

**PAY TO SWITCH OR PAY TO STAY:
PREFERENCE-BASED PRICE DISCRIMINATION
IN MARKETS WITH SWITCHING COSTS**

GREG SHAFFER

*William E. Simon Graduate School of Business
University of Rochester
Rochester, NY 14627
shaffer@simon.rochester.edu*

Z. JOHN ZHANG

*Columbia Business School
Columbia University
New York, NY 10027
zz25@columbia.edu*

In many markets, firms can price discriminate between their own customers and their rivals' customers, charging one price to consumers who prefer their own product and another price to consumers who prefer a rival's product. We find that when demand is symmetric, charging a lower price to a rival's customers is always optimal. When demand is asymmetric, however, it may be more profitable to charge a lower price to one's own customers. Surprisingly, price discrimination can lead to lower prices to all consumers, not only to the group that is more elastic, but also to the less elastic group.

1. INTRODUCTION

In many markets firms can distinguish between their own customers and their rivals' customers. In these markets it may be possible for one or more firms to engage in price discrimination, whereby a firm charges one price to consumers who prefer its own product and another price to consumers who, all else equal, prefer a rival's product. Two important questions that arise in this case are: which group

A previous version of this paper was called "Winners and Losers in a Competitive Coupon Environment: Targeted Promotions Based on Brand Preferences." We thank Joe Harrington, Chakravarthi Narasimhan, Tore Nilssen, Peter Rossi, Michael Waldman, two anonymous referees, and the coeditor for helpful comments. We have also benefited from the comments of workshop participants at MIT, Northwestern, North Carolina, Washington University, Yale, the Conference on Pricing Decision Models in Boston (1994), and the EARIE meetings in Copenhagen (1998).

should receive the firm's lower price, and what is the overall effect on competition?

Invariably the advice given to firms in these markets is that they should charge lower prices to their rivals' customers. For example, Alsop (1985) suggests that firms should mail coupons directly to users of competitive brands, and Blattberg and Neslin (1990) argue that firms should place their discounts in magazines more likely to be read by a rival's customers. Rossi and Allenby (1993, p. 178) suggest that firms should offer discounts to households "that show loyalty toward other brands and yet are price sensitive," and Bester and Petrakis (1996, p. 228) find that firms should offer coupons to attract "some consumers who are otherwise more inclined to buy a competing brand." Chen (1997, p. 877) calls this optimal strategy "paying customers to switch," while Fudenberg and Tirole (1997) call it "customer poaching." Charging a lower price to the rivals' customers is also always optimal in Taylor (1998). The problem is that we sometimes see, to the contrary, firms charging lower prices to their own customers, a strategy that we call "paying customers to stay."

In this paper, we solve a simultaneous-move, differentiated-products model involving two price-setting firms and show, by contrast, that although a strategy of paying customers to switch may sometimes be optimal, at other times it may be optimal for a firm to pay customers to stay. In particular, we find that if only one firm can price-discriminate, or if demand is symmetric, paying customers to switch is always optimal. If both firms can price-discriminate and demand is asymmetric, it may be optimal for both firms to pay customers to switch, or it may be optimal for only one firm to charge a lower price to its rival's customers (in this case, the other firm charges a lower price to its own customers). It is never optimal for both firms to pay customers to stay.

We also consider how competition is affected when firms can discriminate in price between their own customers and their rival's customers. In particular, we ask what happens to prices, who gains and who loses, and whether preference-based price discrimination favors the larger or the smaller firm.

We find that the range of pricing outcomes extends beyond the usual case of third-degree price discrimination in which prices rise to one group of consumers (strong market) and fall to the other group (weak market). For example, depending on parameters, prices may fall to all consumers, leading to all-out competition. That is, in equilibrium, each firm's price to both groups of consumers may be lower than the uniform price it would have charged in the absence of any discrimination. Alternatively, it may be that some prices rise and some prices fall, as in the usual case, but that one firm's weak market is the

other firm's strong market, or the pricing outcomes may be mixed in that the larger firm lowers both of its prices while the smaller firm raises one price and lowers the other. It is never the case that both firms raise their prices to both consumer groups.

In terms of profit outcomes, we find that if only one firm can price-discriminate, that firm always gains at the expense of its rival. If both firms can price-discriminate and demand is symmetric, the outcome is a prisoner's dilemma (as all-out competition ensues). But, if demand is asymmetric, then depending on parameters, both firms may be worse off than they would be without any discrimination, or one or both firms may be better off. However, in all cases, the kind of preference-based price discrimination that we consider favors the firm with the smaller market share.

Markets where consumers incur switching costs provide a natural setting for firms to price-discriminate on the basis of consumer preferences. Casual observations in these markets suggest that some firms use a pay-to-switch strategy and others use a pay-to-stay strategy.¹ In telecommunications, for example, *AT&T* and *MCI* sometimes offer cash payments to induce consumers to switch long-distance carriers. In contrast, *Sprint* sometimes responds by offering a lower price to its own customers.² In financial services, credit-card companies often offer lower interest rates to customers of other bank cards (pay to switch), whereas in the software industry, *Microsoft*, *Symantec*, and *Adobe* often discriminate between consumers who currently use their software and those who use a competitor's version by charging lower upgrade prices to their existing customers (pay to stay).³

More generally, any market in which products are differentiated is subject to preference-based price discrimination if a firm has information on which customers prefer its own product and which prefer its rival's. There are many ways for a firm to obtain such information. One way is to ask consumers in a telemarketing or direct-mail survey.⁴ Another way is to purchase it from one of several companies

1. We abstract from the different sources of switching costs. For example, the switching costs associated with changing long-distance phone carriers may be different from the switching costs associated with having to learn a new software package. See Klemperer (1987) and Nilssen (1992) for more on this distinction.

2. See "Sprint Phone Rebate Program Pays Customers to Stay Put," *Houston Chronicle*, August 29, 1995. "In essence *Sprint* is doing the same thing other carriers do—mail checks to customers—however, *Sprint* will send checks to existing customers to reward loyalty rather than as an inducement to switch" (Richard Klugman, p. B3).

3. However, this is not always the case. In an advertisement in the June 1997 issue of *PC Magazine*, for instance, we found *Adobe Illustrator 7.0* listed at \$89.16 for an upgrade and \$79.80 for a competitive upgrade (after the \$50 rebate). In this instance, *Adobe* is charging a lower price to its competitor's customers.

4. In a recent mailing we were urged to fill out a personal product preference survey: "Fill out this survey now! Major manufacturers really need to know which

specializing in this kind of information.⁵ A firm can also pay to have magazine and newspaper readership data cross-tabulated by product usage,⁶ or infer a customer's preferences from where he or she lives. *Domino's Pizza* and *Stop & Shop*, for instance, often target promotions in residential areas located closer to their rivals' stores.⁷

Information on consumer preferences can also be obtained at the point of sale. Coupon dispensing machines, which reside at the checkout counters in thousands of grocery stores nationwide, allow firms to offer discounts that are based on consumers' current purchases. To encourage consumers to switch sports drinks, for instance, *Gatorade Inc.* can print out a coupon for its own brand whenever a competing sports drink is purchased. Other examples of firms following pay-to-switch strategies include: *Coca-Cola* giving a discount on *Diet Coke* when *Diet Pepsi* is purchased, *Chesebrough Pond* giving a discount on *Mentadent Toothpaste* when *PeroxiCare* is purchased, and *Ortega Foods* giving a discount on its taco shells when *Old El Paso* is purchased. Although less common, some firms follow a pay-to-stay strategy at the checkout counter; for example, *Pepsi* often responds to *Coca-Cola's* promotion by giving discounts to its own customers (purchasers of *Diet Pepsi*).

The pay-to-stay strategies that one observes in reality seem to be counterintuitive. Why then do some firms adopt this strategy? Our results suggest that firms do so to compete effectively for the consumer group that, *from its perspective*, has the higher price elasticity of demand. Sometimes this group is the same for both firms, as when one firm pays customers to switch and another firm pays customers to stay. At other times, this group differs by firm, as when each firm pays customers to switch. Thus, for example, our results suggest that *Pepsi's* profit-maximizing strategy, given that *Coca-Cola* has a pay-to-switch strategy, is to give discounts to its own customers (*Diet Pepsi* customers) if, from its perspective, its own customers have a higher price elasticity of demand.

products you like and actually use so they can offer you savings on the right kind of products. This new survey of your personal preferences will provide that kind of vital information."

