

Behavior-Based Price Discrimination and Customer Recognition*

Drew Fudenberg
(Harvard University)

J. Miguel Villas-Boas
(University of California, Berkeley)

September 2005

ABSTRACT

When firms are able to recognize their previous customers, they may be able to use their information about the consumers' past purchases to offer different prices and/or products to consumers with different purchase histories. This article surveys the literature on this "behavior-based price discrimination."

*Forthcoming in the Handbook on Economics and Information Systems, Elsevier. We thank Mark Armstrong, Michael Baye, Assaf Ben-Shoham, Yongmin Chen, Terry Hendershott, John Morgan, Jae Nahm, Jean Tirole, and Juanjuan Zhang for insightful comments on an earlier draft of this paper, and for helpful conversations. NSF grant SES-04-26199 provided financial support for some of this work.

1. Introduction

When firms have information about consumers' previous purchases, they may be able to use this information to offer different prices and/or products to consumers with different purchase histories. This sort of "behavior-based price discrimination" (BBPD) and use of "customer recognition" occurs in several markets, such as long-distance telecommunications, mobile telephone service, magazine or newspaper subscriptions, banking services, credit cards, labor markets; it may become increasingly prevalent with improvements in information technologies and the spread of e-commerce and digital rights management.

This article focuses on models of "pure" BBPD, in which past purchases matter only for their information value, and do not directly alter consumers' preferences. We do make some comparisons with switching-cost models, where past purchases do have a direct effect, but we say almost nothing about traditional models of third-degree price discrimination, where firms can base their prices on observable and exogenous characteristics of the consumers.

One recurrent theme throughout the article is the possibility that firms may face a commitment problem: although having more information helps the firm extract more surplus with its current prices, consumers may anticipate this possibility, and so alter their initial purchases. Thus, as in the related literatures on bargaining, durable-good monopoly, and dynamic mechanism design,¹ the seller may be better off if it can commit to ignore information about the buyer's past decisions. A second theme is that, as in traditional models of third-degree price discrimination,² more information may lead to more intense competition between firms. Thus, even if each firm would gain by being the only one to practice BBPD, industry profits can fall when all of the firms practice it. Third, and related, firms would often gain from using long-term contracts when they are able to do so as, for example, in the market for cell-phone services. The last implication is somewhat unfortunate from the analyst's perspective: The welfare implications of BBPD seem to be ambiguous, and to depend on many aspects of the market structure.

¹See, for example, Stokey (1981), Bulow (1982), Fudenberg and Tirole (1983), Baron and Besanko (1984), Freixas et al. (1985), Sobel (1991).

² See, e.g., Thisse and Vives (1988).

Section 2 examines a monopoly supplier of a single, non-durable good. We start with a simple two-period model, and then consider the infinite-horizon models of Hart and Tirole (1988) and Schmidt (1993), where all consumers are infinite lived, and Villas-Boas (2004), where there are overlapping generations of consumers who live only two periods. We compare this situation with that of a durable-goods monopolist. Then we consider the use of long-term contracts, and relate the resulting outcome again to that in models of durable-good monopoly. We also discuss the case where the consumer's preferences vary over time, as in Kennan (2001) and Battaglini (2004). Finally, we consider the situation where the monopolist sells more than one good, which we use as a benchmark when studying BBPD with multiple firms; we also compare this with a monopolist seller of two goods in a model of switching costs.

Section 3 studies BBPD with two firms, each still selling a single good. In these models, firms can try to "poach" their rivals' customers by giving new customers special "introductory" prices. We begin with Fudenberg and Tirole (2000)'s analysis of a two-period model of competition in short-term contracts, and its extension by Chen and Zhang (2004) and Esteves (2004) to other distributions of consumer types, where other insights emerge. Next, we discuss Villas-Boas' (1999) model of poaching in an infinite-horizon model with overlapping generations of consumers, each of whom lives only for two periods, where firms cannot distinguish between new consumers and old ones who bought from their rival. We then return to the two-period setting to present Fudenberg and Tirole's analysis of competition in simple long-term contracts, meaning that firms offer both a "spot" or one-period price and also a long-term commitment to supply the good in both periods. Finally, we compare the predictions of these models to models of switching costs, where past decisions are directly payoff relevant, and may also provide information, as in Chen (1997) and Taylor (2003).

Section 4 discusses models where each firm can produce multiple versions of the same product. We begin with Fudenberg and Tirole (1998), and Ellison and Fudenberg (2000), who study the provision of "upgrades" by a monopolist in a setting of vertical differentiation, where all customers agree that one good is better than the other. We then consider the work of Zhang (2005) on endogenous product lines in a Hotelling-style duopoly model of horizontal differentiation. Finally we discuss the papers of Levinthal

and Purohit (1989), Waldman (1996), and Nahm (2004) on the introduction of a new product in models with anonymous consumers and a frictionless second-hand market. Although these papers do not consider behavior-based pricing, the analysis of the anonymous case is an important benchmark for the effects of behavior-based pricing.

Section 5 briefly discusses three related topics: privacy, credit markets, and customized pricing. We discuss the work of Taylor (2004a) and Calzolari and Pavan (2005) on consumer privacy. If consumers are not myopic, they will realize that information revelation can reduce their future surplus; in some cases, this can give firms an incentive to try to protect consumer privacy. In credit markets, lenders may learn about the ability of their borrowers, their customers, to repay loans; this information can then be used by the firms in future loans to those customers. In this case what a firm learns about its previous customers relates to the cost of providing the customer with a given contract, as opposed to the customer's willingness to pay, which is the focus of most of the work we discuss. Our presentation here is based in large part on Dell'Ariccia et al. (1999), and Dell'Ariccia and Marquez (2004); we also discuss Pagano and Jappelli (1993), and Padilla and Pagano (1997, 2000). Finally, for completeness, we briefly present the case of competition when firms already have information about consumer tastes, starting from the initial work of Thisse and Vives (1988). Section 6 presents concluding remarks.

2. Monopoly

We begin with the case of a monopolist who can base prices to its consumers on their past purchase history. For example, in newspaper or magazine subscriptions, firms with market power may offer different rates depending on the consumers' past purchase behavior.³ We start by considering a two-period model to illustrate some of the effects that can be present, discussing the role of commitment, and of forward-looking consumers. Then, we consider the case of overlapping generations of consumers (Villas-Boas 2004), and discuss the case when consumers are long lived (Hart and Tirole 1988, Schmidt 1993). We consider the effect of long-term contracts and the relationship to the

³ See, for example, "Publications Are Trying New Techniques to Win over Loyal Readers," *The New York Times*, January 4, 1999, p. C20.

durable-goods and bargaining literature in that setting (Hart and Tirole). We also discuss the case where the consumer's preferences vary over time, as in Kennan (2001) and Battaglini (2004), who study short-term and long-term contracts, respectively. Finally, we consider the situation where the monopolist sells more than one good, as in Section 5 of Fudenberg and Tirole (2000), which will be an important benchmark case for the next section on competition, and discuss the differences between purely informational behavior-based price discrimination and price discrimination when previous purchases have a direct impact on consumer preferences as in models of switching costs.

2.1. Two-Period Model

2.1.1. Base Model

Consider a monopolist that produces a non-durable good at zero marginal cost in each of two periods. A continuum of consumers with mass normalized to one is in the market in each of the two periods. In each period each consumer can use one unit or zero units of the good.; no consumer has any additional gain from using more than one unit in each period. The consumer preferences are fixed across the two periods. The consumers' valuation for the good is represented by a parameter q distributed in the line segment $[0,1]$ with cumulative distribution function $F(q)$ and density $f(q)$. We assume throughout that $p[1-F(p)]$ is strictly quasi-concave in p (which is the condition necessary for the existence of a unique local maximum in the static monopoly case). The assumption on the support of the distribution is without loss of generality relative to any compact interval. Hart and Tirole (1988) and Villas-Boas (2004) consider the case of the two-point distribution. Schmidt (1993) considers the case of any discrete number of types.⁴ Here, we present the case of a continuum of consumer types, and note differences

⁴ We restrict attention to the case in which the consumers are the only parties with private information. It would also be interesting to investigate what happens when the monopolist has also some private information, and the consumers may learn what price offers they will get in the future from the offers made by the firm in the past. From the literature on "reputation effects" we expect that this could allow the firm to obtain higher profits.

with the two-type case when they arise. In order to obtain some of the sharper results we will sometimes restrict attention to the uniform distribution, with $f(\mathbf{q}) = 1, \forall \mathbf{q}$.

Each consumer is endowed with the same \mathbf{q} in both periods. This valuation \mathbf{q} represents the gross utility the consumer enjoys from using the good in one period. Therefore, the net utility per period of a consumer of type \mathbf{q} purchasing the good at price p in one period is $\mathbf{q} - p$. The lifetime utility of a consumer is the discounted sum of the net utilities of the two periods the consumer is in the market with discount factor \mathbf{d}_C with $0 \leq \mathbf{d}_C < 1$. In the first period the monopolist chooses one price a to be charged to all consumers (the monopolist cannot distinguish among them, and all consumers prefer a lower price). In the second period the monopolist chooses two prices: a price \mathbf{a}_p to be charged to the previous customers of the firm, and a price \mathbf{a}_n to be charged to the consumers that did not buy in the first period, the new customers.

The monopolist wants to maximize the expected discounted value of its profits, using a discount factor \mathbf{d}_F with $0 \leq \mathbf{d}_F < 1$. Except where expressly noted we restrict attention to the case in which $\mathbf{d}_F = \mathbf{d}_C$, and then, the discount factor is denoted by \mathbf{d} .

Given that there is a continuum of consumers, each of them realizes that his decision does not affect the prices charged by the monopolist in the next period. Then a consumer of type \mathbf{q} just entering the market decides to buy in the first period if $\mathbf{q} - a + \mathbf{d}_C \max[\mathbf{q} - \mathbf{a}_p, 0] \geq \mathbf{d}_C \max[\mathbf{q} - \mathbf{a}_n, 0]$. From this inequality one can then obtain directly that given $\mathbf{d}_C < 1$, if a type $\hat{\mathbf{q}}$ chooses to buy in the first period then all the types $\mathbf{q} > \hat{\mathbf{q}}$ also choose to buy in the first period. That is, the consumers that buy for the first time in the second period value the product by less than any of the consumers that buy in the first period.

In order to compute the type of the marginal consumer it is helpful to consider the pricing decision of the monopolist with respect to its previous customers. Define $p^* \equiv \operatorname{argmax}_p p[1 - F(p)]$, the price that maximizes the profit in one period when the consumers do not have any reason to refrain from buying, that is, they buy if their valuation \mathbf{q} is greater than the price charged. This is the monopoly price in the static

case, or if the monopolist is not able to recognize its previous customers or price differently to them.

Denoting \hat{q} as the type of the marginal consumer in the first period, if $\hat{q} > p^*$ the monopolist sets $a_p = \hat{q}$. If, on the other hand $\hat{q} < p^*$, the monopolist sets $a_p = p^*$. That is, $a_p = \max[\hat{q}, p^*]$, the marginal consumer in the first period does not get any surplus in second period. This is the “ratchet effect” of the consumers being hurt (i.e., being charged a higher price) by revealing, even if partially, their types (Freixas et al. 1985).

The marginal consumer in the first period is then determined by

$$(2.1) \quad \hat{q} - a = d_c \max[\hat{q} - a_n, 0],$$

which results in

$$\begin{aligned} \hat{q} &= a \text{ if } a \leq a_n \\ \hat{q} &= \frac{a - d_c a_n}{1 - d_c} \geq a \text{ if } a > a_n. \end{aligned}$$

This expression for \hat{q} shows an important aspect of the market dynamics: If prices are expected to increase, each consumer does not have any reason to behave strategically and buys if his valuation is above the current price. If, on the other hand, prices are expected to decrease, some consumers will behave strategically, not being identified in the first period, and being able to get a better deal in the second period.

2.1.2. No Customer Recognition

Consider first as a benchmark the case of no customer recognition, in which the monopolist cannot price discriminate in the second period between the consumers that bought, and did not buy, in the first period. The optimal price charged in each period is then $p^* \equiv \operatorname{argmax} p[1 - F(p)]$, generating a profit in each period equal to $p^*[1 - F(p^*)]$. Note that, obviously, there is no price variation through time. For the uniform distribution example we have $p^* = 1/2$, a profit per period of $1/4$, and a total profit of $(1 + d)/4$.

2.1.3. Customer Recognition and Behavior-Based Price Discrimination

Consider now the case in which the monopolist is able to recognize the previous customers, as in Hart and Tirole (1988), Schmidt (1993), and Villas-Boas (2004).⁵ For example, an internet store may be able to recognize returning customers through cookies installed in their computer, and charge them different prices. In this setting the monopolist can identify in the second period two different groups of consumers: those who have bought in the first period, and those who have not bought in the first period. In the second period the monopolist can charge two different prices. The price paid by the monopolist's previous consumers, \mathbf{a}_p , and the price paid by consumers who have not bought previously, \mathbf{a}_n .

Given that the marginal consumer buying the product in the first period is $\hat{\mathbf{q}}$, the optimal prices in the second period are $\mathbf{a}_p^*(\hat{\mathbf{q}}) = \max[p^*, \hat{\mathbf{q}}]$ and $\mathbf{a}_n^*(\hat{\mathbf{q}}) = \operatorname{argmax}_{\mathbf{a}_n} \mathbf{a}_n [F(\hat{\mathbf{q}}) - F(\mathbf{a}_n)]$.

The marginal consumer in the first period, $\hat{\mathbf{q}} = \hat{\mathbf{q}}(a)$, is determined by $\hat{\mathbf{q}}(a) = \frac{a - d_c \mathbf{a}_n^*(\hat{\mathbf{q}}(a))}{1 - d_c}$. In order to obtain the optimal first period price, a^* , the monopolist then maximizes

(2.2)

$$\max_a a[1 - F(\hat{\mathbf{q}}(a))] + d_F \{ \max_{\mathbf{a}_p} \mathbf{a}_p [1 - F(\max[\mathbf{a}_p, \hat{\mathbf{q}}(a)])] + \mathbf{a}_n^*(\hat{\mathbf{q}}(a)) [F(\hat{\mathbf{q}}(a)) - F(\mathbf{a}_n^*(\hat{\mathbf{q}}(a)))] \},$$

where the first term represent the first-period profit, and the second term represents the second-period profit, both from the consumers who bought in the first period and from the consumers who did not buy in the first period. Under the assumption that $\hat{\mathbf{q}} > p^*$, which is satisfied in equilibrium, the first order condition that defines the optimal a^* is then

⁵ See also Acquisti and Varian (2005) for results focusing on the role of commitment (see below) and the effect of being able to offer enhanced services. The possibility of enhanced services is also covered in Section 4.

$$(2.3) \quad 1 - F(\hat{q}) - a^* f(\hat{q})\hat{q}' + d_F \hat{q}' [1 - F(\hat{q}) - f(\hat{q})\hat{q}' + f(\hat{q})\mathbf{a}_n^*(\hat{q})] = 0.$$

Note that for $\mathbf{d}_C = \mathbf{d}_F = \mathbf{d}$ the marginal consumer buying the product in the first period has a higher valuation than if there were no customer recognition. To see this note that the first derivative of the objective function above evaluated at $\hat{q}(a) = p^*$ is equal to $f(p^*)p^*[1 - (1 - \mathbf{d})\hat{q}']$ after substituting for $1 - F(p^*) - p^*f(p^*) = 0$ and $p^*(1 - \mathbf{d}) = a - \mathbf{d}\mathbf{a}_n^*(p^*)$. Given that $\hat{q}' = 1/(1 - \mathbf{d} + \mathbf{d}\mathbf{a}_n^*)$ and $\mathbf{a}_n^* > 0$, that derivative is positive, which means that the monopolist should increase a , which implies a higher valuation of the marginal consumer than p^* . One can also obtain for $\mathbf{d}_C = \mathbf{d}_F$ that the present value of profits is $\hat{q}[1 - F(\hat{q})] + \mathbf{d}\mathbf{a}_n^*[1 - F(\mathbf{a}_n^*)]$, which is strictly below the present value of profits under no customer recognition, as p^* uniquely maximizes $p[1 - F(p)]$. The intuition of this result is that the marginal consumers refrain from buying in their first period in the market because they know that they can get a lower price in the next period. This result of lower profits with customer recognition does not hold if the consumers are myopic while the monopolist is forward looking (or \mathbf{d}_F large as compared to \mathbf{d}_C).

For the uniform distribution example one can obtain $\mathbf{a}_n^*(\hat{q}) = \hat{q}/2$, $\hat{q}(a) = 2a/(2 - \mathbf{d})$, and $a^* = (4 - \mathbf{d}^2)/(8 + 2\mathbf{d})$. One can also easily check that, as argued above, the present value of profits is lower than in the no customer recognition case for all \mathbf{d} . One can also get that $2/(4 + \mathbf{d})$ consumers buy in both periods, while $(2 + \mathbf{d})/(8 + 2\mathbf{d})$ consumers only buy in the second period. As consumers become more strategic (greater \mathbf{d}) the number of consumers buying in both periods decreases, as the consumers wait for future deals, and consequently, the number of consumers that only buy in the second period increases.

The main ideas from these results can also be obtained with a two-type distribution as presented in the references listed above.

2.1.4. The Role of Commitment

A crucial feature in the previous section is that the monopolist could not commit in the first period to its second-period price. This led the consumers to refrain from buying in the first period, because the marginal consumers knew that if they bought in the first period they would get zero surplus in the second period. One could then wonder what would happen if the monopolist were able to commit in the first period to its second-period prices. For example, in the market for cellular phone services firms are sometimes able to commit to prices for some future periods. In this case one can then apply the revelation principle, giving incentives for consumer types to reveal themselves in the first period. That is, we suppose that consumers announce their valuations in the first period, and are then assigned a price and a consumption plan for the two periods, such that consumers announce their valuation truthfully. Without commitment, the firm could change the utility (or consumption) a consumer gets in the second period given what the firm learns in the first period.

In a durable-good context Stokey (1979) shows that when firms can commit to the time path of prices, and $\mathbf{d}_C = \mathbf{d}_F$, the monopolist commits to having the same price in all periods, which ends up being the static monopoly price. Hart and Tirole (1988) show that the same conclusion applies when the firm can engage in behavior-based price discrimination: the optimal policy is to forgo the ability to price discriminate and simply charge the static monopoly price in every period.⁶ Villas-Boas (2004) shows that the result also applies when there are overlapping generations of consumers.

To see this in the model presented here, note that if the monopolist is able to commit to the second-period prices for the consumers who bought in the first period, \mathbf{a}_p , and who did not buy in the first period, \mathbf{a}_n , the most that it can get is $\mathbf{a}_p[1 - F(\mathbf{a}_p)] + \mathbf{d}\mathbf{a}_n[1 - F(\mathbf{a}_n)]$ which is maximized when $\mathbf{a}_p = \mathbf{a}_n = p^*$, with a first-period price $a = p^*$, no price discrimination. Note also that commitment allows the monopolist to be better off.

⁶Acquisti and Varian (2005) derive the same result. The result can also be seen as the same as in Baron and Besanko (1984) who show that in a dynamic single-principal single-agent relationship with constant types over time the optimal long-term contract under full-commitment consists in a sequence of static optimal contracts.

Note that when the monopolist is more forward-looking than the consumers, $d_F > d_C$, the firm may then choose to price discriminate, cutting prices through time.

