

The Perils of Behavior-Based Personalization

Juanjuan Zhang

MIT Sloan School of Management, Cambridge, Massachusetts 02142, jjzhang@mit.edu

“Behavior-based personalization” has gained popularity in recent years, whereby businesses offer personalized products based on consumers’ purchase histories. This paper highlights two perils of behavior-based personalization in competitive markets. First, although purchase histories reveal consumer preferences, competitive exploitation of such information damages differentiation, similar to the classic finding that behavior-based price discrimination intensifies price competition. With endogenous product design, there is yet a second peril. It emerges when forward-looking firms try to avoid the first peril by suppressing the information value of purchase histories. Ideally, if a market leader serves all consumers on day 1, purchase histories contain no information about consumer preferences. However, knowing that their rivals are willing to accommodate a market leader, firms are more likely to offer a mainstream design at day 1, which jeopardizes differentiation. Based on this understanding, I investigate how the perils of behavior-based personalization change under alternative market conditions, such as firms’ better knowledge about their own customers, consumer loyalty and inertia, consumer self-selection, and the need for classic designs.

Key words: behavior-based personalization; behavior-based price discrimination; revealed preference; segmentation; targeting; competition; customer relationship management; endogenous information generation

History: Received: May 19, 2008; accepted: August 20, 2010; Eric Bradlow served as the editor-in-chief and Ganesh Iyer served as associate editor for this article. Published online in *Articles in Advance* December 10, 2010.

1. Introduction

“Behavior-based personalization” has gained popularity in recent years, whereby businesses offer personalized products based on consumers’ purchase histories. Firms that have adopted this strategy include Amazon, Barnes and Noble, Digg, eBay, iTunes, Netflix, and YouTube, to name a few. The main argument for behavior-based personalization is the notion of revealed preferences; because consumers’ past choices reflect their tastes, firms can serve consumer needs better by modifying product offerings based on choice histories (Arora et al. 2008).

However, I argue that behavior-based personalization can hurt the profits of competing firms. Suppose a consumer subscribes to Netflix. Blockbuster can infer that this consumer likely appreciates the convenience of online DVD rentals. By offering her an online rental option, Blockbuster serves this consumer better. But this move only diminishes the differentiation between Netflix and Blockbuster to both firms’ detriment. I label this effect the “first peril of behavior-based personalization,” which echoes the classic result that behavior-based *price discrimination* intensifies price competition (Villas-Boas 1999, Fudenberg and Tirole 2000).

With endogenous *product design*, there is yet another peril. It arises as forward-looking firms try to strategically avoid the first peril. One obvious way is to have a market leader serve all consumers at day 1.

For example, let us suppose Netflix adopts a more “mainstream” service mode convenient for everyone and captures the entire market. Subscription to Netflix then contains no information about a consumer’s relative preference for online rentals, thus preventing both firms from offering personalized services. In other words, the mere incentive to avoid the first peril would compel Blockbuster to accommodate Netflix as a market leader. However, the same argument holds true had Blockbuster become the market leader in the first place. Indeed, knowing that their rivals are willing to welcome a market leader, both firms will offer a more mainstream service on day 1, which jeopardizes differentiation. I label this effect the “second peril of behavior-based personalization.”

I formalize the above intuition with a two-period duopoly model. Consumers have heterogeneous tastes and are uniformly distributed on a Hotelling line. Firms cannot collude, or commit to ignoring consumer purchase histories, in making product and pricing decisions. I start with a simple way to conceptualize consumers’ purchase histories: each firm offers one standard product in the first period and recognizes its customers. In the second period, a prisoner’s dilemma outcome arises as each firm chooses to offer different personalized designs to its customers and the rival firm’s patrons even though simultaneous personalization by both firms reduces differentiation. This effect leads to the first peril of behavior-

based personalization, where both firms are worse off in period 2 than if personalization were impossible. In addition, the more symmetric first-period market shares, the more informative purchase histories are in the “entropy” sense (Shannon and Weaver 1949), and the more severe the first peril.

For a detailed view of the market force leading to the second peril, suppose firms were sharing the market evenly in period 1. Now, firm *A* unilaterally moves its product closer to the market center. Firm *B* will normally respond by lowering prices to recoup some of its lost market share. However, firm *B* also realizes that its current loss in market share helps attenuate the first peril in the future and will thus respond with a smaller price cut than if it were myopic. Anticipating firm *B*’s mild response in prices, firm *A* will be more aggressive with its product strategy and will offer a design closer to the market center. However, the same reasoning applies to firm *B* as well. The net result is reduced differentiation and lower profits in period 1, although in this symmetric equilibrium firms still fail to avoid the first peril.

The emergence of the second peril contrasts findings from the behavior-based price discrimination literature. In particular, when analyzing the same market setting, Fudenberg and Tirole (2000) find that price discrimination based on purchase histories hurts profits in period 2, similar to the first peril, but *benefits* firms in period 1 by softening price competition. The benefit comes from a demand-side effect, that forward-looking consumers are less price sensitive in period 1, anticipating period 2 price discrimination. Notably, forward-looking firms’ incentive to avoid the first peril has no impact on period 1 competition. This is because a change in period 1 market shares from an even split has zero first-order effect on period 2 profits—each firm’s marginal gain from one clientele is offset by the marginal loss from the other. Firms thus have no incentive to vary period 1 prices to influence market shares.

This last result no longer holds with endogenous product design. In a symmetric equilibrium, a change in period 1 market shares still has zero first-order effect on period 2 profits. However, forward-looking firms’ incentive to avoid the first peril affects how they respond to their rivals’ period 1 product design in the pricing subgame. As a result, the first-order conditions of firms’ period 1 product choices shift with their degree of patience, which gives rise to the second peril. The second peril can outweigh the effect of decreased period 1 demand elasticity so that period 1 profits can be even lower than if behavior-based personalization were impossible, in contrast to Fudenberg and Tirole (2000).

I extend the main model to explore factors that affect the perils of behavior-based personalization.

I find that the attenuation of one peril often exacerbates the other. For example, one might conjecture that consumer preference information will be more beneficial to a firm if the rival cannot access this information: a firm may observe its customers’ exact preferences besides purchase histories; alternatively, a new generation of consumers may enter the market in period 2 such that the rival cannot perfectly identify the firm’s previous customers. In both cases, knowing more about their own customers does improve firms’ profits in period 2, thus mitigating the first peril, but hurts profits in period 1 as firms compete more aggressively for customers. The same result holds if firms can exploit consumer loyalty and inertia. Behavior-based personalization becomes more profitable in period 2 when customers are reluctant to switch. However, firms in period 1 will again compete more intensively for market shares.

Nevertheless, there do exist market conditions under which firms can reduce the perils of behavior-based personalization. If consumers are able to self-select among all personalized designs, firms will abandon personalization altogether to avoid intra-firm cannibalization. Alternatively, if firms are committed to providing a “classic design” for their old customers, it will help mitigate both perils. Finally, asymmetric patience between firms attenuates the first peril, whereas a larger segment of forward-looking consumers reduces the second peril.

I continue in §2 with a literature review. Section 3 introduces the model setup. Section 4 presents the main analysis, and §5 extends the main model in several ways. Section 6 discusses a set of testable empirical implications. Section 7 concludes the paper. All proofs are collected in the electronic companion, available as part of the online version that can be found at <http://mktsci.pubs.informs.org>.

2. Related Literatures

This paper is related to the theoretical literature on “behavior-based price discrimination,” meaning price discrimination based on customers’ purchase histories.¹ A well-known result is that conditioning prices on purchase histories can damage profits. Villas-Boas (1999) and Fudenberg and Tirole (2000) show that because purchases reveal preferences, firms will poach their rivals’ customers with lower prices than if purchase histories were unobservable. Villas-Boas (2004) finds that behavior-based price discrimination hurts a monopolist seller’s profits because strategic consumers foresee future price

¹ See Fudenberg and Villas-Boas (2007) for a comprehensive review of this literature. There is also an empirical literature that investigates the efficacy of conditioning prices on purchase histories (see, for example, Rossi et al. 1996, Besanko et al. 2003).

Table 1 Does Behavior-Based Discrimination Benefit or Hurt Firms?