5. *Computerized Marketing Technologies* claims to have surveyed millions of consumers through product preference surveys. According to its managing partner, "We know who are current users, who are competitive users, and who are non-users." See "Stalking the Customer," *Business Week*, August 28, 1989, pp. 54-62.

6. There are several marketing companies that run these cross-tabulations. See Lehmann (1985, pp. 263-265).

7. In a conversation with an employee at *Domino's Pizza*, we were told that management often instructs employees to deliver coupons door to door, targeting only those apartment complexes that are located closer to their *Pizza Hut* rival. Zeldis (1987) reports a similar strategy by *Stop & Shop* stores.

We arrive at these conclusions by allowing for asymmetry in switching costs,⁸ a factor that is not modeled in the above-cited literature on preference-based price discrimination.⁹ In particular, we allow for the possibility that it might on average be more (or less) costly for firm 1's customers to switch in the current period to firm 2 than vice versa, and we show that it is the ratio of these switching costs that determine the relative price elasticities of demand between groups, and hence whether a firm wants to adopt a pay-to-stay or a pay-to-switch strategy.

There is some evidence to support our finding that a firm should charge a lower price to the consumer group that has the higher elasticity of demand from its perspective. According to Information Resources, Inc. (IRI), consumers of *Diet Coke* tend to be more loyal than those of *Diet Pepsi*, implying that the latter consumers have a higher elasticity of demand.¹⁰ This is consistent with *Coca-Cola's* strategy of giving discounts to consumers of *Diet Pepsi* and *Pepsi* responding by also giving discounts to consumers of *Diet Pepsi*. In addition, IRI reports that consumers of *Mentadent* toothpaste are more loyal than those of *PeroxiCare*, consumers of *Ortega Foods* taco shells are more loyal than those of *Old El-Paso*, and consumers of *Gatorade* are more loyal than those of other sports drinks, which is consistent with *Gatorade Inc*, *Chesebrough Pond*, and *Ortega Foods* using pay-to-switch strategies. In the case of *Microsoft's* pay-to-stay strategy, we note that it is perhaps much less costly for *Microsoft's* customers to switch to a competitor than vice versa. The low switching costs can result from the fact that there is a large installed base of *Microsoft* software, which encourage *Microsoft's* competitors to develop very good translators that allow easy importation of *Microsoft* files, whereas *Microsoft* may have relatively poor translators (less non-*Microsoft* software suggests reduced benefit from good translators). Alternatively, it may be that *Microsoft* perceives its own customers to have a much higher elasticity of demand because it thinks that current users will receive a smaller

8. We define "switching costs" broadly to include *ex ante* product differentiation. Thus, if a consumer prefers brand 1 over brand 2 at the same price, the reduction in utility from consuming brand 2 is a switching cost; if a consumer lives closer to store A than store B, the additional travel to store B is a switching cost.

9. In this literature, switching costs are assumed symmetric across consumers; regardless what product was purchased in the preceding period, the distribution of switching costs in the current period is the same.

10. IRI measures a product's brand loyalty by its share-of-category-requirements. This measure quantifies the "inertia" in consumers' purchases over the course of a year, and is commonly used by both marketing scholars and practitioners. The measure is reported annually in *The Marketing Fact Book* for major brands in over 700 categories of consumer goods products, and is calculated based on the households panel data that IRI maintains.

incremental benefit from the upgrade than will non-*Microsoft* customers (who are using a product presumed by *Microsoft* to be inferior).

In addition to the above-cited literature, our work is related to that of Lederer and Hurter (1986), Thisse and Vives (1988), and Shaffer and Zhang (1995, 1997), all of whom consider situations in which firms can tailor their promotions to each individual consumer. In this literature, each firm knows the exact location of each consumer along a Hotelling line and can customize prices one to one. We differ in this paper in that we assume that firms are unable to distinguish among individual consumers, identifying them only as either one's own customers or one's rival's customers.

Our work is also related that of Corts (1998), who considers a general model of oligopolistic price discrimination with asymmetric demands. Corts derives necessary conditions for when all-out competition (prices fall to all consumers) may occur, and shows that these conditions do not preclude a situation in which both firms offer a discount to the same group in the discriminatory equilibrium [as opposed to the earlier literature, e.g., Bester and Petrakis (1996) and Chen (1997), which assumes symmetric demands and where all-out competition is always associated with both firms adopting a pay-to-switch strategy]. Although Corts does not explicitly consider the case of preference-based price discrimination, his example of third-degree price discrimination involving firms with products of different qualities can be interpreted as such if one thinks of the choosy (cheap) consumers as being the high-quality (low-quality) firm's consumers. However, whereas in his example the low-quality firm follows a pay-to-stay strategy and prices fall to all consumers, we show that equilibria in which one firm adopts a pay-to-stay strategy are not always associated with lower prices for all consumers. In specifying a particular functional form of demand, we are also able to provide insight into what determines a firm's optimal pricing strategy and how the equilibrium prices and profits of the two firms depend on firm size and relative switching costs or brand loyalties.

Finally, our work is related to the literature on endogenously created switching costs (e.g., Banerjee and Summers, 1987; Caminal and Matutes, 1990), which points out that it may be in a firm's interest to reward continuing patronage by offering frequent-flyer bonuses, coupons, and rebates. In this literature, loyalty-inducing arrangements work over time. Here, we focus on the properties of price discrimination in a static model. To illustrate, suppose a firm has initiated some reward program in the past that is still in effect at the time the consumer must decide whether to switch products. If the program was successfully crafted, it will have created an opportunity cost to each buyer of switching brands. Our analysis takes the distribution of

these switching costs as given and then asks how competition today is affected if the firms can also price-discriminate.

Section 2 describes the model and defines notation. Section 3 solves a benchmark game in which each firm is constrained to charge a single price. It then considers games in which at least one firm can price-discriminate. Sections 4 and 5 discuss the model's implications for prices and profits. Section 6 concludes. The proofs of all lemmas and propositions can be found in the appendix.

2. MODEL AND NOTATION

Suppose two firms—*A* and *B*—sell competing brands of a good that is produced at constant marginal cost c ; individual consumers have unit demands, and no consumer buys from both firms. To simplify the exposition, we assume that competition between the two firms will cause equilibrium prices to be below the level at which individual consumers would drop out of the market.¹¹ Thus, we assume that the market elasticity of demand is zero. Nevertheless, as we will show, price discrimination will still be profitable, because cross-price elasticities will in general differ among consumers. Our results extend to the case of nonzero market demand elasticity, as we discuss in the conclusion.

We partition consumers into two groups of unequal size. Let the larger size group, group α , denote those consumers (a fraction $\theta \in [\frac{1}{2}, 1]$ of all consumers) who would buy from firm *A* if prices for the two brands were the same. Let group β denote the remaining consumers, those who would buy from firm *B* at equal prices. We refer to θ and $1 - \theta$ as the baseline market shares of firms *A* and *B* respectively. These are the market shares that would obtain if prices for the two brands were the same. Actual market shares, however, may differ from these baseline market shares, depending on the prices set by each firm and the distribution of switching costs.

2.1 ASYMMETRIC DEMANDS

We define a consumer's switching cost as the minimum price differential necessary to induce the consumer to switch to the competing brand. One can also think of the switching cost as a measure of the consumer's brand loyalty, and thus we do not distinguish

11. These assumptions best approximate demand in the service industries, e.g., the telecommunications and credit-card industries. See also the discussion in Fudenberg and Tirole (1997).

between markets in which the products are *ex ante* homogeneous and only become differentiated once a product is purchased, and markets in which, for instance, consumers have locational preferences about where to shop.

