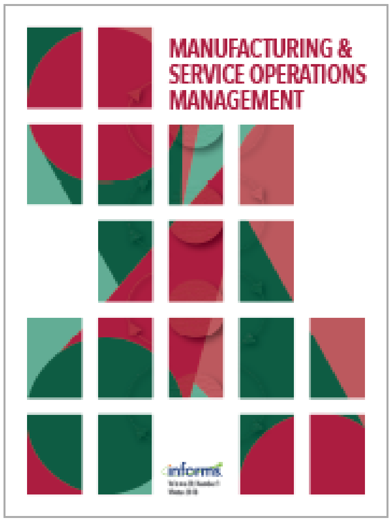

Manufacturing & Service Operations Management

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

Does Loss Aversion Preclude Price Variation?

Ningyuan Chen, Javad Nasiry

To cite this article:
Ningyuan Chen, Javad Nasiry (2019) Does Loss Aversion Preclude Price Variation?. Manufacturing & Service Operations Management

Published online in Articles in Advance 14 May 2019

. https://doi.org/10.1287/msom.2018.0743

informs
http://pubsonline.informs.org/journal/msom/
MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–13
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Does Loss Aversion Preclude Price Variation?

Ningyuan Chen,[a] Javad Nasiry[b]

[a] Department of Industrial Engineering and Decision Analytics, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong; [b] School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

Contact: nychen@ust.hk, http://orcid.org/0000-0002-3948-1011 (NC); nasiry@ust.hk, http://orcid.org/0000-0002-3768-8676 (JN)

Received: September 1, 2017
Revised: March 7, 2018; June 25, 2018
Accepted: June 26, 2018
Published Online in Articles in Advance: May 14, 2019

https://doi.org/10.1287/msom.2018.0743

Abstract. ***Problem definition***: In modern retailing, frequent discounts are seemingly at odds with the idea that price variation antagonizes loss-averse consumers and hence, diminishes their demand for products and services. Our research question then is whether (or not) loss aversion rules out price variation—and in particular, cyclic pricing. ***Academic/practical relevance***: Pricing and revenue management subject to behavioral considerations are key research areas in operations management. Our approach contributes to this research by highlighting the importance of incorporating heterogeneity in consumers' behavioral responses in pricing models. ***Methodology***: We model a monopolist selling a product over time to loss-averse consumers who differ in their sensitivity to gains/losses. Although the market is thus segmented, the firm cannot price discriminate among consumers based on that sensitivity. We then characterize the structural properties of the firm's optimal pricing policy. ***Results***: We show that charging a long-run constant price may be suboptimal and then derive conditions under which the optimal policy is cyclic. These findings establish that loss aversion does not preclude price variation and thereby, underscore the importance of incorporating consumer heterogeneity into pricing policies. ***Managerial implications***: For operations management scholars, our model highlights the importance of heterogeneity in consumers' behavioral responses to firms' policies and shows that structurally different insights are obtained from pricing models if this heterogeneity is appropriately accounted for. This approach offers new avenues in pricing and revenue management research. For managers, our model suggests that they could vary prices, under certain conditions, without worrying that price variation will antagonize consumers. Our model offers insights on what these conditions are, which managers may incorporate in devising the pricing policies.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2018.0743.

Keywords: behavioral pricing • loss aversion • cyclic pricing • markdown management • retailing

1. Introduction

Cyclic pricing, a practice followed in marketing and revenue management, is prevalent in certain industries. Promotional sales account for nearly one-half of all purchases in U.S. supermarkets (Nielsen 2009). Retailers often maintain a regular price and offer seasonal sales, holiday (post-Thanksgiving, "Cyber Monday," etc.) sales, and year-end sales. The year 2016 saw \$3.39 billion in Cyber Monday sales in the United States alone (Adobe 2016), which shows the effectiveness of such pricing policies.

Despite its prevalence and practical success, cyclic pricing seems to be inconsistent with findings in economic psychology (and especially, prospect theory) that price variation antagonizes consumers and reduces their demand for products and services. Prospect theory posits that individuals are loss averse and therefore, react more strongly when their payoffs fall short of (than exceed) a reference point. In a pricing context, loss aversion implies that consumers dislike prices that are higher than expected (perceived surcharges) more than they like prices that are lower than expected (perceived discounts). It follows that consumers are less likely to buy a product or service if its price exceeds expectations (i.e., the reference price). However, stimulating consumer demand by offering discounts has the negative long-term consequence of eroding the consumer's reference point, which makes restoring prices to their former levels costly.

We, therefore, seek to establish whether (or not) loss aversion rules out price variation—and in particular, cyclic pricing. Evidence is building in the operations management literature that loss aversion does not preclude price variation. A demand function subject to strong seasonality effect (Chen et al. 2016), varying consumer shopping schedules (Wang 2016), and demand aggregation (Hu and Nasiry 2017) may lead to price variation, even if all consumers are loss averse. We add to this evidence by suggesting an alternative mechanism. We develop a model of a monopolist selling to loss-averse consumers over time. The model has two key ingredients. The first is *consumer heterogeneity* with respect to perceptions of gain and loss, which we capture by assuming that the market consists of two segments that are each characterized by different sensitivities to reference effects. This type of heterogeneity has strong support in the empirical and experimental literature, and there have been calls to investigate its normative implications; a review is in Kopalle et al. (2012) and the references therein. The second is that the firm may use price to shut down the demand from a consumer segment. This is a key role of pricing in managing demand, especially when the firm cannot price discriminate between the segments.

Considerable evidence supports consumers' heterogeneity in gain/loss perceptions. Bell and Lattin (2000) use scanner panel data to measure loss aversion in nondurable grocery products and find that using a single-segment model tends to significantly overestimate loss aversion. Moon et al. (2006) estimate the reference effect in a consumer product category and find that the market consists of three segments: a segment without gain/loss sensitivity, a segment with a memory-based (i.e., internal) reference point, and one with a stimulus-based (i.e., external) reference point. They show that price sensitivities are significantly different across the segments. Dayaratna and Kannan (2012) provide an empirical estimation for their memory-based reference price reformulation using a three-segment latent class model, across which gain/loss sensitivities differ significantly. We further refer to Neumann and Böckenholt (2014) for a recent comprehensive review.

In our model, each segment consists of a continuum of consumers with uniformly distributed valuations (i.e., willingness to pay [WTP]). This implies that the "base demand" of each segment (that is, the demand in absence of behavioral considerations) is linear. We allow one segment (which we call "segment 1") to have (stochastically) higher WTP. We show that, if consumers in the segment with lower WTP (i.e., "segment 2") exhibit a strong reaction to perceived gains, then the firm's optimal pricing policy is a cyclic one. A common form of cyclic pricing is *markdown pricing*, whereby—once in each cycle—the firm targets segment 2 consumers by offering a discount. If the discount is set optimally, then its stimulation effect on segment 2 demand dominates its dilution effect on the profit margin of products sold to segment 1 consumers. The cycle length and price range of the optimal cyclic policy are volatile and nonmonotone in the model's behavioral parameters and especially, the strength of gain/loss perception. Finally, we show that ignoring consumer heterogeneity in gain/loss perception leads to constant pricing, which significantly underperforms the optimal pricing policy.

We use an exponentially weighted average of past prices as the reference point formation mechanism in our base model. This mechanism is recursive: the new reference point is a function of the previous reference point and the previous piece of information. Recursivity makes a dynamic model amenable to analytical analysis, and hence, exponential smoothing is widely applied in the dynamic pricing and revenue management literature (e.g., Popescu and Wu 2007). However, from an experiment, Baucells et al. (2011) observe that reference prices are not recursive. Our key results do not depend on the particular assumption of exponential smoothing or the recursivity of the reference point. Our results generalize, under mild conditions on the reference point formation mechanism, to contexts in which consumers may remember only a few price points or follow a peak-end rule.

Consumers in this model are loss averse, and our assumptions ensure that the overall demand (of each consumer segment) is also more sensitive to losses than to gains. In that sense, both the research question and our approach to studying it differ from the literature. A few papers have established the optimality of cyclic (or "hi-lo") pricing under reference effects (Greenleaf 1995, Kopalle et al. 1996, Popescu and Wu 2007, Chen et al. 2016, Hu et al. 2016, Wang 2016). However, these papers (except for those of Chen et al. 2016 and Wang 2016) assume that at least some consumers are gain seeking—in other words, they are actually more sensitive to gains (when the firm charges a price below the reference price) than to losses (when the firm charges a price higher than the reference price). In this case, the firm applies hi-lo pricing as follows: after increasing the reference point by charging a series of high prices, it then exploits that higher reference point by offering a discount that significantly boosts the demand from gain-seeking consumers.