Study	Main findings	Mechanism
Villas-Boas (1999)	Hurts firms when firms and consumers are patient	In an infinite-horizon duopolist market with overlapping generations of consumers, firms price below static levels to poach rivals' customers; patient consumers are more indifferent to which product to buy first and more sensitive to current prices, which intensifies price competition
Fudenberg and Tirole (2000)	Hurts firms in period 2; benefits firms in period 1	In period 2, duopolist firms price below static levels to poach their rival's customers, who have revealed their relative preference for the rival; in period 1, firms price above static levels because consumers, who anticipate second-period poaching, are less price sensitive
Villas-Boas (2004)	Hurts firms	In an infinite-horizon monopolist market with overlapping generations of consumers, prices cycle in equilibrium; the firm is worse off than without customer recognition because consumers foresee future prices cuts
Acquisti and Varian (2005)	Generally hurts firms; can benefit firms if they provide enhanced services to returning consumers	If consumers who hold higher valuation of the product also hold higher valuation of the enhanced services provided to returning consumers, the firm can profitably target high-valuation consumers with a high price conditional on initial purchase
Pazgal and Soberman (2008)	Generally hurts firms who add the same values to past customers	Because duopolist firms can lock in their past customers by adding value in period 2, they compete more aggressively for customers in period 1
Chen and Zhang (2009)	Can benefit firms in both periods if the market consists of loyal vs. price-sensitive consumers	In period 2, duopolist firms profit from being able to target loyal vs. price-sensitive consumers; in period 1, firms want to charge a higher price than the rival to identify the loyal consumers, thus softening competition
Shin and Sudhir (2010)	Can benefit firms in both periods with stochastic preferences and heterogeneous purchase quantities	In period 2, duopolist firms compete less aggressively if some consumers automatically prefer their product because of stochastic preferences, or if low prices disproportionately attract low-value customers; in period 1, consumers anticipate second-period poaching and are less price sensitive, which softens price competition
This study	Hurts firms in period 2 (the first peril); can hurt firms in period 1 (because of the second peril) with endogenous product design	In period 2, duopolist firms offer less differentiated designs and lower prices than static levels to poach rivals' customers, who have revealed their relative preference for the rival; in period 1, prices may drop below static levels as firms become less differentiated, knowing that the rival is more likely to accommodate mainstream designs

Note. The static benchmark refers to the case in which consumer purchase histories are unavailable.

cuts offered to low-valuation consumers. Acquisti and Varian (2005) consider a seller's ability to commit to a pricing policy and find that conditioning prices on purchase histories is generally unprofitable. Pazgal and Soberman (2008) explore the possibility of adding value to past customers. Although firms can lock in their customers in this way, they compete more aggressively for customers on day one.

This literature also uncovers a number of reasons why behavior-based discrimination might benefit firms, including different enhanced services to high-versus low-valuation past customers (Acquisti and Varian 2005), the coexistence of loyal versus price-sensitive consumers (Chen and Zhang 2009), and stochastic consumer preferences and heterogeneous

purchase quantities (Shin and Sudhir 2010). Notably, Fudenberg and Tirole (2000) find that the anticipation of behavior-based price discrimination makes consumers less price sensitive, which helps soften early-stage price competition. Table 1 presents a detailed summary of these studies on behavior-based price discrimination, comparing their main findings and underlying market forces.²

This paper extends the behavior-based price discrimination literature by endogenizing firms' product

² For comparability, in presenting the main findings of Fudenberg and Tirole (2000), I focus on the setting closest to mine, which features a two-period market with uniformly distributed tastes and short-term contracts.

design decisions. Practically the product design perspective is important today for at least two reasons. First, the revolution in flexible manufacturing has significantly lowered the costs of personalization (Dewan et al. 2003). Second, price discrimination may raise regulatory concerns and cause customer antagonism (Anderson and Simester 2010), forcing many businesses to obfuscate price disparities by varying product features. Therefore, product design has frequently been an endogenous variable even for the short run. Theoretically, endogenizing product design reveals the second peril of behavior-based personalization; firms can be even worse off in the early stage of competition than if behavior-based discrimination were impossible, a result opposite to the findings of Fudenberg and Tirole (2000).

In modeling firms' product decisions, I draw on the literature of spatial competition (e.g., Hotelling 1929, Wernerfelt 1986, Moorthy 1988, Thisse and Vives 1988, Desai 2001, Xiang and Sarvary 2007). In particular, this paper is related to the research on product customization. Dewan et al. (2003) analyze the strategic implications when firms can offer a continuous spectrum of perfectly customized products. Syam et al. (2005) show that competing firms will customize only one of two product attributes to soften price competition. Syam and Kumar (2006) find that offering customized products in addition to standard products can expand demand and improve profits. In the customization literature, firms are often assumed to exogenously know at least the distribution of consumer preferences. This paper, on the other hand, emphasizes consumer purchase history as the endogenous information basis of product design.³

This paper is also connected with the literature on targetability, which refers to a firm's ability to predict customer characteristics such as brand loyalty (Chen et al. 2001). Another related literature is customer addressability, where database technologies allow firms to know the tastes of a fraction of consumers and offer customized prices accordingly (Chen and Iyer 2002). This study contributes to these literatures by highlighting firms' endogenous production of customer preference information and sophisticated consumers' strategic release of preferences information.

3. Model Setup

Two horizontally differentiated firms, denoted A and B , compete in a nondurable goods market. Both

firms incur the same marginal cost of production, which is normalized to zero. In each period consumers have unit demands. I assume that the intrinsic value of the product v is sufficiently high so that all customers are served in equilibrium. Assuming full market coverage allows us to focus on the competition effects.

A unit mass of consumers are uniformly distributed along the Hotelling interval $[0, 1]$, where consumer's location x represents her ideal product attribute. A consumer incurs a disutility of taste mismatch, or a "transportation cost," by consuming a product away from her ideal point. I assume that transportation costs are quadratic in a consumer's distance to the product. Specifically, if consumer x buys a product located at a , she incurs transportation costs of $t(x - a)^2$, where t measures the degree of consumer taste heterogeneity. The quadratic assumption is appealing for Hotelling models with endogenous location choices. It ensures that for any product locations a pure-strategy price equilibrium exists. This property may not hold for other functional forms. For example, with linear transportation costs, a pure-strategy price equilibrium does not exist if the two firms are located relatively close to each other (see d'Aspremont et al. 1979 for a detailed discussion).

Key to the analysis is the dynamic revelation of consumer taste information, and firms' and consumers' influences over this revelation process. I model the dynamics with a two-period game. In period 1, firms A and B simultaneously choose the location, or "design," of their products on the Hotelling line, denoted as a and b , respectively.⁴ After observing both location choices, firms then simultaneously determine the prices p_a and p_b , respectively. It is common in the spatial competition literature to assume that firms' price decisions occur subsequent to product decisions, because prices are typically faster to change than product design. Consumers choose between products a and b after observing their locations and prices. Each firm recognizes consumers' choices. However, firms do not observe each consumer's *exact* taste (i.e., her exact location on the Hotelling interval). In other words, I make a minimum-information assumption about consumer purchase histories; a consumer's product choice reveals her relative preference between the two firms but not the precise strength of her preference.⁵

³ Arora et al. (2008) distinguish between customization and personalization: customization refers to a consumer's own specification of product features to purchase, whereas personalization refers to a firm's tailored product offerings to an individual consumer based on its data about that consumer. I follow this terminology and use the word "personalization" for the strategy I analyze.

⁴ To identify the market force of interest in a clean and tractable way, I assume that each firm offers one product in the first period. I will discuss this assumption in §5.1.1.

⁵ Often adopted in the behavior-based price discrimination literature, this minimum-information assumption best describes industries where firms do not have data on consumer preferences other than a snapshot of consumers' previous choices. Theoretically,

In period 2, each firm can offer either a *standard product design* to all consumers or *personalized product designs* to its own customers and the rival's customers, respectively.⁶ If both firms choose standard product design, second-period competition reduces to a static location-price game. If a firm chooses to personalize its product design, I use the subscript "o" to refer to the design for its own customers and "n" to denote the design for its new customers. After observing each others' location choices, firms simultaneously set prices. For the main analysis, I allow each firm's second-period product designs to be different from its first-period design, which is plausible if flexible manufacturing enables firms to update product designs frequently at negligible costs. I also assume that firms can target a consumer segment with a specific design. That is, if both firms have selected personalization, then firm A's old customers will choose between designs a_o and b_n , priced at p_{a_o} and p_{b_n} , respectively; firm B's old customers will choose between a_n and b_o , priced at p_{a_n} and p_{b_o} . To focus on the information role of purchase histories, I assume that consumers incur zero switching costs.

