Conditioning Prices on Purchase History

Alessandro Acquisti
Hal R. Varian*
University of California, Berkeley

August 2001
Revised: May 11, 2003

Abstract

The rapid advance in information technology now makes it feasible for sellers to condition their price offers on consumers' prior purchase behavior. In this paper we examine when it is profitable to engage in this form of price discrimination.

Our baseline model involves rational consumers with constant valuations for the good being sold and a monopoly merchant who can commit to a pricing policy. Applying results from prior literature, we show that although it is *feasible* to price so as to distinguish high-value and low-value consumers, the merchant will never find it *profitable* to do so.

We then consider various generalizations of this model, such as allowing the seller to offer enhanced services to previous customers, making the merchant unable to commit to a pricing policy, and allowing competition in the marketplace. In these cases we show that sellers will, in general, find it profitable to condition prices on purchase history.

Keywords: Internet Marketing, Personalized Marketing, Price Discrimination, Targeting.

---

*Research support from NSF grant 9979852 is gratefully acknowledged. Email contact: `acquisti@sims.berkeley.edu, hal@sims.berkeley.edu`.

1 Introduction

In 1988 the cost of a gigabyte of hard disk storage was about $11,500. By 2000 that cost was $13, roughly 900 times cheaper. Today, a gigabyte of storage costs less than a dollar. This remarkable reduction of the cost of storing information has led firms to capture and analyze detailed information about transactions with their customers.

Many industries, including supermarkets, airlines, and credit cards, have compiled vast databases of individual consumer transactions which have been used both to study purchase behavior and to make specific offers to individual consumers, via direct mail or other forms of targeted marketing. Many companies have become expert in using tracking tools to refine marketing strategies (see Bailey [1998], and Dayal et al. [2001]).

Since so many transactions are now computer mediated, and these computers can easily be networked to data centers, sellers now have the ability to access databases of past purchases in real time. This allows them to condition current offers to consumers on their previous purchase behavior. Sellers can offer each individual a different price, a particular prize or coupon, or personalized recommendations. With computer mediated transactions, price discrimination on an individual basis becomes quite feasible.

Collecting and analyzing such information is particularly easy in the online world. Though the HTTP protocol used by Web servers is stateless, browsers typically accept "cookies" from servers that contain information about the current transaction.[1] These cookies persist after the session has ended, so that the next time the user accesses the server (using the same account) the server can retrieve identification which can be matched with details of past interactions. Even without cookies, static IP addresses, credit card numbers, direct user authentication, and a variety of other mechanisms can be used to identify individual users.

There is considerable evidence that Internet merchants have attempted to use "dynamic pricing" in order to price discriminate (see Taylor [2002] and Streifield [2001]). However, there is a long literature in economics, dating at least back to Stokey [1979], that shows that, under certain conditions, a seller facing strategic customers cannot do better than committing to optimal single period pricing. Why, then, do both online and offline merchants invest so much on tracking technologies?

[1] See Schartz [2001] for a history of Web cookies.

Our answer is based on considering the buyers' behavior. Although sellers can now easily use price conditioning strategies, consumers are far from defenseless. No one is forced to join a loyalty program, and it is relatively easy to set one's browser to reject cookies or to erase them after a session is over. Consumers can use a variety of credit cards or more exotic anonymous payment technologies to make purchases hard to trace. In addition, consumers can voice their outrage for pricing policies perceived as discriminatory, as happened after the famous Amazon.com price experiment (Streifield [2001]).

Thus, even though sellers can post prices, observe choices, and condition subsequent price offers on observed behavior, buyers are also able to hide the fact that they bought previously. Hence, it is likely that sellers will have to offer buyers some benefits in order to induce them to reveal their identities.

In this paper we develop models of this strategic interaction between buyer and seller in order to determine conditions under which sellers will find it profitable to condition prices on purchase history. We find that in the simplest model, where consumers' valuations of the good being sold are constant, sellers do not want to condition current price offers on past behavior. However, we also show that if the consumer's value for the good changes as he or she makes more purchases, the seller will find it profitable to condition prices on past behavior. Both of these results have been observed in the previous literature, although we believe that we provide a particularly simple treatment.