Let P_i , $i = A, B$, denote the price charged by firm i to group α consumers. Then the brand choice of a consumer in group α with a loyalty of l is given by

$$P_A \leq P_B + l \Rightarrow \text{choose brand } A,$$

$$P_A > P_B + l \Rightarrow \text{choose brand } B.$$

Similarly, let \tilde{P}_i denote the price charged by firm i to group β consumers. Then the brand choice of a consumer in group β with a loyalty of l is given by

$$\tilde{P}_B \leq \tilde{P}_A + l \Rightarrow \text{choose brand } B,$$

$$\tilde{P}_B > \tilde{P}_A + l \Rightarrow \text{choose brand } A.$$

Define $F_k(x)$ as the fraction of consumers in group k , $k = \alpha, \beta$, with loyalty less than or equal to x . We assume the distribution of loyalties is uniform over $[0, l_k]$; thus $F_k(x)$ can be written as

$$F_k(x) = \begin{cases} 0 & \text{if } x < 0, \\ x/l_k & \text{if } 0 \leq x \leq l_k, \\ 1 & \text{if } x > l_k. \end{cases}$$

Given our assumptions, $F_k(x)$ has a simple interpretation. For example, the fraction of consumers in group α who will buy firm A 's brand if firm A charges a price premium x over firm B is $1 - F_\alpha(x)$, and the fraction of consumers in group α who will buy firm B 's brand is $F_\alpha(x)$.¹²

Normalizing the total number of consumers in the market to one, we can express each firm's overall demand as a weighted average of its demand from each consumer group:

$$D_A(P_A, \tilde{P}_A, P_B, \tilde{P}_B) = \theta[1 - F_\alpha(P_A - P_B)] + (1 - \theta)F_\beta(\tilde{P}_B - \tilde{P}_A), \quad (1)$$

$$D_B(P_A, \tilde{P}_A, P_B, \tilde{P}_B) = \theta F_\alpha(P_A - P_B) + (1 - \theta)[1 - F_\beta(\tilde{P}_B - \tilde{P}_A)]. \quad (2)$$

12. In contrast, in previous work it is often assumed either that all consumers are identical (all have the same cost of switching), or that the distribution of switching costs is symmetric for the consumers of each product.

Note that each firm can induce some amount of brand switching among consumers who would otherwise prefer its rival's brand if and only if its price to them is less than its rival's price. The amount of switching that occurs will depend on the price differential and average loyalty within each group; the larger is l_k , the less brand switching there will be for any given difference in price.

Our framework encompasses the Hotelling (1929) model of product differentiation as a special case. In that model, the parameter t often denotes the unit transportation cost. In the model here, an analogous role is played by the loyalty parameters l_α and l_β . For example, if demand is symmetric, so that $l_\alpha = l_\beta = t$ and $\theta = \frac{1}{2}$, then the demands simplify to the familiar Hotelling demands of $(t + P_B - P_A)/2t$ for firm A and $(t + P_A - P_B)/2t$ for firm B . More generally, though, our framework extends the Hotelling model by allowing for asymmetries in demand.

2.2 GAME FORM

To complete the model, we now describe the games we consider and make some assumptions about each firm's information set. In particular, we consider both a game in which neither firm is able to price-discriminate and a game in which both firms can price-discriminate. We adopt Nash equilibrium as our solution concept and refer to the equilibrium prices as the Bertrand prices.¹³

Formally, the set of strategies available to firm i if it can identify individual consumers by group is $\Sigma_i = \{(P_i, \tilde{P}_i) | P_i, \tilde{P}_i \geq 0\}$. One can think of each firm as literally specifying two prices, or as specifying a regular price and simultaneously offering a discount to one group. Under the latter interpretation, a firm's regular price is defined as $\max(P_i, \tilde{P}_i)$, and its discounted price is defined as $\min(P_i, \tilde{P}_i)$. The size of the firm's discount is then given by $|P_i - \tilde{P}_i|$. The set of strategies available to firm i if consumers are indistinguishable by group is $\Sigma_i = \{(P_i, \tilde{P}_i) | \tilde{P}_i = P_i \geq 0\}$.

13. With the exception of Nilssen (1992) and Taylor (1998), who allow for an arbitrary number of periods, much of the literature looks at two-period models in which consumers purchase one unit of a product each period. In these models, consumers' purchases in the first period reveal their preferences and determine market shares at the start of the second period; in the second period firms can choose to offer discounts to either their own or their rival's first-period customers. In this paper, we simplify by examining a one-period model only, thus abstracting from any prior competition. One can think of firms as inferring a customer's preferences either from where he or she lives, or through his or her self-revelation via customer surveys.

Each firm's profit as a function of all possible strategy combinations is

$$\begin{aligned} \Pi_A(\sigma_A, \sigma_B) = & \theta(P_A - c)[1 - F_\alpha(P_A - P_B)] \\ & + (1 - \theta)(\tilde{P}_A - c)F_\beta(\tilde{P}_B - \tilde{P}_A). \end{aligned} \quad (3)$$

$$\begin{aligned} \Pi_B(\sigma_A, \sigma_B) = & \theta(P_B - c)F_\alpha(P_A - P_B) \\ & + (1 - \theta)(\tilde{P}_B - c)[1 - F_\beta(\tilde{P}_B - \tilde{P}_A)], \end{aligned} \quad (4)$$

where $\sigma_i \in \Sigma_i$ denotes firm i 's strategy, $i = A, B$. Note that group α consumers are more loyal on average than consumers in group β if and only if $l_\alpha > l_\beta$. We place no restrictions on l_α, l_β , and θ other than what is required below to ensure the existence of pure-strategy Nash equilibria.

The expressions in (3) and (4) implicitly contain two simplifying assumptions, which we make for ease of exposition. The first is that all consumers who buy from a given firm will avail themselves of any discounts offered to them, i.e., all consumers within group who buy a firm's product will pay the same price.¹⁴ The second is that, conditional on price discrimination being feasible, there is no cost in specifying two different prices, or in specifying a regular price and offering a discount. This implies that firms can price-discriminate without regard to their relative sizes.¹⁵

3. SOLVING THE MODEL

We begin with the benchmark game in which neither firm can price-discriminate. Thus, we solve for a Nash equilibrium in which $\tilde{P}_k = P_k$ is imposed. To simplify the profit expressions in (3) and (4), it is useful to determine what must be true of prices in equilibrium. There are three possibilities to consider: Firm A has a higher Bertrand price, the same Bertrand price, or a lower Bertrand price than its rival. Intuitively, we might expect firm A to have the higher Bertrand price because of its larger baseline market share. This intuition is confirmed in the following lemma.

14. Alternatively, we could keep track of which consumers take advantage of discounts and which do not. It would then be optimal for each firm to choose a regular price (for those consumers in the two groups who do not take advantage of discounts) and *two* promotional prices (for those consumers in each group who do take advantage of discounts). The main qualitative results would be unchanged.

15. Alternatively, each firm might incur a fixed cost of offering discounts that is proportional to a group's baseline market share, i.e., we could assume a cost $z(\theta)$ to target consumers in group α , where $z(0) = 0$ and $z' \geq 0$. However, as long as these fixed costs were small enough, the main qualitative results would be unchanged.

LEMMA 1: *In any pure-strategy Nash equilibrium to the game in which neither firm can price-discriminate, firm A has the weakly higher Bertrand price (strictly higher if $\theta > \frac{1}{2}$).*

From Lemma 1, we know that in any pure-strategy Nash equilibrium, profits can be written as

$$\Pi_A = \theta(P_A - c) \frac{l_\alpha + P_B - P_A}{l_\alpha}, \tag{5}$$

$$\Pi_B = (1 - \theta)(P_B - c) + \theta(P_B - c) \frac{P_A - P_B}{l_\alpha}. \tag{6}$$

Differentiating (5) and (6) with respect to each firm’s price and then setting the resulting expressions equal to zero, we obtain the candidate Nash equilibrium prices and corresponding profits¹⁶:

$$\begin{aligned} P_A^* &= \frac{1 + \theta}{3\theta} l_\alpha + c, & P_B^* &= \frac{2 - \theta}{3\theta} l_\alpha + c, \\ \Pi_A^* &= \frac{(1 + \theta)^2}{9\theta} l_\alpha, & \Pi_B^* &= \frac{(2 - \theta)^2}{9\theta} l_\alpha. \end{aligned} \tag{7}$$

It must still be verified, using the profit functions in (3) and (4), that neither firm can profitably deviate. In Appendix B, we show that this requires an upper bound, $h^*(\theta)$, on l_β/l_α to ensure that firm B does not want to deviate from its strategy of pricing lower than firm A, and it requires a lower bound, $g^*(\theta)$, on l_β/l_α to ensure that firm A does not want to undercut firm B.

PROPOSITION 1: *A pure-strategy Nash equilibrium to the game in which neither firm can price-discriminate exists if and only if θ and l_β/l_α are such that*

$$g^*(\theta) = \frac{(1 - \theta)(2 - \theta)^2}{9\theta^3} \leq \frac{l_\beta}{l_\alpha} \leq \frac{(1 + \theta)^2}{9(1 - \theta)\theta} = h^*(\theta).$$

The unique equilibrium prices are given by (P_A^*, P_B^*) with corresponding profits (Π_A^*, Π_B^*) .