Firms and consumers maximize their total discounted payoffs over the two periods. Let β and δ denote firm and consumer discount factors, respectively. I treat β and δ as free parameters to trace the different impact of firm versus consumer patience on market outcomes.

4. Main Analysis

I derive the equilibrium of the two-period game through backward induction, using perfect Bayesian equilibrium as the solution concept (see also Fudenberg and Tirole 2000). In each period, I first establish consumers' choices given firm decisions, and then I derive firms' optimal decisions anticipating consumer responses.

4.1. The Second Period

4.1.1. Consumers' Second-Period Choices. In the second period, each consumer is offered two products given her first-period choice and firms' personalization strategies. For example, if both firms adopt standard product design, consumers choose between these two products as in a static game. If both firms

offer personalized designs, firm A's old customer x can either stay with firm A and purchase product a_o , or switch to firm B and buy b_n . Consumer x will prefer product a_o if and only if $p_{a_o} + t(x - a_o)^2 < p_{b_n} + t(x - b_n)^2$. A similar decision rule applies to firm B's old customers. Therefore, if both firms offer personalized designs, they will compete over two markets in the second period, each market composed of one firm's customer base.

4.1.2. Firms' Second-Period Decisions. Suppose each firm had a positive market share in period 1. Consumers' period 1 choices thus reveal their relative preferences, providing firms with a new segmentation variable. The question then is whether firms in period 2 will indeed base their product and pricing decisions on this segmentation variable. I state the answer in the following lemma (see the electronic companion for the proof).

LEMMA 1. *Having segmented consumers based on their first-period product choices, in the second period both firms (weakly) prefer personalized product designs to standard product design and (weakly) prefer to offer different prices to different segments.*

Lemma 1 reflects the well-established result that firms without commitment power would unilaterally prefer the extra degree of freedom from a discriminatory policy (Thisse and Vives 1988, Shaffer and Zhang 1995). However, it is worth noting that the unilateral profitability of a discriminatory policy does not come from a larger strategy space per se (i.e., being able to offer multiple designs as opposed to one) but from market segmentation as its information basis. Indeed, without segmentation information, each firm will offer only one product even if it has the option to offer multiple products. Lemma 2 states this important result formally (see the electronic companion for the proof).

LEMMA 2. *If no segmentation information is available over a uniformly distributed mass of consumers, in a static equilibrium both firms adopt standard product design even if offering additional products is costless.*

The intuition behind Lemma 2 is as follows. When a firm offers multiple products without segmentation information, it cannot target a specific product at a specific consumer. Instead, it has to allow consumers to self-select among all products in the market, a practice called "product proliferation" (Arora et al. 2008). Proliferation creates cannibalization within a firm's own product line, and the decrease in profit margins more than offsets any gain in market share.⁷

this assumption highlights the endogenous nature of consumer preference information: that it can only be revealed through actual purchases. In §5.1.1 I ask what happens if a firm observes the exact preferences of its own customers.

⁶ Given the model setup, each firm offers at most two personalized designs as an equilibrium outcome (Lemma 2). In practice, however, if firms have finer-grained preference information, personalization may imply a wider array of products further tailored to individual consumer preferences.

⁷ From the firms' perspective, uniform taste distribution represents the least demand information in the entropy sense, as every taste

Lemmas 1 and 2 together imply that, supposing firms have shared the first-period market, in the second period there will be four personalized products: products a_o and b_n for firm A’s old customers, and products b_o and a_n for firm B’s old customers. Assume that firm A locates to the left of firm B in any static location-price game without loss of generality. I show in the electronic companion that there exists a consumer \hat{x} who is indifferent between the two firms in period 1, such that the segment of consumers over $[0, \hat{x}]$ are firm A’s old customers, and the segment $(\hat{x}, 1]$ forms firm B’s clientele. Therefore, period 2 competition over either segment is a standard duopolist location-price game. The following lemma summarizes the equilibrium of this class of games (see the electronic companion for the proof).

LEMMA 3 (EQUILIBRIUM OF STATIC LOCATION-PRICE GAMES). Consider an arbitrary segment $[Z, Z + L] \subset \mathfrak{R}$ of uniformly distributed consumers of mass m . Two firms first simultaneously choose locations and then simultaneously set prices observing both location choices. Consumers’ transportation costs are quadratic with coefficient t . The unconstrained equilibrium locations are $Z - L/4$ and $Z + 5L/4$; firms charge the same price of $3tL^2/2$, each serving a demand of $m/2$ and earning a profit of $3tL^2m/4$.

For the rest of the analysis I shall leave product locations unconstrained, which means firms can locate outside the Hotelling interval (see Tyagi 2000 for the same assumption). By doing so, I admit the possibility that firms may locate “off the market” to serve the market. For example, many shopping malls choose remote locations, and all their customers have to travel. More generally, products may contain features that all consumers find undesirable; firms may even introduce “nuisance attributes” for differentiation purposes (Gerstner et al. 1993). Nevertheless, the key intuition of this paper remains unchanged if firms must locate within the Hotelling interval (see §A-6 of the electronic companion).

It follows from Lemma 3 that when $\hat{x} = 0$ or 1, the second-period game degenerates to a static location-price game with the two standard designs located at $-1/4$ and $5/4$ in equilibrium. Correspondingly, each firm charges an equilibrium static price of $3t/2$ and

earns an equilibrium static profit of $3t/4$. In subsequent analyses I will frequently compare market outcomes to this static benchmark.

When $\hat{x} \in (0, 1)$, the equilibrium second-period product locations are

$$\begin{aligned} a_n^*(\hat{x}) &= \frac{5}{4}\hat{x} - \frac{1}{4}, & a_o^*(\hat{x}) &= -\frac{1}{4}\hat{x}, \\ b_n^*(\hat{x}) &= \frac{5}{4}\hat{x}, & b_o^*(\hat{x}) &= -\frac{1}{4}\hat{x} + \frac{5}{4}. \end{aligned} \quad (1)$$

Figure 1 illustrates the equilibrium product locations in the second period given that firms shared the first-period market. The smaller a segment, the more tailored the personalized products are for consumers in that segment. This result is intuitive as a smaller segment represents finer knowledge of consumer tastes. Meanwhile, through personalization each firm offers its old customers a design better tailored to their tastes, compared to the case of standard design (because $-1/4 < a_o^*(\hat{x}) < 0$ and $1 < b_o^*(\hat{x}) < 5/4$ for any $\hat{x} \in (0, 1)$). Among the four new personalized designs, firm A’s products are located to the left of firm B’s products. That is, personalization maintains firms’ core image although part of its purpose is to poach the rival’s customers. Consumers with a strong preference for one firm over the other stay with the same firm in period 2, whereas the relatively indifferent consumers switch to the rival firm.


It also follows from Lemma 3 that for $\hat{x} \in (0, 1)$, the second-period equilibrium prices are $p_{a_o}^*(\hat{x}) = p_{b_n}^*(\hat{x}) = 3t\hat{x}^2/2$ and $p_{a_n}^*(\hat{x}) = p_{b_o}^*(\hat{x}) = 3t(1 - \hat{x})^2/2$. These prices are all lower than the equilibrium price of $3t/2$ in a static location-price game without personalization. In addition, the equilibrium prices of personalized designs all increase with the size of the target segment, because serving a more diverse set of tastes increases differentiation. Combining these results, firms’ second-period profits are

$$\pi_{A2}^*(\hat{x}) = \pi_{B2}^*(\hat{x}) = \frac{3}{4}t[\hat{x}^3 + (1 - \hat{x})^3], \quad \hat{x} \in [0, 1]. \quad (2)$$

The analysis so far reveals a difference between behavior-based price discrimination and personalization. In the former case, a firm is able to charge a higher price yet occupy a larger market share among

is equally possible (Shannon and Weaver 1949). Therefore, offering multiple personalized products, even if production is free, is unattractive because each product lacks a solid target market. In practice, with demand knowledge sufficiently better than a uniform distribution, a firm may offer a proliferation of multiple products. For example, Wernerfelt (1986) analyzes a market where consumers are concentrated at the two ends of the Hotelling interval and finds that firms may each offer two products. In this sense, Lemma 2 best describes new markets where firms initially have little demand information.