We then go on to point out that the seller can induce the necessary change in consumer valuations by offering enhanced services to prior users, such as discount coupons (common in supermarket loyalty clubs), prizes or awards (common with airlines and credit cards), lowered transactions costs (such as one-click shopping), or personalized services (such as recommendations). We derive conditions describing when the use of such technologies is profitable to sellers.

This analysis takes place in the context of rational, fully informed consumers. We also find that if enough customers are myopic, or the costs of using anonymizing technologies are too high, sellers will want to condition pricing on purchase history.

We then briefly consider scenarios where the seller cannot commit to prices; in this case, buyers will want to randomize their purchases in order to hide their type, leading to mixed strategy equilibria.

Finally, we examine the case of competition and find that if the availability of purchase history allows firms to offer higher value, customized offers to consumers, we will end up with various kinds of lock-in equilibria in which firms first invest in acquiring consumer information, and then exploit this information to provide personalized, albeit high priced, service to some consumers.

2 Literature review

The earliest relevant contribution to our research is Stokey [1979]. She shows that intertemporal price discrimination is never optimal for a monopoly seller who can commit to future prices. Salant [1989] extends Stokey's result to the case of multiple types and shows that it follows from the linearity of the constraints in the problem. Riley and Zeckhauser [1983] find a "no haggling" result that is also related to linearity issue examined in Salant [1989].

Several authors have examined intertemporal price discrimination in models in which the seller is unable to commit to future prices. These authors often use the term "ratchet effect" to describe situations in which consumers who have signaled higher willingness to pay for a product tend to receive higher prices, even though the consumers behave strategically so as to avoid this outcome. See, for example, Hart and Tirole [1988], who study the strategies for sellers who are unable to commit to pricing policies in a repeated purchase game with incomplete information.

The above papers consider a monopoly seller. Subsequent work such as Fudenberg and Tirole [1998], Villas-Boas [1999, 2001], Fudenberg and Tirole [2000], and Chen and Zhang [2001] investigates generalizations in both monopolistic and oligopolistic markets.

Fudenberg and Tirole [1998] ask what happens when the ability to identify particular consumers may vary across goods. In particular, they consider a model of goods upgrades and buy-backs where customers may be anonymous or "semi-anonymous." In the latter case, customers may be able to prove that they have purchased a previous version of a product, so as to take advantage of upgrade offers, for example. Fudenberg and Tirole [1998] find that if the new good is a significant improvement over the previous good, a monopolist will prefer to charge a higher price for upgrades than for sales to new consumers.

Fudenberg and Tirole [2000] analyze a duopoly in which some consumers remain loyal and others defect to the competitor, a phenomenon they refer to as "customer poaching." Villas-Boas [1999] finds that two firms in a duopoly can compete by lowering prices to attract the competitor's previous customers. Villas-Boas [2001] shows that targeted pricing by a monopolist who cannot commit to future prices may make it worse off.

Chen and Zhang [2001] demonstrate a similar result in the case of a duopoly. To escape this impasse, Chen and Zhang [2001] propose a "price for information" strategy, with firms pricing less aggressively in order to learn more about their customers, a point related to the lock-in model we develop below.

The most recent work in the economics literature has to do with concerns about consumer privacy. Calzolari and Pavan [2001] and Taylor [2002] have examined the implications of information exchange among sellers.

Calzolari and Pavan [2001] study contracting environments with two principals (e.g., two merchants) that interact sequentially with one common agent (the consumer). They assume that the consumer's tastes for the goods sold by the merchants are perfectly correlated and find that the transmission of information about the consumer between merchants may lead to increases in welfare and consumer surplus when the goods are complements. On the other hand, if consumer valuations for the two goods is additive, then a firm may benefit from committing not to sell customer information.

Taylor [2002] also studies the market for customer information. He considers two settings: an anonymity regime where the sale of customer information is not possible, and a recognition regime where firms can sell customer information. Taylor finds that the welfare implications of the various technologies depend on the sophistication of the consumers. He finds that consumers fare poorly and firms fare well under an "open privacy" regime (where sale of customer information is permitted) when consumers are myopic. However, when customers are more sophisticated, firms benefit from keeping their customers' information private.