At $\theta = \frac{1}{2}$, the Bertrand prices simplify to $l_\alpha + c$, as in the Hotelling model, and at $\theta = 1$, the Bertrand prices simplify to $\frac{2}{3}l_\alpha + c$ for the higher-quality firm and $\frac{1}{3}l_\alpha + c$ for the lower-quality firm, as in Tirole

16. The functions in (5) and (6) are strictly concave in their own prices, thus satisfying the second-order conditions. Uniqueness of the candidate Nash prices is ensured, since $|\partial^2 \Pi_i / \partial P_i^2| - |\partial^2 \Pi_i / \partial P_i \partial P_j| \geq 0$ for all $P_i, P_j, i, j = A, B, i \neq j$.

(1989, p. 296). More generally, the equilibrium prices (and profits) for both firms are decreasing in firm A 's baseline market share. Intuitively, firm B is indifferent at the margin between setting a higher price to extract more from the group β consumers and setting a lower price to attract additional consumers from group α . The fewer group β consumers there are to exploit, the lower firm B 's price will be. As a result, firm A 's price will be lower as well.

Equilibrium prices and profits are increasing in l_α . Intuitively, as l_α increases, it becomes more difficult for firm B to attract additional group α consumers. This decreases firm B 's incentive to lower its price, causing both firms to compete less aggressively. It may be surprising that equilibrium prices and profits are independent of l_β , but this is a necessary consequence of the fact that the marginal consumer belongs to group α , i.e., firm A does not sell to group β consumers.

3.1 SOLVING THE MODEL WHEN BOTH FIRMS CAN PRICE-DISCRIMINATE

We now solve the model for the game in which both firms can price-discriminate. There are four possibilities to consider: Firm A has the higher Bertrand price in (a) both groups, (b) neither group, (c) group β only, or (d) group α only. Intuitively, we would expect each firm to have the higher Bertrand price in its home market. This intuition is confirmed in the following lemma.

LEMMA 2: *In any pure-strategy Nash equilibrium of games in which at least one firm can price-discriminate, each firm will sell to both consumer groups. That is, $P_A > P_B$ and $\tilde{P}_B > \tilde{P}_A$.*

Lemma 2 implies that *each* firm makes strictly positive sales to consumers in both groups in any pure-strategy Nash equilibrium.¹⁷ This allows us to write the profit functions in (3) and (4) as

$$\Pi_A = \theta(P_A - c) \frac{l_\alpha + P_B - P_A}{l_\alpha} + (1 - \theta)(\tilde{P}_A - c) \frac{\tilde{P}_B - \tilde{P}_A}{l_\beta}, \quad (8)$$

$$\Pi_B = \theta(P_B - c) \frac{P_A - P_B}{l_\alpha} + (1 - \theta)(\tilde{P}_B - c) \frac{l_\beta + \tilde{P}_A - \tilde{P}_B}{l_\beta}. \quad (9)$$

Differentiating (8) with respect to firm A 's two prices (P_A, \tilde{P}_A), and (9) with respect to firm B 's two prices (P_B, \tilde{P}_B), and then setting the

17. Thus, brand switching occurs in equilibrium. As pointed out by others, e.g., Fudenberg and Tirole (1997), this is inefficient, since some consumers end up buying from a firm whose product they like less.

resulting four expressions equal to zero, we obtain the candidate Nash equilibrium prices and corresponding profits¹⁸:

$$\begin{aligned} (P_A^{**}, \tilde{P}_A^{**}) &= \frac{2}{3}l_\alpha + c, \frac{1}{3}l_\beta + c, & (P_B^{**}, \tilde{P}_B^{**}) &= \frac{1}{3}l_\alpha + c, \frac{2}{3}l_\beta + c, \\ \Pi_A^{**} &= \frac{4}{9}\theta l_\alpha + \frac{1}{9}(1 - \theta)l_\beta, & \Pi_B^{**} &= \frac{1}{9}\theta l_\alpha + \frac{4}{9}(1 - \theta)l_\beta. \end{aligned} \quad (10)$$

Since the two markets are completely separated when both firms can price-discriminate, it remains only to check that neither firm would want to deviate by having a strictly lower price in its home market. This is easy to show, implying that the price pairs in (10) form a Nash equilibrium.

PROPOSITION 2: *There exists a unique pure-strategy Nash equilibrium to the game in which both firms can price-discriminate. Equilibrium prices are $(P_A^{**}, \tilde{P}_A^{**})$ and $(P_B^{**}, \tilde{P}_B^{**})$. Firm A's profit is Π_A^{**} , and firm B's profit is Π_B^{**} .*

Solving the game in which both firms can price-discriminate is similar to solving two separate games with vertically differentiated products in which each firm has the higher-quality product in one of the markets. Brand switching occurs in equilibrium because the firm with the higher-quality product exploits its advantage in quality by charging a price in excess of marginal cost, thus providing an opportunity for its rival to charge less and capture additional sales. Having the ability to set two prices is key to this result, since it enables each firm to charge a price to its competitor's customers that is independent of what it charges its own customers.

Note that the complete separation between markets means that although equilibrium prices within each group are increasing in that group's average loyalty, they do not depend on the average loyalty of consumers outside the group or on baseline market shares. If $l_\alpha \geq l_\beta$, firm A earns strictly higher profit because of its greater baseline market share. If l_β is sufficiently large relative to l_α , however, it is possible for firm B to earn the higher profit despite its smaller market share.

Depending on the actual values of l_α and l_β , profits may be increasing or decreasing in θ . If $l_\alpha > 4l_\beta$, profits for both firms are increasing in θ . If $l_\beta/4 < l_\alpha < 4l_\beta$, only firm A's profit is increasing in θ . If $l_\alpha < l_\beta/4$, profits for both firms are decreasing in θ . Intuitively, each firm's profit is a weighted average of its profit from the

18. The profit functions in (8) and (9) are strictly concave in their own prices, thus satisfying the second-order conditions. Uniqueness of the candidate Nash prices is ensured, since the two markets are separable, and within each market, $|\partial^2 \Pi_i / \partial P_i^2| - |\partial^2 \Pi_i / \partial P_i \partial P_j| \geq 0$ and $|\partial^2 \Pi_i / \partial \tilde{P}_i^2| - |\partial^2 \Pi_i / \partial \tilde{P}_i \partial \tilde{P}_j| \geq 0$ for all $P_i, P_j, \tilde{P}_i, \tilde{P}_j, i, j = A, B, i \neq j$.

two consumer groups. As the relative weight changes, one or both firms may gain. To illustrate, suppose group α were to increase in size. Clearly, firm A would be more likely to benefit from such a change, as its market share would increase (recall prices are unaffected by base-line market shares). But it is possible for firm B to gain as well, if the relative brand loyalty of group α consumers is sufficiently large. In this case, the decrease in firm B 's market share would be outweighed by the higher per-unit profit it could make from sales to group α consumers (assuming average brand loyalty doesn't change). The converse is that both firms are worse off if the relative brand loyalty of group α consumers is sufficiently small. Over a large region of parameter values, however, an increase in θ will benefit firm A and harm firm B .

4. EQUILIBRIUM PRICING STRATEGIES

In this section, we consider each firm's optimal pricing strategy: should it charge a lower price to its own customers or its rival's customers? We say that firm A 's optimal pricing strategy is pay-to-switch if $P_A^{**} > \tilde{P}_A^{**}$ and pay-to-stay otherwise. Similarly, we say that firm B 's optimal pricing strategy is pay-to-switch $\tilde{P}_B^{**} > P_B^{**}$ and pay-to-stay otherwise. Table I summarizes our results.