Figure 1 First-Period Consumer Choices and Second-Period Product Equilibrium



its old customers than its rival firm (Villas-Boas 1999, Fudenberg and Tirole 2000). This incumbent advantage comes from the fact that a firm's clientele naturally consists of consumers whose tastes are more aligned with the firm. However, this advantage dissipates when cost-effective personalization becomes available. On the one hand, the rival firm can offer a product better designed for the incumbent firm's old customers. On the other hand, the incumbent firm will offer a design less suited to its old customers in order to differentiate. That is, because of the personalization capacity, incumbent firms can no longer commit to offering the same product designs that attracted their clientele in the first place. As a result, firms earn equal profits over each other's customer base, as Equation (2) suggests.

It is worth noting that consumer preference information hurts both firms in the second period. The only payoff-relevant variable for the second period is \hat{x} , the location of the first-period indifferent consumer. From the information theory perspective, \hat{x} also measures the informativeness of purchase histories. Given uniform taste distribution over $[0, 1]$, the closer \hat{x} is to $1/2$, the lower the Shannon entropy, and the more information it contains (Shannon and Weaver 1949). In the extreme case of \hat{x} being 0 or 1, purchase histories contain no information beyond the prior distribution of consumer tastes. However, both firms' second-period profits are the lowest when $\hat{x} = 1/2$ and the highest when \hat{x} equals 0 or 1. This is because if purchase history information is available, by Lemma 1, both firms will exploit it as the basis for personalization, but simultaneous personalization "localizes" competition to both firms' detriment. Therefore, a prisoner's dilemma outcome emerges in period 2, where firms are worse off than if personalization were infeasible (i.e., the static benchmark). I call this effect the first peril of behavior-based personalization and summarize this result in the following proposition.

PROPOSITION 1 (THE FIRST PERIL OF BEHAVIOR-BASED PERSONALIZATION). *In the second period, both firms will offer personalized designs based on consumer purchase histories. By doing so, both firms are worse off than if personalization were infeasible.*

To escape the first peril, firms would ideally want to avoid consumer preference information altogether. At the extreme, they can suppress preference information by having one firm sell to all consumers in period 1, such that firms cannot differentiate consumers based on their past choices. In less extreme scenarios, firms can weaken the informativeness of purchase history by splitting the first-period market more asymmetrically. However, strategic consumers may change their

first-period choices to influence the amount of preference information they reveal and thus the product offers they receive in period 2. In the following subsection, I analyze how these opposing incentives affect the first-period market outcome.

4.2. The First Period

4.2.1. Consumers' First-Period Choices. In the first period, consumers choose between products a and b , knowing that this decision also determines their choice set in period 2.⁸ If consumer x buys a , in period 2 she will choose between products a_o and b_n . The discounted total cost she pays to buy a is therefore

$$p_a + t(x-a)^2 + \delta \min[p_{a_o} + t(x-a_o)^2, p_{b_n} + t(x-b_n)^2]. \quad (3)$$

Similarly, if consumer x buys b in the first period, she incurs a discounted total cost of

$$p_b + t(x-b)^2 + \delta \min[p_{a_n} + t(x-a_n)^2, p_{b_o} + t(x-b_o)^2]. \quad (4)$$

By definition, these two costs are equal when $x = \hat{x}$. Meanwhile, because switching always happens in the second period, if the indifferent consumer \hat{x} chooses a she will switch and buy b_n in period 2, and if she chooses b she will subsequently purchase a_n . Therefore, \hat{x} solves $p_a + t(x-a)^2 + \delta[p_{b_n} + t(x-b_n)^2] = p_b + t(x-b)^2 + \delta[p_{a_n} + t(x-a_n)^2]$. Plugging in the equilibrium values of a_n , b_n , p_{a_n} , and p_{b_n} , which in turn are functions of \hat{x} , I derive the first-period indifferent consumer as

$$\hat{x} = \frac{a+b}{2} + \frac{p_b - p_a}{25\delta t/8 + 2t(b-a)} + \frac{\delta(1-a-b)}{2\delta + 32(b-a)/25}. \quad (5)$$

Equation (5) reveals the underlying mechanism by which consumers' strategic reactions influence the market outcome. The first term on the right-hand side, $(a+b)/2$, captures the direct demand-expansion effect of firms positioning toward the market center. The second term is the price effect on first-period demand, which is attenuated by not only product differentiation $(b-a)$ and taste heterogeneity t , but also by consumer patience δ . To see why, suppose firms share the market evenly in period 1. A unilateral price cut helps a firm attract a larger clientele, but consumers in a larger clientele will receive less tailored products at higher prices in period 2. Therefore, the more forward looking these marginal consumers are, the less sensitive they are to current prices. The third term is more subtle. It has the opposite sign of $(a+b)/2 - 1/2$ whenever $\delta > 0$, but it is

⁸ Because there is a continuum of consumers, each consumer's first-period decision alone does not alter firms' strategies in the second period.

smaller in magnitude (unless $a = b$). In other words, consumer patience weakens the demand-expansion effect; if product a is closer to the center of the market than product b (i.e., $a + b > 1$), the third term partially offsets this location advantage. To understand this result, suppose consumers are myopic, and the corresponding indifferent consumer is $\hat{x}_{\text{myopic}} > 1/2$. By choosing product b over a , this consumer will be strictly better off in the second period because of more localized competition over the smaller segment $[\hat{x}_{\text{myopic}}, 1]$. Had this consumer been patient, she would have strictly preferred product b in anticipation of receiving better offers in the second period. As a result, $\hat{x}_{\text{patient}} < \hat{x}_{\text{myopic}}$ when $\hat{x}_{\text{myopic}} > 1/2$. By the same logic, $\hat{x}_{\text{patient}} > \hat{x}_{\text{myopic}}$ if $\hat{x}_{\text{myopic}} < 1/2$. That is, patient consumers are less responsive to aggressive designs that are closer to the market center. In fact, absent the price effect, the indifferent consumer will be closer to the center of the market when consumers are patient than when they are myopic.

In summary, consumer patience makes it more likely that firms split the first-period market symmetrically, an outcome representing the most preference information and the least second-period firm profits. Therefore, firms and consumers have conflicting interests regarding the production of consumer taste information. Below I investigate firms' first-period decisions and the market equilibrium.

4.2.2. Firms' First-Period Decisions and the Market Equilibrium. In period 1, firms first simultaneously choose locations a and b and then simultaneously set prices p_a and p_b , observing both location choices. In doing so, firms maximize their discounted total profits

$$\begin{aligned}\pi_A &= \pi_{A1} + \beta\pi_{A2}^*(\hat{x}) = p_a\hat{x} + \beta\pi_{A2}^*(\hat{x}), \\ \pi_B &= \pi_{B1} + \beta\pi_{B2}^*(\hat{x}) = p_b(1 - \hat{x}) + \beta\pi_{B2}^*(\hat{x}),\end{aligned}\quad (6)$$

where $\pi_{A2}^*(\hat{x})$ and $\pi_{B2}^*(\hat{x})$ are given by Equation (2) and \hat{x} is given by Equation (5). These profit functions suggest another source of conflicts in the market place. Not only do firms and consumers have opposite preferences regarding personalization, each firm itself also faces an intertemporal trade-off. Should firms only care about second-period profits, they would have one firm sell to all consumers in period 1. However, if firms were completely myopic, they would play a static location-price game in period 1 and split the market evenly. The equilibrium outcome is therefore jointly shaped by firm patience and consumer patience. I first present the results and then discuss the intuition (see the electronic companion for the proof).

PROPOSITION 2 (THE SECOND PERIL OF BEHAVIOR-BASED PERSONALIZATION). *There exists a symmetric*

pure-strategy equilibrium in which firms split the market evenly in period 1. In this equilibrium, forward-looking firms' incentive to avoid the first peril of behavior-based personalization further erodes profits: first-period differentiation and prices both decrease with firm patience. Meanwhile, first-period differentiation and prices both increase with consumer patience.


The condition for the symmetric equilibrium to arise holds for the majority of the parameter space, including the case where firms and consumers are equally forward looking (see Figure A-2 in the electronic companion). The rest of this section will focus on describing this symmetric equilibrium. Figure 2 presents how consumer patience and firm patience affect equilibrium first-period product differentiation and prices.