Chen et al. (2016) show numerically that a strong seasonality effect on the demand function of loss-averse consumers causes a cyclic pricing pattern. Wang (2016) studies the dynamic pricing problem of a monopolist in a market of consumers who have *different shopping schedules*. In his model, consumers' reference points are based only on the prices that they encounter when visiting the store; therefore, consumers with different shopping schedules have different reference points. Wang (2016) proves that the firm's optimal policy for loss-neutral consumers is cyclic, and he shows numerically that the same insight holds also for loss-averse consumers. Our paper complements his work by accounting for consumer heterogeneity when pricing for reference-dependent consumers. However, the results reported here contradict his claim that, "if customers have homogenous arrival times, then even if they have different demand functions and memory factors, the optimal pricing policy must be an asymptotically constant one" (Wang 2016, p. 291). We confirm the optimality of cyclic pricing in the absence of heterogeneous arrival times by incorporating the heterogeneity in consumers' *sensitivity to reference effects*. We discuss Wang (2016) more thoroughly in Section 2.

Besbes and Lobel (2015) study the dynamic pricing problem of a monopolist in a market of strategic consumers who are heterogeneous with respect to both WTP and willingness to wait. These authors assume that the monopolist commits to a price sequence and then show that the optimal price policy is cyclic, with length no more than twice the maximum of the customer population's willingness to wait.

Hu and Nasiry (2017) argue that behavioral biases, such as loss aversion, are individual-level phenomena and therefore, need not be "inherited" by such aggregate variables as market demand. That is, it may be that aggregate demand is more sensitive to gains than to losses, although all consumers in the market are loss averse. These authors suggest that, instead of adding a gain/loss component to a conventional demand function, it would be preferable to derive individual demand subject to loss aversion and then aggregate it while accounting for the consumer population's valuation heterogeneity. Hu and Nasiry (2017) show that gain/loss preferences are inherited by the aggregate demand function if and only if consumer valuations are distributed uniformly. The consumers in our model are loss averse, and consumer valuations in each of the two segments are uniformly distributed. This implies that the overall demand in each market segment is more sensitive to losses than to gains. However, the two segments still have different sensitivities, and therefore, if the firm could price discriminate, then the price path would converge to a steady state (Popescu and Wu 2007, Hu and Nasiry 2017). This type of price discrimination is not plausible in the retail context, where each product carries only one price at any given time. We show that, when the firm cannot discriminate, its optimal pricing policy may not admit a steady state and could, in particular, be cyclic.

Periodic markdowns as a pricing and demand management method have been extensively studied in the operations management and economics literature. Özer and Zheng (2015) show that a firm may prefer markdown pricing to an "everyday low price" strategy when consumers are prone to regret and probability weighting. Yin et al. (2009), Mersereau and Zhang (2012), Ovchinnikov and Milner (2012), and Cachon and Feldman (2015), among others, study the optimality of markdown pricing in the presence of strategic consumers. Our own focus is not on markdown policies per se, and also, we do not incorporate strategic consumer behavior. Instead, we endeavor to prove that loss aversion—even at the aggregate demand level—does not rule out price variation if we account for the heterogeneity in consumer perceptions of gains and losses.

From a managerial perspective, our work suggests that the effect of consumer loss aversion on pricing policies is mediated by consumer heterogeneity in the strengths of gain/loss perceptions. In particular, attempts to fit a single gain/loss parameter to market demand data will result in adopting suboptimal pricing policies (e.g., uniform instead of cyclic policies) and profit loss.

2. Model

In this section, we introduce the general setup of our pricing model. A monopolist sells a product to a market of loss-averse consumers over an infinite horizon. We assume that the market consists of two segments $i \in \{1, 2\}$. There is a continuum of a_i consumers in segment i whose valuations are distributed uniformly in the interval $[0, a_i/b_i]$.[1] The parameter a_i/b_i is the maximum WTP in segment i. The firm cannot price discriminate between the two segments, and therefore, the consumers in each segment encounter the same price at time t. In deciding whether to buy at that time, consumers compare p_t (the price at time t) with a reference point r_t. A price higher (lower) than the reference point is perceived as a surcharge (discount).

We assume that the reference price is an exponentially weighted average of past prices:

Assumption 1. *The reference price satisfies* $r_t = \theta r_{t-1} + (1-\theta)p_{t-1}$ *for some* $\theta \in (0, 1)$.

The parameter θ is a memory parameter that captures how well consumers recall past prices. Memory-based reference prices are popular in both empirical and modeling research in the marketing and revenue management literature (e.g., Greenleaf 1995, Popescu and Wu 2007, Kopalle et al. 2012; a review is in Mazumdar et al. 2005). In Section 3, we will show that our key results generalize to a broader set of reference point formation mechanisms, of which exponential smoothing is a special case.

Segment i's demand function at time t is

$$D_i(p_t, r_t) = (a_i - b_i p_t + \gamma_i (r_t - p_t)^+ - \lambda_i (p_t - r_t)^+)^+, \quad (1)$$

where γ_i and λ_i are the demand function's sensitivities to discounts and surcharges, respectively.

We choose uniform consumer valuation distributions for two reasons. First, this assumption implies that the "base" demand function of segment i, $D_i(p_t) = a_i - b_i p_t$, is linear (e.g., chapter 3 of Phillips 2005). Linear demand functions are commonly applied and perform sufficiently well in modelling demand in dynamic pricing settings (Besbes and Zeevi 2015).

Second, because consumer valuations in segment i are distributed uniformly, their biases toward gains and losses are inherited by the segment's aggregate demand (Hu and Nasiry 2017). In other words, if all individual consumers in segment i are loss averse or loss neutral, then the overall demand of segment i is more sensitive to losses ($\lambda_i > \gamma_i$) or equally sensitive to losses and gains ($\lambda_i = \gamma_i$), respectively. This property is important, because our goal is to investigate whether loss aversion at the aggregate demand level precludes price variation. If loss aversion is present at the individual level but absent at the aggregate level, which Hu and Nasiry (2017) show is possible, then the optimal price policy may not admit a steady state (Hu and Nasiry 2017).

If there is only one segment in the market, then in the long run, the firm's optimal policy converges to a steady state. The reason is that the overall demand function is more sensitive to losses than to gains (Fibich et al. 2003, Popescu and Wu 2007, Hu and Nasiry 2017). However, we show that this generalization no longer holds when there are two market segments with heterogeneous sensitivities to losses and gains. We shall, therefore, assume without loss of generality that $a_2/b_2 \le a_1/b_1$.[2] Obviously, at prices above a_1/b_1, the demand from both segments is zero unless customers perceive substantial gains. Because that is unlikely to happen,[3] we assume that p_t (and hence, r_t) is no larger than a_1/b_1. This assumption, however, allows the firm to shut down the demand from segment 2 by charging prices above the "choke price" of segment 2 (i.e., $p_t \ge a_2/b_2$). Thus, our setup differs from that of Wang (2016), who assumes the firm can charge only those prices at which the demand from every segment is positive. That assumption rules out a key function of prices in managing demand, and it is the main reason why cyclic pricing cannot be optimal in his setup when arrival times are homogeneous. Popescu and Wu (2007) make the same assumption and prove that, in the long run, it is optimal for the firm to charge a uniform price if there are multiple market segments with loss-averse consumers. Cyclic pricing will emerge in their model of heterogeneous consumers if the monopolist can price out one or more segments.

In this paper, we mainly focus on the case in which $\gamma_1 = \lambda_1 = 0$ and $\gamma_2 \le \lambda_2$. That is, demand from segment 1 is not subject to reference effects, whereas demand from segment 2 is either more sensitive to losses or equally sensitive to losses and gains. Assumption 2 facilitates the analytics in our paper. In Section 4.3, our results generalize to a setting where both segments are subject to reference effects. In summary, we make the following assumption in this section.