In the marketing literature, identification of specific customers is often called "customer addressability," and is viewed as a prominent feature of "interactive marketing." See Blattberg and Deighton [1991], McCulloch et al. [1996] and Rossi and Allenby [1999] for a general overview of the issues raised by these capabilities.

The term "customer addressability" is usually applied to situations where customers can be identified by a characteristic such as a zip code, which is a predictor of the valuation that they might place on a good. Examples of this literature include Ulph and Vulkan [2000, 2001], Chen and Iyer [2002], and Chen et al. [2003].

The difference between "customer addressibility" and "conditioning prices on purchase history" comes in the fact that the latter case considers the strategic response of consumers to the pricing strategies set by the seller. "Customer addressibility" is related to first-degree price discrimination, which is based on identity as a signal of value, while "conditioning on purchase history" is a form of second-degree price discrimination, with previous behavior as a signal of value. In fact, Fudenberg and Tirole [2000] use the term "behaviorially based price discrimination" to describe what we call "conditioning prices on purchase history."

Perhaps the easiest way to clarify the difference between "addressibility" and "behaviorially based" or "conditioning" effects would be to observe the following: few consumers would change their zip code to get a lower price on a book they purchase online, but many may well delete a cookie or delay purchase in order to get a better price.

Our contribution to the literature described above is as follows.

- We offer a simple algebraic and geometric treatment of the classic results that sellers will not want to engage in intertemporal price discrimination when consumers are sophisticated.
- We show that sellers will want to condition prices on past behavior if a large enough fraction of the consumer population is myopic.
- We show that sellers will want to condition prices if they can offer enhanced services that high-value and low-value consumers value differently.
- In this case, we compare consumer and producer welfare in conditioning and flat-price regimes, shedding some light on the consequences of regulating price discrimination of this sort.
- We compare our results to the literature where firms cannot commit to future prices, and examine the impact of consumers delaying purchase versus making anonymous purchase on equilibrium outcomes.
- We develop a model of competitive sellers in which the purchase history of consumers provides a competitive advantage by allowing firms to tune offers to individual tastes. This imposes a switching costs on consumers, since they would lose their personalized services by switching to a competitor, leading to a form of lock-in equilibrium.

3 The baseline model

We begin with a simple model of a single profit-maximizing seller of a good that can be provided at zero marginal cost.

We assume that the seller has a mechanism for recording purchase histories of customers. This could be based on technologies such as loyalty program identification, credit card numbers, static Internet addresses, or other such devices, but we will refer to it as a "cookie."

New customers come to the seller and are offered a price. Their decision about whether to purchase at this price is observed. The second time they come to the seller the price they are offered can be conditioned on their earlier behavior. The following list summarizes the possible actions the seller can take depending on whether or not a cookie is present and what it indicates about prior behavior.

No cookie. The seller offers a price and records whether or not the customer purchases. It sets a cookie indicating whether or not purchase took place at the offered price.

Cookie shows customer bought before at price offered. The seller offers a price which may depend on the details of the previous purchase.

Cookie shows customer did not buy before at price offered. The seller offers a possibly different price.

We approach this problem from the perspective of mechanism design. For simplicity we restrict ourselves to two types and two periods. Let v_H be the high-value type's willingness to pay for one unit of consumption, and let v_L be the low-value type's willingness to pay for one unit of consumption. Let π indicate the fraction of the population that has the high value.

We will assume that if the consumer is indifferent, he will act in the manner preferred by the seller, since the seller could always cut a price by a penny if it were profitable to do so. For simplicity, we also assume a zero discount rate.

Obviously, if the seller sets a flat price of v_H each period, it will make a profit of $2\pi v_H$ and if it sets a flat price of v_L each period, it will make a profit of $2v_L$. The maximum profit available from flat pricing is therefore $\max\{2\pi v_H, 2v_L\}$.

We are interested in whether the seller can do better by some form of individual conditioning that will allow price discrimination so that the high-value person pays more than the low-value person.

3.1 All consumers myopic

Myopic consumers are those who base their purchase decision on the price that they see today, not recognizing that the price they face on the next purchase may depend on today's behavior. We present this case as a baseline for comparisons to more interesting scenarios below.