First-period differentiation, as measured by $b^* - a^*$, increases as consumers become more forward looking. As we have seen, consumer patience makes \hat{x} sticky to the market center (the third term in Equation (5)). The more forward looking consumers are, the less likely it is for a firm to become the market leader by offering a design close to the market center, thereby decreasing firms' incentive to do so. Similarly, first-period prices increase with consumer patience because forward-looking consumers are less price sensitive (the second term in Equation (5)). This price effect is well known in the behavior-based price discrimination literature (Villas-Boas 1999, Fudenberg and Tirole 2000). However, with endogenous product decision, the effect is amplified as consumer patience increases first-period differentiation, which further softens price competition.

Notably, first-period prices decrease with firm patience, contrary to findings from the behavior-based price discrimination literature. In particular, Fudenberg and Tirole (2000), when analyzing a similar two-period duopolist model with uniformly distributed demand, show that firm patience does not affect first-period prices. In their model, behavior-based discrimination only affects first-period prices by lowering consumers' price sensitivity. Even though firms are forward looking, the first-order effect of shifting the indifferent consumer equals zero around the market center—a firm's marginal gains in profit over one segment are exactly canceled out by losses over the other in period 2. This is an intriguing result because it implies that forward-looking firms' incentive to avoid the unprofitable use of purchase history information has no impact on the first-period market equilibrium.

I investigate why this classic result changes with endogenous product locations. Given any a and b , the interior subgame price equilibrium is specified by Equation (A-3) of the electronic companion. Note

Figure 2 Consumer Patience, Firm Patience, and First-Period Equilibrium



Notes. Each graph presents a set of curves for consumer patience $\delta \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The arrow indicates the direction along which δ increases. Each curve is plotted over $\beta \in [0, \bar{\beta}(\delta)]$, where $\bar{\beta}(\delta)$ is the maximal value of β for the symmetric equilibrium to exist given δ .

that whenever first-period products are symmetrically located ($a + b = 1$), these equilibrium prices become

$$p_a^*(a, b) = p_b^*(a, b) = t(b - a) + \frac{25}{16} \delta t. \tag{7}$$

With exogenously fixed a and b , as in most behavior-based price discrimination models, Equation (7) replicates the result that first-period prices increase with consumer patience δ but are invariant to firm patience β . Therefore, the only way by which firm patience affects first-period prices is through its impact on first-period differentiation. As products become less differentiated, price competition escalates. The key question then is why does first-period differentiation decrease with firm patience.

The reason is subtle. Consider firm A 's location decision as an example. By the chain rule, the derivative of firm A 's discounted total profit (Equation (6)) with respect to location a is

$$\frac{\partial \pi_A}{\partial a} = \frac{\partial \pi_{A1}}{\partial a} + \beta \frac{d\pi_{A2}^*(\hat{x})}{d\hat{x}} \frac{\partial \hat{x}}{\partial a}. \tag{8}$$

However, we also know from Equation (2) that the marginal effect of first-period market shares on second-period profits is zero in a symmetric equilibrium:

$$\frac{d\pi_{A2}^*(\hat{x})}{d\hat{x}} \Big|_{\hat{x}=1/2} = 0. \tag{9}$$

The removal of this last term from Equation (8) implies that product location has no impact on second-period profits in equilibrium. One might then conclude that firm patience does not affect first-period product locations. However, there is a less obvious

effect. Suppose the two firms were symmetrically located, but firm A unilaterally relocates a closer to the market center. This perturbation shifts the indifferent consumer to the right. Firm B will then react by cutting prices to regain some market share. However, with sufficient foresight, firm B should realize that its current loss of market share improves its second-period profits and will hence price less aggressively in response. Mathematically, it can be verified that in the subgame price equilibrium $\partial^2 p_b / \partial a \partial \beta > 0$. Anticipating less resistance in prices from a forward-looking rival, firm A is then more inclined to locate closer to the market center. Indeed, it can be shown that

$$\frac{\partial^2 \pi_A}{\partial a \partial \beta} = \frac{\partial^2 \pi_{A1}}{\partial a \partial \beta} > 0. \tag{10}$$

The same reasoning applies to firm B , who will also move closer to the market center with greater firm patience. In other words, firm patience reduces first-period differentiation not because firms are able to influence second-period profits via first-period locations; it is because, as firms try to improve second-period profits, firm patience shifts the first-order conditions of their period-one location choices, making aggressive designs (those close to the market center) unilaterally more profitable in the first period.

In contrast, for the exogenous product design setting of Fudenberg and Tirole (2000), the first-order condition of firm A 's period 1 price decision is $\partial \pi_A / \partial p_a = \partial \pi_{A1} / \partial p_a$ in a symmetric equilibrium because of the same effect of Equation (9). However, $\partial \pi_{A1} / \partial p_a = p_a \partial \hat{x} / \partial p_a + \hat{x}$, where \hat{x} as specified in Equation (5) does not depend on firm patience β . Therefore, $\partial^2 \pi_A / \partial p_a \partial \beta = \partial^2 \pi_{A1} / \partial p_a \partial \beta = 0$. Intuitively, firms'

incentive to avoid the first peril does not improve period 2 profits in a symmetric equilibrium; without a separate product design stage, neither does such incentive change period 1 profits. As a result, in equilibrium firm patience has no effect on first-period competition.

It will be useful to compare equilibrium first-period differentiation and prices to the benchmark values from a static location-price game (as given by Lemma 3). As Figure 2 shows, first-period differentiation is lower than the static value of $3t/2$ except when firms are myopic ($\beta = 0$). First-period prices can be either higher or lower than the static level of $3t/2$. This is because two countervailing forces affect first-period prices: whereas consumer patience lowers price sensitivity and increases first-period prices directly, firm patience reduces differentiation and decreases prices indirectly. The overall comparison with the static level depends on which effect dominates. Notably, behavior-based discrimination can intensify competition in the first period beyond the static benchmark, which differs from the finding of Fudenberg and Tirole (2000) that first-period prices are higher than the static level. I summarize these results in the following corollary.

COROLLARY 1. *In the first-period equilibrium, products are less differentiated than in a static location-price game unless firms are myopic; prices can be even lower than the static level.*

To complete the analysis, I derive the second-period equilibrium that follows from the symmetric first-period equilibrium. The first-period indifferent consumer is $\hat{x}^* = 1/2$, which produces the most consumer preference information and leads to the most unprofitable personalization. There are four personalized designs in the second period:

$$a_o^* = -\frac{1}{8}, \quad a_n^* = \frac{3}{8}, \quad b_n^* = \frac{5}{8}, \quad b_o^* = \frac{9}{8}.$$

The second-period prices for all four products equal $3t/8$. Consumers are all better off in the second period compared with the static case: every consumer receives a better design at a lower price. Each firm's second-period total profit is $3t/16$, lower than the static-game value of $3t/4$, which reflects the first peril. Moreover, all consumers are better off in period 2 than in period 1, but firms are worse off because of lower prices (see the electronic companion for the proof). I state this last result below.

COROLLARY 2. *Firms are worse off and consumers are better off in period 2 than in period 1 because of behavior-based personalization.*

In summary, the main analysis identifies the double perils of behavior-based personalization: firms cannot

help but use consumer purchase information to personalize products, a move that hurts second-period profits; but trying to suppress consumer purchase information hurts firms in the first period.

5. Extensions

In this section, I investigate how alternative market conditions affect the two perils of behavior-based personalization. In particular, I will relax the following assumptions of the main analysis: (1) a firm does not have better information about its own clientele than its rival, (2) there is no switching cost, (3) firms can target a design to a specific segment, and (4) the first-period designs are abandoned in period 2.

5.1. When a Firm Knows More About Its Customers Than Its Rival

In the main model, the amount of consumer preference information is fully captured by the indifferent consumer \hat{x} . Therefore, a firm has no information advantage over its rival. Two model features have contributed to this symmetric information outcome. First, a firm observes no further consumer preference information besides purchase histories. Second, firms can infer who the rival's customers are by recognizing their own. I revisit these assumptions in order.

5.1.1. Perfect Discrimination Within a Firm's Clientele. Consumer database technologies may help a firm gather more information about its customers besides purchase histories. Firms can collect customer demographic information through user accounts, conduct post-purchase satisfaction surveys, or gauge customer preferences by analyzing their product search behaviors (Hauser et al. 2009). To model this type of information asymmetry between firms, I assume that each firm observes the exact locations of its old customers. I continue to assume that personalization is free, which is plausible if investments in flexible manufacturing are sunk costs. As a result, in the second period, firms will implement first-degree discrimination by offering each old customer a personalized design that perfectly suits her taste (Thisse and Vives 1988), a strategy called "mass personalization" in practice.