It seems intuitive that firms should always pay customers to switch, since, by definition, a rival's customers value a firm's product less than the firm's own customers. However, as Table I shows, this strategy is optimal for both firms if and only if $\frac{1}{2} \leq l_\beta/l_\alpha \leq 2$, i.e., if the average loyalties of the two groups are not too dissimilar. If consumers in group α have an average loyalty that is more than twice as large as the average loyalty of the consumers in group β , then firm B should pay customers to stay and firm A should pay customers to switch. Conversely, if the consumers in group β have an

TABLE I.
OPTIMAL PRICING IF BOTH FIRMS CAN
PRICE-DISCRIMINATE

Relative Loyalty l_β/l_α	Regular Price		Amount of Discount		Group with Discount	
	Firm A	Firm B	Firm A	Firm B	Firm A	Firm B
$l_\beta/l_\alpha \geq 2$	$\frac{1}{3}l_\beta + c$	$\frac{2}{3}l_\beta + c$	$\frac{1}{3}(l_\beta - 2l_\alpha)$	$\frac{1}{3}(2l_\beta - l_\alpha)$	Group α	Group α
$\frac{1}{2} \leq l_\beta/l_\alpha < 2$	$\frac{2}{3}l_\alpha + c$	$\frac{2}{3}l_\beta + c$	$\frac{1}{3}(2l_\alpha - l_\beta)$	$\frac{1}{3}(2l_\beta - l_\alpha)$	Group β	Group α
$l_\beta/l_\alpha < \frac{1}{2}$	$\frac{2}{3}l_\alpha + c$	$\frac{1}{3}l_\alpha + c$	$\frac{1}{3}(2l_\alpha - l_\beta)$	$\frac{1}{3}(l_\alpha - 2l_\beta)$	Group β	Group β

average loyalty that is more than twice as large as the average loyalty of the consumers in group α , then firm A 's optimal strategy is pay-to-stay and firm B 's optimal strategy is pay-to-switch. Intuitively, each firm's equilibrium price to each group is increasing in that group's average loyalty. If one group's average loyalty is very low relative to the other group's average loyalty, then the consumers of the group whose loyalty is low can be much more easily switched, and the home firm may have no choice but to defend its home market more aggressively.

Note that Table I implies that at least one firm will pay customers to switch, but the identity of this firm cannot be determined by firm size alone. The smaller firm, for example, may instead find it optimal to charge a lower price to its own customers, a finding that runs contrary to the conjecture in Bester and Petrakis (1996, p. 238) that "smaller firms have a higher incentive to distribute coupons because they can gain more by attracting customers from other market segments."

In order to forge a deeper understanding of the difference in results between our paper and previous literature, and why we find that paying customers to stay is sometimes optimal, it is useful to consider the differences in (cross-price) elasticities of demand between the two groups.

Formally, let ϵ_B and $\tilde{\epsilon}_B$ denote firm B 's elasticity of demand among group α consumers and group β consumers respectively, according to the standard definition (percentage change in sales divided by the percentage change in price):

$$\epsilon_B \equiv \left| \frac{dF_\alpha(x)}{dP_B} \frac{P_B}{F_\alpha(x)} \right|, \quad \tilde{\epsilon}_B \equiv \left| \frac{d[1 - F_\beta(x)]}{d\tilde{P}_B} \frac{\tilde{P}_B}{1 - F_\beta(x)} \right|. \tag{11}$$

Since Lemma 2 requires $\tilde{P}_B > \tilde{P}_A$ and $P_A > P_B$, firm B 's demand in each group can be written as

$$F_\alpha(x) = \frac{P_A - P_B}{l_\alpha}, \quad 1 - F_\beta(x) = \frac{l_\beta + \tilde{P}_A - \tilde{P}_B}{l_\beta}.$$

Substituting $F_\alpha(x)$ and $1 - F_\beta(x)$ into the definitions in (11) and simplifying yields

$$\epsilon_B = \frac{P_B}{P_A - P_B}, \quad \tilde{\epsilon}_B = \frac{\tilde{P}_B}{l_\beta + \tilde{P}_A - \tilde{P}_B}. \tag{12}$$

Substituting firm A 's Bertrand prices ($P_A^{**}, \tilde{P}_A^{**}$) into the elasticity expressions in (12), we obtain

$$\epsilon_B^{**} = \frac{-P_B}{\frac{2}{3}l_\alpha + c - P_B}, \quad \tilde{\epsilon}_B^{**} = \frac{-\tilde{P}_B}{\frac{4}{3}l_\beta + c - \tilde{P}_B}. \quad (13)$$

Comparing ϵ_B^{**} in (13) with $\tilde{\epsilon}_B^{**}$, we see that firm B 's price elasticity of demand at equal prices is higher among group α consumers ($\epsilon_B^{**} > \tilde{\epsilon}_B^{**}$) if and only if $2l_\beta > l_\alpha$. If this condition holds, profit maximization requires firm B to pay customers to switch. On the other hand, if $2l_\beta < l_\alpha$, demand is more elastic among its own consumers, and thus firm B should pay customers to stay. Note that these cutoff levels are the same as in Table I. The analysis for firm A proceeds similarly.

PROPOSITION 3: *In the game in which both firms can price-discriminate, each firm should charge a lower price to the consumer group in which it faces the more elastic demand. Both firms will pay to switch if and only if $\frac{1}{2} \leq l_\beta/l_\alpha \leq 2$. If $l_\beta/l_\alpha > 2$, firm A will pay to stay. If $l_\beta/l_\alpha < \frac{1}{2}$, firm B will pay to stay. It is never optimal for both firms to give discounts to their own customers.*

The common intuition for why a firm should always charge a lower price to its rival's customers is that, in doing so, it can generate profitable incremental sales without giving up any profit margin on its own customers. That is, intuition suggests that a firm can costlessly increase its market share, and hence profit, by following a pay-to-switch strategy. However, this intuition, and the literature on which it is based, either (a) assumes that demands are symmetric or (b) implicitly assumes that the rival firm cannot itself price-discriminate. We consider each in turn.

Suppose demands are symmetric, so that the distribution of loyalties is the same in each group ($l_\alpha = l_\beta$). Then we can see from (13) that firm B 's elasticity of demand at equal prices is always higher among group α consumers ($\epsilon_B^{**} > \tilde{\epsilon}_B^{**}$). Analogous reasoning holds for firm A . This implies that, under symmetry, each firm should indeed always charge a lower price to its rival's customers.

The symmetry that is required here is defined with respect to the aggregate demand (across both groups) that each firm faces. Note that the demands within each group are not symmetric. This within-group asymmetry is what allows the firms to differ in their ranking of consumer groups by elasticity of demand. From (12) and its analog, we see that, at a given rival's price, the strong market of firm A is always the weak market of firm B and vice versa. Thus, to use

Corts's (1998) terminology, our model exhibits *best-response asymmetry*, which turns out to be necessary for price discrimination to have unambiguous price effects (see Corts's Proposition 4, p. 315).

The intuition that a firm should always charge a lower price to its rival's customers also holds when only one firm can price-discriminate. To see this, suppose that only firm *B* can price-discriminate. Then, for this firm, the claim is that charging a lower price to group α consumers is always optimal. Since Lemma 2 applies to all games in which at least one firm can price-discriminate (we prove this in Appendix C), the substitutions we made to derive (12) hold, and the result follows immediately from noting that $P_A = \hat{P}_A$ when firm *A* cannot price-discriminate. Analogously, a pay-to-switch strategy is also optimal in the game in which only firm *A* can price-discriminate.

The general result that each firm should charge a lower price to the consumer group in which it faces the more elastic demand holds independently of the size of the two consumer groups (baseline market shares) and relative brand loyalties. We can best illustrate its implications with a concrete example involving the two pizza competitors, *Pizza Hut* and *Domino's Pizza*. Suppose *Pizza Hut* has the ability to price-discriminate between its own customers and *Domino's* customers, but that *Domino's Pizza* does not yet have this ability. Then, we have just shown that it is optimal for *Pizza Hut* to target *Domino's* customers with discounts. Now suppose that both firms have the ability to price-discriminate and that all consumers purchase delivered pizza. In this case, we might expect switching costs (increased delivery time for acquiring a more-distant pizza) to be symmetrical, implying that each firm should pay customers to switch. On the other hand, suppose *Domino's* delivers, but *Pizza Hut* does not. Then, because *Domino's* delivers, there is minimal cost for customers close to *Pizza Hut* to switch to *Domino's Pizza*, while the cost for *Domino's* customers (who live far from *Pizza Hut*) to switch to dining in at *Pizza Hut* is significant. In this case, it may be that *Pizza Hut* should pay customers to stay, while *Domino's Pizza* should pay customers to switch.

Our model has three additional implications for a firm's optimal pricing strategy. First, in comparing the first two columns in Table I, we see that firm *A*'s regular price exceeds firm *B*'s regular price if and only if $l_\alpha > l_\beta$. Similarly, in comparing the first and third columns in Table I, we see that the same is true of the size of each firm's profit-maximizing discount, i.e., $|P_A^{**} - \hat{P}_A^{**}| > |P_B^{**} - \hat{P}_B^{**}|$ if and only if $l_\alpha > l_\beta$. Thus, Table I implies that the firm with the higher regular price will offer the larger discount (e.g., *AT&T* will offer a larger discount than *MCI*).