2.2. *Overlapping Generations of Consumers*

The two-period model above is able to highlight some of the effects under customer recognition and behavior-based price discrimination, but since it focuses on the effects of the beginning of the market (in the first period) and the end of the market (in the second period), it potentially may not get at some of the effects in an on-going market.

Consider then a market where there is an infinitely lived monopolist facing overlapping generations of consumers as in the previous section (Villas-Boas 2004). Each generation lives for two periods, and in each period there are two generations of consumers in the market (each of mass one for a total mass in the market of two in each period), one in its first period in the market, the other in its second and final period in the market. Assume further that $1 - F(p) - 2pf'(p) = 0$ has only one solution in the set of real numbers. This last assumption is not necessary but simplifies the presentation of the results.⁷

Note first that if the monopolist is not only able to recognize whether a consumer bought in the past, but also his “age,” all the results that we obtained in the previous section (including the equilibrium prices) apply directly, the monopolist charging three prices in each period: One price for the customers that are just arriving into the market; one price for the consumers who are in their second period in the market and bought the product in the previous period; and finally one price for the consumers who are in their second period in the market and did not buy the product in the previous period.

However, in many situations, a firm may not be able to recognize a consumer’s “age,” and therefore have to charge the same price to both consumers that are just entering the market and consumers that have been in the market in the previous period, but did not buy the product. Note also that this has the realistic feature of the monopolist

⁷This assumption is implied by the condition $3f(p) + 2pf'(p) > 0$ which is satisfied for distributions close to the uniform or the truncated normal with sufficiently large variance.

knowing more about the consumers that bought the product in the previous period than about the new customers. In terms of the notation of the previous section, not recognizing the customers age means that $a = \mathbf{a}_n$.

In order to concentrate on the dynamic effects of customer recognition we focus the analysis on the Markov perfect equilibria (MPE; Fudenberg and Tirole (1991), p. 513) of this game, i.e., equilibria in which the actions in each period depend only on the payoff-relevant state variables in that period. In this particular game the payoff-relevant state variable in each period is the stock of previous customers of the monopolist in each period.

From the analysis in the previous section, we know that in each period the consumers just arriving in the market who buy the product in that period are the ones with the highest valuation. That is, in a period t , the payoff-relevant state variables can be summarized by the type of the marginal consumer entering the market in period $t-1$ who chooses to buy in period $t-1$, denoted by $\hat{\mathbf{q}}_t$. The computation of $\hat{\mathbf{q}}_t$ is exactly as in the previous section. In what follows, let a_t be the price charged to new customers in period t , and \mathbf{a}_t be the price charged to previous customers in period t .

Denoting as $\hat{\mathbf{q}}_{t+1} = \hat{\mathbf{q}}(a_t)$ the marginal consumer purchasing in period t given price a_t , and $V(\hat{\mathbf{q}}_t)$ the net present value of the monopolist's profits from period t onwards if the marginal consumer purchasing in period $t-1$ had valuation $\hat{\mathbf{q}}_t$ we can write the monopolist's problem as

$$(2.4) \quad \begin{aligned} V(\hat{\mathbf{q}}_t) = & \max_{\mathbf{a}_t} \mathbf{a}_t [1 - F(\max[\mathbf{a}_t, \hat{\mathbf{q}}_t])] + \max_{a(\hat{\mathbf{q}}_t)} a(\hat{\mathbf{q}}_t) [1 - F(\hat{\mathbf{q}}(a(\hat{\mathbf{q}}_t)))] \\ & + \max[F(\hat{\mathbf{q}}_t) - F(a(\hat{\mathbf{q}}_t)), 0] + dV(\hat{\mathbf{q}}(a(\hat{\mathbf{q}}_t))). \end{aligned}$$

The function $a(\hat{\mathbf{q}}_t)$ is the price to charge the new customers in period t if the marginal consumer purchasing in period $t-1$ has valuation $\hat{\mathbf{q}}_t$. The right hand side of (2.4) is composed of three terms. The first term is the profit from repeat buyers. The second term is the profit from first-time buyers which are either new in the market,

$1 - F(\hat{\mathbf{q}}(a(\hat{\mathbf{q}}_t)))$, or in their second period in the market, $\max[0, F(\hat{\mathbf{q}}_t) - F(a(\hat{\mathbf{q}}_t))]$. The set of new buyers which are in their second period in the market has only positive measure if $a(\hat{\mathbf{q}}_t) < \hat{\mathbf{q}}_t$. The third term represents the net present value of profits from the next period on.

The MPE is then characterized by the functions $V(\hat{\mathbf{q}}_t)$, $a(\hat{\mathbf{q}}_t)$ satisfying (2.4) and $\hat{\mathbf{q}}(a_t)$ satisfying $\hat{\mathbf{q}}(a_t) = \max[a_t, \frac{a_t - \mathbf{d}a(\hat{\mathbf{q}}(a_t))}{1 - \mathbf{d}}]$. Note also that if $a(\hat{\mathbf{q}}_t) \geq \hat{\mathbf{q}}_t$ then $a(\hat{\mathbf{q}}_t)$ is a constant (the case of $\hat{\mathbf{q}}_t$ small) because the maximization in (2.4) is independent of $\hat{\mathbf{q}}_t$. This also means that if for a certain $\hat{\mathbf{q}}_t$ the optimal $a(\hat{\mathbf{q}}_t) \geq \hat{\mathbf{q}}_t$ then $a(x) \geq \hat{\mathbf{q}}_t, \forall x \leq \hat{\mathbf{q}}_t$. If, on the other hand $a(\hat{\mathbf{q}}_t) < \hat{\mathbf{q}}_t$ then $a(\hat{\mathbf{q}}_t)$ is increasing in $\hat{\mathbf{q}}_t$ because the objective function is supermodular in $\hat{\mathbf{q}}_t$ and $a(\hat{\mathbf{q}}_t)$.

2.2.1. No Constant Prices in Equilibrium

We now show that in general prices are not constant through time. Suppose that we are in the steady-state, with the monopolist charging the price \bar{a} to the new customers in every period. Then, because prices are not going to decrease and the marginal consumer gets always zero surplus in the second period, all consumers with valuation above \bar{a} buy in the current period, that is, $\hat{\mathbf{q}}(\bar{a}) = \bar{a}$. Then, we also know that $a(x) = \bar{a}, \forall x \leq \bar{a}$. Let \hat{a} be defined by $1 - F(\hat{a}) - 2\hat{a}f(\hat{a}) = 0$, and note that $\hat{a} < p^*$.

If $\bar{a} > \hat{a}$, a small price cut da from \bar{a} attracts all consumers with valuation $\mathbf{q} \geq \bar{a} - da$, and the effect on the present value of profits is $-\{1 - F(\bar{a}) - 2\bar{a}f(\bar{a}) + \mathbf{d} \min[1 - F(\bar{a}) - \bar{a}f(\bar{a}), 0]\}da$, which is always positive. Then, $\bar{a} > \hat{a}$ cannot be an equilibrium. The intuition is that if the candidate constant price is not low enough the monopolist gains from cutting prices in the next period to attract the consumers of the older generation that have a lower valuation for the good.

Consider now $\bar{a} \leq \hat{a}$, and a deviation where the monopolist chooses in the current period t , $a_t = \mathbf{d}\bar{a} + (1 - \mathbf{d})p^*$, followed by $a_{t+1} = \bar{a}$. That is, in period t the monopolist

charges a price above \bar{a} and in period $t+1$ returns to the equilibrium price \bar{a} . Once \bar{a} is charged, the consumers believe that no lower price is charged in the future, and all the consumers with the valuation above \bar{a} buy the product. In period t , under the deviation, the marginal consumer buying the product can be computed to be $\hat{q}_{t+1} = p^*$. The present value of profits from this deviation is then

$$p^*[1-F(p^*))(1+d)+[d\bar{a}+(1-d)p^*][1-F(p^*)]+d\bar{a}[1-F(\bar{a})+F(p^*)-F(\bar{a})]+d^2V(\bar{a})$$

while the equilibrium present value of profits can be represented by $p^*[1-F(p^*))(1+d)+\bar{a}[1-F(\bar{a})](1+d)+d^2V(\bar{a})$. The difference between the former and the latter can then be obtained to be $(1-d)\{p^*[1-F(p^*)]-\bar{a}[1-F(\bar{a})]\}$ which is greater than zero because p^* maximizes $p[1-F(p)]$. Then, this deviation is profitable and the monopolist charging always $\bar{a} < \hat{a}$ cannot also be an equilibrium. That is, if the monopolist charges a sufficiently low price that it does not have the incentive to cut prices in the next period (to attract the consumers of the older generation that have a lower valuation for the good) then it would gain from deviating and charging a high price for one period in order to identify the consumers that value more the good in the incoming generation. This shows that there are going to be price fluctuations in any MPE.⁸

Let us briefly note that if the analysis is not restricted to MPE one can obtain subgame perfect equilibria in which prices are constant through time (as in Ausubel and Deneckere 1992) at the level obtained when future price commitments are possible, which is also the case with no customer recognition. In such a case, a deviation by the monopolist is “punished” with the equilibrium path in the MPE.

2.2.2. Price Cycles in Equilibrium

⁸ Villas-Boas (2004) shows that this same argument also goes through in a two-type distribution for some parameter values. However, because in a two-type distribution, continuous deviations may not be possible, under some parameter values, there are equilibria with constant prices through time.

Let us now present an equilibrium with price cycles, for the particular case where $f(\mathbf{q}) = 1$, the uniform case. We restrict attention to smooth equilibria - equilibria where in steady-state the prices being chosen by the monopolist result from the maximization of a smooth concave function. As noted below, there are equilibria of this type if \mathbf{d} is sufficiently small.⁹ In the steady state, the monopolist alternates between high and low prices for the new customers, denoted by a^h and a^ℓ , respectively.

If in period t the marginal consumer from the previous generation, $\hat{\mathbf{q}}_t$, is high, the monopolist charges a low price in order to attract not only the new generation consumers but also the old generation consumers who did not buy in the previous period. If, on the other hand, in period t the marginal consumer from the previous generation, $\hat{\mathbf{q}}_t$, is low, the monopolist charges a high price targeted only at the new generation of consumers. In this case we can see that $V(a^\ell) = V^\ell$ is independent of a^ℓ .

One can then obtain that for \mathbf{d} small there is a MPE where the behavior of the games settles in steady state into a price cycle alternating between $(a_t = \frac{8 - \mathbf{d}^2}{16 + 2\mathbf{d}}, \mathbf{a}_t = \frac{1}{2})$ and $(a_t = \frac{6 + \mathbf{d}}{16 + 2\mathbf{d}}, \mathbf{a}_t = \frac{4 + \mathbf{d}}{8 + \mathbf{d}})$. The prices for the new customers are always lower than the prices to the previous customers. However, both prices fluctuate in opposite directions: The price for the new customers is high when the price for the previous customers is low, and vice versa. The monopolist charges a high price to the new customers when it had in the previous period a high demand of new customers. Then, it has relatively small demand from new customers of $4/(8 + \mathbf{d})$ (all from the new generation), and a large demand from the previous customers, $1/2$. In the next period the monopolist charges a low price to the new customers attracting all the customers from the new generation that have a valuation higher than the price being charged (with mass $(10 + \mathbf{d})/(16 + 2\mathbf{d})$), plus the consumers from the previous generation that waited for the low price in this period, with mass $(2 + \mathbf{d})/(16 + 2\mathbf{d})$, for a total demand of new customers of $(6 + \mathbf{d})/(8 + \mathbf{d})$. The demand from the previous customers is equal to all the

⁹ When $\mathbf{d} \rightarrow 0$ all the equilibria converge to the equilibrium presented here.

new customers of the previous generation, $4/(8+d)$. Profits in each of the alternating periods can also be immediately obtained.¹⁰

It is also interesting to check the effect of the discount factor on prices, demands, and profits. In the periods in which the monopolist charges a high price to the new customers, an increase in the discount factor decreases that price, the demand from new customers, and therefore profits from new customers. In the periods in which the monopolist charges a low price to the new customers, an increase in the discount factor increases that price, the price to the previous customers, the demand from new customers, and profits from new customers, and decreases the demand from the previous customers. The average per period profit decreases with an increase in the discount factor.

An increase in the discount factor makes the customers more willing to wait for price cuts. This means that in periods in which the monopolist charges a high price to new customers, the monopolist has less overall demand, which makes it lower its price, and results in lower profits. Given that the marginal customer buying the product has now a greater valuation, in the next period the profits are greater, and the monopolist chooses to charge a greater price to the new customers. However, if one computes the effect of a higher discount factor on the normalized discounted profit (the constant profit that would yield the same net present value of profits), one finds that profits decrease in the discount factor. This is because with a higher discount factor, consumers are “more strategic”, i.e., in the periods in which the monopolist charges a high price more consumers refrain from buying.

It is also interesting to compare the equilibrium profits with the case in which the monopolist is not able to recognize its customers from the current period on. One can

¹⁰The condition on d being small is important because if d were high, more consumers would wait for the lower prices in the future, which means that there is less advantage for the monopolist to charge a high price. That is, if d were high, after charging supposedly the lowest price (in steady-state), $(6+d)/(16+2d)$, the monopolist would gain from cutting the price even further (and “surprising” some of its previous customers). One can check that if $d < 1/2$ there is no such profitable deviation. One can also check that when d is high there is an equilibrium with prices alternating between high and low prices for the new customers, with similar properties to the ones of the equilibrium presented here, and where the low price is such that the monopolist does not want to cut the price even further (for $d \rightarrow 1$ the low price converges to $1/3$).

then obtain, as in the previous section, that the average per period profit without customer recognition is higher than if the monopolist were able to recognize its customers.

Comparing the equilibrium profits with the case in which the monopolist is able to recognize both the previous customers and the consumers' age one obtains that the monopolist is hurt by being able to recognize the consumers' age in addition to recognizing its previous customers. The result is interesting because it reinforces the idea that the monopolist having more information (in this case the consumers' age) ends up hurting the monopolist. The intuition is that when the monopolist recognizes the consumers' age in the market, it offers an even lower price to the consumers that do not buy the product in their first period in the market, which makes consumers refrain even more from buying in the first period.

2.3. Long Lived Consumers

The longer consumers are in the market, the more information they potentially can give about their preferences through their decisions to buy or not to buy at different prices. This means that the firm's policy with respect to its previous customers is exponentially more complicated with the number of periods that a consumer has been in the market. Hart and Tirole (1988) consider the perfect Bayesian equilibrium of this case of long lived consumers with a two-type distribution, $\{q_1, q_2\}$ with $q_1 < q_2$, and only one generation of consumers. They find that in equilibrium, if $d > 1/2$, there is no price discrimination when the horizon tends to infinity, with the monopolist always charging the low price (the valuation of the low type). The intuition for this result is that if a high-valuation consumer q_2 were to buy the product at a higher price, he would reveal that he has high valuation and will have zero surplus from that period onwards. If there were a price strictly above the lowest valuation q_1 for which the high valuation consumer would buy the product with positive probability (such that after that price, if there were no purchase, the monopolist would charge a price q_1 forever), a high valuation consumer buying the product would be better off deviating, not buying the product, and getting a low valuation price from then on. By buying the product the high valuation consumer

would get a surplus of at most $q_2 - q_1$, while if the high valuation consumer waited for one period (and not be identified as a high valuation consumer) he would get a surplus approaching $[d/(1-d)](q_2 - q_1)$, which is greater than $q_2 - q_1$ for $d > 1/2$.

Schmidt (1993) considers the case with any discrete number of consumer types,¹¹ $\{q_1, q_2, \dots, q_n\}$ with $q_1 < q_2 < \dots < q_n$, while restricting attention to the MPE.¹² He finds, as in Hart and Tirole, that, for $d > 1/2$, there is no price discrimination when the horizon tends to infinity, with the monopolist always charging the low price q_1 (the valuation of the low type).

The method of proof used in Schmidt allows us to better understand the relation of this result with the general results on reputation (e.g., Kreps et al. 1982, Fudenberg and Levine 1989). The proof is similar to the one in Fudenberg and Levine (1989) on the reputation of a long-term player facing a sequence of short-term players. Schmidt first shows that if there is a price strictly above q_1 on the equilibrium path, then there is a strictly positive minimum probability of that price being accepted and revealing a consumer type with valuation strictly above q_1 . He then shows that because types $q > q_1$ can build a reputation for being of type q_1 , they will do so. That is, the no-discrimination equilibrium can be seen as a response of the monopolist to the consumers' threat to build a reputation that they have the lowest valuation for the product if the price is above q_1 . In Fudenberg and Levine's model, the type that a consumer would like to be seen as is type q_1 . Given the greater structure of the game considered here (in comparison to the general class of games considered in Fudenberg and Levine) Schmidt is able to extend the results of Fudenberg and Levine to the case of two long-run players, and characterize the equilibrium actions (while Fudenberg and Levine only characterize payoffs). Schmidt looks at a long finite horizon game using backward induction, which is what allows him to show that q_1 acts like a Fudenberg-Levine "commitment type" and rejects all prices above q_1 .

¹¹He considers that it is the monopolist that is the party that has private information (on her costs). We present here the result in terms of private information of the consumers.

¹²The Markov assumption is necessary for the case of any n to guarantee that the continuation payoffs are the same for a price equal or below q_1 (with $n = 2$ this can be shown without the Markov assumption).

Kennan (2001) considers the case in which the consumer types can change randomly through time but with positive serial correlation. He then finds that we can then have stochastic price cycles because (no) purchases indicate a high (low) consumer valuation and are followed by a high (low) price.

It is interesting to discuss in this context of long-lived consumers what happens if the firm is allowed to offer long-term contracts, and the relationship of behavior-based price discrimination with the results from the durable-goods and bargaining literatures.

2.3.1. Long-Term Contracts

Suppose that the firm would be able to offer a contract to a customer committing itself to a sequence of prices for the future to be charged to that consumer. Note that the effect of this possibility is that a consumer would know now that he would not be taken advantage of in the future for revealing his high valuation. Hart and Tirole (1988) consider this situation with the possibility of renegotiation, such that the firm might be able to offer different contracts in the future.¹³ For example, in the market for cellular phone service firms can offer long-term contracts, and can change which long-term contracts to offer in the future. Hart and Tirole show that in such a setting with two consumer types, the firm might now be able to sell to the high valuation consumers at a price above the lowest price. The intuition is that with a long-term contract the monopolist has greater ability to price discriminate. It can get the high valuation consumer to buy the product at an average price per period above the lowest price (low type valuation), because it commits to this average price with a long-term contract.

For example, if the monopolist offers a long-term contract at an average per-period price $p > q_1$, the surplus for the high-valuation consumer if he accepts the contract is $(q_2 - p)/(1 - d)$. If this consumer decides not to buy in this period, the most he consumer is able to get is $d(q_2 - q_1)/(1 - d)$, if the monopolist offers in the next period a contract with an average per-period price of q_1 (the monopolist will never offer a lower

¹³ Laffont and Tirole (1990) consider a two-period version of such contracts with continuous consumption per period in the context of procurement.

average per-period price). Then, if $p = d\mathbf{q}_1 + (1-d)\mathbf{q}_2$ the high valuation consumers accepts the contract, and the monopolist is able to sell to such consumer at a price strictly above \mathbf{q}_1 . As shown in Hart and Tirole, the equilibrium long-term contract is for the monopolist to offer a contract in a number of initial periods with average per-period price strictly above $d\mathbf{q}_1 + (1-d)\mathbf{q}_2$, such that type \mathbf{q}_2 randomizes between accepting and not accepting the contract, and then, after a certain number of periods, the monopolist offers a contract with average per-period price \mathbf{q}_1 , and both types accept the contract.