In period 2, a firm's information advantage about its customer base translates into a competitive edge. Take firm A 's clientele $[0, \hat{x}]$ as an example. Although firm B can still serve this segment with product b_n (or even multiple products), with mass personalization, firm A can undercut firm B for every consumer x in $[0, \hat{x}]$ by offering her a price slightly lower than $p_{b_n} + t(x - b_n)^2$. Therefore, firm B cannot make a positive profit from firm A 's clientele. For simplicity I assume that firm B does not offer product b_n , as any infinitesimal manufacturing cost will strictly prevent

firm B from doing so. It follows that firm A will charge each of its previous customers a price equal to her reservation price v ; so will firm B for its own customers. The second-period prices are higher than the static level and hence higher than those in the main model, which negates the first peril of behavior-based personalization.⁹

Note that in period 2 each firm benefits from having a larger customer base, in contrast to the main model where both firms benefit from an asymmetric split of the market. Specifically, firms' discounted total profits are

$$\pi_A = p_a \hat{x} + \beta v \hat{x}, \quad \pi_B = p_b(1 - \hat{x}) + \beta v(1 - \hat{x}), \quad (11)$$

where the indifferent consumer \hat{x} is obtained by substituting $\delta = 0$ into Equation (5). This is because, knowing that they will receive zero surplus in period 2 regardless of their purchase histories, consumers in period 1 will maximize their current utility as if they were myopic. It follows that first-period competition is equivalent to a static location-price game with an additional profit margin of βv . We know from §A-2.1 of the electronic companion that this extra margin does not affect firms' period 1 equilibrium product designs, which are the same as in the static case: $a^* = -1/4$ and $b^* = 5/4$. However, it intensifies the price war, with the first-period equilibrium prices being lower than the static level of $3t/2$:

$$p_a^* = p_b^* = \frac{3}{2}t - \beta v. \quad (12)$$

Moreover, first-period prices are lower than those in the main model, and the difference widens with firm patience.¹⁰ That is, firms' ability to perfectly discriminate among their customers exacerbates the second peril of behavior-based personalization. The more forward looking firms are, the deeper the price cut they take. First-period prices may even be negative if consumers' intrinsic value v is sufficiently high. In practice, a firm may recruit customers by paying them cash incentives. An example is cell phone carriers who compete to sign up users by offering free phones and cash-back bonuses.

Because firms in period 1 "compete away" any extra consumer surplus they can extract through perfect discrimination, they are ex ante worse off compared with the static benchmark (unless $\beta = 0$ in

which case firms are indifferent). In addition, it can be shown that firms' discounted total profit of $3t/4$ is lower than that in the main model for a nonempty set of parameters. That is, firms can also be worse off than if they can only recognize consumers' purchase histories. The following proposition summarizes the results. The proof holds by construction.

PROPOSITION 3. *Firms' ability to perfectly discriminate among their own customers negates the first peril of behavior-based personalization but exacerbates the second peril. Firms are ex ante (weakly) worse off than in the static case and can be worse off than if they only recognized consumers' purchase histories.*

These results suggest that finer discrimination does not always help. Technologies that allow firms to collect individual customers' taste information may end up eroding industry profits. The results are also useful in understanding whether firms, in a broader context, want to offer multiple products in the first period. If a firm privately observes its own customers' choices, by offering multiple products in period 1 it can gain better knowledge about its customers than its rival. As the above analysis suggests, firms might then have greater incentive to compete for a larger market share, which may intensify first-period competition. Future research can formally model this trade-off.

5.1.2. New Generation of Consumers. New generations of consumers may enter the market over time, especially in expanding product categories (Villas-Boas 1999, 2004). Firms may not be able to distinguish between new-generation consumers and the rival's previous customers. Formally, I assume that in the second-period market a fraction γ of consumers comes from the new generation. That is, I normalize the mass of new-generation consumers as $\gamma/(1 - \gamma)$, whereas the mass of old-generation consumers continues to be 1. I also assume that new-generation consumers are uniformly distributed along the Hotelling interval. Firms offer personalized designs a_o and b_o to their old customers and a_n and b_n to whoever did not buy their products in period 1.


I begin by analyzing the second period. Let \hat{x}_n denote the location of the new-generation consumer who is indifferent between the two firms. Such an indifferent consumer exists, with \hat{x}_n solving $p_{a_n} + t(x - a_n)^2 = p_{b_n} + t(x - b_n)^2$. Let \hat{x}_A denote the consumer in firm A 's customer base who is indifferent between buying a_o and b_n , and let \hat{x}_B denote the consumer in firm B 's clientele indifferent between a_n and b_o . (If firm A 's old customers all prefer a_o , let $\hat{x}_A = \hat{x}$. Other boundary cases are similarly specified.) Firms' second-period profits are

$$\begin{aligned} \pi_{A2} &= p_{a_n} \hat{x}_n \gamma / (1 - \gamma) + p_{a_n} (\hat{x}_B - \hat{x}) + p_{a_o} \hat{x}_A, \\ \pi_{B2} &= p_{b_n} (1 - \hat{x}_n) \gamma / (1 - \gamma) \\ &\quad + p_{b_n} (\hat{x} - \hat{x}_A) + p_{b_o} (1 - \hat{x}_B). \end{aligned} \quad (13)$$

⁹ As proof, a necessary condition for full market coverage is that $v > 3t/2$, the price of the static location-price game.

¹⁰ The proof is that in the main model, v must be no smaller than $p_a^* + t(1/2 - a^*)^2$ to ensure that the marginal consumer at $1/2$ is willing to buy, where p_a^* is given by Equation (A-7) in the electronic companion. It is then straightforward to show that p_a^* of the main model is greater than $3t/2 - \beta v$ and that the difference increases with β .

Figure 3 Second-Period Equilibrium Differentiation and Prices with the Entry of New-Generation Consumers ($t = 1, \hat{x} = 1/2$)




Given these profit functions, I can derive the second-period equilibrium in the same way as in the main model. The results involve high-order polynomials, which I report graphically in Figure 3 to facilitate interpretation. For illustrative purpose, I set $t = 1$ and $\hat{x} = 1/2$.

As a larger fraction of period 2 consumers are from the new generation, the personalized designs for new customers (a_n and b_n) become more differentiated. Intuitively, serving the new generation whose tastes are broadly distributed increases differentiation. Meanwhile, serving the new generation diverts a firm’s focus from poaching the rival’s clientele, thus allowing the rival to offer a more tailored design to its old customers. When the fraction of the new generation γ is sufficiently high, the products for firms’ old customers (a_o and b_o) become less differentiated than those for the new customers. This result is the opposite of the main model, where old customers receive more “outlandish” designs in the second period. Finally, the prices of all personalized designs increase with γ ; the new generation mitigates the competition between the products for new customers, which in turn softens the competition over each firm’s clientele. Therefore, the entry of a new generation mitigates the first peril of behavior-based personalization by raising second-period prices above the level in the main model.

I next ask how the new generation affects the first-period equilibrium. Figure 4 presents first-period product differentiation and prices where $t = 1, \beta = 0.5$, and $\delta = 0.5$. Other values of β and δ suggest similar patterns. A larger fraction of new-generation consumers increases differentiation and reduces prices in the first period. These results can be interpreted as follows. As discussed above, when the fraction of the new generation increases, firms become less interested in poaching and are consequently more likely to

retain their old customers. This shifts forward-looking firms’ first-period imperative from suppressing preference information to growing the customer base. The implications are twofold: aggressive designs are less likely to be accommodated by the rival, leading to greater differentiation compared to the main model; however, firms compete more intensively in prices for market share. The latter effect is further aggravated by the fact that old-generation consumers’ price sensitivity increases with γ . This is because as more new-generation consumers enter the market, old-generation consumers are more likely to get poaching offers that are independent of their first-period choices and are thus more responsive to period 1 prices compared to the main model. The net result is that first-period prices and profits are lower than that in the main model. I summarize these findings below. The proof holds by construction.

Figure 4 First-Period Equilibrium Differentiation and Prices with the Entry of New-Generation Consumers ($t = 1, \beta = 0.5, \delta = 0.5$)



PROPOSITION 4. *The entry of a new generation of consumers mitigates the first peril of behavior-based personalization but exacerbates the second peril.*

This finding parallels the result of §5.1.1, where firms can perfectly discriminate within their clientele: when there is gain from having served a larger clientele, firms will compete more intensively in the first period for market share, which dissipates first-period profits.