Assumption 2. *We have* $a_2/b_2 \le a_1/b_1$, $\gamma_1 = \lambda_1 = 0$, $\gamma \triangleq \gamma_2 \le \lambda \triangleq \lambda_2$, *and* $p_t, r_t \in [0, a_1/b_1]$.

Moon et al. (2006) provide empirical support for this assumption. They estimate consumers' responses to price changes using a data set on toilet tissue purchase records. They find that some consumers are not prone to reference effects (i.e., $\lambda_1 = \gamma_1 = 0$), whereas others are loss averse (i.e., $\gamma_2 < \lambda_2$). Moreover, loss-averse consumers constitute a larger proportion of the market ($a_2 > a_1$) and are more price sensitive ($b_2 > b_1$).

By Assumption 2, the demand of segment 1 is $D_1(p_t) = a_1 - b_1 p_t$, and that of segment 2 is $D_2(p_t, r_t) = (a_2 - b_2 p_t + \gamma(r_t - p_t)^+ - \lambda(p_t - r_t)^+)^+$. Let $\Pi_1(p_t) = p_t D_1(p_t)$ and $\Pi_2(p_t, r_t) = p_t D_2(p_t, r_t)$ denote the single-period revenue functions corresponding to segments 1 and 2, respectively. The aggregate demand from both segments is $D(p_t, r_t) \triangleq D_1(p_t) + D_2(p_t, r_t)$, and the corresponding single-period revenue is $\Pi(p_t, r_t) \triangleq p_t D(p_t, r_t)$.

2.1. Myopic Pricing Policy

We first study the firm's myopic pricing policy, under which the firm focuses on maximizing current period revenue and ignores this strategy's possible consequences on the firm's long-term revenues. That is, given the reference price r in the current period, the firm uses the price $p_m(r) = \operatorname{argmax}_p\{\Pi(p, r)\}$. As we shall see, the myopic pricing policy has an analytical solution and thus, provides useful insights into the benefit of cyclic pricing.

To give an intuition for why myopic firms may prefer cyclic pricing, consider a *periodic markdown* policy: in a pricing cycle of n periods, the firm targets only segment 1 for $n-1$ periods and then, offers a discount to "skim" segment 2 once in a cycle. More precisely, from period 1 to $n-1$, the firm charges a regular price $\frac{a_1}{2b_1}$, which is maximizing the revenue $\Pi_1(p)$ from segment 1 and ignores segment 2. In period n, the firm charges a price (which is characterized in Proposition B.1 in the online appendix) that is less than $\frac{a_1}{2b_1}$ to appeal to segment 2 customers and meanwhile, optimize its revenue from both segments $\Pi_1(p) + \Pi_2(p, r_n)$. Then, from period $n+1$ to $2n-1$, the same regular price is used and so on. Seasonal sales are an example of periodic markdown policies commonly used in practice.

To understand why periodic markdown can be optimal for a myopic firm, we track how the reference price of segment 2 customers changes in a pricing cycle. In period 1, the reference price is relatively low. To make a profit from segment 2, the firm has to set a low price and forgo a substantial amount of revenue from segment 1. Therefore, the optimal myopic pricing policy is to set $p_m(r_1) = \operatorname{argmax}_p \Pi_1(p) = \frac{a_1}{2b_1}$, which prices out segment 2. Because this price is higher than the reference price r_1, by Assumption 1, the reference price in the next period r_2 increases slightly because of customers' memory in period 1. However, r_2 is still not large enough, and the firm chooses to price out segment 2 in period 2 as well. The same happens until period n, in which customers in segment 2 have a sufficiently large reference price r_n based on their memory from periods 1 to $n-1$. It is now optimal for the firm to set a low price so that segment 2 customers perceive gains, and the boosted demand from segment 2 outweighs the lost revenue from segment 1 as the price deviates from $\frac{a_1}{2b_1}$. In period $n+1$, however, the reference price drops because of the discount in period n, and the pricing cycle starts over again. Proposition B.1 in the online appendix provides the condition under which a cycle n periodic markdown is optimal for a myopic firm.

Having explained the intuition, we study the analytical solution for the myopic pricing policy (i.e., maximizing the revenue function $\Pi(p,r)$ for given r). In general, $\Pi(p,r)$ is piecewise, quadratic, nonsmooth, and nonconcave, making the optimization quite challenging. It has one of the following forms:

$$\Pi(p,r) = \Pi_1(p) = p(a_1 - b_1 p) \quad \text{Segment 2 is priced out.}$$

$$\Pi(p,r) = p((a_1+a_2) - (b_1+b_2)p + \gamma(r-p)) \quad \text{Segment 2 perceives gains.}$$

$$\Pi(p,r) = p((a_1+a_2) - (b_1+b_2)p + \lambda(r-p)) \quad \text{Segment 2 perceives losses.}$$

Therefore, we have to identify $p_m(r)$ from potential local maxima $\{\frac{a_1}{2b_1}, p^\gamma(r), p^\lambda(r), r\}$, where we define $p^\gamma(r) \triangleq \frac{a_1+a_2+\gamma r}{2(b_1+b_2+\gamma)}$, $p^\lambda(r) \triangleq \frac{a_1+a_2+\lambda r}{2(b_1+b_2+\lambda)}$. Although the first three local maxima in the set correspond to the three quadratic functions above, the last local maximum is produced by the nonsmoothness at $p = r$. Next, we compare the local maxima and derive conditions under which one of them becomes the global maximum. This results in Proposition 1.

Define

$$r^\gamma \triangleq \frac{a_1+a_2}{\gamma + 2(b_1+b_2)},$$
$$r^\lambda \triangleq \frac{a_1+a_2}{\lambda + 2(b_1+b_2)},$$
$$\bar{r}^\lambda \triangleq \frac{a_1}{\lambda}\sqrt{\frac{b_1+b_2+\lambda}{b_1}} - \frac{a_1+a_2}{\lambda},$$
$$\bar{r}^\gamma \triangleq \frac{a_1}{\gamma}\sqrt{\frac{b_1+b_2+\gamma}{b_1}} - \frac{a_1+a_2}{\gamma}, \text{ and}$$
$$\bar{r}^0 \triangleq \frac{(a_1+a_2) - \sqrt{(a_1+a_2)^2 - (b_1+b_2)a_1^2/b_1}}{2(b_1+b_2)}.$$

Proposition 1. *One and only one of the following cases emerges.*

1. *If* $\bar{r}^0 \le r^\lambda$, *then* (i) $p_m(r) = a_1/2b_1$ *for* $r \in [0, \bar{r}^\lambda]$, (ii) $p_m(r) = p^\lambda(r)$ *for* $r \in (\bar{r}^\lambda, r^\lambda]$, (iii) $p_m(r) = r$ *for* $r \in (r^\lambda, r^\gamma]$, *and* (iv) $p_m(r) = p^\gamma(r)$ *for* $r \in (r^\gamma, a_1/b_1]$.
2. *If* $\bar{r}^0 \in (r^\lambda, r^\gamma]$, *then* (i) $p_m(r) = a_1/2b_1$ *for* $r \in [0, \bar{r}^0]$, (ii) $p_m(r) = r$ *for* $r \in (\bar{r}^0, r^\gamma]$, *and* (iii) $p_m(r) = p^\gamma(r)$ *for* $r \in (r^\gamma, a_1/b_1]$.
3. *If* $\bar{r}^0 > r^\gamma$, *then* (i) $p_m(r) = a_1/2b_1$ *for* $r \in [0, \bar{r}^\gamma]$ *and* (ii) $p_m(r) = p^\gamma(r)$ *for* $r \in (\bar{r}^\gamma, a_1/b_1]$.

In each case, the myopic policy $p_m(r)$ is piecewise linear and discontinuous. Take case 1 for an example. When the reference price is low, the firm simply maximizes the revenue from segment 1 and prices out segment 2 (case 1, i). As the reference price increases, segment 2 becomes more profitable, and the firm targets both segments. It may cause segment 2 to perceive losses (case 1, ii), gains (case 1, iv), or neither (case 1, iii) under the myopic policy. Cases 2 and 3 are similar to case 1, whereas one or more subcases become nonexistent. Proposition 1 is illustrated in Figure 1.