However, this possibility of selling to the high valuation consumers with an average per-period price strictly above \mathbf{q}_1 is not possible without a long-term contract. Without a long-term contract a high-valuation consumer gets zero surplus after revealing his type, and therefore, must be offered a price below the low-type valuation to accept buying.¹⁴ But then the low valuation consumer would also buy the product, and, therefore, no information would actually be revealed about the type of the customer buying the product. Hart and Tirole then show that, because of this greater ability to price discriminate a firm is better off when it has the ability to offer a long-term contract.

It turns out that this effect of long-term contracts does not occur if the consumer lives only for two periods, with the second period being the last period. In the two-period model presented above it turns out that the introduction of long-term contracts does not have any effect, and the equilibrium with long-term contract is exactly the same as the equilibrium without long-term contracts. This is because the zero surplus obtained by the marginal consumer after revealing his type only lasts for one period.

Battaglini (2004) considers the case of infinitely lived consumers where the preferences change through time following a Markov process, as in Kennan (2001), but allowing for continuous consumption. A consumer's per-period utility in period t is $\mathbf{q}_t q - p$, for q units bought at price p . The monopolist's cost of selling q units is $\frac{1}{2}q^2$. For future reference note that the efficient quantity to be sold in period t is $q^e(\mathbf{q}_t) = \mathbf{q}_t$. The marginal benefit \mathbf{q}_t in period t is private information of the consumer, can only take

¹⁴One can see this as a high valuation consumer maintaining the reputation that he may have a low valuation. See the discussion above.

one of two values, $\{\underline{q}, \bar{q}\}$, with $\bar{q} > \underline{q}$, and evolves through according to a Markov process. The transition probabilities between states are in $(0,1)$, and are denoted by $\Pr(\mathbf{q}_{t+1} | \mathbf{q}_t)$. Types are assumed to be positively correlated over time, $\Pr(\bar{q} | \bar{q}) \geq \Pr(\bar{q} | \underline{q})$ and $\Pr(\underline{q} | \underline{q}) \geq \Pr(\underline{q} | \bar{q})$. At date 0 the monopolist has a prior m that the consumer's type is \bar{q} and a prior $1 - m$ that the consumer's type is \underline{q} .

Battaglini computes the optimal long-term contract. First, he shows that under commitment the optimal contract always involves the efficient quantity being supplied if in the history of the relationship (including the current period) there has been a period in which the marginal benefit has been equal to \bar{q} . That is, with varying types we have the result that a long-term contract supply is at the efficient level in finite time (which is not the case for fixed types). The intuition for this result has to do with the role of the quantity distortions in the contract. Distortions are introduced only to extract more surplus from higher types, and therefore, there is no reason not to offer the highest type the efficient quantity. After any history the rent that must be paid to a high type to reveal himself is independent of the future quantities. That is, the monopolist is the residual claimant on the surplus generated on histories after a high type report, and therefore the quantities that follow such report are the efficient ones. In addition, Battaglini finds that if the history has never had a period where the buyer had type \bar{q} , the quantity distortion vanishes through time as the initial state has less and less information about the current buyer's type.

Battaglini then considers the case in which the contract can be renegotiated, and shows that under general conditions the contract with commitment is renegotiation-proof, and when these conditions fail, the contract is renegotiation-proof after a finite amount of time.

Battaglini's analysis relies heavily on the assumption that there are only two types. As noted in the paper, with n types, then the conditional distribution for each type is represented by a $n-1$ vector, each type has $n-1$ characteristics, and we would need to solve a multidimensional screening problem. It would be interesting to investigate further this n -type case, even if putting some structure on the conditional distribution for each type.

2.3.2. Relationship to Durable Goods and Bargaining

The strategic behavior of consumers when firms practice behavior-based price discrimination is related to settings where a monopolist sells a durable-good, or to settings where two parties bargain, and in which one of the parties has private information. Here, we first briefly discuss some of the forces present in a market where a monopolist sells a durable good, or in a bargaining situation between two parties, in which one party has private information. Then, we relate the durable-goods setting with the behavior-based price discrimination model. For some discussion of the durable-good monopoly literature see, for example, Chapter 1.5 in Tirole (1988). For a survey of the bargaining literature see, for example, Chapter 10 in Fudenberg and Tirole (1991).¹⁵

Durable Goods and Bargaining

Consider the two-period model above, but suppose now that the monopolist sells a product in the first period that lasts for the two periods. Let A be the price of the durable in the first period. Denoting \hat{q} as the type of the marginal consumer in the first period, the surplus of this consumer is $(1 + d_c)\hat{q} - A$ when buying in the first period, and is $d_c(\hat{q} - a_n)$ if waiting for the second period.

The marginal consumer in the first period is then determined by

$$(2.5) \quad \hat{q} = A - d_c a_n$$

Given that the marginal consumer buying the product in the first period is \hat{q} , the optimal price in the second period is $a_n^*(\hat{q}) = \operatorname{argmax}_{a_n} a_n [F(\hat{q}) - F(a_n)]$. Using this we then have that (2.5) defines \hat{q} as a function of A , $\hat{q}(A)$.

In order to obtain the optimal first period price, A , the monopolist then maximizes

$$(2.6) \quad \max_A A[1 - F(\hat{q}(A))] + d_f a_n^*(\hat{q}(A))[F(\hat{q}(A)) - F(a_n^*(\hat{q}(A)))],$$

¹⁵ For early work on the durable-goods monopolist problem see, for example, Stokey (1981) and Bulow (1982).

where the first term represent the first-period profit, and the second term represents the second-period profit. The first order condition that defines the optimal A^* is then

$$(2.7) \quad 1 - F(\hat{q}) - A^* f(\hat{q}) \hat{q}' + d_F \hat{q}' f(\hat{q}) a_n^*(\hat{q}) = 0.$$

Note that for $d_C = d_F$ the marginal consumer buying the product in the first period has a higher valuation than if the firm were selling a non-durable good. To see this note that the first derivative of the objective function above evaluated at $\hat{q}(A) = p^*$ is equal to $f(p^*) p^* [1 - \hat{q}']$ after substituting for $1 - F(p^*) - p^* f(p^*) = 0$ and $p^* = A - d a_n^*(p^*)$. Given that $\hat{q}' = 1/(1 + d a_n^*)$ and $a_n^* > 0$, that derivative is positive, which means that the monopolist should increase A , which implies a higher valuation of the marginal consumer than p^* . One can also obtain for $d_C = d_F$ that the present value of profits is $\hat{q}[1 - F(\hat{q})] + d a_n^*[1 - F(a_n^*)]$, which is strictly below the present value of profits under no customer recognition, as p^* uniquely maximizes $p[1 - F(p)]$. The intuition of this result is that the marginal consumers refrain from buying in their first period in the market because they know that they can get a lower price in the next period.

For the uniform distribution example one can obtain $a_n^*(\hat{q}) = \hat{q}/2$, $\hat{q}(A) = 2A/(2 + d)$, and $A^* = (2 + d)^2 / (8 + 2d)$. One can also get that $2/(4 + d)$ consumers buy in the first period, while $(2 + d)/(8 + 2d)$ consumers buy in the second period.

The model above can also represent a bargaining situation where now there is a single buyer, and if the buyer does not take the first offer A^* then he is offered a_n^* in the second period. In such a setting one can then obtain that the private information of the buyers leads to an inefficient outcome for some consumer types (if rejection occurs in the first period).

In a durable-goods setting, if new generations of consumers come into the market in every period, there are incentives for the monopolist to raise its price in order to try to extract more surplus from the consumers who have a high valuation and who have entered the market most recently. This can then generate price cycles in which prices come down to clear the demand from low valuation consumers, and then go up to better

extract the surplus from the consumers with high valuation who just entered the market. This setting is analyzed in Conlisk et al. (1984), and Sobel (1984, 1991). Although having the flavor of the results in subsection 2.2 for overlapping generations of consumers and behavior-based price discrimination, and as also discussed below, the results are different, as we can have price cycles in the behavior-based price discrimination model, but constant prices in the corresponding durable goods model.

In some situations the monopolist may also have some private information regarding its costs, so that the price offers can potentially reveal some information about the monopolist's costs. Fudenberg and Tirole (1983), with a bargaining set-up, characterize the set of equilibria in two-period games when the monopolist and the buyer each have two potential types (two-sided incomplete information). They show that this additional private information may lead to a continuum of perfect Bayesian equilibria. Ausubel and Deneckere (1992) consider the infinite horizon version of this two-sided incomplete information model, and show that we may have (stationary) equilibria in which prices stay high, and the seller tries to maintain a reputation of having a high cost.

Relationship of Durable Goods to Behavior-Based Price Discrimination

When a monopolist is selling a durable good through time, consumers refrain from buying in the initial periods because they foresee that the monopolist may cut its price in future periods. In such a setting consumers may prefer to forsake the benefits of the product if buying earlier, with the lower price if buying later. On the other hand, with customer recognition and behavior-based price discrimination for a non-durable, consumers refrain from buying in the initial periods because they foresee that the monopolist may cut its price in future periods, and therefore they will be identified as lower valuation consumers and get lower prices for the future. This difference between durable-goods and behavior-based price discrimination for non-durables leads to different consumer surplus effects from purchasing the product, and therefore different market implications.

When buying a durable-good the consumer pays a price and gets a benefit of using the product for the duration of the product's life. Consumers for whom the present

value of future benefits is greater than that price may be willing to purchase the product. However, under customer recognition and behavior-based price discrimination for a non-durable good, the marginal consumer buying the product pays the price and gets a benefit in the current period, but then gets zero surplus in all future periods. Therefore, in order for a consumer to be willing to buy, the initial price must be so low, that even consumers with very low valuation may be willing to buy the product. For an infinite horizon with two types, Hart and Tirole (1988) show then that the durable-good case is better for the monopolist than the case of a non-durable with the ability to recognize customers and price discriminate according to past behavior (Hart and Tirole consider this possibility in terms of rental of the durable good).

In the long lived consumers with two consumer types model that they consider, Hart and Tirole also find that the durable-good case is exactly the same as when the monopolist can offer long-term contracts (and different from short-term contracts), as the separation between types can be done ex-ante. In the two-period model considered above the durable-good case results in exactly the same outcome as the long-term contract case, and generates exactly the same outcome as the customer recognition case for a short-term sales of a non-durable, or equivalently with short-term rentals of a durable. This is because, in a two-period model, the consumer surplus effects of purchasing a durable-good are the same as purchasing a non-durable with customer recognition, as the zero surplus of the marginal consumers under customer recognition lasts only for one period.

In the case of overlapping generations of consumers, with consumers living for two periods (and without the ability to recognize the customer's age) selling a durable good may not generate price cycles, as selling the durable good for a consumer that only uses the product for one period requires a much lower price than selling the durable good for a consumer who uses the product for two periods (Villas-Boas 2004). That is, with overlapping generations of consumers, selling a durable good does not yield the same outcome as selling a non-durable with customer recognition (with or without long-term contracts).

Thus, in general, the sale of a durable good is not the same as a sequence of short-term rentals. Although the price falls over time, the price a consumer faces is not based

directly on its own past behavior. Loosely speaking, the commitment involved in selling a durable good lets the monopolist commit to not use behavior-based pricing.

2.4. Two-Good Monopoly

In order to serve as an introduction to the next section on competition, and to serve as a benchmark, consider now the case of a monopoly selling two goods, A and B . The presentation here follows closely part of Section 5 of Fudenberg and Tirole (2000). To focus on the interaction between the two goods we set up preferences such that consumers buy a unit of one of the goods in every period. Indexing the relative preferences for B over A as \mathbf{q} , let the valuation per period of a consumer of type \mathbf{q} be $v - \mathbf{q}/2$ if the consumer buys good A , and $v + \mathbf{q}/2$ if the consumer buys good B , with v “large” and \mathbf{q} distributed in $[\underline{\mathbf{q}}, \bar{\mathbf{q}}]$, where $\underline{\mathbf{q}} = -\bar{\mathbf{q}} < 0$, with cumulative distribution function $F(\mathbf{q})$, strictly positive density $f(\mathbf{q})$, and $F(\mathbf{q})$ is symmetric about zero and satisfies the monotone hazard rate (MHR) condition that $f(\mathbf{q})/[1 - F(\mathbf{q})]$ is strictly increasing in \mathbf{q} . The parameter v is assumed large, such that the monopolist chooses prices such that all consumers buy one unit of one of the goods in every period. For this reason, the monopolist’s production will be constant across the pricing regimes we consider, so that the costs of production are a constant that can be ignored.

Let a and b be the prices charged in the first period for products A and B , respectively, \mathbf{a} and \mathbf{b} be the prices charged in the second period for products A and B , respectively, for consumers that bought the same product in the previous period, and $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ be the prices charged in the second period for products A and B , respectively, for consumers that bought a different product in the previous period.

Consider first the case in which long-term contracts are not available. Then the firm will charge $a = \mathbf{a} = b = \mathbf{b} = v$, for a present value of profits of $(1 + \mathbf{d})v$. Note that consumers do not switch products.

Consider now the case in which long-term contracts are available (with commitment not to renegotiate). Then, in the first period the monopolist can offer four product consumption experiences: product A in both the periods, product A followed by

product B , product B followed by product A , and product B in both periods. By symmetry it is enough to analyze the interval $[\mathbf{q}, 0]$. As argued in Fudenberg and Tirole, incentive compatibility requires that consumers in an interval $[\mathbf{q}, \hat{\mathbf{q}}]$ choose product A in both periods (which we call AA), and consumers in the interval $[\hat{\mathbf{q}}, 0]$ choose product A followed by product B (which we call AB). In order for the type 0 to be indifferent between buying and not buying, and between the “switching” bundles AB and BA , it must be that the price for each of these bundles is $(1+d)v$. Indifference for type $\hat{\mathbf{q}}$ between bundles AA and AB requires that the price of AA be above the price of AB by $-d\hat{\mathbf{q}}$. Thus, it is as if the monopolist first sold all consumers a “switching” bundle at price $(1+d)v$, and then offered an “upgrade” to AA or BB for a premium of $-d\hat{\mathbf{q}}$. The present value of profits is then $(1+d)v - 2d\hat{\mathbf{q}}F(\hat{\mathbf{q}})$, where the optimal $\hat{\mathbf{q}}$ satisfies $F(\hat{\mathbf{q}}) + \hat{\mathbf{q}}f(\hat{\mathbf{q}}) = 0$. Note that the optimum has some consumers switching products across periods. Since consumer preferences are the same in both periods, this switching is inefficient; it is used to extract rents for the privilege of not switching. For the uniform distribution one can obtain $\hat{\mathbf{q}} = \mathbf{q}/2$, so that one half of the consumers switch products from the first to the second period.

Fudenberg and Tirole also show that the monopolist can do better than the above deterministic menu, with a randomized menu where consumers in $[\hat{\mathbf{q}}, -\hat{\mathbf{q}}]$ get a $(1/2, 1/2)$ randomization between products A and B . This allows the monopolist to extract a greater surplus from the consumers that get no “switching”. Again, as in the deterministic menu, we have some inefficient (stochastic) switching by some consumers.

When we reach the second period, as stated above, “switching” consumers are consuming a product that is not the best for them. This means that there are gains to be made from the monopolist renegotiating the product that is offered to those consumers. This renegotiation may then affect the choices of consumers in the first period (and the monopolist’s offers). It would be interesting to investigate whether we would still have inefficient switching in equilibrium if the monopolist can offer long-term contracts subject to renegotiation.

In order also to compare with the next section consider now the case of switching costs, where a monopolist sells two products in each of two successive periods, consumers have the same valuation v for each product per period and incur in a cost s if they change products from the first to the second period. It is clear that in this situation the best the monopolist can do is extract v per period per consumer (with a price equal to v), and there is no switching products from the first to the second period. This can be accomplished either with short or long-term contracts.

Consider now in the model above (with heterogeneous consumers) the role of the introduction of switching costs s (suppose s small). The price of the switching bundles can then be at most $(1+d)v - ds$ and indifference for type \hat{q} between bundles AA and AB requires that the price of AA be above the price of AB by $d(s - \hat{q})$. Thus, comparing with the no switching costs case, the price of the switching bundle is now lower, but the premium to upgrade to the non-switching bundle became now greater. The present value of profits is now $(1+d)v - 2d\hat{q}F(\hat{q}) - ds[1 - 2F(\hat{q})]$, where the optimal \hat{q} satisfies $F(\hat{q}) + (\hat{q} - s)f(\hat{q}) = 0$. Note that the optimum has some consumers switching products across periods, but the number of switching consumers is decreasing in the switching cost s . Note also that switching is inefficient, and that the monopolist profit and welfare are decreasing in the switching cost s . For the uniform distribution one can obtain $\hat{q} = (q + s)/2$, so that less than one half of the consumers switch products from the first to the second period.

3. Competition

Several new issues arise in models of behavior-based price discrimination with multiple firms. Starting with the most obvious, firms can try to “poach” their rivals’ customers by giving them special “introductory” prices.¹⁶ This raises the questions of how much switching we should expect to occur, and of its efficiency consequences. At a

¹⁶In 1994, about 20% of all U.S. households changed their provider of long-distance telephone services (Schwartz 1997).

more theoretical level, we have already seen that in equilibrium a monopolist without commitment power can be made worse off by the ability to condition the price it charges a customer on that customer's past decisions, because consumers will foresee this condition and adjust their earlier behavior. The same sort of foresight can operate in models with multiple firms, but now its impact on profit is *a priori* ambiguous, because of the interactions between the customers' behavior (basically the elasticity of demand) and the equilibrium in the pricing decisions of the firms. Furthermore, while a monopolist with commitment power can always do at least as well when behavior based discrimination is possible (by committing itself to ignore past behavior in setting prices), a group of oligopolists with commitment power can all be worse off if all of them become able to discriminate based on past customer behavior, as the better information may lead to more intense price competition (see subsection 5.3 below).¹⁷ For this reason, while each firm has a dynamic incentive to adjust its prices so that it learns more about the consumers and can better segment the market, the firm also has an incentive to reduce the information that is obtained by its rivals.

The way that these various effects combine to determine equilibrium prices and allocations depends on the nature of preferences and on the form of market competition. The first part of this section considers Fudenberg and Tirole (2000)'s analysis of a two-period model of competition in short-term contracts, and some variations on the distributions of consumer types studied by Chen and Zhang (2004) and Esteves (2004). The second part discusses Villas-Boas' (1999) extension of the two-period model to an infinite horizon with overlapping generations of consumers, each of whom lives only for two periods. We then return to the two-period setting to present Fudenberg and Tirole's analysis of competition in simple long-term contracts, meaning that firms offer both a "spot" or one-period price and also a long-term commitment to supply the good in both periods. Finally, we compare the predictions of these models of "pure price discrimination," where the past matters only for the information it provides about preferences, to models of switching costs, where past decisions are directly payoff relevant, and may also provide information, as in Chen (1997) and Taylor (2003).

¹⁷Of course, a single oligopolist with commitment power who is given the ability to condition prices on customer history cannot be made worse off, provided that none of the other firms are allowed to have this ability.