If consumers are myopic, the seller can offer a price of v_H in the first meeting with the consumer. If the consumer does not purchase at this price, the seller can offer a price of v_L the second time.

This strategy results in sales of 2 units to the high-value population and 1 unit to the low-value population, yielding revenues of

$$2\pi v_H + (1-\pi)v_L.$$

How does price conditioning affect overall welfare? There are two cases.

Case 1. $\pi v_H > v_L$. If conditioning were not possible, the seller would sell only to the high-value consumers. Allowing conditioning doesn't change the price the high-value consumers face, but does allow the seller to identify the low-value consumers. This allows the merchant to sell to the low-value consumers at v_L, giving the them zero surplus from the purchase. Hence overall welfare (producer plus consumer surplus) rises, but this is entirely due to the increased profit received by the seller.

Case 2. $\pi v_H < v_L$. If conditioning were not possible, the seller would sell to everyone at price v_L. If conditioning is feasible, it will be profitable when

$$\pi > \frac{v_L}{v_H}\left(\frac{1}{2 - v_L/v_H}\right).$$

Thus, there will be a set of values of π determined by

$$\frac{v_L}{v_H} > \pi > \frac{v_L}{v_H}\left(\frac{2}{1 - v_L/v_H}\right),$$

for which the seller would sell to everyone if it didn't have a way to condition, but chooses to restrict output when a conditioning technology is available.

It follows that the high-value consumers loose under conditioning and the low-value consumers are no worse off (though they consume in one period rather than two.) The seller is better off due to selling at a higher price to the high-value consumers but worse off from losing one period of revenue from the low-value consumers. If the seller voluntarily chooses to condition, it must be better off, but overall welfare is reduced due to price conditioning.

3.2 All consumers sophisticated

High-value consumers may eventually come to recognize that purchasing at a high price is not the best strategy, since it guarantees that they will face a high price in the future.

Let us suppose that consumers can delete cookies, use an anonymous payment system, delay purchase, or take some other steps to avoid establishing a purchase history. The literature on intertemporal price discrimination discussed above has shown that there are conditions under which dynamic pricing will *not* be optimal for a monopolistic seller when sellers cannot identify individual buyers (see, for example, Stokey [1979] and Hart and Tirole [1988]). Here we ask if the same result holds when sellers can track individual buyers and condition pricing on their purchase history.

Let p_H be the present value (in this case, the sum) of the prices charged to the high-value person, and p_L the present value of the prices charged to the low-value person. Let x_H be the total amount consumed by the high-value type and x_L the total amount consumed by the low-value type.

The optimization problem facing the seller is:

$$\begin{aligned}
\max_{x_H, x_L, p_H, p_L} \quad & \pi p_H + (1-\pi) p_L && (1)\\
v_H x_H - p_H \geq\ & v_H x_L - p_L && (2)\\
v_H x_H - p_H \geq\ & 0 && (3)\\
v_L x_L - p_L \geq\ & v_L x_H - p_H && (4)\\
v_L x_L - p_L \geq\ & 0. && (5)
\end{aligned}$$

It is clear that due to the linearity of the problem x_L and x_H can only take on the values $\{0, 1, 2\}$. Table 1 lists the maximum revenue associated with the eight possible cases.

Note that the last three cases dominate the others, so that there are relatively few interesting cases. The case $(x_H, x_L) = (2, 0)$ is the case, already analyzed, where the pricing plan induces only the high-value consumers

x_H	x_L	Maximum revenue
0	0	0
0	1	Not incentive compatible
0	2	Not incentive compatible
1	0	πv_H
1	1	v_L
1	2	Not incentive compatible
2	0	$2\pi v_H$
2	1	$\pi v_H + v_L$
2	2	$2v_L$

Table 1: Payoffs and profits.

to purchase: the optimal strategy is therefore pricing $p = v_H$. The case $(x_H, x_L) = (2, 2)$ is the one where all consumers are induced to purchase in both periods: the price must therefore be $p = v_L$, to avoid that certain customers delay their purchase (in the hope of bargaining down the price in another period) or delete their cookies after a low initial price (in the hope that they will get the same low price also in later periods). We will examine two of the other cases in the table to get a flavor for the analysis.