2.1.1. Steady States and Cyclic Pricing.

This section focuses on the behavior of the myopic pricing policy and investigates the conditions under which it is cyclic—in other words, when it does not admit a steady state. Given a reference price r_t at time t, the myopic firm charges $p_m(r_t)$ (as specified in Proposition 1), and therefore, the reference price at time $t+1$ becomes $r_{t+1} = \theta r_t + (1-\theta)p_m(r_t)$. A *myopic steady-state price* is a price from which the myopic firm has no incentive to deviate; it is a price $\tilde{r}$ that solves $p_m(\tilde{r}) = \tilde{r}$. If $r_t = \tilde{r}$ at time t, then $p_m(r_t) = r_t$, and $r_{t+1} = \theta r_t + (1-\theta)p_m(r_t) = r_t$, which imply that r_t and the optimal myopic price $p_m(r_t)$ remain $\tilde{r}$ over time.

We can use Figure 1 to illustrate the existence of steady states. The steady state corresponds to the intersection of an identity map (dotted lines in Figure 1) and the function $p_m(\cdot)$. Therefore, cases 1 and 2 in Proposition 1 (which correspond to panels (a) and (b) of Figure 1) admit a range of steady states.

In case 3 of Proposition 1, there are two possibilities. If $a_1/2b_1 > \bar{r}^\gamma$, then there is no steady state (panel (c) of Figure 1). However, if $a_1/2b_1 \le \bar{r}^\gamma$, then there is a single steady state: $a_1/2b_1$ (panel (d) of Figure 1). The following proposition formally characterizes the myopic steady states.

Proposition 2.

1. *If* $\bar{r}^0 \le r^\lambda$, *then any* $r \in [r^\lambda, r^\gamma]$ *is a steady state.*
2. *If* $\bar{r}^0 \in (r^\lambda, r^\gamma]$, *then any* $r \in [\bar{r}^0, r^\gamma]$ *is a steady state.*
3. *If* $\bar{r}^0 > r^\gamma$ *and* $\bar{r}^\gamma \ge a_1/2b_1$, *then there is one steady state:* $a_1/2b_1$.
4. *If* $\bar{r}^0 > r^\gamma$ *and* $\bar{r}^\gamma \in (r^\gamma, a_1/2b_1)$, *then there is no steady state.*

Two observations follow from this proposition. First, the steady states from cases 1 and 2 are essentially different from the steady state in case 3. In cases 1 and 2, the firm sells to both segments in the steady state, whereas segment 2 is priced out in case 3. Second, the emergence of a cyclic pricing policy (case 4) depends on γ but not on λ, because $\bar{r}^0$, r^γ, and $\bar{r}^\gamma$ are not functions of λ. As $\gamma \to \infty$, we have $r^\gamma = O(\gamma^{-1})$ and $\bar{r}^\gamma = O(\gamma^{-1/2})$. These imply that, when the sensitivity to perceived gains increases, the condition in case 4 is more likely to be met; in that event, hi-lo pricing becomes preferable (from the firm's perspective) to uniform pricing. Our next result

Figure 1. (Color online) Myopic Pricing Policy $p_m(r)$ for the Three Cases Listed in Proposition 1

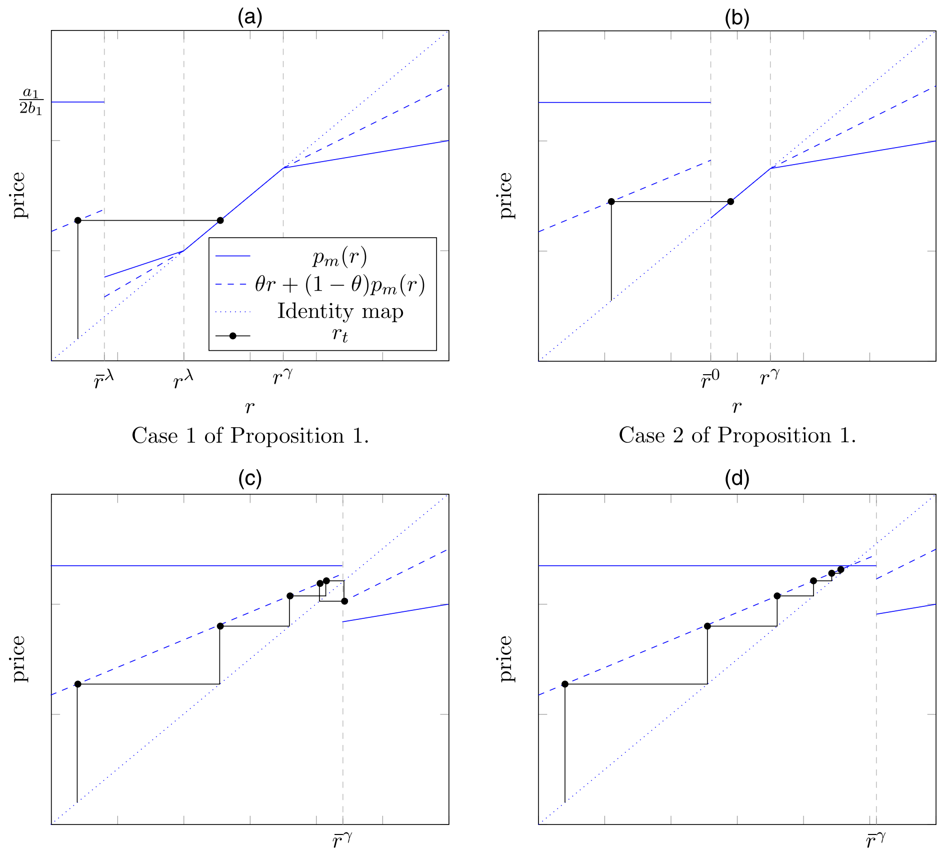

Case 1 of Proposition 1. Case 2 of Proposition 1.

Case 3 of Proposition 1: no steady state. Case 3 of Proposition 1: steady state.

shows that, if a steady state exists, then in the long run, the paths of the optimal myopic price charged by the firm and the reference price of customers converge to that steady state—although the convergence is not monotonic.

Proposition 3.

1. *For all* $r_0 \in [0, a_1/b_1]$, *if the myopic pricing policy admits steady states* (*cases* 1–3 *in Proposition* 2), *then* r_t *and* $p_m(r_t)$ *converge to the steady state.*

2. *If there is no steady state* (*case* 4 *in Proposition* 2), *then* r_t *and* $p_m(r_t)$ *are cyclic in the long run.*

Figure 1 plots the reference price path. If the reference price at time t is r_t, then at time $t+1$, the reference price is $r_{t+1} = \theta r_t + (1-\theta)p_m(r_t)$. In all panels of Figure 1, the horizontal line from $(r_t, \theta r_t + (1-\theta)p_m(r_t))$ intersects the identity map (the 45° line) at (r_{t+1}, r_{t+1}). The bold dots in Figure 1 indicate how r_t evolves over time. The reference price clearly converges to a steady state in panels (a), (b), and (d) of Figure 1, whereas in panel (c) of Figure 1, the path cycles.

However, the cycle length is not monotone in any of the parameters. Instead, it displays the so-called *bifurcation phenomenon* of dynamical systems theory (Rajpathak et al. 2012 and the references therein): between two regions of the parameter space corresponding to cycle lengths n_1 and n_2, there exists a region with cycle length $n_1 + n_2$. This behavior implies that the cycle length is extremely discontinuous in the parameters; the phenomenon is consistent with the fact that widely different promotion patterns are observed for products that are but slightly differentiated.

It is worth mentioning that, unlike in Popescu and Wu (2007) and Nasiry and Popescu (2011), convergence to the steady state is not order preserving; that is, a larger r_0 does not guarantee a larger steady state. For example, in panel (a) of Figure 1, if $r_0 = r^{\lambda}$, then r_0 is already a steady state; if r_0 is set as illustrated by the dots in panel (a) of Figure 1 and thus, less than r^{λ}, then r_0 reaches a steady state that is greater than r^{λ}. This is because the evolution of the reference price is discontinuous and nonmonotone in r.