Second, in comparing a firm's regular price with its optimal pricing strategy, the second and fourth columns of Table I, we see that the firm with the higher regular price always pays customers to switch (or, alternatively stated: if a firm's optimal pricing strategy is pay-to-stay, it must have the lower regular price). This is not surprising, since otherwise the firm with the higher regular price would not have positive sales in both markets, contradicting Lemma 2. However, the converse is not true; depending on parameters, the firm with the lower regular price may either want to pay customers to switch (this is *MCI's* strategy) or pay customers to stay (this is *Sprint's* strategy).

Third, in comparing each firm's discount and its optimal pricing strategy, the third and fourth columns of Table I, we see that if each firm offers a discount to the same consumer group, the firm that is paying customers to switch will have the higher discount (e.g., in the *Diet Coke vs. Diet Pepsi* example, the discount on *Diet Coke* was \$.50 and the discount on *Diet Pepsi* was \$.25). This observation is the result of two factors. First, we know that the firm that is paying customers to switch must have the higher regular price, and second, we know that its discounted price must be low enough to attract brand switchers. These two factors imply that the gap between this firm's prices must exceed the gap between the rival's prices, i.e., that its discount must be larger.

5. IMPLICATIONS FOR PRICES AND PROFITS

We have shown that preference-based price discrimination is privately profitable for each firm; otherwise it would not arise as part of the firm's equilibrium strategy. This does not mean, however, that both firms are collectively better off as a result. To address whether they are or are not, we now compare equilibrium prices and profits of the game in which both firms can price-discriminate with the equilibrium prices and profits of the benchmark game in which neither firm can discriminate.

Table II summarizes our findings. Note that the entries in this table make sense only for those parameter values that satisfy $g^*(\theta) \leq l_\beta/l_\alpha \leq h^*(\theta)$, the region of parameter space in which a pure-strategy Nash equilibrium exists. However, it is easy to show that this requirement does not rule out any of the cases in the first column. In the remaining columns, a plus signifies an increase in the equilibrium value of the indicated variable and a minus signifies a decrease.

The first thing to notice from Table II is that, for each case, the equilibrium prices of both firms decrease to group α consumers. This is a direct consequence of the increase in competition that occurs over this group. Intuitively, in the absence of price discrimination,

TABLE II.
CHANGES IN PRICES AND PROFITS WITH PRICE
DISCRIMINATION

Relative Loyalty l_β/l_α	Firm A's Price		Firm B's Price		Profit	
	Group α	Group β	Group α	Group β	Firm A	Firm B
$\frac{l_\beta}{l_\alpha} \geq \frac{1+3\theta}{\theta}$	-	+	-	+	+	+
$\frac{1+\theta}{\theta} \leq \frac{l_\beta}{l_\alpha} < \frac{1+3\theta}{\theta}$	-	+	-	+	-	+
$\frac{1}{\theta} \leq \frac{l_\beta}{l_\alpha} < \frac{1+\theta}{\theta}$	-	-	-	+	-	+
$\frac{2-\theta}{2\theta} \leq \frac{l_\beta}{l_\alpha} < \frac{1}{\theta}$	-	-	-	+	-	-
$\frac{l_\beta}{l_\alpha} < \frac{2-\theta}{2\theta}$	-	-	-	-	-	-

firm B trades off charging a higher price to exploit group β consumers and a lower price to attract more group α consumers (recall that firm A faces no such trade-off, since in the absence of discrimination it sells only to group α consumers). The ability to price-discriminate effectively breaks this linkage between groups and enables firm B to lower its price to consumers in group α without sacrificing profit from sales to consumers in group β . In anticipation of firm B's lower price, it is a best response for firm A also to lower its price to these consumers (reaction functions are upward sloping in this model).

In contrast, the prices to group β consumers may increase or decrease, depending on the ratio of brand loyalties. Consider first the latter case. In the last row of Table II, we see that both firms' prices to group β consumers are lower. This leads to the following counter-intuitive result.

PROPOSITION 4: *If $l_\beta/l_\alpha < (2 - \theta)/2\theta$, then each firm will lower its price to both groups of consumers relative to the common price each would charge if price discrimination were not possible.*

In the usual case of third-degree price discrimination (e.g., Holmes, 1989), both prices rise in the strong (lower-elasticity) market and fall in the weak (higher-elasticity) market.¹⁹ The strong market is

19. As Nahata et al. (1990) show, however, this result need not extend to profit functions that exhibit multiple maxima in each market. In our model, the profit functions in each market are concave (one maximum).

simply assumed to be the same for both firms. In this model, at a given rival's price, the strong market of firm *A* is always the weak market of firm *B* and vice versa. This difference leads to a wider range of pricing outcomes. For example, Proposition 4 suggests that all prices may decrease relative to the case of no discrimination. Curiously, all prices may fall even if, in equilibrium, the two firms rank the two markets similarly. For instance, a *sufficient* condition for all consumers to be better off with price discrimination is $l_\beta/l_\alpha < \frac{1}{2} < (2 - \theta)/2\theta$ (a subset of row 5 in Table II). In this case, we know from Table I that group *B* consumers form each firm's more elastic group and therefore that the weak market of both firms is the same when evaluated at the prices in the discriminatory equilibrium. However, in other instances, $\frac{1}{2} < l_\beta/l_\alpha < (2 - \theta)/2\theta$, consumers are unambiguously better off, and the weak markets of the two firms differ (e.g., consider the symmetric case $l_\alpha = l_\beta$, $\theta = \frac{1}{2}$).

These results suggest that third-degree price discrimination that is based on consumers' brand preferences is qualitatively different from other forms of price discrimination. In addition to the usual segmentation issues that arise from grouping consumers with differing elasticities of demand, there is also an increase in competition that takes place when each firm can identify the other's customers. The joint interplay of these two forces leads to the wider range of possible outcomes.

Lower prices for all consumers imply lower profits for both firms (row 5).²⁰ And, as we can see from row 4, the region in which both firms are worse off with price discrimination also includes part of the region in which firm *B* increases its price to group β consumers. In contrast, in the first three rows, it happens that the gain to firm *B* from charging a higher price to group β consumers outweighs the loss it incurs from more intense price competition over group α consumers. If the average loyalty of group β consumers is sufficiently large relative to the average loyalty of group α consumers (row 1), even firm *A*'s profit can increase. Thus, the decrease in profit from group α consumers may be outweighed, for both firms, by the increase in profit from group β consumers.

PROPOSITION 5: *Both firms' profits decrease relative to no price discrimination if and only if $l_\beta/l_\alpha < 1/\theta$. If $l_\beta/l_\alpha \geq 1/\theta$, firm *B* earns higher profit than it would earn without discrimination, and if $l_\beta/l_\alpha \geq (1 + 3\theta)/\theta$, both firms earn higher profit than they would earn if price discrimination were not possible.*

20. This region of parameter space corresponds to a situation in which both firms would like to collude *ex ante* to limit price discrimination and eliminate it if possible. For more on this, see Winter (1997), who derives general conditions under which symmetric firms may want to limit price discrimination by colluding on relative prices.

This contrasts with the results in Borenstein (1985) and Holmes (1989), where price discrimination always leads to higher profits for both firms when market demand is fixed. Although it is possible in our model for both firms' profits to increase, this can only happen if group β consumers are sufficiently less price-sensitive than group α consumers (row 1 in Table II). In particular, it must be that $l_\beta/l_\alpha \geq (1 + 3\theta)/\theta$, which is not possible in our model, for instance, if demands are symmetric.

The results in Propositions 4 and 5 also contrast with the results in Bester and Petrakis (1996) and Chen (1997). In those articles, the equilibrium prices of both firms *always* decrease to both groups of consumers, thereby reducing each firm's profit. Although this scenario is possible in our model, it can nonetheless only happen if l_α is sufficiently large relative to l_β (row 5 of Table II). In particular, it must be that $l_\beta/l_\alpha < (2 - \theta)/2\theta$, which is satisfied, for instance, when demands are symmetric.

Because our model allows for both asymmetric loyalties and asymmetric baseline market shares, it has elements from each of the aforementioned four articles. For example, if, as in row 3 in Table II, $1/\theta \leq l_\beta/l_\alpha < 1 + \theta/\theta$, the qualitative changes in firm B 's two prices and their impact on its profits are similar to what happens to the firms in Holmes's and Borenstein's models, whereas the qualitative changes in firm A 's two prices and their impact on its profits are similar to what happens to the firms in Bester and Petrakis's and Chen's models. Intuitively, the increase in competition over group α consumers by itself lowers prices and profits for both firms. If group β consumers are sufficiently more loyal on average than group α consumers, we have Holmes and Borenstein's result. If group α consumers are sufficiently more loyal on average than group β consumers, we have Bester and Petrakis and Chen's result. Otherwise, we have a mixture of the results from the four articles.