3.1. Two Periods, Short-Term Contracts

Following FT, suppose that there are two firms, A and B who produce non-durable goods A and B , respectively, at constant marginal cost c . There are two periods, 1 and 2; each period a consumer can either consume a unit of good A or a unit of good B or neither, but not both. There is a continuum of consumers, whose preferences are quasi-linear in money and are indexed by $\theta \in [\underline{\theta}, \bar{\theta}]$, where $\underline{\theta} = -\bar{\theta} < 0$. The consumption utility from goods A and B is $v - \theta/2$ and $v + \theta/2$, respectively, so that θ measures the consumer's preference for good B over good A . There is a known distribution F on θ , which is assumed to be symmetric about 0. Fudenberg and Tirole assume that F is smooth, with density f , and that F is symmetric and that it satisfies the monotone hazard rate (MHR) property that $f(\theta)/[1 - F(\theta)]$ is strictly increasing in θ ; their sharpest results are for the special case of the uniform distribution. Esteves (2004) considers the case where F has a two-point support;¹⁸ Chen and Zhang (2004) assume that F is concentrated on the three points $\underline{\theta}, 0, \bar{\theta}$. Fudenberg and Tirole assume that all agents use a common discount factor δ ; the other papers suppose that firms use discount factor δ_F while consumers use the possibly different discount factor δ_C .

With simple short-term contracts, and no commitment power, each firm will offer a single first-period price, which we denote a and b , respectively. In the second period, each firm can offer two prices, one to its own past customers and another price to all others. (We will assume that the reservation value is high enough that all consumers purchase in the first period, so that agents who didn't purchase from firm A must have purchased from firm B.¹⁹) Note that if firms do not observe the identities of their customers, there is no link between the periods, and the equilibrium reduces to two repetitions of the static equilibrium. Our question is how the prices and efficiency of the

¹⁸Esteves supposes that the two mass points are in the interval $[0,1]$, symmetric about the point $1/2$; to map her notation to ours suppose that the mass points are at $y_A = t(2x_A - 1)$ and $y_B = -y_A$, and that $v = v' - t(1 - x_A)$, where v' is the reservation value in her notation.

¹⁹Chen and Zhang consider an extension of their model to the case where agents with $\theta = 0$ have lower reservation values; in this case not all agents purchase in the first period.

equilibrium with short-term contracts and customer poaching compare to that of the static benchmark.

Under FT's assumptions, the static one-period problem is well behaved: each firm's objective function, $\pi^j = F(p^j - p^i)(p^j - c)$, is strictly quasi-concave, so that firms are never willing to randomize, and the game has a unique equilibrium, namely

$$p^A = p^B = \frac{F(0)}{f(0)} + c.$$

In the case of a uniform distribution, this simplifies to $p = c + \frac{\bar{\theta} - \theta}{2} = c + \bar{\theta}$, so that each firm's profit is $\frac{\bar{\theta}}{2}$. Moreover, in the uniform case the dynamic equilibrium is also in pure strategies, and can be characterized with first-order conditions. With the discrete supports specified in the other two papers, the static equilibrium is in mixed strategies, which makes the calculations more complex and the intuition more subtle. For this reason we use the FT case for exposition, and try to explain informally the effects of the other distributional assumptions.

3.1.1. Analysis of the Two-period Model under the MHR Assumption

A standard argument shows that at any pair of first-period prices such that all consumers purchase and both firms have positive sales, there is a cut-off θ^* such that all consumers with types $\theta < \theta^*$ purchase from firm A in the first period.²⁰ Given this cut-off, the second period game is as depicted in Figure 1: consumers to the left of θ^* lie in "firm A's turf" and the consumers on the right lie in firm B's. On firm A's turf, firm A offers price α , while firm B offers price $\hat{\beta}$; on B's turf B charges β and A charges $\hat{\alpha}$. Thus a consumer on firm A's turf will stick with good A if $-\theta < \alpha - \hat{\beta}$, and otherwise

²⁰To deal with out-of-equilibrium beliefs, we suppose that if first period prices are such that no consumer is expected to buy from firm A, a consumer who unexpectedly does purchase from A is assumed to have type $\underline{\theta}$, and similarly a consumer who unexpectedly purchases from B is assumed to have type $\bar{\theta}$.

will switch to good B. If θ^* is very near the endpoint $\underline{\theta}$, then A's turf is very small, and consists only of agents with a strong preference for A, and firm A can charge the monopoly price in this market and not lose any sales to firm B. The paper shows that this occurs when $\theta^* < \theta^-$, where θ^- is the "isoelastic point" where $F(\theta^-) + f(\theta^-)\theta^- = 0$ so that marginal revenue equals 0. In this case firm A sets $\alpha = v - \theta^* / 2$ and sells to everyone on its turf, while firm B sets $\hat{\beta} = c$. Otherwise, both firms will have positive sales in each market, which implies that the "poacher's" price in a market must be lower than the incumbent's.

The intuition for this result comes from the fact that on the interior of firm A's turf, its second-period reaction function reflects a trade-off between losing marginal customers at θ^A and inframarginal rents on types below θ^A , and so the reaction function does not depend on the first-period cut off θ^* , while decreasing θ^* decreases B's sales on A's turf, and so makes firm B price more aggressively, as shown in Figure 2, where the curves R are the reaction curves on firm A's turf when it had the entire first period market (which is why they intersect on the 45 degree line) and \hat{R} is firm A's reaction curve on its turf as a function of the first-period cut-off θ^* .

The next step is to work backwards and determine the equilibrium first period prices. Before presenting the analysis, we can identify some general considerations to keep in mind:

- 1) If consumers are forward looking (as assumed by FT) they realize that they will be offered a "poaching" offer in the second period. FT show that this can lead to a less elastic first-period demand and hence higher first period prices.
- 2) Firms may distort first period price to increase second period profit. Specifically, each firm would rather that its opponent have less information about consumer preferences, and is willing to distort first period prices for that purpose. Moreover, this preference is large enough that firms do better when neither of them has any information about consumer's identities. The impact of this consideration depends on the way that changes in price change what is learned about consumer demand, which in turn depends on the distribution of types.

- 3) If customers buy from their preferred firm in the first period (as they do in FT) then second-period switching lowers welfare.

To explore this second point in more detail, we present a more detailed analysis of second-period competition in the uniform case than is given in FT. Solving for the intersection of the second period reaction curves (corresponding to equation (6) and (7) in

FT) shows that $\theta^A = \frac{\theta^* + \underline{\theta}}{3}$, $\theta^B = \beta - \hat{\alpha} = \frac{\bar{\theta} + \theta^*}{3}$. In its home turf, firm A sells to

types below θ^A ; this is mass $\frac{\theta^* - 2\underline{\theta}}{3(\bar{\theta} - \underline{\theta})} = \frac{\theta^* + 2\bar{\theta}}{6\bar{\theta}}$. On B's turf A sells to types between

θ^* and θ^b ; this has mass $\frac{\bar{\theta} - 2\theta^*}{3(\bar{\theta} - \underline{\theta})} = \frac{\bar{\theta} - 2\theta^*}{6\bar{\theta}}$.

So the second period profit of firm A is

$$\frac{(\theta^* - 2\underline{\theta})^2}{18\bar{\theta}} + \frac{(\bar{\theta} - 2\theta^*)^2}{18\bar{\theta}} = \frac{5\theta^{*2} + 5\bar{\theta}^2}{18\bar{\theta}},$$

provided that $\theta^* > \underline{\theta} = \underline{\theta}/2$, so that there is poaching in both markets. By symmetry this is also the second period profit of firm B.

Note that the symmetric outcome $\theta^* = 0$ is the global minimum of firm A's second period profits; it does better not only with a larger first period market share, but also with a smaller one! Specifically, when $\theta^* = 0$ the second period profit is

$\frac{5\bar{\theta}^2}{18\bar{\theta}} = \frac{5\bar{\theta}}{18}$. As θ^* increases to $\bar{\theta}/2$ profit increases to

$\frac{(\bar{\theta}/2 + 2\bar{\theta})^2}{18\bar{\theta}} = \frac{25\bar{\theta}^2/4}{18\bar{\theta}} = \frac{25\bar{\theta}}{72}$. From this point on, there is no competition in firm B's

second period market. Firm A's profit is $\frac{(\theta^* - 2\underline{\theta})^2}{18\bar{\theta}}$, which converges to the static

equilibrium value of $\frac{\bar{\theta}}{2}$ as θ^* goes to $\bar{\theta}$.

This shows that both firms do best when neither has first-period information. When θ^* is near the endpoints, firms have less precise information in the larger market, and hence competition there is less intense.

Perhaps surprisingly, in the uniform case this second-period consideration has no impact on first period pricing. This is because the first-period equilibrium will have equal market shares, i.e., $\theta^* = 0$, and because $\theta^* = 0$ leads to the lowest level of second-period profit, there is no first-order effect when it changes. For this reason, the only reason that first-period prices differ from the static equilibrium is that consumer demand differs. In the static case, the cut-off θ^* shifts one-for-one with prices, while in the dynamic setting $\theta^* = \frac{b - a + \delta_C(\hat{\alpha}(\theta^*) - \hat{\beta}(\theta^*))}{1 - \delta_C}$, because type θ^* must be

indifferent between the different plans to switch and get the “poaching price” next period, and so must be indifferent between buying good A now at price a and then buying B tomorrow at price $\hat{\beta}$, or buying B now at price b and then buying A tomorrow at price $\hat{\alpha}$. In the uniform case this leads to a less elastic first period demand ($|\partial\theta^*/\partial a| < 1$) and hence higher prices; with zero production costs and consumer types distributed on the unit interval, the first-period price is $1 + \delta/3$ and the second period prices (on the equilibrium path) are $2/3$ to the firm’s old customers and $1/3$ to the customers it is trying to “poach.”

This finding for the uniform case leaves open the possibility that for other distributions the second-period-profit effect could have an impact on first-period pricing. However, it seems plausible that $\theta^* = 0$ is the global minimum of second period profits for general symmetric distributions, so that the effect of second period profit on first period decisions vanishes, provided that the first-order approach is valid. However, the fact that firms would do better in the second period with a less symmetric first-period outcome suggests a possible non-concavity in the problem. The MHR assumption makes the static optimization problem concave, which implies that the firms’ first-period objective functions are concave for discount factors close to 0 and any distribution that satisfy MHR; FT show that they are also concave under the uniform distribution for all

(common) discount factors (that is, $\delta_f = \delta_c = \delta$). However, concavity does not seem to be implied by the MHR condition, and when it fails there can be mixed-strategy equilibria. To investigate this possibility it may be interesting to abandon the first-order approach altogether, and work with discrete types, as in Esteves and Chen and Zhang.

3.1.2. Discrete Distributions

In Esteves' model, whenever the difference in price is less than $y_B - y_A$ each firm buys from their preferred firm, while if the difference is larger than this all consumers buy from the same firm and no information is revealed, which corresponds to the case $\theta^* = \pm\bar{\theta}$ in FT. Again as in FT, the second-period profits are symmetric in the information: firms do better when the first period prices are very different, but as far as second period prices go they are indifferent between having a large turf or a small one. To simplify the analysis, Esteves assumes that consumers are completely myopic. The first-period equilibrium is in mixed strategies, and she shows that the probability that both firms have positive first-period sales decreases as they become more patient. Moreover, she shows that first period prices tend to fall as the discount factor increases.

Chen and Zhang suppose that there are three types. A mass γ is loyal to A (they get 0 utility from B, so they buy A whenever the price is below their reservation value v), a mass γ is loyal to B, and a mass of $1 - 2\gamma$ who are exactly indifferent. Neither firm can hope to sell to the loyalists of the other, so what each firm wants to do is distinguish its loyalists from the neutrals. Starting from equal first period prices, a small increase in firm A's price shifts all of the neutrals to firm B, and results in an asymmetric knowledge about the consumers: firm A knows who its loyalists are, but firm B does not. Thus, in contrast to the previous two papers, the firm with the smaller first-period sales has strictly higher second period profits. They show that this leads to prices that are, on average, higher than in the static equilibrium, even when consumers are myopic.

We should point out some unusual features of the assumed demand distribution. Specifically, second period profits when consumer types are known are exactly the same as in the static model, while in general we may expect that known types could lead to

fiercer competition and lower profit. This suggests that competition in the static model is particularly fierce, which observation may help explain why equilibrium profits here are higher than when firms lack information on purchase history.

3.1.3. Welfare

Finally we compare the welfare effects of price discrimination in the three models. In FT, the first-period outcome is efficient, so all second-period switching lowers welfare. In Esteves, both the static equilibrium and the first-period equilibrium of the two period price discrimination game are not fully efficient, due to the randomized nature of the equilibrium. Moreover, when the first period prices reveal the customers' types, the second period outcome is efficient, and there is no switching, even though firms offer lower second-period prices to their opponents' customers. This stems from the two-point distribution of demand, and would not extend to a discrete model with more types. Combining these two observations, we see that price discrimination can increase efficiency provided that it doesn't lower first-period efficiency too much, and she shows that this is indeed the case. In the Chen and Zhang model, efficiency considerations are moot, as the only consumers whose purchases change when price discrimination is allowed are those who are completely neutral. There can however be efficiency implications of price discrimination when the reservation value of the neutrals is less than the other players, as price discrimination allows the firms to offer the neutrals a second-period price that is below their reservation value without losing sales to the loyalists.

3.2. Infinite Lived Firms, Overlapping Generations of Consumers, and Short-Term Contracts

Villas-Boas (1999) extends the FT model to the case of two infinite-lived firms facing overlapping generations of consumers. Each consumer lives for two periods, and each generation has unit mass. Each firm knows the identity of its own past customers, but not those of its opponent, and it does not observe the consumer's "age," so it cannot distinguish young consumers from old ones who bought from the opponent last period.

The basic setup of the model, and the notation, are the same as in FT, with θ uniform on $[-1/2, 1/2]$ and zero production costs. The timing of the game is a bit different, as in each period the firms first simultaneously set the price for new customers, and then set the prices to existing customers after observing the price the competitor charges to new ones.

In order to focus on the dynamics of price discrimination, and abstract from (possibly important) repeated game aspects, the paper restricts attention to the state-space or Markov perfect equilibria (MPE) of the game. Given the linear-quadratic nature of the model, there are MPE in which the strategies are piecewise affine in the state variable, and these are the ones considered in the paper.²¹ As a benchmark case, note that the MPE here would be exactly the outcome in FT if, as in FT, firms can recognize both their own and the opponent's customers, and all prices are set simultaneously. If firms can recognize both types of old customers, but prices are set sequentially as specified above, timing, the prices will be $1 + \delta_C - \delta_F / 4$ to new customers, and the prices will be $3/4$ and $1/2$ to the firm's and the competitor's old customers, as opposed to $2/3$ and $1/3$ with simultaneous price setting. (Prices are higher with sequential moves because the reaction curves slope up, this is a form of the "puppy dog effect" (Fudenberg and Tirole 1984).)

We now turn to the MPE of the game where firms only recognize their own customers. If the reservation value is high enough that all consumers purchase every period, Villas-Boas shows that the equilibrium is again characterized by cut-offs θ_t^* such that each new consumer arriving in period t purchases from firm A iff their type $\theta < \theta_t^*$. Thus the payoff-relevant state in each period is simply the previous period's cutoff.

The easiest part of the model to solve is the prices firms charge to their old customers. Since these consumers will leave the market at the end of the period, neither they nor the firm need to consider future periods in making their decision, and since prices are set after observing the rival's poaching price, the firm faces a simple static maximization. In contrast, the price set to unrecognized consumers must take into

²¹The reason to consider piecewise affine strategies instead of affine ones is that there are "kinks" in the value functions corresponding to the states where a firm completely retains all of its clientele; these kinks are roughly analogous to the points $\pm\theta^*$ in FT.

account that some of these are new agents who will purchase again in the next period, and the demand of new customers must also take the future into account.

Neither of these complications is present in the case of complete myopia, $\delta_F = \delta_C = 0$. Here the cutoff converges to the steady state with equal market shares. Except possibly in the first period, the convergence is monotone, and customers sort themselves as in FT: those with strong preference for one firm buy from that firm in each period, while those with more intermediate preferences switch. As in FT, prices to the recognized consumers are lower than in the static case. Prices to the unidentified consumers are also lower than the static prices, while in FT the first period price equals the static price when firms are myopic; this is because the pool of unidentified consumers here contains both new consumers (as in the first period of FT) and old consumers who prefer the other firm.

Villas-Boas then considers the case of myopic firms but patient consumers.; this differs from the previous analysis in that consumers take into account the prices they will be charged next period; it differs from FT because a consumer who buys A in the first period is offered a second-period price for B that is tailored to a mixture of “A-preferrers” (i.e. $\theta < \theta_t^*$) and new agents, as opposed to a “poaching price” for A-preferrers alone. This mixed price will in general be less responsive to changes in θ^* than is the poaching price, which makes the marginal new customers more responsive to changes in price. For this reason, the price to new consumers is lower than in FT, and in fact it goes to 0 as $\delta_C \rightarrow 1$.

Finally Villas-Boas considers the case where δ_F and δ_C are both non-zero. As in Esteves and Chen and Zhang, patient firms have an incentive to shift their prices in a way that softens future competition, which here leads to higher prices. In the case $\delta_C = \delta_F = \delta \rightarrow 1$, the price charged to new consumers converges to 0 with δ , while the price charged to old ones converges to $1/2$. Thus firms are worse off than when they could credibly share their information. We discuss the issue of information sharing in Section 5.2 on credit markets.

3.3. Long-Term Contracts

As we remarked in Section 2, long term contracts are used in a variety of consumer markets. This section considers the impact of competition in simple long term contracts in the setting of the two-period FT model. Specifically, we suppose that in the first period firms A and B offer to sell their good this period at spot prices a and b , and that they also offer long-term contracts to supply the good in both periods for A and B. In the second period, firms know the first-period prices announced by their rival, and they also know from whom each consumer purchased, but do not observe the contracts chosen by their rivals' customers.

If a firm chooses to only sell long-term contracts, it would prevent poaching by its rival; but the fact that a monopolist with commitment power induces switching suggests that the complete lock-in will not be optimal here either. And indeed, Fudenberg and Tirole show that the equilibrium has the form depicted in Figure 3: consumers who most prefer A buy a long-term contract from A; this is the interval $[\underline{\theta}, \underline{\theta}^A]$. The next interval $[\underline{\theta}^A, \theta^A]$ purchases A in each period on the spot market, interval $[\theta^A, \theta^*]$ buys from A in the first period and B in the second, and so on.²² Thus, as in the case of short-term contracts, there is inefficient switching.

A key fact in determining the equilibrium outcome is that when firm A locks in more of its customers with long-term contracts (increases $\underline{\theta}^A$), it becomes more aggressive on its turf in the second period, as cuts in its second-period price α do not reduce revenue from locked-in consumers.²³ Since changes in $\underline{\theta}^A$ do not change firm B's

²²Because this is a deterministic model, equilibrium prices must satisfy the no-arbitrage condition $A = a + \delta\alpha$, so that all consumers who plan to purchase from A in both periods are indifferent between purchasing the long term contract or a sequence of short-term ones. The results reported here rely on the tie-breaking assumption that when the no-arbitrage condition holds, it is the customers who most prefer A who choose the long-term contract. Intuitively, there is an option value to the sequence of short-term contracts, and this value is increasing in the probability that the customer decides to purchase B instead of A in the first period. It seems plausible that this option value is higher for consumers with higher values of θ , and indeed this tie-breaking rule corresponds to taking the limit of models where the second-period valuation is imperfectly correlated with first period value, and the distributions are ranked by first-order stochastic dominance in the first-period valuation. Some sort of tie-breaking rule is needed in any deterministic model where there are multiple ways of purchasing the same consumption stream.