2.2. Optimal Dynamic Pricing

In this section, we study the firm's dynamic pricing problem. This problem differs from the myopic pricing problem, in which the firm considers only current period profits. Given the initial reference price r_0, the firm maximizes its infinite horizon problem as $V(r_0) = \max_{p_t} \sum_{t=0}^{\infty} \beta^t \Pi(p_t, r_t)$, where $\beta < 1$ is the discount factor. Because $0 \le p_t \le a_1/b_1$, it follows that per-period revenue is bounded; hence, the value function is the unique solution to the Bellman equation:

$$V(r) = \max_{p\in[0,a_1/b_1]} \{\Pi(p,r) + \beta V(\theta r + (1-\theta)p)\}. \quad (2)$$

Let $p^*(r)$ denote the optimal pricing policy. In general, the Bellman equation (2) does not admit a closed form solution. To obtain structural results, we first derive a performance bound—that is, an upper bound for the revenue gap between the myopic and the optimal pricing policies; then, we can show that the optimal dynamic policy does not admit steady states when the *gains*-related reference effect is strong.

2.2.1. A Performance Bound. We use $V_m(r)$ to denote the discounted revenues corresponding to the myopic pricing policy when the current reference price is r. By definition, $V_m(r) = \Pi(p_m(r), r) + \beta V_m(\theta r + (1-\theta)p_m(r))$. We can obtain the following result.

Proposition 4.

$$V(r) - V_m(r) \le \frac{\beta(1-\theta)\lambda a_1^2}{(1-\beta)(1-\beta\theta)b_1^2}.$$

Two key observations can be made with regard to this theorem. First, the bound approaches zero when $\beta \to 0$ or $\theta \to 1$. That is to say, the less weight that is given to future revenues (i.e., the more myopic that the firm is), the tighter the bound. The myopic case is equivalent to $\beta = 0$ in (2). When $\theta = 1$, consumers never update their reference prices. Thus, the dynamic problem is equivalent to a sequence of identical single-period pricing problems, and therefore, the myopic pricing policy is optimal.

Second, the performance bound increases with the loss-aversion coefficient λ. This outcome follows, because a higher λ increases the weight of future revenues in the tradeoff between them and current revenues. The myopic pricing policy completely ignores the effect of current pricing on future revenues; as a consequence, it is significantly outperformed by optimal dynamic pricing when γ or λ is large.

2.2.2. Cyclic Pricing and Reference Effects. We now establish that, if the reference effect for perceived gains γ is sufficiently large, then a long-run constant price (steady state) is never optimal in the dynamic case. This result is at odds with theorems 2 and 4 in Popescu and Wu (2007), which state that there always exists a steady state for loss-neutral or loss-averse customers. That is, no matter how strong reference effects may be, the firm should avoid hi-lo pricing and instead, charge a constant long-run price if consumers are loss neutral or loss averse. Their results differ, because their model does not allow a market segment to be priced out. Therefore, although the perceived gain may be substantial when using hi-lo pricing, the perceived loss is also large. In the model of Popescu and Wu (2007), then the negative effects of these losses dominate the positive effects of gains, and hence, adopting a constant price is always optimal in the long run.

In contrast, our model considers the more realistic case, in which gain/loss preferences are heterogeneous. The loss effect is capped, because the worst case scenario is to have zero demand from segment 2. The firm can use a regular high price that shuts down the demand from segment 2 and focus only on segment 1, but they will occasionally reap some segment 2 revenue when the reference price of those consumers is sufficiently high. Such a strategy is practical and frequently observed when the market is segmented but the firm cannot price discriminate. This insight is formalized in our next result.

Proposition 5. *When γ is sufficiently large, the optimal dynamic pricing policy that solves* (2) *does not have any steady states.*

By Assumption 2, we must have $\lambda \ge \gamma$, and therefore, λ is also large for large γ. Nevertheless, it is the magnitude of γ that leads to cyclic pricing, regardless of λ. This may seem counterintuitive given that, when the firm charges a high price, loss aversion ($\lambda > \gamma$) restricts the revenue earned from segment 2; hence, a sufficiently large λ should discourage the firm from hi-lo pricing and lead to uniform pricing. However, this is not the case. There is indeed very little demand from segment 2 when the firm uses hi-lo pricing and charges a high price. The revenue in this phase is mainly generated from segment 1, with demand that is not subject to reference effects. In other words, no matter how large λ is, its negative effect on demand in the high-price phase is limited, because the demand from segment 2 cannot be less than zero. However, the profitability of hi-lo pricing (compared with uniform pricing) is reflected in the price-cut phase—that is, when the price changes from "hi" to "lo." The more sensitive the demand of segment 2 to perceived gains, the more profitable hi-lo pricing will be. It follows that the emergence of a steady state depends mainly on γ. This is consistent with the discussion after Proposition 2.

3. A General Reference Point Formation Mechanism

So far, we have assumed that consumers' price expectation in each time period is an exponentially weighted average of past prices. This framework lends itself to analytical analysis and is widely applied in pricing literature. Under this assumption, the reference point formation mechanism is recursive (that is, the new reference point is a function of the previous reference point and the previous piece of information (Baucells et al. 2011): $r_t = f(r_{t-1}, p_{t-1})$). Exponential smoothing also implies implicitly that consumers recall all past prices paid for a product.

In this section, we introduce a general reference price formation mechanism, of which exponential smoothing is a special case. Our goal is to show that, under mild conditions on this general mechanism, our key insights on a monopolist's pricing policy (i.e., Proposition 5) continue to hold. That is, cyclic pricing is optimal when the heterogeneity in the reference effects of the two segments is sufficiently large.

Given a sequence of prices, p_t and the initial reference price r_0, consider an array of weights $w_{t,i} \in [0,1]$, where $t = 0, 1, \ldots$ and $i = -1, 0, \ldots, t-1$, and $\sum_{i=-1}^{t-1} w_{t,i} = 1$. The reference price in period t then is

$$r_t = w_{t,-1} r_0 + \sum_{i=0}^{t-1} w_{t,i} p_i. \tag{3}$$

In other words, the reference price in period t is a weighted average of all past prices and the initial reference price. Because we do not specify $w_{t,i}$, it is more general than most reference formation models in the literature.

As we shall see, the reference point in (3) encompasses exponential weighting, the model proposed in equation 7 of Baucells et al. (2011), and extrapolative expectations, where the reference price depends on most recent (e.g., the last two) prices (e.g., Jacobson and Obermiller 1990). Moreover, Equation (3) is not necessarily recursive.

To proceed, we define two properties for a reference point formation mechanism: *retentiveness* and *asymptotic obliviousness*.

Definition 1. The reference point in (3) is (K, δ) *retentive* if there exist $K \in \mathbb{Z}_+$ and $\delta > 0$ such that, for all $t \geq 1$, we have $\sum_{i=(t-K)^+}^{t-1} w_{t,i} \geq \delta$.

If a reference formation process is (K, δ) retentive, then the customer always puts a positive weight on the most recent K prices, regardless of t. Retentiveness is not restrictive, because the choice of K and δ is free. Customers with a strong memory may have a large K and a small δ. For example, a customer may have a $1/3$ weight on the prices in the last 10 periods and the remaining $2/3$ on the historical prices more than 10 periods ago. For forgetful customers, we may have a small K. For example, a customer forming her reference price based on the average of the prices in the last three periods is $(3, 1)$ retentive.

Definition 2. The reference point in (3) is *asymptotically oblivious* if, for all $k \geq 1$, we have $\lim_{t\to\infty} \sum_{i=-1}^{k\wedge(t-1)} w_{t,i} = 0$.

Asymptotic obliviousness implies that customers' memory of the prices in the first few periods fades away in the long run. If a reference formation process is not asymptotically oblivious, then the customer always assigns some weight to the prices in the first few periods. As a result, even if the firm charges a constant price p in the long run, the reference price may not converge to p as $t \to \infty$. This is unlikely to happen in practice.

Example 1 (Exponentially Smoothed Adaptive Expectations Process). The reference price model that we use in Section 2 has a recursive form: $r_t = \theta r_{t-1} + (1-\theta) p_{t-1}$. This is a special case of the formation mechanism in (3), because we can write $r_t = (1-\theta) \sum_{i=0}^{t-1} \theta^{t-1-i} p_i + \theta^t r_0$. It is straightforward to observe that this reference point formation mechanism is $(1, 1-\theta)$ retentive and asymptotically oblivious.

Example 2 (Extrapolative Expectations). In this framework, customers' reference point is a weighted average of k most recent prices: $r_t = \sum_{i=1}^{k} a_i p_{t-i}$ and $\sum_{i=1}^{k} a_i = 1$. For instance, the model in Jacobson and Obermiller (1990) is a special case with $k = 2$. Clearly, this reference price formation mechanism is $(k, 1)$ retentive and asymptotically oblivious. However, it is not recursive.