Consider case $(x_H, x_L) = (1, 2)$. Among the self-selection constraints are

$$v_H - p_H \geq 2v_H - p_L \tag{6}$$

$$2v_L - p_L \geq v_L - p_H. \tag{7}$$

Rearranging these inequalities gives us the contradiction

$$v_L \geq p_L - p_H \geq v_H. \tag{8}$$

Now consider case $(x_H, x_L) = (2, 1)$, where the pricing strategy is such to induce the high-value consumers to purchase twice, and the low-value consumer to purchase only once. The self-selection constraints are

$$2v_H - p_H \geq v_H - p_L \tag{9}$$

$$2v_H - p_H \geq 0 \tag{10}$$

$$v_L - p_L \geq 2v_L - p_H \tag{11}$$

$$v_L - p_L \geq 0. \tag{12}$$

It is easily seen that the solution to these inequalities is $p_L = v_L$ and $p_H = v_H + v_L$, which yields a profit of $\pi v_H + v_L$. That is, in order to have the high-value consumers buy twice, the price must be discounted in one period. When does this exceed the profit from flat pricing? That is, when do we have

$$\pi v_H + v_L > \max\{2\pi v_H, 2v_L\}? \tag{13}$$

The following result shows the answer is "never."

Fact 1 (Conditioning is not optimal.) *If consumers are sophisticated, the seller does at least as well by setting a flat price as it does by condition prices on past behavior.*

Proof. Writing out the necessary inequalities in 13 we have

$$\begin{aligned} \pi v_H + v_L &> 2\pi v_H & (14) \\ \pi v_H + v_L &> 2v_L. & (15) \end{aligned}$$

Adding these together gives a contradiction. □

We can illustrate the essence of the argument geometrically. The most interesting case of price conditioning in the previous section was case (2,1), in which the high-value type consumed both periods and the low-value type consumed only in the second period.

The self-selection constraints for this case are given in inequalities 9-12, which can be written as

$$\begin{aligned} v_H + p_L &\geq p_H & (16) \\ 2v_H &\geq p_H & (17) \\ p_H &\geq v_L + p_L & (18) \\ v_L &\geq p_L. & (19) \end{aligned}$$

We have plotted these inequalities in Figure 1. It is clear by inspection that $(p_H, p_L) = (v_H + v_L, v_L)$ is the only candidate for profit-maximization; this is the point illustrated by the black dot.

But does this solution dominate flat pricing? Refer to Figure 2 where we have plotted the demand curve for total consumption for two groups of consumers. The shaded part of the curve equals the revenue extracted by the seller under these prices. It is easy to see that if $A > B$, charging only v_H

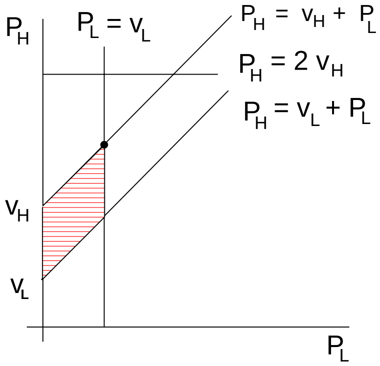

Figure 1: Self-selection constraints.

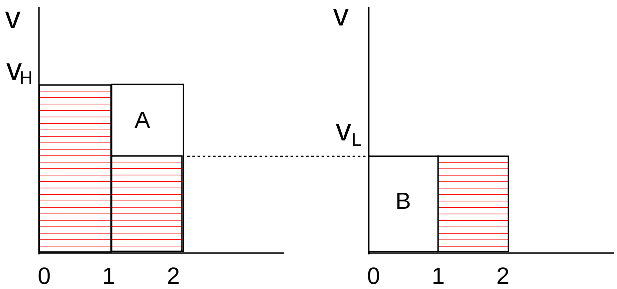

Figure 2: Demand curves, with shaded area indicating revenue.

yields more revenue, and if $B > A$ charging the v_L yields more revenue. Hence one of these two flat pricing strategies must dominate differential pricing.

Intuitively, if selling to the high-value customers is more profitable than selling to both high- and low-value customers, then the seller wants to always sell to those customers. There is no advantage to cutting its price to sell to the low-value customers.