5.2. Consumer Loyalty and Inertia

Firms may benefit from a large customer base not only by knowing their own customers better, but also through customer brand loyalty or inertia (see Chen 1997, Taylor 2003). In §A-7 of the electronic companion, I extend the main model by assuming that consumers incur a cost when switching to another firm in the second period. I find that a firm has a competitive advantage on its turf in period 2: it is able to offer a design better tailored to its customers and charge a higher price than its rival. In doing so, firms also earn a higher second-period equilibrium profit than in the main model, thus mitigating the first peril of behavior-based personalization.

In period 1, firms will try to expand their turf to increase their period 2 profits. Indeed, it can be shown that the first-order effect of market share on a firm's second-period profit is positive around $\hat{x} = 1/2$ and is increasing in the switching cost. This result is different from the main model, where the same first-order effect is zero, and firms are interested in asymmetric market shares rather than large market shares. Consistent with this intuition, first-period equilibrium differentiation and prices both decline with switching cost, exacerbating the second peril. I summarize these results below (see the electronic companion for the proof).

PROPOSITION 5. *Consumer loyalty and inertia mitigates the first peril of behavior-based personalization but exacerbates the second peril.*

These results are in line with common findings of the switching cost literature: firms can charge a higher price to consumers who are “locked in” because of the difficulty of switching; however, to compete for the surplus from locked-in consumers, firms escalate price wars early on (see Fudenberg and Villas-Boas 2007 for a review of this literature).

5.3. Consumer Self-Selection

The main model assumes that firms have perfect targeting abilities in the sense that they can limit a consumer's access to a specific product. Firms can then implement third-degree discrimination in period 2 based on consumers' past choices. This assumption is common in the literature and is relevant to some

industries. A familiar example is Amazon, whose personalized product recommendations can be thought of as a form of targeted offers—with search costs, a consumer may not access the products not recommended to her. Pazgal and Soberman (2008) note similar practices—Air Canada offers double frequent-flyer miles exclusively to newly registered members, and Scandinavian Airlines provides automatic flight information to travelers who sign up for wireless access at specific airports. In general, targeted offers are more frequently seen in categories such as travel, telecommunications, credit cards, catalogue retail, and Internet retail.

In other industries, the norm might be to provide a full menu of products to all consumers and let them self-select. For the market setting I analyze, self-selection means consumers will have equal access to all four personalized designs in the second period. However, this is equivalent to the case of proliferation in which firms compete over the whole market without the aid of segmentation information. By Lemma 2, each firm will offer a standard design to all consumers to avoid intrafirm cannibalization. That is, consumer self-selection offsets the effect of personalization by making competition more “global,” so that each firm earns the static profit of $3t/4$ in the second period. Because purchase histories no longer affect firm profits and consumer surplus in period 2, first-period competition reduces to the static location-price game. I summarize the results in the following proposition. The proof holds by construction.

PROPOSITION 6. *Consumer self-selection obviates both perils of behavior-based personalization. Firms solve a static location-price game in both periods.*

Some industries' intrinsic characteristics endow firms with weak targeting power. Meanwhile, targeted offers may antagonize consumers and may spur arbitrage activities. Interestingly, Proposition 6 suggests that firms could be better off under these seemingly adverse market conditions. The general message is that finer discrimination does not always benefit firms, either in the form of first-degree discrimination (see §5.1) or third-degree discrimination (i.e., targeted offerings). However, it should be noted that if the targeting ability can be endogenously acquired, then firms in equilibrium may do so and offer targeted personalized designs (Lemma 1). In this case, the prisoner's dilemma uncovered in the main model will again emerge, and firms will again incur the perils of behavior-based personalization.

5.4. Classic Designs


In some industries it is infeasible to abandon old product designs for every consumption cycle. In

particular, brand image concerns may require the provision of a timeless classic design to the existing clientele. Firms' personalization problem in this situation becomes a product line extension problem. To capture this market feature, I modify the main model by assuming that in period 2, each firm's old customers continue to receive the product they bought in period 1, either a or b . In addition, each firm offers a new design, a_n and b_n , to poach its rival's customers. However, because prices are often easier to change than product design, I allow the firms to charge their old customers a different price than what they paid in the first period. I denote the first-period prices of a and b as p_{a1} and p_{b1} and the second-period prices as p_{a2} and p_{b2} . In each period firms set prices simultaneously after observing each other's design choices.

In the second period, firms choose the new designs taking the classic designs as given, anticipating the price equilibrium that follows. Firms' commitment to the classic designs thus allows them to be the incumbent over their customer base. Intuitively, this incumbent advantage is maximized if a firm is located at the center of its turf. This effect might lead firms to locate aggressively toward the market center in period one. However, there is a countervailing effect, that aggressive designs will trigger deep price cuts from the rival firm. To see this, suppose firm A unilaterally moves a slightly closer to the market center, which also increases \hat{x} . Firm B will react by cutting prices to regain some market share even if it is myopic. If firm B is forward looking, it will cut prices even more. This is because when the classic design a is closer to the market center, firm B will earn lower poaching profits in period 2 and therefore will prefer to reduce \hat{x} and shrink the size of the unprofitable poaching market (mathematically, it can be verified that $\partial^2 \pi_{B2} / \partial a \partial \hat{x} < 0$). As a result, aggressive classic designs can ignite price wars worse than in the static case. Interestingly, this force is the opposite of the main model, where forward-looking firms accommodate aggressive designs to reduce the first peril.

Figure 5 presents equilibrium classic designs and their first-period prices as a function of firm patience β . The differentiation between the classic designs is higher than the static level of $3/2$ and increases with firm patience. This outcome suggests that the cost of aggressive designs outweighs the benefit as forward-looking firms resort to intensive price wars to protect their second-period profits. First-period prices are also above the static level of $3t/2$ and increase with firm patience. The high prices are partly attributed to the high differentiation between the classic designs, and partly attributed to patient consumers' lower price sensitivity as discussed in the main analysis. Overall, first-period profits are higher than in the main model and increase with

Figure 5 Differentiation Between Equilibrium Classic Designs and First-Period Equilibrium Prices ($t = 1, \delta = 1$)




firm patience. This effect negates the second peril of behavior-based personalization.

Last, Figure 6 presents the second-period equilibrium. The new designs are less differentiated than the classic designs, and the degree of differentiation between new designs decreases with firm patience. Naturally, as forward-looking firms position the classic designs far from the market, in period 2 they can poach their rival's customers by offering a new design better tailored to their tastes. In doing so, firms charge higher prices for the new designs than for the classic designs and earn the majority of their second-period profits from poaching new customers. Moreover, it can be verified that firms' total second-period profits are higher than the value of $3t/16$ in the main model, thus mitigating the first peril of behavior-based personalization. I summarize these results below.

PROPOSITION 7. *Firms' commitment to retaining a classic design for their old customers mitigates the first peril of behavior-based personalization and negates the second peril.*

In summary, §5 explores whether firms can circumvent the perils of behavior-based personalization in more general settings. I find that attenuating one peril often exacerbates the other. In particular, although firms can improve their second-period profits by gaining better information about their customers or by exploiting consumer loyalty and inertia, to grow a larger clientele they compete more intensively in period 1. Nevertheless, there do exist market conditions that help firms reduce the perils of behavior-based personalization, such as consumer self-selection and commitment to classic designs for previous customers.

Finally, because the symmetric equilibrium marks the lowest second-period profits, various forms of

Figure 6 Second-Period Equilibrium Differentiation, Prices, and Profits When Firms Retain a Classic Design for Their Old Customers ($t = 1, \delta = 1$)

market asymmetry may mitigate the first peril of behavior-based personalization and thus affect the second peril. In particular, I investigate asymmetric patience between firms (for example, because of unequal access to the credit market) in §A-8.1 of the electronic companion. I find that asymmetric firm patience does attenuate the first peril, although the more patient firm fares worse in period 1. I also consider heterogeneous consumer patience in §A-8.2 of the electronic companion. Firms are better off with a larger segment of forward-looking consumers because consumers' effort to induce behavior-based personalization counters firms' incentive to avoid it, thus weakening the second peril, a result consistent with Proposition 2.