Next, we analyze the optimal dynamic policy for the firm given the reference point formation in (3). The firm's optimal pricing policy $\{p_t\}_{t=0}^{+\infty}$ solves the following optimization problem:

$$\begin{aligned} \max \quad & \sum_{t=0}^{+\infty} \beta^t \Pi(p_t, r_t) \\ \text{subject to} \quad & r_t = w_{t,-1} r_0 + \sum_{i=0}^{t-1} w_{t,i} p_i, \\ & p_t \in [0, a_1/b_1],\ t = 0, 1, \ldots \end{aligned}$$

Because the reference formation mechanism in (3) does not necessarily have a recursive form, the optimization problem cannot be transformed into a dynamic program and is thus analytically intractable. Nevertheless, we are able to characterize its cyclic behavior in Theorem 1 using a perturbation approach. The result applies to Examples 1 and 2.

Theorem 1. *If the reference formation mechanism in* (3) *is* (K, δ) *retentive and asymptotically oblivious, then for a sufficiently large* γ*, the optimal pricing policy* p_t^* *for* $t = 1, 2, \ldots$ *does not admit a steady state. That is, there exists* $\epsilon > 0$ *such that, for all* $T \geq 1$, $\max_{t\geq T} p_t^* - \min_{t \geq T} p_t^* > \epsilon$.

Retentiveness and asymptotic obliviousness are not restrictive conditions, and the mechanism in (3) encompasses widely applied reference points. Thus, Theorem 1 establishes the general validity of cyclic pricing when the low-valuation segment has sufficiently strong reference effects.

We finally remark that some reference point structures are not covered by (3). Nonetheless, our key insights may generalize to such structures as well. For example, the "peak-end rule" is a recursive reference point formation mechanism that assumes that consumers anchor on the most recent and the minimum prices paid for a product (Nasiry and Popescu 2011) or $r_t = \theta m_{t-1} + (1-\theta)p_{t-1}$, where $m_{t-1} = \min(m_{t-2}, p_{t-1})$. We show in the online appendix that, with the peak-end rule and for sufficiently large γ, the firm's optimal policy does not admit a steady state.

4. Numerical Studies

In this section, we conduct several numerical studies. We first compare the optimal dynamic pricing policy (which may be cyclic) with a uniform price policy, in which the firm charges an optimally set constant price over time. A uniform price policy could arise if the firm ignores heterogeneity and decides to charge a constant price to loss-averse consumers. We show that this policy can lead to a substantial revenue loss (as high as 46%) compared with cyclic pricing.

In the second study, we use parameters estimated by Dahana and Terui (2006) and Moon et al. (2006) to compute the optimal pricing policy. Our goal in this study is to show that cyclic pricing is optimal for empirical estimates of the key parameters in our model. In the third study, we relax Assumption 2 so that both segments are loss averse. This example further extends the insight derived in Theorem 1: cyclic pricing is optimal if the loss effect in the high-valuation segment (segment 1) is much lower than the gain effect in the low-valuation segment (segment 2; i.e., $\gamma_2 \gg \lambda_1$, whereas λ_2 and γ_1 are not critical for the emergence of cyclic pricing).

4.1. What Is the Revenue Loss If the Firm Ignores Heterogeneity?

We built the numerical example under Assumptions 1 and 2. In particular, we set $a_1 = 1, b_1 = 0.3, a_2 = 4, b_2 = 2$, $\beta = 0.99$, and $r_0 = 0$. We let $\lambda = \gamma$ (i.e., consumers in segment 2 are gain/loss neutral) and vary γ. This is for simplicity, because by Proposition 2 and the discussion following Theorem 1, λ does not significantly affect cyclic pricing behavior. To solve the Bellman equation numerically, we discretize the state r and use value iteration and linear interpolation to compute the optimal value function $V(r)$ and the optimal pricing policy $p^*(r)$. Table 1 shows the structure of the optimal pricing policy for various combinations of γ and θ. Table 2 compares the optimal pricing policy with the optimal uniform pricing policy. A cycle length of one implies that $p^*(r_t)$ and r_t converge to a steady state.

When combined with the previous analysis, the numerical results allow us to make three conclusions. First,

Table 1. The Cycle Length, Price Range, and Total Revenue of the Optimal Dynamic Pricing Policy

γ	$\theta = 0.1$	$\theta = 0.3$	$\theta = 0.5$	$\theta = 0.7$	$\theta = 0.9$
Cycle length					
0.2	1	1	1	1	1
0.4	1	1	1	1	1
0.6	4	5	7	4	1
0.8	3	4	4	6	3
1.0	3	3	4	5	3
1.2	3	3	3	4	3
1.4	4	3	3	3	3
1.6	4	3	3	3	3
1.8	2	4	3	3	3
Price range					
0.2	[1.67, 1.67]	[1.65, 1.65]	[1.65, 1.65]	[1.65, 1.65]	[1.64, 1.64]
0.4	[1.67, 1.67]	[1.67, 1.67]	[1.64, 1.64]	[1.64, 1.64]	[1.62, 1.62]
0.6	[1.31, 3.33]	[1.48, 3.33]	[1.49, 3.33]	[1.54, 1.71]	[1.59, 1.59]
0.8	[1.37, 3.33]	[1.38, 3.33]	[1.54, 3.33]	[1.55, 3.33]	[1.59, 3.33]
1.0	[1.34, 3.33]	[1.46, 3.33]	[1.45, 3.33]	[1.52, 3.33]	[1.57, 3.33]
1.2	[1.32, 3.33]	[1.45, 3.33]	[1.53, 3.33]	[1.52, 3.33]	[1.58, 3.33]
1.4	[1.33, 3.33]	[1.44, 3.33]	[1.52, 3.33]	[1.62, 3.33]	[1.61, 3.33]
1.6	[1.31, 3.33]	[1.43, 3.33]	[1.51, 3.33]	[1.57, 3.33]	[1.58, 3.33]
1.8	[1.67, 3.33]	[1.41, 3.33]	[1.49, 3.33]	[1.56, 3.33]	[1.55, 3.33]
Total revenue					
0.2	166	166	166	165	162
0.4	166	166	165	164	158
0.6	173	169	166	164	157
0.8	194	183	176	170	160
1.0	216	201	190	181	167
1.2	238	220	205	193	175
1.4	260	238	221	206	185
1.6	283	257	237	220	196
1.8	308	277	254	234	207

Notes. A cycle length of one implies that $p^*(r_t)$ and r_t converge to a steady state. The price range is evaluated for a cycle while ignoring the initial burn-in period.

Table 2. The Optimal Uniform Pricing Policy and Its Percentage Loss Compared with the Optimal Dynamic Pricing Policy Reported in Table 1

γ	$\theta = 0.1$	$\theta = 0.3$	$\theta = 0.5$	$\theta = 0.7$	$\theta = 0.9$
Optimal uniform price					
0.2	1.67	1.65	1.65	1.65	1.64
0.4	1.67	1.67	1.64	1.64	1.62
0.6	1.66	1.66	1.65	1.64	1.59
0.8	1.66	1.66	1.66	1.64	1.59
1.0	1.66	1.65	1.65	1.65	1.59
1.2	1.66	1.66	1.65	1.64	1.60
1.4	1.66	1.66	1.66	1.65	1.59
1.6	1.66	1.66	1.66	1.64	1.60
1.8	1.66	1.66	1.65	1.64	1.60
Percentage revenue loss relative to the optimal policy					
0.2	0%	0%	0%	0%	0%
0.4	0%	0%	0%	0%	0%
0.6	4%	2%	1%	1%	1%
0.8	4%	10%	7%	5%	4%
1.0	23%	18%	14%	11%	9%
1.2	30%	25%	20%	16%	14%
1.4	36%	31%	26%	22%	20%
1.6	41%	36%	31%	27%	24%
1.8	46%	40%	36%	31%	29%

cyclic behavior tends to emerge for sufficiently large γ—a dynamic that is consistent with Proposition 5 and Theorem 1. In the numerical example, a long-run uniform pricing policy is optimal only for $\gamma \leq 0.4$ (with one exception: $\gamma = 0.6$ and $\theta = 0.9$), which is roughly one-fifth the price sensitivity of customers in the same segment (b_2). Second, the cycle length is not monotone in any of the parameters, including γ and θ (see also the discussion after Proposition 3); the values reported in Table 1 confirm the highly irregular behavior of the cycle length. Third, the potential loss when a firm ignores consumer heterogeneity can be substantial. If the firm assumes that consumers are *homogeneous* in their perceived gains and losses (i.e., the market is not segmented), then the firm will charge the long-run uniform price shown in Table 2 (theorem 4 of Popescu and Wu 2007). When customers differ significantly in the strength of their reference effects ($\gamma \geq 1.4$), the potential revenue loss can be as much as 46%, as shown in Table 2.