We conclude this section by considering whether the kind of preference-based price discrimination that we consider favors the firm with the smaller or the larger market share. The relevant comparison can be made from the profit expressions in (7) and (10). It turns out that the firms do not share equally in the gains (or losses) from price discrimination. If there is a gain, the smaller firm benefits more; if there is a loss, the smaller firm suffers less. Intuitively, the reason firm B is better off (or less worse off) than firm A with price discrimination is that its profits may increase among group α consumers whereas firm A always loses there (the ability to price-discriminate enables firm B to compete more aggressively for firm A 's customers), and among group β consumers, firm B is better able than firm A to exploit its own consumers when l_β is large.

6. CONCLUSION

We have studied preference-based third-degree price discrimination when two firms compete in prices for consumers with varying loyalties. Our main innovation is in extending the existing literature to encompass asymmetric markets. Interestingly, for such markets some standard conclusions are reversed. In a symmetric market the firms always offer a discount to induce the competitor's customers to switch; here it may happen that a firm offers a discount to its own customers. In addition, we find that price discrimination may lessen competition rather than intensify it.

Our results also differ from the usual third-degree price discrimination in that firms may or may not rank groups of consumers in the same way on the basis of their price elasticity of demand. When evaluated at the discriminatory prices in equilibrium, we find that the more elastic group of one firm may or may not be the more elastic group of the other. These situations require different pricing strategies. In the first case, the equilibrium requires that only one firm pay customers to switch. In the second case, the equilibrium requires that both firms pay customers to switch.

We find that when preference-based price discrimination is possible a firm should charge a lower price to the consumer group that, from its perspective, is more price-elastic, which may or may not be the competitor's customers. If its own customers are not very loyal, for instance, a firm may want to pay customers to stay. We derived this result when the market elasticity of demand was zero. It is straightforward to show, however, that this result extends to situations of nonzero market elasticity, i.e., where raising prices causes some consumers to drop out of the market.

Extending the model in this way enhances its applicability. To illustrate, consider the *Microsoft* example of pricing upgrades. If all consumers value the upgrade by the same amount, but differ significantly in their costs of switching, then our model applies directly. But the general principle of when to pay customers to stay vs. when to pay customers to switch holds even when switching costs are symmetrical. For example, suppose *Microsoft* expects that current users will receive a smaller incremental benefit from the upgrade than will non-*Microsoft* customers who switch. In this case, although the cross-price derivatives may not differ between consumer groups, *Microsoft's* current customers will nonetheless have a higher elasticity of demand for *Microsoft's* upgrade, implying that they—and not the competitive users—should receive *Microsoft's* lower price.

Our paper has several limitations. First, the asymmetry in switching costs of the order necessary to induce a firm to adopt a pay-to-stay

strategy is large. However, this may be more a result of the special characteristics of the model than a fact about an actual market environment. For example, our model does not allow for possible differences in the market elasticities of demand between the consumer groups. If these are taken into account, then pay-to-stay strategies might be optimal even if there are no differences in switching costs. On the other hand, our model suggests that pay-to-stay strategies will not be as common as pay-to-switch strategies, a prediction that seems to be borne out in practice. Indeed, we would expect a bias towards this outcome even in a more general model. For instance, a pay-to-switch strategy is always optimal in the absence of any asymmetry (market demand elasticities or switching costs) between the two consumer groups.

A second limitation of the model arises from the fact that we abstract from prior competition; the firms in our model inherit their baseline market shares and switching costs. In contrast, the rest of the literature on preference-based price discrimination (which imposes symmetry) derives endogeneously the baseline market shares of each firm. One consequence of our omission is that in a multiple-round setting firms may in reality have more uncertainty about consumer preferences than what we assume here. For instance, even myopic consumers might cause inference difficulties (consumers who switched last period might in reality have preferences for the other brand).

One way of relaxing our assumption of the complete separability of markets is to introduce some randomness in the accuracy of a firm's information. That is, for any given consumer, it might be assumed that firms know only with some probability (less than one) whether he/she belongs to group α or group β . Alternatively, firms might not know the identity of only some consumers, in effect, creating a third group of unknowns.²¹ Of these two ways, the second approach can more easily be handled within our existing framework. In this second approach, each firm would then choose three prices: one for the known group α consumers, one for the known group β consumers, and one for the group of unknown consumers. As long as the same information is available to both firms, so that the identified groups are the same for each, our analysis and insights would essentially be unchanged. The first approach—that the firms know only

21. It may be that there is a flow of new consumers entering the market every period about which firms have little or no information. Incorporating this factor into the analysis might help to explain recent developments in the automotive industry, in which *General Motors* initiated a program called "Loyalty First," mailing millions of coupons to owners of GM vehicles for discounts off the purchase of new GM cars and trucks. Contrary to what our model predicts, Ford and Chrysler then followed with their own customer loyalty programs and agreed to accept GM certificates. GM did not agree to accept Ford and Chrysler discounts. See *Detroit Free Press*, April 23, 1998, pp. F-1, 2.

with some probability whether a given consumer belongs to group α or group β —is more complicated and left for future research.

However, if there are leakages between groups, so that some consumers are misidentified, we conjecture that the changes in equilibrium prices and profits relative to the benchmark game would move in the same direction but have a smaller magnitude than when there are no leakages. We would also expect the leakages between groups to mitigate any asymmetry in switching costs, thus reducing the incidence and relative importance of pay-to-stay strategies—which require substantial asymmetries in switching costs when the market elasticities of demand are the same across groups.

APPENDIX A

LEMMA 1: *In any pure-strategy Nash equilibrium to the game in which neither firm can price-discriminate, firm A must have the weakly higher Bertrand price (strictly higher if $\theta > \frac{1}{2}$).*

Proof. The proof consists of two parts. In part 1, we show that a pure-strategy Nash equilibrium to the game in which group identities are unknown does not exist in which firm B's price exceeds firm A's price. This establishes that if a pure-strategy Nash equilibrium is to exist, firm A must have the weakly higher Bertrand price. In part 2, we show that a pure-strategy Nash equilibrium exists in which firm B's price equals firm A's price if and only if $\theta = \frac{1}{2}$. This establishes that in any pure strategy Nash equilibrium, firm A's price must be strictly higher if $\theta > \frac{1}{2}$.

- *Part 1.* Suppose a pure-strategy Nash equilibrium exists in which (\bar{P}_A, \bar{P}_B) are the equilibrium prices and $\bar{P}_B > \bar{P}_A$. It follows that

$$\bar{P}_A \in \arg \max \left(\theta(P_A - c) + (1 - \theta)(P_A - c) \frac{\bar{P}_B - P_A}{l_\beta} \right),$$

$$\bar{P}_B \in \arg \max \left((1 - \theta)(P_B - c) \frac{l_\beta + \bar{P}_A - P_B}{l_\beta} \right).$$

Solving yields

$$\bar{P}_A = \frac{1 + \theta}{3(1 - \theta)} l_\beta + c, \quad \bar{P}_B = \frac{2 - \theta}{3(1 - \theta)} l_\beta + c.$$

It is easily seen that $\bar{P}_B > \bar{P}_A$ if and only if $\theta < \frac{1}{2}$. A contradiction.

- *Part 2.* Suppose a pure-strategy equilibrium exists in which (\bar{P}_A, \bar{P}_B) are the equilibrium prices and $\bar{P}_A = \bar{P}_B$. It follows that neither

firm should have any incentive to deviate unilaterally by either increasing or decreasing its price. This implies that the left and right derivatives of both firms' profit functions must have proper signs. Hence, four constraints must be satisfied:

$$\bar{P}_A \geq l_\alpha + c,$$

$$\bar{P}_A \leq \frac{\theta}{1-\theta}l_\beta + c,$$

$$\bar{P}_B \geq l_\beta + c,$$

$$\bar{P}_B \leq \frac{1-\theta}{\theta}l_\alpha + c.$$

Since $\bar{P}_A = \bar{P}_B$ by assumption, these inequalities imply $\theta = \frac{1}{2}$. □

APPENDIX B. COMPETITION WHEN NEITHER FIRM CAN PRICE-DISCRIMINATE

To verify that (P_A^*, P_B^*) form a Nash equilibrium, we must show that neither firm can profitably deviate. There are two cases to check. First, we derive a condition for when firm *A* cannot earn higher profit by choosing a price less than or equal to P_B^* . Second, we derive a condition for when firm *B* cannot earn higher profit by choosing a price greater than or equal to P_A^* . The intersection of these two conditions forms a region of parameter space where neither firm wants to deviate, and hence, where (P_A^*, P_B^*) form a Nash equilibrium. We verify that this region is not empty.

