.* 41(8):1371–1391.

Fudenberg D, Levine D (1989) Reputation and equilibrium selection in games with a patient player. *Econometrica* 57(4):759–778.

Gallego G, van Ryzin GJ (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.

Gaur V, Park Y-H (2007) Asymmetric consumer learning and inventory competition. *Management Sci.* 53(2):227–240.

Hann I-H, Terwiesch C (2003) Measuring the frictional costs of online transactions: The case of a name-your-own-price channel. *Management Sci.* 49(11):1563–1579.

Harris M, Raviv A (1981) A theory of monopoly pricing schemes with demand uncertainty. *Amer. Econom. Rev.* 71(3):347–365.

Iyer AV, Bergen ME (1997) Quick response in manufacturer-retailer channels. *Management Sci.* 43(4):559–570.

Kuksov D, Villas-Boas MJ (2010) When more alternatives lead to less choice. *Marketing Sci.* 29(3):507–524.

Lariviere M, van Mieghem J (2004) Strategically seeking service: How competition can generate Poisson arrivals. *Manufacturing Service Oper. Management* 6(1):23–40.

Liu Q, Shum S (2013) Pricing and capacity rationing with customer disappointment aversion. *Production Oper. Management* 22(5):1269–86.

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

Masta DA (2011) Competition and product quality in the supermarket industry. *Quart. J. Econom.* 126(3):1539–1591.

Miller GA (1956) The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psych. Rev.* 63(2):81–97.

Ovchinnikov A, Milner JM (2012) Revenue management with end-of-period discounts in the presence of customer learning. *Production Oper. Management* 21(1):69–84.

Özer Ö, Zheng Y (2015) Markdown or everyday low price? The role of behavioral motives. *Management Sci.*, ePub ahead of print May 21, http://dx.doi.org/10.1287/mnsc.2014.2147.

Radnor R (1986) Repeated partnership games with imperfect monitoring and no discounting. *Rev. Econom. Stud.* 53(1):43–57.

Raman A, Fisher M (1996) Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* 44(4):87–99.

Roberts JH, Lattin JM (1997) Review of research and prospects for future insights. *J. Marketing Res.* 34(3):406–410.

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.

Swinney R (2011) Selling to strategic consumers when product value is uncertain: The value of matching supply and demand. *Management Sci.* 57(10):1737–1751.

Swinney R (2012) Inventory pooling with strategic consumers: Operational and behavioral benefits. Working paper, Duke University, Durham, NC.

Talluri KT, van Ryzin GJ (2004) *The Theory and Practice of Revenue Management* (Springer, New York).

Tereyagoglu N, Veeraraghavan S (2012) Selling to conspicuous consumers: Pricing, production, and sourcing decisions. *Management Sci.* 58(12):2168–2189.

Whang S (2015) Demand uncertainty and the Bayesian effect in markdown pricing with strategic customers. *Manufacturing Service Oper. Management* 17(1):66–77.

Yin R, Aviv Y, Pazgal A, Tang C (2009) Optimal markdown pricing: implications of inventory display formats in the presence of strategic consumers. *Management Sci.* 55(8):1391–1408.