²³Note that firm A does not directly set $\underline{\theta}^A$, instead, this switchpoint is determined by the condition that equilibrium prices satisfy the no-arbitrage conditions $A = a + \delta a$ and $B = b + \delta b$.

prices on firm A's turf, increases in $\underline{\theta}^A$ lead both firms to set lower prices. Moreover, the monotone hazard rate condition implies that the slopes of the second-period reaction curves are less than 1, so increases in $\underline{\theta}^A$ move the switchpoint θ^A to the right, which means fewer agents switch. Hence, if the firms use any long term contracts at all, there will be less switching than with short term contracts.

Fudenberg and Tirole show that on the path of a symmetric equilibrium, firms do use some long-term contracts, so there is less switching (and more efficiency) than with short term contracts. The intuition for this is as follows: by locking in some of its customers, a firm can commit itself to more aggressive second-period pricing on its own turf, which induces a lower second-period poaching price from firm B. The marginal first-period A purchaser plans to switch in the second period, so lowering B's poaching price lets firm A charge a higher first-period price, which raises its profit.

Conversely, a firm always uses some short term contracts. Indeed, using only short-term contracts dominates using only long-term ones whenever first period sales exceed the isoelastic point θ^\sim . To see why, suppose that all customers in the interval $[\theta, \theta^*]$ buy a long-term contract from firm A, and that $\theta^* > \theta^\sim$. Now suppose that firm A deviates and offers only a short-term contract in the first period, where the price a is set so that $a = A - \delta\hat{\beta}(\theta^*) + \delta\theta^*$, where $\hat{\beta}(\theta^*)$ is firm B's poaching price when none of firm A's customers have a short-term contract. This price has been chosen so that a consumer of type θ^* gets exactly the same utility from purchasing A in the first period at price a and then buying B at the poaching price as it received from purchasing the long-term contract from A, and since the change does not affect competition on firm B's turf it leads to the same first-period cutoff.²⁴ Moreover, firm A would receive exactly the same payoff as with the long-term contract by offering a second-period price on its turf of $\alpha'' = \hat{\beta}(\theta^*) - \theta^*$, as this price will induce all of its first period customers to purchase from it again. However, when $\theta^* > \theta^\sim$, this pricing is more aggressive than is optimal, and firm A does strictly better by raising its second-period price, even though this leads some customers to switch.

²⁴It is easy to see that there is an equilibrium with the same cutoff. Fudenberg and Tirole prove that (under their tie-breaking rule) any profile of first period contracts leads to a unique first period cutoff.

Fudenberg and Tirole go on to show that the equilibrium they construct remains an equilibrium when more general contracts are allowed, but they do not discuss uniqueness, and it is an open question whether more general contracts can lead to qualitatively different outcomes. Moreover, as with the analysis of short term contracts, the MHR condition does involve a loss of generality; the effect of long-term contracts with the sorts of distributions studied by Esteves (2004) and Chen and Zhang (2004) is open as well.

3.4. Switching Costs

To conclude this section we return to the case of short-term contracts to compare the impact of purely information-based duopoly poaching with price discrimination in the presence of switching costs. These costs are real social costs in, e.g., complementary equipment or in learning how to use the product; as such they differ from “pecuniary” switching costs such as cancellation fees.

Before addressing price discrimination, we briefly discuss the forces present in models of switching cost without price discrimination. (For extended surveys of the switching costs literature see Klemperer 1995, and Farrell and Klemperer 2004.²⁵) In two-period models such as Beggs (1989) and Klemperer (1987a), all consumers are locked-in in the second period, while none are in the first. Second-period lock-in leads second-period prices to be higher than without switching costs, while first-period prices are lower, as firms compete for the rents from locked-in customers. Finally, consumers in the first period foresee being locked-in in the second period, and become less price sensitive, which is a force towards higher prices in the first period.

To illustrate these forces, we will use a simple two-period model. Each firm sells a fixed and given product in each of the two periods. Each consumer buys at most one unit in each period. Consumers are uniformly distributed along a Hotelling segment, whereas firms are located at the extremes of the segment (as in the previous subsection).

²⁵ For early papers on switching costs see also, for example, von Weizsacker (1984), Klemperer (1987b), Farrell and Shapiro (1988, 1989). For a recent survey on information technology and switching costs see Chen and Hitt (2005).

Transportation costs are t per unit and production costs are zero. A fraction s of the consumers who buy from firm i in the first period incur a high switching cost if buying from firm $j \neq i$ in the second period (so that they never switch firms in the second period). The parameter s can then be seen as an index of switching costs. The remaining consumers, in fraction $1-s$, have zero switching costs.

Given these assumptions, we can start by determining second-period demand for each firm. Let q_1^i be the distance to Firm i of a consumer with switching costs that is indifferent in the first period between the two firms (note $q_1^i = 1 - q_1^j$). Then Firm i is guaranteed a demand of $s q_1^i$ in the second period from the consumers that have switching

costs. The total demand in the second period for Firm i is then $s q_1^i + (1-s) \frac{t + p_2^j - p_2^i}{2t}$;

the unique second period equilibrium prices are $p_2^i(q_1^i) = t[1 + s(2q_1^i - 1)/3]/(1-s)$, and the second period equilibrium profit for Firm i as a function of q_1^i is

$\pi_2^i(q_1^i) = \frac{t}{2(1-s)} (1 + s \frac{2q_1^i - 1}{3})^2$. This illustrates a first effect of switching costs.

Consumers that bought initially from one firm would continue to prefer that firm, and in addition have now a more intense preference due to the switching costs. This would then decrease the demand own-price sensitivity in the second period, which would lead to greater prices and profits in the second period.

Working backwards to the first period, consumers without switching costs behave exactly as in the static case, because their decisions do not affect what happens in the second period. Consider now the decisions of the consumers that have switching costs.

For the marginal consumer buying product i , denoted by q_1^i , the total cost of buying product i is $p_1^i + t q_1^i + \mathbf{d}(p_2^i(q_1^i) + t q_1^i)$, while the total cost of buying product j is $p_1^j + t(1 - q_1^i) + \mathbf{d}(p_2^j(q_1^i) + t(1 - q_1^i))$. Indifference between buying product i and j leads

then to $q_1^i = \frac{1}{2} + \frac{3(1-s)(p_1^j - p_1^i)}{2t[3(1+\mathbf{d}-s) - \mathbf{d}s]}$ and a total demand in the first period of

$$(3.1) \quad \bar{q}_1^i = \frac{1}{2} + s \frac{3(1-s)(p_1^j - p_1^i)}{2t[3(1+\mathbf{d}-s) - \mathbf{d}s]} + (1-s) \frac{p_1^j - p_1^i}{2t}.$$

If consumers are myopic, or there are no switching costs, this reduces to the static Hotelling demands. Equation (3.1) illustrates a second effect of competition with switching costs: Switching costs and forward-looking consumers make the first-period demands less price sensitive because the marginal consumers realize that by buying one product they will be locked-in and pay a higher price in the next period. This is a force towards higher equilibrium prices in the first period.

The firms set first period prices to maximize the total value of profits, $p_1^i \bar{q}_1^i + \mathbf{d} p_2^i(q_1^i)$. This maximization illustrates a third effect of switching costs. In order to get higher profits in the second period, firms charge lower prices in the first period to increase q_1^i .

This is a force towards lower prices and lower profits. In this particular problem this maximization by each firm yields unique first period equilibrium prices $p_1^i = p_1^j = t(1 + \mathbf{d})/[1 + \mathbf{d}(1 - s/3)]$.

In general, which effects dominate (for lower or higher profits) will depend on the particular characteristics in the market. In the particular example above, equilibrium profits are higher with switching costs. An example where it goes the other way can be obtained if consumers have small switching costs, change preferences from period to period, and are not too patient. Beggs and Klemperer (1992) look at the impact of large switching costs on the MPE of an infinite horizon duopoly model with uniform pricing. Each period, a fraction v of new consumers enter the market with horizontally differentiated preferences that are fixed over time. Once a consumer purchases from a firm it is unable to purchase from its rival in the future.²⁶ In this model, firms use a single price both to exploit locked-in consumers and to attract new ones, so the effects of switching costs on prices are less obvious; Beggs and Klemperer show that switching costs increase prices in symmetric equilibria of the affine MPE that they consider.

In some markets switching costs can be created endogenously by the competing firms by putting incompatibility features in its products. This possibility may end up making all firms worse off in equilibrium (e.g., Cabral and Villas-Boas 2005). Nilssen

²⁶As in Taylor (2003), discussed below, the model abstracts from the determination of initial market shares, and takes these as exogenous.

(1992) distinguishes between switching costs that are incurred each time a consumer changes supplier, and “learning” costs that are incurred each time a consumer uses a supplier for the first time. Nilssen argues that a greater relative size of switching to “learning” costs leads to higher prices for the loyal consumers, and lower introductory prices.

Turning to our main interest of behavior-based pricing, we focus on the model of Chen (1997), which is a two-period, two-firm model that is very similar to that of Section 3.1, except that all consumers are identical in the first period, and that after making their first-period purchases, each consumers privately observes a switching cost s . As we will see, the main difference with the work discussed above is that second period prices on the two “turfs” are independent of the relative sizes of these two markets. We will then discuss Taylor (2003) who extends Chen to oligopoly, multiple periods, and switching costs that are correlated over time, and conclude with a brief mention of some other related work.

Following Chen, assume that all consumers have common value v for each of the two goods, and that their switching costs are distributed uniformly on an interval $[0, \bar{s}]$. In the second period, a consumer will switch from firm A to firm B if the difference in prices $\alpha - \hat{\beta}$ is greater than his switching cost, so sales on firm A’s turf will be $x(1 - G(\alpha - \hat{\beta}))$ and $xG(\alpha - \hat{\beta})$ for firms A and B, respectively, where x is the size of firm A’s turf and G is the cumulative distribution function for the switching costs. Since the size of firm A’s turf simply has a multiplicative effect on second period profits, it clearly has no impact on second period pricing or sales, at least at interior equilibria where both firms have sales on A’s turf.²⁷ Intuitively, the fact that a customer bought from firm A last period tells us nothing at all about his preferences, except that the customer must now pay the switching cost to use B, so the size of firm A’s turf has no bearing on second-period competition. This is in contrast to the models of horizontal differentiation we considered earlier, where if firm A has a larger first-period market share it knows that the consumers in B’s turf have a stronger preference for B, and so firm A is more aggressive on firm B’s turf as firm B’s first-period sales decrease. For

²⁷Chen shows that the equilibrium is interior; Taylor extends this finding to distributions G such that both G and $1-G$ satisfy the MHR condition.

this reason, we suspect that adding a small amount of horizontal differentiation to the switching cost model would make the second-period prices respond to market shares.

With the uniform distribution, each firm charges second-period prices $c + 2\bar{s} / 3$ and $c + \bar{s} / 3$ on its own and the rival's turf respectively, where v is assumed larger than $c + \bar{s}$; firms sell to $2/3$ of their old consumers and $1/3$ of their rivals, so second period profits are $\frac{4\bar{s}}{9}x + \frac{1}{9}\bar{s}(1-x) = \frac{\bar{s}}{3}(x + \frac{1}{3})$ and $\frac{\bar{s}}{3}(1-x + \frac{1}{3})$, for firms A and B, respectively. Because the first period product is completely homogenous, and second-period profit is increasing in market share, the first-period prices will be below cost: at the profile where both firms charged marginal cost, and so have second-period profit of $5\bar{s} / 18$, either firm would gain by undercutting slightly, capturing the whole market, and having second-period profit $4\bar{s} / 9$. In fact, Chen shows that the unique subgame perfect equilibrium has first period prices of $c - \delta\bar{s} / 3$; at this point cutting price would incur a large enough first period loss to offset the second period gain. Thus the conclusion that prices rise over time extends from switching-cost models without targeted pricing to switching-cost models with behavior-based pricing. This prediction is in contrast to that of the FT model of short-term contracts, where prices rise over time.²⁸

If firms can not observe the consumers' past purchases, then firms with larger first period sales will price less aggressively in the second period. Chen shows that this would lead to less aggressive first period pricing, so that, as in FT, firms are made worse off when they can both engage in targeted pricing.²⁹ Moreover, consumers need to forecast first period sales to know second period prices, and the assumption of homogenous consumers means that the model may have multiple equilibria.

As noted above, Taylor extends Chen's analysis in several ways. To simplify the analysis, he also assumes that consumers are already "assigned" to one of the firms at the start of the first period. For this reason, first-period demand is very different than in Chen's model, and maintaining the rest of Chen's set-up, first period prices are now

²⁸ Of course the dynamics of prices are different in stationary infinite-horizon models such as Villas-Boas (1999).

²⁹ Chen analyzes one of the equilibria for the uniform-price model, we do not know whether there are others.

above marginal cost, and second-period market shares depend on the initial conditions; prices in the second period, being independent of market share, are the same as in Chen.

Taylor extends this analysis to multiple periods, finding that prices in the two markets are constant over time until the last period. This is intuitive: only the most recent purchase matters for the evolution of switching costs, so all periods before the last are strategically similar (given the assumption that consumers enter the game already assigned to a firm). More surprisingly, moving from two firms to three makes a substantial qualitative difference: when there are at least three firms, at least three of them offer marginal cost pricing to other firm's customers. The reason that three is the key number here is that with three firms, there are two firms competing to get customers from each other firm, so that there is Bertrand competition for the switchers. This insight suggests that it would be interesting to study information-based price discrimination in models with three or more firms; this will be complicated by the need to consider a richer specification of preferences, with a two-dimensional taste parameter θ . As usual with differentiated products, we would not expect prices to be driven to marginal cost, but new and interesting features could emerge.

Finally, Taylor considers a two-period model with two types of consumers, those whose switching costs tend to be low and those whose costs tend to be high. Here a customer who "switches" in the first period is thought on average to have lower switching costs, so that agents who switch will be offered a lower price by their first-period supplier than agents who buy from that supplier without switching. It would be interesting to extend this analysis to more than two periods. In that case, consumers will be all the more concerned about their "reputations," and the impact of being known as a low-cost switcher may be ambiguous, as firms may wish to avoid "recruiting" consumers who are likely to soon move on to another brand.

In addition to these papers, we should mention the paper by Schaffer and Zhang (2000) which looks at a static game corresponding to the last period of the sort of two-period model studied above, with the additional feature that switching may be more costly in one direction than in the other. With symmetric switching costs, firms always charge a lower price to their rival's consumers, but this need not be true when switching costs are sufficiently asymmetric. More recently, Dobos (2004) analyzes a model that

combines horizontal differentiation in the first period, switching costs in the second, and network externalities in both; he finds that profits are decreasing in the size of the network effect, as this effect leads to more aggressive first period pricing.³⁰

4. Behavior-Based Pricing with Multiple Products, and Product Design

So far we have been assuming, for the most part, that each firm produces a single good. We now consider cases where each firm may produce multiple versions of the same product. Even in the case where the set of goods is fixed, this leads to interesting forms of behavior-based pricing, such as price discounts for consumers who are upgrading as opposed to new purchasers. In addition, there are the questions of how many different goods a firm will choose to sell, and (assuming it has this choice) what their characteristics will be.³¹

The literature on behavior-based pricing and multiple goods has studied two rather different sorts of goods and demand structures. Fudenberg and Tirole (1998), and Ellison and Fudenberg (2000) study “upgrades” in models of vertical differentiation, where all customers agree that one good is better than the other; these models study only the monopoly case. Thus these papers are most closely related to the literature we discussed in Section 2. In contrast, Zhang (2005) studies endogenous product lines in a Hotelling style duopoly model of horizontal differentiation that is similar to the model of Fudenberg and Tirole (2000) except for the assumption of quadratic “transportation costs.” We focus on these two sorts of models, and do not discuss the related literature on the monopolist’s profit-maximizing menu of goods and prices in a static model.³² We do however discuss the papers of Levinthal and Purohit (1989), Waldman (1996), and Nahm (2004), which study the introduction of a new product in models with anonymous consumers and a frictionless second-hand market. Although behavior-based pricing is not

³⁰His model is an extension of Doganoglu and Grzybowski (2004) who consider the same preferences but without price discrimination.

³¹This latter question can also be asked when each firm is only allowed to produce a single good, but that question does not seem to have been explored in the literature on behavior-based pricing.

³²See Mussa and Rosen (1978) and Deneckere and McAfee (1996) for discussions of the way the monopolist’s desire to extract surplus leads to distortions in the product line.

considered in these papers, the analysis of the anonymous case is an important benchmark for the effects of behavior-based pricing.

4.1 Upgrades and Buybacks with an Anonymous Second Hand Market

In this subsection and the next we discuss the two-period model of Fudenberg and Tirole (1998). We begin with the case of anonymous consumers and a frictionless second-hand market, which corresponds to the market for textbooks, and is also a useful benchmark for evaluating the impact of behavior-based pricing. As noted above, behavior-based pricing is impossible when consumers are anonymous, just as it is in the durable-good models of Section 2. Indeed those models can be viewed as a special case of this one, because whether or not there is a second-hand market makes no difference given that there is a single perfectly durable good and all consumers enter the market at the beginning and remain until the end.

In period 1, the monopolist produces a low-quality version of a durable good; this good is denoted L . In period 2, the monopolist can produce both L and an improved version H . These goods are produced under constant returns to scale, with cost c_L for L and $c_H = c_L + c_\Delta$ for good H , where $c_\Delta \geq 0$.³³ There is a continuum of consumers, indexed by $\mathbf{q} \in [0, 1]$; a type- \mathbf{q} consumer has utility $\mathbf{q}V + I$, where I is her net income, and $V = V_L$ or $V_H = V_L + V_\Delta$, $V_\Delta > 0$ depending on whether she consumes L or H . This is a fairly standard demand structure, and it is easy to work with, but involves some loss of generality, as can be seen from the fact that in a static model the monopolist will not offer both goods if their costs are the same.³⁴

Following the paper, we assume that $V_L > c_L$ and $V_\Delta > c_\Delta$. To simplify, we also assume that the distribution of types is uniform; the paper assumes that the distribution has a continuous density that satisfies the monotone hazard rate condition. The firm and the consumers use the common discount factor \mathbf{d} .

Because the monopolist lacks commitment power, we solve the problem by working backwards from the second period. The solution here depends on the stock x_1 of

³³ The production cost of upgrading an L unit to H is the same as that of making H from scratch.

³⁴ This can be seen by considering equation (4.1) when $x_1 = 0$.

L that is already in the hands of the consumers, but the assumptions of anonymity and a frictionless second-hand market mean that we do not need to worry about which consumers bought the product, and indeed we can suppose that all old units are sold in the second-hand market, with some of them possibly repurchased by their original owners. The form of the utility function implies that there will be three (not necessarily non-empty) segments of consumers in the second period: types in the interval $[0, \mathbf{q}_L]$ do not consume; types in $[\mathbf{q}_L, \mathbf{q}_H]$ consume good L , and types in $[\mathbf{q}_H, 1]$ consume good H . The market price of good L is then the value $p_L = \mathbf{q}_L V_L$ that makes \mathbf{q}_L indifferent between purchasing L and not purchasing, while the price of H makes \mathbf{q}_H indifferent between purchasing H or purchasing L , so $p_H = p_L + \mathbf{q}_H V_\Delta$.