The best firm *A* can do by pricing less than or equal to P_B^* is

$$\max_{P_A \leq P_B^*} \Pi_A^d(P_A, P_B^*) = \max_{P_A \leq P_B^*} \left(\theta(P_A - c) + (1-\theta)(P_A - c) \frac{P_B^* - P_A}{l_\beta} \right).$$

Let P_A^d be the solution and define $\Delta \Pi_A^d = \Pi_A^* - \Pi_A^d(P_A^d, P_B^*)$. The constraint $P_A \leq P_B^*$ binds if and only if $l_\beta/l_\alpha > (1-\theta)(2-\theta)/(3\theta^2)$. When the constraint does not bind, we have

$$P_A^d = \frac{2-\theta}{6\theta}l_\alpha + \frac{\theta}{2(1-\theta)}l_\beta + c,$$

$$\Delta \Pi_A^d = \frac{[(1-\theta)l_\alpha - \theta l_\beta][(\theta-1)(2-\theta)^2l_\alpha + 9\theta^3l_\beta]}{36(1-\theta)\theta^2l_\beta}.$$

The denominator of $\Delta \Pi_A^d$ is strictly positive. Moreover, the first bracketed factor in the numerator is also strictly positive—we know this

because it is implied throughout the region where the constraint does not bind. Therefore, $\Delta\Pi_A^d \geq 0$ if and only if the second bracketed factor is nonnegative:

$$\frac{l_\beta}{l_\alpha} \geq \frac{(1-\theta)(2-\theta)^2}{9\theta^3} = g^*(\theta).$$

Since it is not profitable for firm A to deviate if and only if $\Delta\Pi_A^d \geq 0$, we can conclude that P_A^* is firm A 's best response to P_B^* if and only if $l_\beta/l_\alpha \geq g^*(\theta)$.

The best firm B can do by pricing greater than or equal to P_A^* is

$$\max_{P_B \geq P_A^*} \Pi_B^d(P_A^*, P_B) = \max_{P_B \geq P_A^*} \left((1-\theta)(P_B - c) \frac{l_\beta + P_A^* - P_B}{l_\beta} \right).$$

Let P_B^d be the solution, and define $\Delta\Pi_B^d = \Pi_B^* - \Pi_B^d(P_A^*, P_B^d)$. The constraint $P_B \geq P_A^*$ binds if and only if $l_\beta/l_\alpha < (1+\theta)/(3\theta)$. When the constraint does not bind, we have

$$P_B^d = \frac{1+\theta}{6\theta} l_\alpha + \frac{1}{2} l_\beta + c,$$

$$\Delta\Pi_B^d = \frac{[\theta l_\beta - (1-\theta)l_\alpha][(1+\theta)^2 l_\alpha - 9\theta l_\beta(1-\theta)]}{36\theta^2 l_\beta}.$$

The denominator of $\Delta\Pi_B^d$ is strictly positive. Moreover, the first bracketed factor in the numerator is also strictly positive—we know this because it is implied throughout the region where the constraint does not bind. Therefore, $\Delta\Pi_B^d \geq 0$ if and only if the second bracketed factor is nonnegative:

$$\frac{l_\beta}{l_\alpha} \leq \frac{(1+\theta)^2}{9\theta(1-\theta)} = h^*(\theta).$$

Since it is not profitable for firm B to deviate if and only if $\Delta\Pi_B^d \geq 0$, we can conclude that P_B^* is firm B 's best response to P_A^* if and only if $l_\beta/l_\alpha \leq h^*(\theta)$.

In summary, a pure-strategy Nash equilibrium exists if and only if $g^*(\theta) \leq l_\beta/l_\alpha \leq h^*(\theta)$.

APPENDIX C

LEMMA 2: *In any pure-strategy Nash equilibrium of games in which at least one firm can price-discriminate, each firm will sell to both consumer groups. That is, $P_A > P_B$ and $\tilde{P}_B > \tilde{P}_A$.*

Proof. The proof is by contradiction. Suppose a pure-strategy Nash equilibrium does exist in which $P_B \geq P_A$. Then firm A 's profit from group α consumers is $(P_A - c)\theta$, and firm B 's profit from group α consumers is zero. For this to be true in equilibrium, it must be that firm A cannot price-discriminate, for otherwise, it could increase its profit by charging something slightly more in group α . It also must be that firm B cannot price-discriminate, for otherwise, it could charge something slightly less than P_A in group α and earn positive profit—we know P_A must strictly exceed marginal cost in any equilibrium. But if neither firm can price-discriminate, the hypothesis of Lemma 2 is violated. This establishes a contradiction. Using similar reasoning, it can be shown that no pure-strategy Nash equilibrium exists in which $\hat{P}_A \geq \hat{P}_B$.

REFERENCES

- Alsop, R., 1985, "Companies See Ways to Put Coupons Where They'll Count," *Wall Street Journal*, August 8, B-1.
- Banerjee, A.L. Summers, 1987, "On Frequent-Flyer Programs and Other Loyalty-Inducing Economic Arrangements," Working Paper 1337, Harvard University.
- Bester, H. and E. Petrakis, 1996, "Coupons and Oligopolistic Price Discrimination," *International Journal of Industrial Organization*, 14, 227-242.
- Blattberg, R. and S. Neslin, 1990, *Sales Promotion: Concepts, Methods, and Strategies*, Prentice-Hall.
- Borenstein, S., 1985, "Price Discrimination in Free-Entry Markets," *Rand Journal of Economics*, 16, 380-397.
- Caminal, R. and C. Matutes, 1990, "Endogenous Switching Costs in a Duopoly Model," *International Journal of Industrial Organization*, 8, 353-373.
- Chen, Y., 1997, "Paying Customers to Switch," *Journal of Economics & Management Strategy*, 6, 877-897.
- Corts, K., 1998, "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," *Rand Journal of Economics*, 29, 306-323.
- Fudenberg, D. and J. Tirole, 1997, "Customer Poaching and Brand Switching," Mimeo, MIT.
- Holmes, T., 1989, "The Effects of Third Degree Price Discrimination in Oligopoly," *American Economic Review*, 79, 244-250.
- Hotelling, H., 1929, "Stability in Competition," *Economic Journal*, 39, 41-57.
- Klemperer, P., 1987, "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, 102, 375-394.
- Lederer, P. and A. Hurter, Jr., 1986, "Competition of Firms; Discriminatory Pricing and Location," *Econometrica*, 54, 623-640.
- Lehmann, D., 1985, *Market Research and Analysis*, 2nd ed., Homewood, IL: Richard D. Irwin.
- Nahata, B., Ostaszewski, K., and P.K. Sahoo, 1990, "Direction of Price Changes in Third Degree Price Discrimination," *American Economic Review*, 80, 1254-1258.
- Nilssen, T., 1992, "Two Kinds of Switching Costs," *Rand Journal of Economics*, 23, 579-589.
- Rossi, P.E. and G.M. Allenby, 1993, "A Bayesian Approach to Estimating Household Parameters," *Journal of Marketing Research*, 30, 171-182.
- Shaffer, G. and Z.J. Zhang, 1995, "Competitive Coupon Targeting," *Marketing Science*, 14, 395-416.

- Shaffer, G. and Z.J. Zhang, 1997, "One-to-One Pricing: Competitive Implications," Mimeo, U. of Rochester.
- Taylor, C., 1998, "Supplier Surfing: Price-Discrimination in Markets with Repeat Purchases," Mimeo, U. of Rochester.
- Thisse, J.F. and X. Vives, 1988, "On the Strategic Choice of Spatial Price Policy," *American Economic Review*, 78, 122-137.
- Tirole, J., 1989, *The Theory of Industrial Organization*, Cambridge, MA: The MIT Press.
- Winter, R., 1997, "Colluding on Relative Prices," *Rand Journal of Economics*, 28, 359-371.
- Zeldis, N., 1987, "Premiums and Promotions: Targeted Coupons Hit Non-users," *Advertising Age*, 58, 26.