4.2. Does Cyclic Pricing Occur for Empirically Observed Parameter Values?

The parameters that we use below are based on two papers: Dahana and Terui (2006) and Moon et al. (2006). Dahana and Terui (2006) use the scanner panel data from two categories in the Japanese market: curry roux and instant coffee. The analysis in Moon et al. (2006) is based on the toilet tissue purchase records of 341 households over 114 weeks.

We consider two segments of customers in a market. Segment 1 is not subject to reference effects, but segment 2 is. The estimated relative size of segment 1 to segment 2 ranges approximately from 1 : 1 (Figure 3(b) in Dahana and Terui 2006) to 1 : 5 (table 4 in Moon et al. 2006). In the numerical example, we use $a_1 = 1$ and $a_2 = 2$. The estimated gain sensitivity of segment 2 ranges approximately from $\gamma = 0.2$ to $\gamma = 0.9$, whereas the ratio λ/γ ranges from one to five (table 4 of Dahana and Terui 2006 and table 4 of Moon et al. 2006). In this example, we use $\gamma = 0.8$ and $\lambda = 2$. The estimated price sensitivity for segment 1 (i.e., b_1) ranges from 0.4 to 3.3, whereas that of segment 2 (i.e., b_2) ranges from 1.3 to 3.3 (Dahana and Terui 2006, p. 17). In this example, we use $b_1 = 0.5$ and $b_2 = 3$. We also use the memory parameter $\theta = 0.79$ estimated in table 4 of Moon et al. (2006). Because these papers consider products that are purchased relatively frequently, we set the discount factor $\beta = 0.99$.

The firm's optimal pricing policy is illustrated in Figure 2. It is optimal for the firm to vary price over time and profit from the heterogeneity in reference effects of

Figure 2. (Color online) The Optimal Dynamic Pricing Policy and the Optimal Uniform Pricing Policy for Empirically Calibrated Parameters

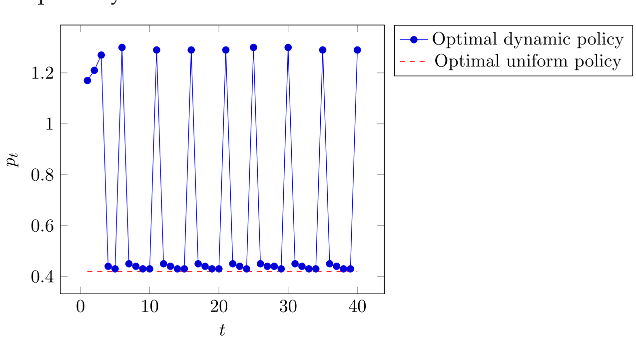

the two segments. The optimal total discounted revenue is 65.3. However, if the firm ignores the heterogeneity, then the optimal uniform price is $p_t \equiv 0.42$, which generates a total discounted revenue of 62.9.

4.3. Is Cyclic Pricing Optimal When Both Segments Are Loss Averse?

Consider a more general model, in which both segments are loss averse. Because the revenue functions are not concave, deriving the optimal dynamic pricing policy is an analytically intractable problem. Therefore, we rely on numerical experiments to showcase that our insights apply more generally.

Consider $a_1 = 1$, $b_1 = 0.3$, $a_2 = 4$, $b_2 = 2$, $\beta = 0.99$, and $\theta = 0.7$. Panel (a) of Figure 3 is the base case, in which we set $\gamma_1 = 0$, $\lambda_1 = 0.1$, $\gamma_2 = 1.5$, and $\lambda_2 = 2$. Because the reference effect in segment 2 is not sufficiently strong, the optimal dynamic price converges to a steady state. A segment 1 customers become less sensitive to losses (λ_1 decreases) or segment 2 customers become more sensitive to gains (γ_2 increases), cyclic behavior eventually emerges. In this example, cyclic pricing dominates uniform pricing (with respect to revenues) when $\gamma_2 \geq 1.75$ or $\lambda_1 \leq 0.05$ (panels (b) and (c) of Figure 3, respectively). However, when segment 2 customers become less sensitive to losses or segment 1 customers

Figure 3. (Color online) The Path of the Optimal Dynamic Pricing Policy and the Reference Price

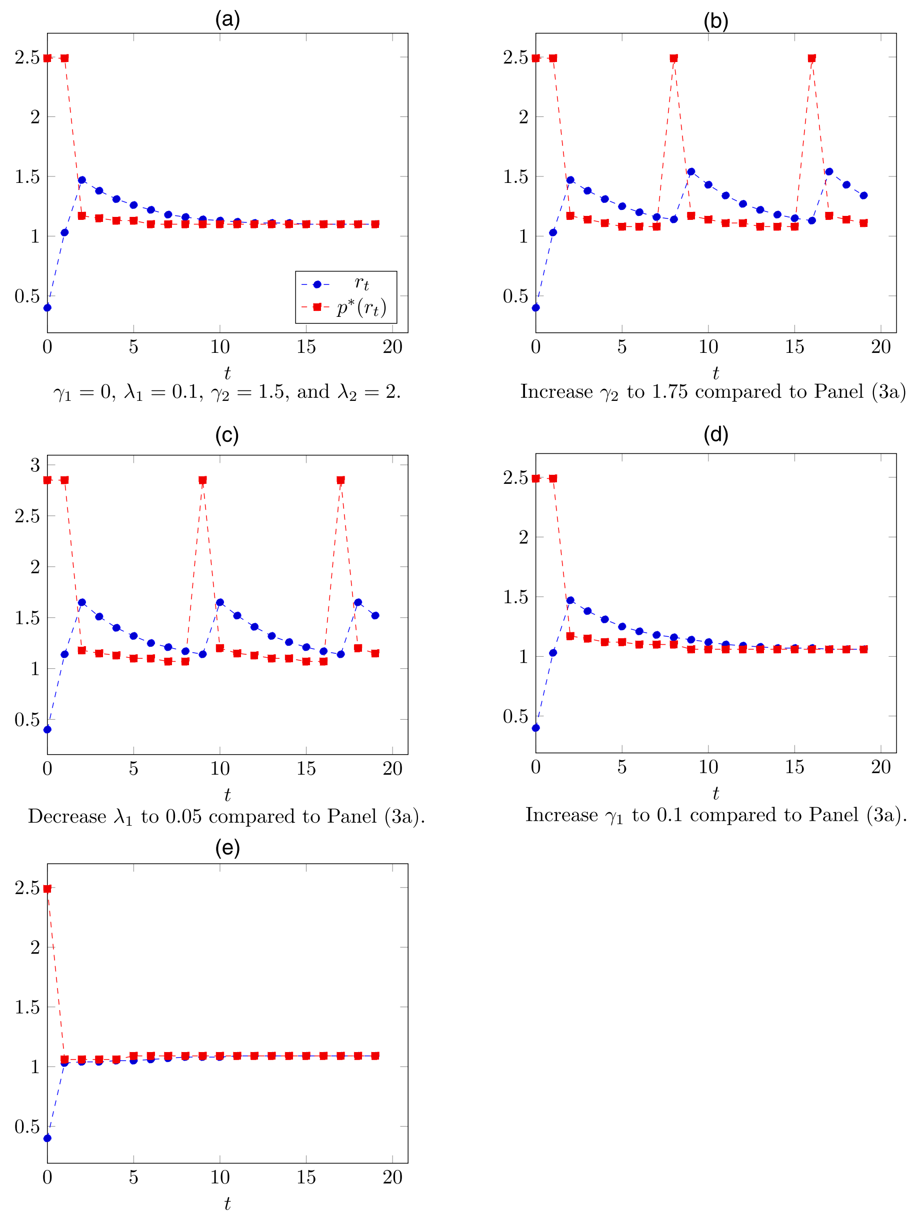

Notes. Increasing γ_2 or decreasing λ_1 results in cyclic pricing. Decreasing λ_2 or increasing γ_1 does not change the pricing behavior.

become more sensitive to gains, cyclic pricing does not emerge. To illustrate, we reduce λ_2 from two to 1.5 (note that $\lambda_2 \geq \gamma_2$ because of loss aversion) in panel (e) of Figure 3 and increase γ_1 from zero to 0.1 (loss aversion means that we must have $\lambda_1 \geq \gamma_1$) in panel (d) of Figure 3. As panels (d) and (e) of Figure 3 show, cyclic pricing is inferior to uniform pricing in such cases.