If the mass of consumers consuming good L is greater than the existing stock, i.e., if $\mathbf{q}_H - \mathbf{q}_L > x_1$, the monopolist is a net seller of L in period 2; when the reverse inequality holds, the monopolist is engaged in “buybacks,” and when $\mathbf{q}_H - \mathbf{q}_L = x_1$, the monopolist is inactive on the L market. Each of these regimes can arise for some values of the first-period stock; moreover, each of these regimes can arise for an open set of parameters in the full equilibrium, where x_1 is determined by the monopolist’s first-period sales.

When $\mathbf{q}_H - \mathbf{q}_L > x_1$, so there are net sales, the monopolist has second-period profit

$$(4.1) \quad \begin{aligned} \Pi_2^{\text{net sales}} &= (\mathbf{q}_H - \mathbf{q}_L - x_1)(\mathbf{q}_L V_L - c_L) + (1 - \mathbf{q}_H)(\mathbf{q}_L V_L + \mathbf{q}_H V_\Delta - c_H) \\ &= (1 - \mathbf{q}_L - x_1)(\mathbf{q}_L V_L - c_L) + (1 - \mathbf{q}_H)(\mathbf{q}_H V_\Delta - c_\Delta). \end{aligned}$$

Thus, it is as if the monopolist faces two separate, unlinked markets in period 2. All consumers above \mathbf{q}_L purchase L , with x_1 of this coming from the pre-existing supply. Separately, the monopolist supplies the ‘upgrade’ to types above \mathbf{q}_H ; this (fictitious) good has incremental cost c_Δ and sells at price $\mathbf{q}_H V_\Delta$. Thus when the net-sales regime prevails, the monopolist sells exactly the same amount of good L as it would if good H did not exist, and sales of the old good follow the standard Coasian path discussed in

Section 2. Similarly, price and sales in the upgrade market are not influenced by x_1 . Thus, the first-order conditions for maximizing (4.1) are given by the standard formulas:

$$\frac{q_L V_L - c_L}{q_L V_L} = \frac{1 - q_L - x_1}{q_L}$$

and

$$\frac{q_H V_\Delta - c_\Delta}{q_H V_\Delta} = \frac{1 - q_H}{q_H}.$$

When $q_H - q_L < x_1$, so there are buy-backs, we suppose that the monopolist has no use for repurchased units. Thus the payoff function in this region is the same as $\Pi_2^{\text{net sales}}$ except that c_L is replaced by 0. That is,

$$\Pi_2^{\text{buybacks}} = (1 - q_L - x_1)(q_L V_L) + (1 - q_H)(q_H V_\Delta - c_H).$$

Note that once again the “upgrade market” decouples from the market for L . However, the price for L (given x_1 and the buy-back regime) is lower than it would have been if H had not been introduced, for now the “effective cost” of L is zero. Thus, while the monopolist’s second-period payoff is continuous at the boundary between net sales and buybacks, it has a kink there, as the effective marginal cost changes from 0 to c_L . For this reason, the “inactive” regime is the equilibrium for a range of values of x_1 . In this regime the constraint $q_H - q_L = x_1$ is binding, and the markets do not decouple.

Fudenberg and Tirole show (in Proposition 2) that there are numbers $0 \leq \underline{x}_1 < \bar{x}_1 < 1$ such that when $x < \underline{x}_1$ the solution has net sales, and is exactly the solution to maximizing $\Pi_2^{\text{net sales}}$ while ignoring the net-sale constraint. For $\underline{x}_1 < x_1 < \bar{x}_1$ the solution that maximizes $\Pi_2^{\text{net sales}}$ has negative sales of L , while the solution that maximizes Π_2^{buybacks} has net sales; here the second-period equilibrium is at the kink. Finally, for $\bar{x}_1 < x_1$ the solution has buybacks. Moreover, p_L is a continuous and weakly decreasing function of x_1 , and $q_L + x_1$ is continuous and weakly increasing.

What we are really interested in is the full equilibrium of the two-period game. Fudenberg and Tirole show that setting a first-period price of p_1 leads to sales to all types above the cutoff value $q(p_1)$, so that the stock on hand at the start of the second

period is $x_1 = 1 - q_1$. The monopolist's problem is thus to maximize the discounted sum of first and second period profits, taking into account the way that first period sales determine the second period regime. The following examples show that each regime can arise for some parameter values, and give a flavor of when they might be expected, but stop far short of a characterization of when each regime prevails.

First, if $c_L = c_H = 0$, then there are always buybacks. To see this, note that in this case Π_2^{buybacks} simplifies to $(1 - q_L - x_1)(q_L V_L) + (1 - q_H)(q_H V_\Delta)$, so that the optimum in the H -market is $q_H = 1/2$, which is the same as the optimum in the L market when $x_1 = 0$.³⁵ Thus, there are buybacks when x_1 is close to zero, and as x_1 increases, q_H is unchanged while $q_L + x_1$ increases, so buybacks (which are $x_1 + q_L - q_H$) increase as well.

Next, net sales occurs whenever $c_L = 0$ and c_Δ is almost as large as V_Δ , so that the new good is sold to only the highest-value consumers. This is true for any value of the discount factor, but it is easiest to see for the case $d = 0$, as here first-period output is the amount sold by a static, zero cost monopolist, which is $1/2$ for the uniform case considered here, while the first-order condition for q_L in the net sales regime simplifies to $q_L(x_1) = \frac{1-x_1}{2}$, so that $q_L(1/2) = 1/4 < 1/2$ and the second-period solution following $x_1 = 1/2$ indeed has net sales.

Finally, the equilibrium will have neither sales nor buybacks if $c_\Delta = 0$ and $c_L = c_H$ is very close to V_L . Intuitively, when $c_\Delta = 0$ there will be no production of the old good in period 2, and because costs are close to V_L , there will be very little production of L in the first period, so x_1 is small, which makes buybacks less likely.³⁶

At this point we should mention the related work of Levinthal and Purohit (1989) and Lee and Lee (1994) on monopolists with an anonymous second-hand market. Levinthal and Purohit consider a model with costless production, where the second-period market is described by a pair of linear demand curves, and the rental prices of each

³⁵This equality does not depend on the uniform distribution but rather on the assumptions that costs are zero and quality enters the demand function multiplicatively.

³⁶The formal argument uses continuity and monotonicity properties.

generation are equally affected by an increase in the output of the new generation.³⁷ In their model, buybacks are only optimal when the firm is sufficiently patient, and otherwise there are net sales.³⁸ Lee and Lee suppose that the monopolist is unable to sell or buy units of the old product in period two.

4.2 Upgrades and Buybacks with Non-Anonymous Consumers

Fudenberg and Tirole go on to consider two other sorts of information structures: “identified consumers,” where the firms know which consumers purchased at date 1, and “semi-anonymous consumers,” where consumers can prove that they purchased if they wish to do so, but can also pretend not to have purchased, which constrains the price to new customers to be no lower than the “upgrade price” offered to old ones. Following the paper, we now assume that $V_L > dV_H$, which implies that any first-period price induces a cut-off $\mathbf{q}_1(p_1)$ such that the consumer of type \mathbf{q} purchases when $\mathbf{q} > \mathbf{q}_1(p_1)$. This assumption is stronger than one would like, but we are not aware of a weaker condition that guarantees a first-period cutoff, nor of related analyses that allow for disjoint sets of consumers to purchase in the first period. We also assume that $c_L = 0$, and that $c_H = 0$ as well; the paper does not make this last assumption.

We begin with the case of identified consumers. Here the monopolist faces two distinct second-period markets, patrons and non-patrons. On the patron’s market the monopolist maximizes $(1 - \mathbf{q}_u)\mathbf{q}_u V_\Delta$ subject to $\mathbf{q}_u \geq \mathbf{q}_1$, so $\mathbf{q}_u = \max[1/2, \mathbf{q}_1]$, and $p_u = \mathbf{q}_u V_\Delta$. On the non-patron’s market, the monopolist will sell good H to consumers with values between \mathbf{q}_H and \mathbf{q}_1 , where \mathbf{q}_H is chosen to maximize $(\mathbf{q}_1 - \mathbf{q}_H)\mathbf{q}_H V_H$; the solution to this is $\mathbf{q}_H = \mathbf{q}_1 / 2$, with price $p_H = V_H \mathbf{q}_1 / 2$. Comparing the objective functions in the two markets lets us identify two competing effects. First, non-patrons of any given type have more to gain from purchasing because they have a lower payoff

³⁷Note that in the Fudenberg and Tirole model, the price of L is $\mathbf{q}_L V_L$ and so depends only on the supply of good L .

³⁸Their results imply that the inactive region never occurs without pre-commitment. They show that in some cases the monopolist can gain from a first-period commitment not to produce L in the second period, just as it could if good H did not exist.

without a purchase; this “reservation utility effect” pushes the upgrade price to be lower than the price to new consumers. On the other hand, former customers have higher types; this “ratchet effect” means that non-patrons should get lower prices. These effects will help us understand when the identified and semi-anonymous cases coincide.

Fudenberg and Tirole show that in equilibrium the monopolist chooses $q_1 > 1/2$, so that all old patrons upgrade, and there is no “leapfrogging” of lower-value consumers past higher-value ones.³⁹ Moreover in this case the second-period upgrade price is $q_1 V_\Delta$.

In the semi-anonymous case, the payoff functions in the two markets are the same as with identified consumers, but the markets are linked by the customers’ incentive compatibility constraint, which requires that $p_u \leq p_H$. The calculations above show that this constraint is slack, and the two solutions coincide, if and only if $V_\Delta \leq V_H/2$, or equivalently, if $V_\Delta \leq V_L$, i.e., if the size of the innovation is not too large. The intuition for this is that for large innovations, upgrading is very attractive to high-value types, so the “ratchet effect” dominates the reservation utility effect; this is true for general distributions and not just the uniform.

Finally, Fudenberg and Tirole show that with costless production the monopolist’s profits are higher under anonymity than with identified consumers. With costless production, when $q_1 \geq 1/2$ (which is the relevant range) the anonymous-market solution is for customers between $q_1/2$ and $1/2$ to consume L , and customers from $1/2$ up to consume H ; with identified consumers, the monopolist sells H in the second period to all types above $q_1/2$. The commitment solution is to sell H to consumers above $1/2$, and nothing at all to the others; the anonymous solution is closer to this outcome, and so yields higher payoffs. The point is that the presence of the second-hand market leads the monopolist to sell less of H in period 2, which helps alleviate the commitment problem in period 1. (Note that this finding does not immediately extend to the semi-anonymous case, except for parameters where it coincides with the solution with identified consumers: The no-arbitrage constraint cannot help the monopolist in the second period,

³⁹Leapfrogging can occur when $c_H > 0$, as here the monopolist will not induce all old patrons to upgrade but it will sell H to non-patrons so long as c_H is not too high.

for any given first period outcome, but the constraint could have an impact on first-period play.)

4.3. *Endogenous Innovation*

Waldman (1996) and Nahm (2004) analyze endogenous innovation in the anonymous case. Waldman suppose that there are only two types, q_L and q_H , with $q_L V_L < c_L$. This means that the firm would not sell to the low types in a one-period model, and moreover in the absence of the new good the firm would not produce in period 2. That is, the assumed demand structure means that the firm would not face the usual Coasian commitment problem. However, the sale price of the low good in period 1 is decreasing in the probability that the firm will introduce an improved good H in the second period, and Waldman shows that the firm does face a commitment problem with respect to introducing the improved good.

Nahm points out that this conclusion relies on the assumed demand structure. In a two-type model with $q_L V_L > c_L$, the price of good L will fall over to $q_L V_L$ in the second period whether or not the new good is introduced, and the firm does not face a commitment problem with respect to introducing the new good. Nahm goes on to investigate the incentives for introducing the new product in a model of section 4.1, where in between period 1 and period 2 the firm spends resources on R&D, which in turn determines the probability that the high quality good is available in period 2.

As we saw above, in the net-sales case, the second-period price of good L is the same whether or not H is introduced, and investment in R&D only influences payoff in the “upgrade” market. Hence the monopolist does not face a time-inconsistency problem with respect to R&D, and it chooses the same level of investment that it would chose if it could commit to the choice in period 1. However, in the inactive and buy-back regimes, the second period price of L is lower if H is introduced than if it is not. Hence to maximize first-period sales and overall profit, the monopolist would benefit from a commitment that limited its R&D.

Ellison and Fudenberg (2000) analyze the semi-anonymous, costless-production case in a model intended to correspond to markets for software. It is very similar to that

discussed above, with one good in period 1 and the possibility of producing an improved version in period 2; the main difference is that their model includes (positive) network externalities. In their model, consumers incur set-up or training costs each time they adopt or upgrade their software, and differing versions of software are backwards but not forwards compatible, so that users of the newest version of the software enjoy the largest network benefits. In their dynamic model,⁴⁰ consumers are ex-ante identical, but not all of them are present in the first period. They show that the monopolist suffers from a commitment problem that can lead it to introduce upgrades that are not only welfare-decreasing but also lower its own overall present value. The idea of this result is simple: in the second period the monopolist may prefer to sell upgraded, higher-value software to new consumers, but this forces the old consumers to either incur the costs of learning to use the new version or settle for smaller network benefits due to incompatibility with new consumers. This can lead to a loss of first-period profits that outweighs the second-period gain.

As it is common in models of network externalities, consumers' purchasing decisions have the flavor of a coordination game, and can have multiple equilibria. Ellison and Fudenberg assume that in the second period, new consumers coordinate on the equilibrium that is best for them, and consider two different equilibrium-selection rules for the old consumers who are deciding whether to upgrade; in either case there is a region of the parameter space where the monopolist introduces the upgrade when the social optimum would be to sell only the old good in both periods.

4.4. Endogenous Location Choice in Duopoly

Waldman, Nahm, and Ellison and Fudenberg consider a monopolist whose innovation decision is whether to introduce or research an improved version whose characteristics are fixed. Zhang (2005) considers endogenous location choice in a two-period poaching model. The idea is that the rise of flexible manufacturing makes it cheaper for firms to customize products to various clienteles, and since purchase

⁴⁰The paper also considers a static model with a continuum of types, and shows that even with commitment the monopolist may introduce socially inefficient upgrades to help it price discriminate. That model is less closely related to the themes of this survey.

decisions convey information, firms might want to design one product for its established customers and another for those they are trying to poach from a competitor.

The information structure and institutional assumptions are the same as in the short-term contracts section of Fudenberg and Tirole (2000), but the payoff functions are different: Consumers are uniformly distributed on the interval $[0,1]$, while firm locations are endogenous, and transportation cost is quadratic in distance: the utility of for type q of consuming a good at location a is $v - t(q - a)^2$, where the reservation utility is assumed high enough that in equilibrium all consumers purchase. At the start of the first period, the two firms simultaneously choose locations a and b respectively, and in the second period, each firm can produce products at two (or more) locations, and offer different prices and goods to consumers whose first-period actions were different.

In the base model, designing new products is costless.⁴¹ If the firms and consumers have the same discount factor, or more generally, if the consumers are sufficiently patient compared to the firms, the equilibrium is for the firms to split the market in the first period, and for each firm to offer two new and distinct models in the second period, with firm A choosing a_o, a_n and firm B choosing b_o, b_n , where “o” and “n” are for old and new consumers respectively. However, as in the poaching models discussed in section 3, firms do better when they have less first-period information, and if firms are sufficiently patient compared to consumers then the first-period pure-strategy equilibria are asymmetric, with one firm capturing all of the market, so that first-period purchases reveal no information.

To understand these results, we explain the outcome in the second-period markets for types who have been revealed to lie in an interval $[Z, Z + L]$, which is the same as in a static model with these types as the single market. It is interesting to note that although introducing varieties is costless, and firms are allowed to introduce as many as they wish, in equilibrium each firm only sells a single product. This fact is closely related to the fact that if each firm can only introduce a single product, they will choose locations outside the support $[Z, Z + L]$ of the distribution of consumer types if such locations are allowed, and at the boundaries of the distribution if it is not. Intuitively, firms face a trade-off

⁴¹ The paper speculates briefly about the case where innovation costs are such that firms introduce a single new product in period 2.

between locating near the center of the distribution, which increases profits holding the opponent's price fixed, and locating towards the edges, which lessens price competition and raises the opponent's equilibrium price. With quadratic transportation costs and the uniform distribution, the strategic effect dominates until the locations are well outside the support of the distribution of types.⁴² The fact that the optimal locations for a single product are outside of the support provides an intuition for why introducing a second variety would not be helpful: If the new variety is to provide an increase in efficiency, it must be closer to the opponent's location, but this would provoke the price competition that the first location was chosen to avoid.

Now consider firms simultaneously choosing locations and prices in two different second-period markets, corresponding to the first-period purchase of the consumers. The previous paragraph explains why each firm will choose a single product for each market; in general, these products will be different, and a better match for the tastes of the market they are designed for.

Now we turn to the consumer's decision in the first period. As in Fudenberg and Tirole (2000), the first-period decisions of consumers will generate a cut-off rule, so that first-period sales identify two intervals of consumers, corresponding to each firm's turf. Also as in that model, the consumers who are near the cutoff in the first period switch suppliers in the second, and increased consumer patience makes first-period demand less elastic. Consumers benefit most when they are identified as being in a small interval, as this leads to intense price competition; the firm's second-period profit is highest when all consumers purchase from the same firm in the first period, so that the purchases reveal no information.

Working backwards to the firm's first-period decisions, Zhang shows that when consumers and firms are equally patient, and more generally if the consumers are sufficiently patient compared to the firms, the first period outcome is symmetric, with firms A and B located equal distance from the market center, and each taking half the

⁴²Economides (1986) studies the Hotelling location-price game where duopolists each offer one product, with a uniform distribution on types, and transportation costs proportional to t^a . He shows that for $a \in [1.26, 1.67]$ the firms locate within the distribution of types, while for $a \in [1.67, 2]$ they locate at the endpoints. (He constrains them not to locate outside of it.) For $a \in [1, 1.26]$ there is no pure strategy equilibrium; see d'Aspremont et al. (1979) for the linear and quadratic cases.

market. In the second period, each firm introduces two new products, one for each segment of the first period market. On the other hand, if firms are patient and consumers are myopic, the firms are able to avoid segmenting the first period market, and their combined profits increase.

Zhang's results on product design seem to reinforce the idea that customer recognition leads to more intense competition, and lower profits. It would be interesting to understand what happens if we have a longer time horizon (possibly with changing consumer tastes), and what would happen under product choice and monopoly, with customized product advertising (and where this customized advertising could also depend on past behavior).

5. Related Topics: Privacy, Credit Markets, and Customized Pricing

This section briefly discusses the issues of consumer privacy protection, pricing in credit markets, and standard third-degree price discrimination that is based on exogenous characteristics. We focus on the work of Calzolari and Pavan (2005), Taylor (2004a), Dell'Araccia et al. (1999), Dell'Araccia and Marquez (2004), and Thisse and Vives (1988), and also discuss Pagano and Jappelli (1993), Padilla and Pagano (1997, 2000), and Taylor (2004b).