6. Empirical Implications

In this section I summarize the findings of this paper in light of their empirical implications. One key prediction is that behavior-based personalization may hurt the profits of competing firms. This prediction contains two aspects. Cross-sectionally, firm profits can be lower in industries where behavior-based personalization prevails than in industries where it is infeasible, other things being equal. Longitudinally, firm profits are lower in mature markets than in new markets, with products becoming less differentiated and prices declining over time. This prediction challenges the common belief that personalization contributes to the bottom line by better meeting consumer needs (see Arora et al. 2008 for a review). I argue that personalization does deliver greater value to consumers but damages profits by intensifying firm competition.

The second result centers on how much firms know about their own customers. If firms can obtain fine-grained preference information about individual customers, there could emerge mass personalized products at high prices during the mature stage of

the market. However, firms compete intensively at the early stage and might offer cash incentives to attract customers. That is, what firms know about their customers may imply different price trends in an industry. I predict prices to increase over time when firms are equipped to analyze their customers at an individual level but to decline when firms can only recognize consumers' purchase histories.

Third, in stable industries with negligible entry of new-generation consumers, I expect that firms will offer their old customers more "extreme" product designs that symbolize the firm's core image and serve their new customers with more mainstream designs that cater to average tastes. However, I expect the opposite in expanding industries with a heavy influx of new-generation consumers. Moreover, the impact of the new generation on prices depends on the stage of the market—the analysis suggests that prices will decrease with the proportion of new-generation consumers in the early days but increase with it during the mature stage.

A fourth prediction of the model is that there will be less personalization if consumers can self-select between personalized designs than if firms can target different consumers with different products. Moreover, firm profits are higher if consumers can self-select. This prediction is counterintuitive because it implies that, other things being equal, firms can be worse off with better targeting technologies such as direct mail, back-of-receipt offers, and personalized product recommendation systems.

Last, whether firms maintain a classic design implies different market dynamics. Early-stage product designs are expected to be more differentiated if firms will retain them for their old customers than if firms are free to subsequently replace them with new designs. People would normally imagine long-lived classic designs to be more nuanced and tailored to mainstream tastes and seasonal designs to be

more avant-garde. However, the analysis predicts that classic designs should be sufficiently differentiated to epitomize their distinct brand personalities, whereas new designs should be moderately positioned to target mainstream buyers.

7. Concluding Remarks

This paper identifies two perils of behavior-based personalization in competitive markets. First, by the same force that behavior-based price discrimination intensifies price wars, competitive use of consumer purchase histories in product design commoditizes the marketplace (the first peril of behavior-based personalization). With endogenous product design, there is yet another peril: had a market leader served all consumers on day 1, purchase histories would have contained no information about consumer preferences, which could have eliminated the first peril. However, knowing that their rivals are willing to accommodate a market leader, firms are more likely to offer a mainstream design on day 1, which damages differentiation (the second peril of behavior-based personalization). Based on this understanding, I explore how alternative market conditions affect both perils. I find that firms' better knowledge about their own customers and switching costs mitigate the first peril but exacerbate the second. On the other hand, consumer self-selection and the need for classic designs help reduce both perils.

This paper suggests a perspective to understanding the era of product personalization. There are a number of paths to extend this research. It would be interesting to investigate consumer coproduction, which will shed light on the optimal mix of firm-supplied personalization and consumer-initiated customization. It would also be interesting to study platforms' incentives to provide purchase-based product recommendations, given that such recommendations may influence the competition between participating sellers. Future research can also analyze the effects of personalized advertising based on consumer purchase histories, and how market expansion and order-of-entry concerns affect firms' personalization strategies.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mktsci.pubs.informs.org/>.

Acknowledgments

This paper is based on Chapter 3 of the author's doctoral dissertation at the University of California, Berkeley. The author thanks Drew Fudenberg, Richard Gilbert, Teck-Hua Ho, Ganesh Iyer, John Morgan, Duncan Simester, Weiwei Tong, Glen Urban, J. Miguel Villas-Boas, and Birger Wernerfelt for helpful comments. The paper has also received

valuable feedback from seminar participants at the Haas School of Business and the Department of Economics at the University of California, Berkeley; and attendees of the 2006 Marketing Science Conference. The author thanks the editor, area editor, and reviewers for their great suggestions.

References

- Acquisti, A., H. R. Varian. 2005. Conditioning prices on purchase history. *Marketing Sci.* **24**(3) 367–381.
- Anderson, E. T., D. I. Simester. 2010. Price stickiness and customer antagonism. *Quart. J. Econom.* **125**(2) 729–765.
- Arora, N., X. Drèze, A. Ghose, J. D. Hess, R. Iyengar, B. Jing, Y. Joshi et al. 2008. Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Lett.* **19**(3–4) 305–321.
- Besanko, D., J.-P. Dubé, S. Gupta. 2003. Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Sci.* **49**(9) 1121–1138.
- Chen, Y. 1997. Paying customers to switch. *J. Econom. Management Strategy* **6**(4) 877–897.
- Chen, Y., G. Iyer. 2002. Consumer addressability and customized pricing. *Marketing Sci.* **21**(2) 197–208.
- Chen, Y., Z. J. Zhang. 2009. Dynamic targeted pricing with strategic consumers. *Internat. J. Indust. Organ.* **27**(1) 43–50.
- Chen, Y., C. Narasimhan, Z. J. Zhang. 2001. Individual marketing with imperfect targetability. *Marketing Sci.* **20**(1) 23–41.
- d'Aspremont, C., J. J. Gabszewicz, J.-F. Thisse. 1979. On Hotelling's "Stability in competition." *Econometrica* **47**(5) 1145–1150.
- Desai, P. S. 2001. Quality segmentation in spatial markets: When does cannibalization affect product line design? *Marketing Sci.* **20**(3) 265–283.
- Dewan, R., B. Jing, A. Seidmann. 2003. Product customization and price competition on the Internet. *Management Sci.* **49**(8) 1055–1070.
- Fudenberg, D., J. Tirole. 2000. Customer poaching and brand switching. *RAND J. Econom.* **31**(4) 634–657.
- Fudenberg, D., J. M. Villas-Boas. 2007. Behavior-based price discrimination and customer recognition. T. Hendershott, ed. *Economics and Information Systems*. Elsevier, Amsterdam, 377–436.
- Gerstner, E., J. Hess, W. Chu. 1993. Demarketing as a differentiation strategy. *Marketing Lett.* **4**(1) 49–57.
- Hauser, J. R., G. L. Urban, G. Liberali, M. Braun. 2009. Website morphing. *Marketing Sci.* **28**(2) 202–223.
- Hotelling, H. 1929. Stability in competition. *Econom. J.* **39**(153) 41–57.
- Moorthy, K. S. 1988. Product and price competition in a duopoly. *Marketing Sci.* **7**(2) 141–168.
- Pazgal, A., D. Soberman. 2008. Behavior-based discrimination: Is it a winning play, and if so, when? *Marketing Sci.* **27**(6) 977–994.
- Rossi, P. E., R. E. McCulloch, G. M. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* **15**(4) 321–340.
- Shaffer, G., Z. J. Zhang. 1995. Competitive coupon targeting. *Marketing Sci.* **14**(4) 395–416.
- Shannon, C. E., W. Weaver. 1949. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana.
- Shin, J., K. Sudhir. 2010. A customer management dilemma: When is it profitable to reward one's own customers? *Marketing Sci.* **29**(4) 671–689.
- Syam, N. B., N. Kumar. 2006. On customized goods, standard goods, and competition. *Marketing Sci.* **25**(5) 525–537.

- Syam, N. B., R. Ruan, J. D. Hess. 2005. Customized products: A competitive analysis. *Marketing Sci.* **24**(4) 569–584.
- Taylor, C. R. 2003. Supplier surfing: Competition and consumer behavior in subscription markets. *RAND J. Econom.* **34**(2) 223–246.
- Thisse, J.-F., X. Vives. 1988. On the strategic choice of spatial price policy. *Amer. Econom. Rev.* **78**(1) 122–137.
- Tyagi, R. K. 2000. Sequential product positioning under differential costs. *Management Sci.* **46**(7) 928–940.
- Villas-Boas, J. M. 1999. Dynamic competition with customer recognition. *RAND J. Econom.* **30**(4) 604–631.
- Villas-Boas, J. M. 2004. Price cycles in markets with customer recognition. *RAND J. Econom.* **35**(3) 486–501.
- Wernerfelt, B. 1986. Product line rivalry: Note. *Amer. Econom. Rev.* **76**(4) 842–844.
- Xiang, Y., M. Sarvary. 2007. News consumption and media bias. *Marketing Sci.* **26**(5) 611–628.