The experiment suggests that the cyclic behavior depends mostly on λ_1 and γ_2. In particular, a sufficiently large γ_2 and a sufficiently small λ_1 together give rise to cyclic pricing, whereas the values of λ_2 and γ_1 are not critical for the emergence of cyclic pricing. This is consistent with previous discussions. Although the loss effect from segment 2 is limited so that λ_2 is not critical, it is the opposite for segment 1. Because the firm does not price out the high-valuation segment 1, the magnitude of λ_1 significantly restricts the profitability of cyclic pricing. From the numerical examples, segment 2 customers have to be more sensitive to gains than segment 1 customers are to losses for cyclic pricing to be optimal.

5. Conclusion

The prevalent practice of offering discounts in consumer markets seems to be at odds with findings in the economic psychology literature that consumers are loss averse and dislike price variations. Those findings reflect the erosion, by discounts, of the internal reference price for a product—after which, a return to the product's "normal" price is perceived as a loss by consumers, reducing their demand for it. Hence, retailers have been advised to keep their prices constant in the long term, a strategy that limits the role that prices can play in managing consumer demand. However, recent developments in the operations management literature offer plausible setups, in which loss aversion does not preclude price variation. We add to this literature and suggest that the optimality of constant (long-run) prices hinges on the assumption that consumers are homogenous in their gain/loss perceptions (i.e., that their response to perceived gains is similar to perceived losses of the same amount). Hence, we model a market consisting of two segments of loss-averse consumers with *heterogeneous* gain/loss perceptions while making the (realistic) assumption that a firm cannot price discriminate between the two segments. We also restore a key role of prices in managing demand from multiple segments and let the firm use price to turn off the demand from a consumer segment. We explain why the firm's optimal dynamic pricing policy may be cyclic. One example is a markdown policy, in which the firm charges a regular high price but occasionally offers a discount, thereby significantly boosting demand by appealing to consumers who are relatively more sensitive to gains.

Our model applies to fast-moving consumer goods, where repeated purchase experiences allow customers to anchor on the product's price history to make purchase decisions. The choice of key parameters in Section 4 reflects this category of products, for which the firm's optimal policy is to vary prices over time and profit from the heterogeneity in reference effects across customer segments. Our results suggest that demand estimation procedures should account for differences in behavioral responses to a firm's prices. Attempts to fit a single gain/loss parameter to market demand data will result in adopting suboptimal pricing policies (e.g., uniform instead of cyclic policies) and profit loss.

Our paper complements a nascent literature that emphasizes the importance of consumer heterogeneity for a firm's pricing policies. Although our focus here is on heterogeneity in consumers' gain/loss perceptions, consumers are also likely to be heterogeneous in their internal (or external) reference prices, shopping habits (e.g., Wang 2016), and search behavior. Competitive markets offer another interesting context in which to study such characteristics. A natural extension to our model is to incorporate forward-looking behavior, because consumers often wait in anticipation of discounts and promotions. However, combining a memory-based reference price and forward-looking behavior may be challenging. We leave these ideas for future research.

Endnotes

[1] It may be that $a_1 = a_2$ and $b_1 = b_2$.

[2] Mathematically, the valuation distribution of segment 1 stochastically dominates that of segment 2.

[3] If the reference point is arbitrarily large, then the demand from either segment may be positive for prices above a_1/b_1. However, such large reference price values can occur only if past prices are consistently above a_1/b_1—a pricing policy that would be impractical for the firm. Even if the firm charges such prices, they are unlikely to be "assimilated" into the consumers' reference point (Narasimhan et al. 2005, pp. 365–366).

References

Adobe (2016) Media alert: Adobe data shows cyber Monday largest online sales day in history with $3.39 billion. Accessed June 8, 2017, http://news.adobe.com/press-release/marketing-cloud/media-alert-adobe-data-shows-cyber-monday-largest-online-sales-day.

Baucells M, Weber M, Welfens F (2011) Reference-point formation and updating. *Management Sci.* 57(3):506–519.

Bell DR, Lattin JM (2000) Looking for loss aversion in scanner panel data: The confounding effect of price response heterogeneity. *Marketing Sci.* 19(2):185–200.

Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1): 92–110.

Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Sci.* 61(4):723–739.

Cachon GP, Feldman P (2015) Price commitments with strategic consumers: Why it can be optimal to discount more frequently than optimal. *Manufacturing Service Oper. Management* 17(3): 399–410.

Chen X, Hu P, Hu Z (2016) Efficient algorithms for the dynamic pricing problem with reference price effect. *Management Sci.* 63(12):4389–4408.

Dahana WD, Terui N (2006) Reference price forma-tion model for heterogeneous consumers. Working paper, Osaka University, Suita, Osaka, Japan.

Dayaratna KD, Kannan PK (2012) A mathematical reformulation of the reference price. *Marketing Lett.* 23(3):839–849.

Fibich G, Gavious A, Lowengart O (2003) Explicit solutions of optimization models and differential games with nonsmooth (asymmetric) reference-price effects. *Oper. Res.* 51(5):721–734.

Greenleaf EA (1995) The impact of reference price effects on the profitability of price promotions. *Marketing Sci.* 14(1):82–104.

Hu Z, Nasiry J (2017) Are markets with loss-averse consumers more sensitive to losses? *Management Sci.* 64(3):1384–1395.

Hu Z, Chen X, Hu P (2016) Dynamic pricing with gain-seeking reference price effects. *Oper. Res.* 64(1):150–157.

Jacobson R, Obermiller C (1990) The formation of expected future price: A reference price for forward-looking consumers. *J. Consumer Res.* 16(4):420–432.

Kopalle PK, Rao AG, Assuncao JL (1996) Asymmetric reference price effects and dynamic pricing policies. *Marketing Sci.* 15(1):60–85.

Kopalle PK, Kannan PK, Boldt LB, Arora N (2012) The impact of household level heterogeneity in reference price effects on optimal retailer pricing policies. *J. Retailing* 88(1):102–114.

Mazumdar T, Raj S, Sinha I (2005) Reference price research: Review and propositions. *J. Marketing* 69(4):84–102.

Mersereau AJ, Zhang D (2012) Markdown pricing with unknown fraction of strategic customers. *Manufacturing Service Oper. Management* 14(3):355–370.

Moon S, Russell GJ, Duvvuri SD (2006) Profiling the reference price consumer. *J. Retailing* 82(1):1–11.

Narasimhan C, He C, Anderson ET, Brenner L, Desai P, Kuksov D, Messinger P, et al. (2005) Incorporating behavioral anomalies in strategic models. *Marketing Lett.* 16(3/4):361–373.

Nasiry J, Popescu I (2011) Dynamic pricing with loss-averse consumers and peak-end anchoring. *Oper. Res.* 59(6):1361–1368.

Neumann N, Böckenholt U (2014) A meta-analysis of loss aversion in product choice. *J. Retailing* 90(2):182–197.

Nielsen (2009) Almost half of U.S. supermarket purchases are sold on promotion. Accessed July 28, 2017, http://www.nielsen.com/us/en/insights/news/2009/almost-half-of-u-s-supermarket-purchases-are-sold-on-promotion.html.

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.* 62(2):326–346.

Phillips RL (2005) *Pricing and Revenue Optimization* (Stanford University Press, Stanford, CA).

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

Rajpathak B, Pillai HK, Bandyopadhyay S (2012) Analysis of stable periodic orbits in the one dimensional linear piecewise-smooth discontinuous map. *Chaos* 22(3):033126.

Wang Z (2016) Intertemporal price discrimination via reference price effects. *Oper. Res.* 64(2):290–296.

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