5.1. Privacy

As we have seen, the efficiency consequences of BBPD are ambiguous, so there is some reason to consider the impact of various regulations and firm-based initiatives that protect consumer privacy.⁴³ One interpretation of consumer privacy is that firms

⁴³This ambiguity should not be a surprise in view of previous results on related issues. Hirshleifer (1971) noted that the efficiency impact of information acquisition is ambiguous when markets are incomplete. This holds in particular for firms acquiring more information about the characteristics of each consumer. For example, Hermalin and Katz (2004) show that third degree price discrimination may be better or worse from a social point of view than second-degree price discrimination. Wathieu (2004) argues that information about consumers may lead to inefficiently many products being produced, each at too low a scale. For a recent survey on the economics of privacy see Hui and Png (2005).

cannot track consumers' past behavior.⁴⁴ Consumers that buy early may be recognized as consumers that value the product highly, and then be charged a higher price in subsequent periods. In this sense consumers are hurt by losing their privacy, they are charged higher prices. As discussed above, consumers, if aware of this loss of privacy, may be strategic in the earlier periods, and refrain from purchasing the product, not to reveal their high valuation. This may give firms an incentive to commit to privacy protection.

Taylor (2004a) uses a variation of the two-period model of Section 2 to focus on the privacy issue. Consumers interact sequentially with each of two firms, and each consumer's valuations for the products of the two firms are positively correlated, so that, if the second firm is able to observe that a consumer bought from the first firm, then the second firm's beliefs about the valuation of that consumer for its product is higher than if the consumer declined to purchase. Taylor assumes that the second firm is unable to commit to its prices until after consumers interact with the first one. Privacy is the case in which the second firm is not able to observe whether a consumer bought or did not buy in the first period. Without privacy, the first firm can sell the list of its customers, and allow the second firm to price discriminate between the consumers that bought and did not buy from the first firm.

If there is no privacy, the first firm sells the customer data to the second, and consumers do not foresee that sale (in the context of Section 2 this is the case when the consumers are myopic), then the first firm has a greater incentive to charge higher prices in order to make the customer data more valuable. If consumers foresee that the first firm is going to sell the customer data to the second firm, then they strategically refrain from buying, which makes the customer data being sold less valuable, and gives incentives for the first firm to lower prices. Firms prefer the no-privacy case when consumers are myopic, but prefer consumer privacy if consumers are able to foresee that under no privacy their purchase information is going to be sold. Taylor shows that welfare can be higher or lower under consumer privacy depending on the demand elasticity.

⁴⁴Upon realizing that Amazon was charging different prices for the same item, possibly based on different purchase histories, some consumers showed concern about shopping there ("Customers Balk at Variable DVD Pricing," *Computerworld*, September 11, 2000, p. 4).

Calzolari and Pavan (2005) consider the case where two principals sequentially contract with a common agent, and where the upstream principal can sell its information to the downstream principal. They assume that the agent's valuations with the two sellers are perfectly correlated, which is more restrictive than Taylor's assumption of imperfect correlation, but otherwise their model is more general.⁴⁵ As in Taylor, the second principal posts its contract after the consumer has already decided whether to accept the contract of the first firm. By selling information to the downstream principal, the upstream principal may get some payment from the downstream principal (possibly due to greater efficiency, or less information rents provided to the agent, in the downstream relationship), or appropriate any rents of the agent in the downstream relationship that are generated by this sale of information. Calzolari and Pavan identify three conditions under which, if the upstream principal can commit not to disclose any information (commitment to privacy) she will choose to do so. The first condition is that the upstream principal is not personally interested in the decisions taken by the downstream principal. In the context of Taylor (2004a) this is just that the profit of the first firm is independent of the decisions taken by the second firm. The second condition is that the agent's exogenous private information is such that the sign of the single crossing condition is the same for both the upstream and downstream decisions. In the context of Section 2 this condition is just that the valuation of a consumer type is the same across products. In Taylor (2004a) this is that the valuation for the product of the first firm is positively correlated with the valuation for the product of the second firm. Finally, the third condition is that the preferences in the downstream relationship are additively separable in the two contractual decisions. In the context of Section 2, or Taylor (2004a), this is immediately obtained because the second-period profit or utility is independent of whether there was a purchase in the first period.

It is interesting to try to informally relate the first condition with the two-period model in Section 2. Denote the first-period profit under disclosure of information as a function of the first-period action a_1 as $\mathbf{p}_1^d(a_1)$, the first-period profit under privacy as a function of the first-period action as $\mathbf{p}_1^p(a_1)$, the second-period profit under disclosure of

⁴⁵ Ben-Shoham (2005) extends the Calzolari and Pavan analysis to allow for imperfect correlation, and also for imperfect (i.e., noisy or partial) revelation of information from the first principal to the second.

information as a function of the first-period action as $\mathbf{p}_2^d(a_1)$, and the second-period profit under privacy as \mathbf{p}^p . Note that in the model of Section 2 the second-period profit under privacy is independent of the first-period action.⁴⁶ In the context of Section 2 the firm chooses its first-period action under disclosure of information to maximize $\mathbf{p}_1^d(a_1) + \mathbf{p}_2^d(a_1)$ (where the discount factor was set to one). In Calzolari and Pavan the upstream principal is able to receive a payment for the disclosure of information from the downstream principal in the amount of $\mathbf{p}_2^d(a_1) - \mathbf{p}_2^p$. The upstream principal chooses then her action under disclosure of information to maximize $\mathbf{p}_1^d(a_1) + [\mathbf{p}_2^d(a_1) - \mathbf{p}_2^p]$, which results in the same optimal action as in the model of Section 2. Finally, note that in the model of Section 2 the firm chooses privacy if and only if $\max_{a_1} \mathbf{p}_1^p(a_1) + \mathbf{p}_2^p \geq \max_{a_1} \mathbf{p}_1^d(a_1) + \mathbf{p}_2^d(a_1)$, while in the context of Calzolari and Pavan the upstream principal chooses privacy if and only if $\max_{a_1} \mathbf{p}_1^p(a_1) \geq \max_{a_1} \mathbf{p}_1^d(a_1) + [\mathbf{p}_2^d(a_1) - \mathbf{p}_2^p]$. It is immediate to see that these are exactly the same conditions, that privacy is chosen in both models in exactly the same conditions (no customer recognition in the model of Section 2). So, even though in Calzolari and Pavan there are two principals, in the case where the upstream principal expropriates the informational rent from the downstream principal, the model corresponds to single-principal models discussed in Section 2.

Calzolari and Pavan (2005) also show that under the second condition, if the upstream principal discloses information to the downstream principal, the increase in the rent that has to be given to the agent always offsets any potential benefit from the sale of information, or from a greater rent of the agent in the downstream relationship. This is because, if information is disclosed, the agent becomes more protective of his type and the upstream principal does not have the possibility of using any distortion of the downstream relationship contractual variable to help the agent reveal his type. This then implies that when the upstream principal is not personally interested in downstream decisions (the first condition), then there is no advantage in disclosing information and

⁴⁶Calzolari and Pavan allow for the second-period profit to be additively separable in the two contractual decisions, and therefore to be also a function of the first-period actions even under privacy. This possibility does not affect the argument above.

the optimal policy is committing to full privacy. The paper then argues that each of these conditions is necessary for the full privacy result, and that if one of the conditions does not hold it may be optimal for the upstream principal to disclose information to the downstream principal. In such cases, there are also situations in which disclosure of information benefits all three players.

Taylor (2004b) considers a market situation in which firms first post prices, and then decide on how much to screen the customers that demand their product. The profit that a firm derives from a customer depends not only from the price charged, but also from the cost of servicing that customer which varies in the population (and that is also not known by the customer). The amount of screening chosen by a firm allows that firm to receive a noisy signal about the cost of servicing a customer. More screening reduces the noise of the signal. In relation to the papers above, this paper can be seen as looking at quantity discrimination, while the papers above looked at price discrimination. Given that the cost of servicing a certain customer depends on the a priori unknown characteristics of the customer, this model matches well the market features of credit markets, discussed below.

Consider the case in which the screening device searches for “bad news”, that is, good news are always identified appropriately as good news, but bad news are only identified as bad news with some probability less than one. Then, one can obtain that competitive firms screen customers too much. A firm’s incentive to screen customers is given by the difference between the cost of servicing the costly customers and the price it is getting as revenue from those customers, while the social incentive is the difference between the cost of servicing the costly customers and the consumers’ valuation. As in a competitive market the price is below the consumers’ valuation, a firm’s incentive to screen customers is greater than the social incentive. If the screening device is not very good, or the social cost of servicing the costly customers is small, then it may be better not to allow firms to screen (customers have privacy) and for firms to service all customers. If rejected customers stay in the market and apply for the other firms, the situation may become worse, with even more equilibrium screening, so that no screening (privacy) is even better from a welfare point of view. Consumers can improve their situation (of too much screening) by reducing the quantity that they purchase.

Another possibility is for firms to offer consumers the option of disclosing their valuation or keeping it private. McAdams (2005) considers this case, in which consumers who do not disclose their valuation pay a “sticker price,” while consumers who allow the firm to learn their valuation pay a fee to get a “customized price,” and where learning a consumer’s valuation is costly to the firms. McAdams shows that there are parameter values such that welfare can increase if the firms are required to offer the same prices to all consumers (and consumers are forbidden to reveal their valuation/give up their privacy).

5.2. Credit Markets

In credit markets, lenders may learn about the ability of their borrowers, their customers, to repay loans; this information can then be used by the firms in the future loans to those customers. In this case what a firm learns about its previous customers relates to the cost of providing the customer with a given contract, as opposed to the customer’s willingness to pay, which has been the focus of the work we have discussed so far. This feature is also present in other markets, such as labor markets (information about employees), rental markets (information about tenants), insurance markets (information about policy holders), and some forms of service contracts (fussy customers take more time to service). Our presentation here is cast in terms of credit markets because the existing literature has used this type of markets as main motivation.

We start by discussing what happens in credit markets when lenders have private information about their own previous borrowers, and then consider the possibility and effects of lenders sharing their information. The presentation is based in large part on Pagano and Jappelli (1993), Padilla and Pagano (1997, 2000), Dell’Ariccia et al. (1999), and Dell’Ariccia and Marquez (2004). Some of the discussion is also related to some of the material presented in the privacy section above, in particular, Taylor (2004b).

Following Dell’Ariccia et al (1999), consider a market with two competing lenders, 1 and 2. Borrowers have to borrow \$1 to invest in a project that pays R with probability q , and zero with probability $1-q$. Borrowers are heterogeneous on the success probability q , with cumulative distribution function $G(q)$ (density $g(q)$) on

[0,1]. Furthermore, the borrowers are, independent of q , in one of three groups: Either they are “new” borrowers, and so no lender knows about the borrower’s q ; or they are “old” borrowers from Lender 1, so that Lender 1 knows q , but this is not known by Lender 2; or they are “old” borrowers from Lender 2, so that Lender 2 knows q , but this is not known by Lender 1. Let I be the proportion of “new” borrowers ($1-I$ of “old” borrowers), and let a_i be the proportion of “old” borrowers from Lender i . Dell’Ariccia et al. (1999) assume that a lender is not able to distinguish between “new” borrowers and “old” borrowers from the other lender, and that, lenders first simultaneously set interest rates r_i for the borrowers for whom they do not know q , and then they set, also simultaneously, the interest rates r_{iq} for the borrowers for whom they know q .⁴⁷ The paper focuses on the analysis of this market interaction, which can be seen as the second period of a two-period model.⁴⁸

Consider first the behavior of a Lender i with respect to its previous “old” borrowers. These borrowers have access to an offer from the other lender at an interest rate r_j . In order for Lender i to attract them, it has to offer them at least an interest rate r_j . The expected profitability of a borrower of type q is then $qr_j - 1$. Lender i then only wants to extend credit to the borrowers that will generate positive expected profit, that is for the borrowers with $q \geq 1/r_j$. Lender i expected profits from its previous “old” borrowers is then

$$a_i(1-I) \int_{1/r_j}^1 (qr_j - 1)g(q)dq .$$

Note that these expected profits from the lender’s previous borrowers are independent of the lender’s interest rate to the “new” borrowers.

Consider now the profit of a Lender i from the borrowers that borrow from that lender for the first time, given interest rates (r_i, r_j) . Lender i gets an expected profit from

⁴⁷ These two assumptions are as in Villas-Boas (1999), discussed in Section 3. Sharpe (1990), in the context of credit markets, and with borrowers choosing investment levels, makes the assumption that lenders make first the offers to the borrowers that they know, and then, after observing the offer policies (but not the actual offers), make offers to the borrowers that they do not know.

⁴⁸ The appendix of the paper presents some analysis on the two-period model (without discussing if forward-looking borrowers would play a role), and argues, as in Sharpe (1990), that the first period competition is more intense because of the informational advantages the lenders enjoy in the second period.

the “new” borrowers of $\mathbf{I}(r_i E(\mathbf{q}) - 1)$ if $r_i < r_j$, of $\frac{1}{2}\mathbf{I}(r_i E(\mathbf{q}) - 1)$ if $r_i = r_j$, and of zero if $r_i > r_j$. The expected profits for Lender i of the “old” borrowers of the other lender, due to the poor quality borrowers that are denied credit by the other lender, are $\mathbf{a}_j(1 - \mathbf{I})G\left(\frac{1}{r_i}\right)\left[r_i E(\mathbf{q} / \mathbf{q} \leq \frac{1}{r_i}) - 1\right]$.

Because of the discontinuity of the expected profits from the “new” borrowers at $r_i = r_j$, by standard arguments (for example, related to Varian 1980), one can show that the market equilibrium involves mixed strategies in the interest rates r_i and r_j . One can also show that the lender with a smaller share of the “old” borrowers, makes zero expected profits from its new customers, while the lender with a greater share makes positive expected profits from this type of customers. This is because the lender with a greater market share of “old” borrowers suffers less asymmetric information, and lends to less poor quality “old” borrowers than the lender with a smaller market share of the “old” borrowers. Dell’Ariccia et al. (1999) go on to show that this equilibrium with two lenders is exactly the same as the equilibrium with a third lender potentially entering the market, as this new lender would prefer to stay out. This is because this potential entrant cannot protect itself from the lower quality “old” borrowers from both firms. As the incumbent smaller market share lender makes zero expected profits, the new entrant would make negative profits if entering the market (have a positive market share), and prefers to stay out. We have then that the ability to recognize previous customers in credit markets leads to blockaded entry.⁴⁹

Dell’Ariccia and Marquez (2004) considers a variation of the model above where only one lender has previous “old” borrowers, this informed lender has higher costs of funds than the competitor, and \mathbf{q} is uniformly distributed on the segment $[0,1]$. The paper fully characterizes the mixed-strategy equilibrium, and analyzes how the existence of this informed lender affects the loan portfolio allocation. Greater information asymmetry leads to higher interest rates as the informed lender takes advantage of its

⁴⁹ Baye et al. (1992) show the existence of a continuity of asymmetric equilibria in the symmetric Varian (1980) model. It would be interesting to investigate the implications of those results for the model above when there are more than two incumbents.

information advantage. Furthermore, as the competitor has lower costs of funds, the informed lender concentrates more on its previous borrowers, as competing for the “new” borrowers requires now lower interest rates.

This problem of a new firm trying to poach some of the “old” customers of an incumbent firm, and having to be aware of the lemons problem associated with it, is also related to auction problems when one of the bidders is better informed (as in e.g., Engelbrecht-Wiggans et al. 1983), and to competition for auditing business, when the incumbent auditor is better informed about the business risk of a client compared to a rival entrant (e.g., Morgan and Stocken 1998).

One issue that is particularly important in credit markets is what happens if the lenders exchange information about the borrowers. Pagano and Jappelli (1993) investigate this issue with two types of borrower quality, where each lender is in a different “town,” and learns about the credit quality of the borrowers in that town in the previous period. Some of the borrowers change towns from period to period, and there is heterogeneity on the return from the borrowers’ projects if successful. Lenders can price discriminate across three types of borrowers: the safe “old” borrowers, the risky “old” borrowers, and the “new” borrowers. If the interest rate to the “new” borrowers is too high, only the risky “new” borrowers apply for credit. Consider first the case in which lenders are local monopolies in their own towns. In this case profits are decreasing in the proportion of “new” borrowers, as the lenders have less ability to price discriminate between the types of borrowers. If there is information sharing across towns, then lenders can distinguish the types of all borrowers, and profits increase. However, the lending volume increases with information sharing if the safe “new” borrowers were not served in the case without information sharing, and decreases otherwise.

Consider now the case of competition where lenders can offer credit to borrowers in neighboring towns, although at a cost disadvantage. “New” borrowers are assumed to come from far away towns. In order to simplify the analysis (to get away from mixed strategy equilibria), Pagano and Jappelli (1993) assume that outside lenders make offers after the offers made by the local lenders. The paper finds that, as above, lenders are able to deter entry given their informational advantages, and that information sharing leads to lower profits, given the greater threat of the potential entrants. The incentives for lenders

to share information depend then on the monopoly effects above for information sharing, and on the competition effects against information sharing. Which effect dominates depends on their relative strength.

Another potential important issue in credit markets is the possibility of borrowers exerting effort to increase the probability of success of their project. This issue is addressed in Padilla and Pagano (1997). In this case, borrowers may be concerned about exerting effort and then being taken advantage of by high interest rates from the informed lenders (hold-up problem). Padilla and Pagano suggest that lenders may be able to correct this incentive problem by committing to share their information about the borrowers with other lenders, such that the borrowers can benefit from interest rate competition. In another paper, Padilla and Pagano (2000) consider the case in which lenders cannot take advantage of their information about the borrowers because they compete away ex-ante any gains from future private information. In this case the paper argues that the lenders may still want to commit to share the borrowers default rate with other lenders as an incentive device for the borrowers to exert more effort to increase the probability of the project success. However, if the lenders share the information about the type of the borrower, the incentives to exert effort are lower than if only defaults are shared, and the borrowers exert the same level of effort as if no information were shared.

5.3. Customized Pricing

In some markets competing firms may have information about the consumer preferences and price discriminate based on consumer preferences. Competition in such a setting may end up being more intense, if this leads to less differentiation in the competition for each consumer.

Thisse and Vives (1988) consider this effect in the Hotelling line with two firms located at the extremes of the segment $[0,1]$. Suppose that consumers are uniformly distributed on this segment, and that a consumer located at x pays “transportation costs” tx , if buying from the firm located at 0, Firm 0, and “transportation costs” $t(1-x)$, if buying from the firm located at 1, Firm 1.

If firms do not know the location of the consumers they have to charge a uniform price for all consumers. Let the price charged by Firm 0 be p_0 , and the price charged by Firm 1 be p_1 . Then, it is well known that the demand for Firm 0 is $D_0(p_0, p_1) = \frac{t + p_1 - p_0}{2t}$, and that the demand for Firm 1 is $D_1(p_0, p_1) = 1 - D_0(p_0, p_1)$. The equilibrium prices are then $p_0 = p_1 = c + t$ (assume constant marginal costs c), and the equilibrium profit for each firm is $t/2$.

Consider now that the firms know the location of each consumer. Then, each firm can charge a price per location x , $p_i(x)$. The price competition in each location x is like competition with a homogeneous good, where the consumer has different valuations for the product. For $x \leq 1/2$ (the case of $x > 1/2$ is symmetric) we have in equilibrium $p_0(x) = c + t(1 - 2x)$, $p_1(x) = c$, and the consumers choose Firm 0's product. The average price received as revenue by a firm is then $c + t/2$, and each firm has a profit of $t/4$, one half of the profit when customized prices were not possible. This result points to a general effect that competition with customized prices is more intense than competition without customized prices, if customization leads to less differentiation in the competition for each consumer. That is, competition with customized prices becomes like competition with no differentiation, in which at the equilibrium prices, an infinitesimal small price cut attracts all the demand. Variations of this result can be seen in Borenstein (1985), Holmes (1989), Corts (1998).⁵⁰ For the case of competition with second degree price discrimination see, for example, Stole (1995), Villas-Boas and Schmidt-Mohr (1999), Armstrong and Vickers (2001), Desai (2001). For a recent survey of competition with price discrimination see Stole (2004).⁵¹

However, as noted by Armstrong (2005), more information about the consumer preferences may not necessarily lead to less differentiation and lower profits. Armstrong notes that if the additional information is about the "transportation costs" parameter in the traditional Hotelling model, additional information leads to significantly higher prices

⁵⁰ See also Katz (1984) for the case of price discrimination in monopolistic competition.

⁵¹ See also Armstrong (2005) for a recent survey on economic models of price discrimination.

for the consumers with the higher transport costs; this may lead to higher equilibrium profits.

One interesting extension of the variation of the Thisse and Vives model above is the case in which we allow firms to only know the locations of some of the consumers in the line (the firm's database), and therefore, can only offer customized prices to those consumers. This case is considered in Chen and Iyer (2002). We then have that at each location some consumers are in the database of both firms, some consumers are in the database of only one of the firms, and some consumers are not in any data base. The databases can be available from the firms' internal sources or from external sources such as syndicated vendors of information.⁵²

Chen and Iyer show that firms may choose to have not all consumers in their database as this alleviates price competition. However, it turns out that allowing firms to offer some degree of customized prices leads to higher profits than no customization at all. That is, there is an intermediate level of price customization that leads to higher profits. The intuition for why having limited databases may alleviate price competition is related to Grossman and Shapiro's (1984), who show, in the context of uniform prices, that decreased advertising costs may reduce profits because it leads firms to increase their advertising. This increased advertising leads to more consumers that can compare prices, which leads to a greater benefit for a firm of cutting prices, and thus to lower equilibrium prices and profits. In Chen and Iyer, larger databases allow firms to do more customized pricing, which we know from Thisse and Vives, may lead to greater price competition. Ulph and Vulkan (2000) consider the incentives for firms to invest in customization capabilities under different transportation cost functions. Ulph and Vulkan (2001) discuss what happens when customization may allow a firm to offer customized products. Iyer et al. (2005) consider the effects of customized advertising (in a model similar to Grossman and Shapiro 1984 for uniform advertising), and show that customized advertising decreases price competition.⁵³

A related but different form of competition with price discrimination is when firms with capacity constraints advance-sell their products, possibly at a discount. Dana

⁵²This can then be seen as a later period of some dynamic interaction where firms learn the complete preferences of some consumers (the consumers in the firm's database).

⁵³ See also Stegeman (1991), and Roy (2000).

(1998) considers this case, and finds that in equilibrium we may have advance-selling discounts that are bought by consumers with lower valuation for the product, but that have a more certain demand.

6. Conclusion

This paper presents a summary of existing research on the effects of firms being able to recognize their previous customers, and behave differently towards them. The importance of understanding the effects of this market practice has increased in the recent past given the development of information technologies and the Internet (for example, web-browser cookies) that allow firms to keep, gather, and process more information about their past customers.⁵⁴ This increase in information has led to the proliferation of customer relationship management practices in most industries. As of now, it seems that many firms collect more information about their customers' behavior than they are able to process. As firms get better at processing this large amount of information, the effects of customer recognition are going to become more and more important. In fact, the Internet allows also firms to interact more directly with their customers, and better respond to this increase in information.

Most of the work until now has been on the firms' pricing decisions, (with the exception of the limited work discussed in Section 4). Firms use consumer behavior to target many other sorts of decisions, including their product offerings and communication policies. As of now we have still very little understanding of how these activities can interact with the ability of firms to recognize customers. This means that research on this problem has so far just uncovered the "tip of the iceberg," and that there is much work to

⁵⁴ See Rossi et al. (1996) for a discussion of available databases of purchase histories and their possible use in direct marketing. Pancras and Sudhir (2005) present an empirical application of personalization activities (for example, offering of coupons) in grocery retailing. Lewis (2005) presents an application to subscriber data of a large metropolitan newspaper of the dynamic issues in pricing using the past consumer purchase behavior.

be done on behavior-based targeting in the future. It would also be interesting to see more empirical work testing for the results presented in this literature.⁵⁵

Research to date has identified several pricing effects in both monopoly and competition. As discussed in Section 2, in monopoly, we have to account for both behavior of the firm anticipating the future gain of having more information, and the strategic behavior of consumers anticipating what firms will do in the future with their information. As discussed there, we may end up having a “ratchet effect,” as consumers realize that they would be hurt by revealing their information, so that they incur costs (forgo utility) to conceal their preferences. Important factors in how these forces play out include the relative discount factors of the firm and the consumers, the feasibility of the firm offering long-term contracts, the effect of new generations of consumers coming into the market, and the effect of consumer preferences changing (with positive correlation) through time.

In markets with multiple firms there is the additional effect of firms poaching each other’s customers with special deals. This generates interesting strategic effects, possibly inefficient switching, and effects on the intensity of competition. In addition to the possibility of firms offering long-term contracts, and the entry of new customers (or customers changing preferences), another effect that can be important in several markets is the presence of switching costs or network externalities.

Allowing firms to recognize customers raises the question of what can firms do with such information, and whether consumers should have the right to privacy in their market interactions. Furthermore, in some markets, the characteristics of consumers may affect profits directly and this may have additional effects on the functioning of the market as discussed in Section 5 in the context of credit markets.

Finally, the possibility of firms recognizing their past customers interact with several market aspects that have been substantially studied in the past such as customized pricing, switching costs, durable-goods markets, and bargaining.

⁵⁵There is already some related empirical work. See, for example, Goldberg (1996) and Guha and Wittink (1996) who show that empirical dealer discounts for new cars are a function of whether it is a first-time purchase and whether there is a trade-in.

References

- Acquisti, A., and H.R. Varian, 2005, "Conditioning Prices on Purchase History," *Marketing Science*, forthcoming.
- Armstrong, M., and J. Vickers, 2001, "Competitive Price Discrimination," *RAND Journal of Economics*, 32, 579-605.
- Armstrong, M., 2005, "Economics Models of Price Discrimination," working paper, University College London.
- Ausubel, L., and R. Deneckere, 1992, "Durable Goods Monopoly with Incomplete Information," *Review of Economic Studies*, Vol. 59, 795-812.
- Baye, M.R., D. Kovenock, and C.G. de Vries, 1992, "It Takes Two-to-Tango:Equilibria in a Model of Sales," *Games and Economic Behavior*, 4, 493-510.
- Beggs, A.W., 1989, "A Note on Switching Costs and Technology Choice," *Journal of Industrial Economics*, 37, 437-440.
- Baron, D., and D. Besanko, 1984, "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1, 447-470.
- Battaglini, M., 2004, "Long-Term Contracting with Markovian Consumers," *American Economic Review*, forthcoming.
- Beggs, A. and P. Klemperer, 1992, "Multi-Period Competition with Switching Costs," *Econometrica*, 60, 651-666.
- Ben-Shoham, A., 2005, "Information and Order in Sequential Trade," working paper, Harvard University.
- Bulow, J.I., 1982, "Durable-Goods Monopolists." *Journal of Political Economy*, 90, 314-332.
- Cabral, L., and J.M. Villas-Boas, 2005, "Bertrand Supertraps," *Management Science*, 51, 599-613.
- Calzolari, G., and A. Pavan, 2005, "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, forthcoming.
- Caminal, R. and Matutes, C., 1990, "Endogenous Switching Costs in a Duopoly Model." *International Journal of Industrial Organization*, 8, 353-373.
- Chen, P.-Y., and L.M. Hitt, 2005, "Information Technology and Switching Costs," this handbook.

- Chen, Yongmin, 1997, "Paying Customers to Switch." *Journal of Economics and Management Strategy*, 6, 877–897.
- Chen, Yuxin and G. Iyer, 2002, "Consumer Addressability and Customized Pricing," *Marketing Science*, 21,197–208.
- Chen, Yuxin, C. Narasimhan, and Z.J. Zhang, 2001, "Individual Marketing and Imperfect Targetability," *Marketing Science*, 20, 23-41.
- Chen, Yuxin and Z.J. Zhang, 2004, "The Benefit of Behavior-Based Targeted Pricing with Strategic Consumers", working paper, New York University and University of Pennsylvania.
- Conlisk, J., E. Gerstner, and J. Sobel, 1984, "Cyclic Pricing by a Durable Goods Monopolist." *Quarterly Journal of Economics*, 99, 489-505.
- D'Aspremont, C., J. Gabszewicz, and J.-F. Thisse, 1979, "On Hotelling's Stability in Competition," *Econometrica*, 17, 1145-1151.
- Dana, J., 1998, "Advance-Purchase Discounts and Price Discrimination in Competitive Markets." *Journal of Political Economy*, 106, 395-422.
- Dell'Ariccia, E. Friedman, and R. Marquez, 1999, "Adverse Selection as a Barrier to Entry in the Banking Industry", *RAND Journal of Economics*, 30, 515-534.
- Dell'Ariccia, and R. Marquez, 2004, "Information and Bank Credit Allocation," *Journal of Financial Economics*, 72, 185-214.
- Deneckere, R., and P. McAfee, 1996, "Damaged Goods," *Journal of Economics and Management Strategy*, 5, 149-174.
- Desai, P., 2001, "Quality Segmentation in Spatial Markets: When Does Cannibalization Affect Product Line Design?," *Marketing Science*, 20, 265-283.
- Dobos, G., 2004, "Poaching in the Presence of Switching Costs and Network Externalities," working paper, University of Toulouse.
- Doganoglu, T., and L. Grzybowski, 2004, "Dynamic Duopoly Competition with Switching Costs and Network Externalities," working paper, University of Munich.
- Economides, N., 1986, "Minimal and Maximal Product Differentiation in Hotelling's Duopoly," *Economics Letters*, 21, 67-71.
- Ellison, G. and D. Fudenberg, 2000, "The Neo-Luddite's Lament: Excessive Upgrades in the Software Industry." *RAND Journal of Economics*, 31, 253-272.

Engelbrecht-Wiggans, R., P. Milgrom, and R. Weber, 1983, "Competitive Bidding and Proprietary Information," *Journal of Mathematical Economics*, 11, 161-169.

Esteves, R.B., 2004, "Pricing under Customer Recognition", working paper, Oxford University

Farrell, J., and C. Shapiro, 1988, "Dynamic Competition with Switching Costs," *RAND Journal of Economics*, 19, 123-137.

Farrell, J., and P. Klemperer, 2004, "Coordination and Lock-In: Competition with Switching Costs and Network Effects," *Handbook of Industrial Organization*, forthcoming.

Freixas, X., R. Guesnerie, and J. Tirole, 1985, "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52,173–191.

Fudenberg, D., and D. Levine, 1989, "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759-778.

Fudenberg, D., and J. Tirole, 1983, "Sequential Bargaining with Incomplete Information," *Review of Economic Studies*, 50, 221-247.

Fudenberg, D., and J. Tirole, 1984, "The Fat Cat Effect, the Puppy Dog Ploy and the Lean and Hungry Look," *American Economic Review, Papers and Proceedings*, 74, 361-368.

Fudenberg, D., and J. Tirole, 1991, *Game Theory*, The MIT Press, Cambridge, Massachusetts.

Fudenberg, D. and J. Tirole, 1998, "Upgrades, Trade-ins, and Buybacks," *RAND Journal of Economics*, 29, 238–258.

Fudenberg, D., and J. Tirole, 2000, "Customer poaching and brand switching," *RAND Journal of Economics*, 31, 634–657.

Goldberg, P.K., 1996, "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 104, 622-654.

Grossman, G.M., and C. Shapiro, 1984, "Informative Advertising with Differentiated Products," *Review of Economic Studies*, 51, 63-81.

Guha, R., and D. Wittink, 1996, "The Relationship Between Replacement Time, Brand Loyalty and Transaction Price for Automobiles," working paper, Cornell University.

- Hart, O.D. and J. Tirole, 1988, "Contract renegotiation and coasian dynamics," *Review of Economic Studies*, 55, 509–540.
- Hermalin, B.E., and M.L. Katz, 2004, "Privacy, Property Rights & Efficiency: The Economics of Privacy as Secrecy," working paper, University of California, Berkeley.
- Hui, K.-L., and I.P.L. Png, 2005, "The Economics of Privacy," this handbook.
- Iyer, G., D. Soberman, and J.M. Villas-Boas, 2005, "The Targeting of Advertising," *Marketing Science*, forthcoming.
- Katz, M.L., 1984, "Price Discrimination and Monopolistic Competition," *Econometrica*, 52, 1453-1472.
- Kennan, J., 2001, "Repeated Bargaining with Persistent Private Information." *Review of Economic Studies*, 68, 719-755.
- Kim, B.D., M. Shi, and K. Srinivasan, 2001, "Reward Programs and Tacit Collusion," *Marketing Science*, 20, 99-120.
- Klemperer, P., 1987a, "The Competitiveness of Markets with Switching Costs," *RAND Journal of Economics*, 18, 138-150.
- Klemperer, P., 1987b, "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, 102, 375-394.
- Klemperer, P., 1995, "Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics and International Trade," *Review of Economic Studies*, 62, 515–539.
- Kreps, D., P. Milgrom, J. Roberts, and R. Wilson, 1982, "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 27, 245-252.
- Laffont, J.-J., and J. Tirole, 1990, "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 57, 597-625.
- Lee, I.H., and J. Lee, 1994, "Durable Goods Monopoly under Technological Innovation," working paper, University of Southampton.
- Lenvinthal, D.A., and D. Purohit, 1989, "Durable Goods and Product Obsolescence," *Marketing Science*, 8, 35-56.
- Lewis, M., 2005, "Research Note: A Dynamic Programming Approach to Customer Relationship Pricing," *Management Science*, 51, 986-994.
- McAdams, D., 2005, "A Theory of Sticker Prices," working paper, M.I.T.

- Morgan, J., and P. Stocken, 1998, "The Effects of Business Risk on Audit Pricing," *Review of Accounting Studies*, 3, 365-385.
- Mussa, M., and S. Rosen, 1978, "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301-317.
- Nahm, J., 2004, "Durable-Goods Monopoly with Endogenous Innovation," *Journal of Economics and Management Strategy*, 13, 303-320.
- Nilssen, T., 1992, "Two Kinds of Consumer Switching Costs," *RAND Journal of Economics*, 23, 579-589.
- Padilla, J., and M. Pagano, 1997, "Endogenous Communications among Lenders and Entrepreneurial Incentives," *Review of Financial Studies*, 10, 205-236.
- Padilla, J., and M. Pagano, 2000, "Sharing Default Information as a Borrower Incentive Device," *European Economic Review*, 44, 1951-1980.
- Pagano, M., and T. Jappelli, 1993, "Information Sharing in Credit Markets," *Journal of Finance*, 48, 1693-1718.
- Pancras, J., and K. Sudhir, 2005, "The Personalization Services Firm: What to Sell, Whom to Sell to and For How Much?," working paper, University of Connecticut.
- Rajan, R.G., 1992, "Insiders and Outsiders: The Choice between Informed and Arm's-Length Debt," *Journal of Finance*, 47, 1367-1400.
- Rossi, P.E., R.E. McCulloch, and G.M. Allenby, 1996, "The Value of Purchase History Data in Target Marketing," *Marketing Science*, 15, 321-340.
- Roy, S., 2000, "Strategic Segmentation of a Market," *International Journal of Industrial Organization*, 18, 1279-1290.
- Schmidt, K., 1993, "Commitment through Incomplete Information in a Simple Repeated Bargaining Game," *Journal of Economic Theory*, 60, 114-139.
- Schwartz, M., 1997, "Telecommunications Reform and the United States: Promises and Pitfalls," in P.J.J. Welfens and G. Yarrow, eds., *Telecommunications and Energy in Systemic Transformation*. New York: Springer.
- Shaffer, G., and Z.J. Zhang, 1995, "Competitive Coupon Targeting," *Marketing Science*, 14, 395-416.

- Shaffer, G., and Z.J. Zhang, 2000, "Pay to Switch or Pay to Stay: Preference-Based Price Discrimination in Markets with Switching Costs," *Journal of Economics and Management Strategy*, 9, 397-424.
- Sharpe, S., 1990, "Asymmetric Information, Bank Lending and Implicit Contracts: A Stylized Model of Customer Relationships," *Journal of Finance*, 45, 1069-1087.
- Sobel, J., 1984, "The Timing of Sales," *Review of Economic Studies*, 51, 353-368.
- Sobel, J., 1991, "Durable Goods Monopoly with Entry of New Consumers," *Econometrica*, 59, 1455-1485.
- Stegeman, M., 1991, "Advertising in Competitive Markets," *American Economic Review*, 81, 210-223.
- Stokey, N., 1979, "Intertemporal Price Discrimination," *Quarterly Journal of Economics*, 93, 355-371.
- Stokey, N., 1981, "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics*, 12, 112-128.
- Stole, L., 1995, "Nonlinear Pricing and Oligopoly," *Journal of Economics and Management Strategy*, 4, 529-563.
- Stole, L., 2004, "Price Discrimination in Competitive Environments," *Handbook of Industrial Organization*, forthcoming.
- Taylor, C.R., 2003, "Supplier Surfing: Price-Discrimination in Markets with Repeat Purchases," *RAND Journal of Economics*, 34, 223-246.
- Taylor, C.R., 2004a, "Consumer Privacy and the Market for Customer Information," *RAND Journal of Economics*, 35, 631-650.
- Taylor, C.R., 2004b, "Privacy and Information Acquisition in Competitive Markets," working paper, Duke University.
- Thisse, J.-F. and X. Vives, 1988, "On the Strategic Choice of Spatial Price Policy," *American Economic Review*, 78, 122-137.
- Tirole, J., 1988, *The Theory of Industrial Organization*, The MIT Press, Cambridge, Massachusetts.
- Ulph, D. and N. Vulkan, 2000, "Electronic Commerce and Competitive First-Degree Price Discrimination," Technical report, University College, London.

Ulph, D. and N. Vulkan, 2001, "E-commerce, Mass Customization and Price Discrimination," Technical report, University College, London.

Varian, H., 1980, "A Model of Sales," *American Economic Review*, 70, 651-659.

Villanueva, J. P. Bhardwaj, Yuxin Chen, and S. Balasubramanian, 2004, "Managing Customer Relationships: Should Managers Really Focus on the Long Term?" working paper, University of California, Los Angeles.

Villas-Boas, J.M., 1999, "Dynamic Competition with Customer Recognition," *RAND Journal of Economics*, 30, 604-631.

Villas-Boas, J.M., and U. Schmidt-Mohr, 1999, "Oligopoly with Asymmetric Information: Differentiation in Credit Markets," *RAND Journal of Economics*, 30, 375-396.

Villas-Boas, J.M., 2004, "Price Cycles in Markets with Customer Recognition," *RAND Journal of Economics*, 35, 486-501.

Waldman, M., 1996, "Planned Obsolescence and the R&D Decision," *RAND Journal of Economics*, 27, 583-595.

Wathieu, L., 2004, "Marketing and the Privacy Concern," working paper, Harvard University.

Zhang, Juanjuan, 2005, "Dynamic Segmentation and Product Design," working paper, University of California, Berkeley.

FIGURE 1




FIGURE 2




FIGURE 3