As we have indicated, the result that price discrimination is not profitable when the seller can commit to prices has been analyzed by Stokey [1979], Salant [1989], and Riley and Zeckhauser [1983]. However, note that we allow the seller to identify individual buyers and condition pricing on purchase history of individual consumers, whereas the Stokey-Salant treatment only considered posted prices. Nevertheless, we find that with sophisticated consumers, our model has the same "reduced form" as the Stokey model. This is basically a consequence of the revelation principle: both the Stokey-Salant model and the model we examine are equivalent to the same mechanism design problem.

This result that sellers will not want to condition prices on purchase history is somewhat disconcerting since sellers have invested many millions of dollars in computer systems to allow them to collect data to allow them to do exactly that. Though experimentation with such systems has only gone on for a few years in the online world, loyalty programs for airlines travelers and supermarket shoppers have been around for years. Such programs commonly offer special prices to consumers with different purchase histories. But the results described above show that, at least in the simplest model, such behavior is not profitable. What is missing from this model?

4 Profitable conditioning

We saw earlier that if all consumers are myopic, price conditioning could be profitable, so it is not surprising that if a large enough fraction of the population is myopic, the same result will hold. We will discuss the relevant conditions in the next section.

Another, perhaps more interesting answer, is customer resistance. However, frequent flyer programs and supermarket loyalty cards have been extremely popular. No one likes to think that they have been charged more than anyone else, but everyone likes to get a price break. Structuring a personalized pricing program in a palatable way is important, but the airline and supermarket examples show that it can be done.

Consider again Figure 2. The geometric argument for no conditioning depended heavily on the fact that the value of the first purchase was the same as the value of the second purchase. It appears likely that price discrimination could be optimal when the value of the second unit of consumption has a different value than the first.

We could, for example, assume that marginal utility of consumption is decreasing so that the value of the second unit of consumption is less than that of the first. This is equivalent to the standard analysis of quality or quantity discrimination, in which utility is assumed to be a concave function of quality/quantity. See Mussa and Rosen [1978] and Maskin and Riley [1984] for early treatments and Tirole [1988] and Varian [1992] for textbook analysis. Salant [1989] establishes conditions for price discrimination to dominate flat pricing when marginal utilities decrease with consumption.

A more interesting assumption, in our context, is to examine the case where the second unit of consumption is *more* valuable than the first. This could arise because the second visit to the merchant is more efficient or pleasant than the first one, which might occur because the seller offered enhanced services of some form, which could be enabled by the information the customer has revealed during the first purchase. Examples could be targeted recommendations, personalized service or content, one-click shopping, prizes, or a variety of other enhanced services. This case has not been examined in the previous literature, but is relatively easy to handle in our discrete framework.

4.1 Some consumers are myopic

Suppose that a fraction m of each type is myopic, with a fraction $1-m$ being sophisticated. This case is undoubtedly realistic, but relatively straightforward in terms of analysis, so we will conduct only a cursory examination.

Assume that the seller conditions prices on purchase history by first charging a high price to everyone and then offering a low price to those who did not purchase.

The low-value consumers and the sophisticated high-value customers will wait for second period to buy at the low price. The myopic high-value consumers pay the high price each period. The revenue the seller receives is therefore

$$2m\pi v_H + (1 - m\pi)v_L.$$

This will exceed the revenue from flat pricing when

$$m\pi > \max\left\{\frac{2\pi v_H - v_L}{2v_H - v_L}, \frac{v_L}{2v_H - v_L}\right\}.$$

Hence if the fraction of myopic consumers is large enough, the seller will want to condition prices on purchase history.

Note that, as usual, the presence of unsophisticated consumers can exert a negative externality on the sophisticated consumers, by creating an equilibrium that makes the sophisticated consumers worse off.

Of course, the myopic consumers could just be consumers who had a particularly high cost engaging in anonymizing behavior. We can make the choice between being myopic and sophisticated endogenous by considering a cost to deleting cookies. If the seller's strategy is to offer a low price first, the customer may want to buy at the low price, but then delete her history in order to present herself as a new customer to the same seller on the next encounter. The inconvenience cost of deleting cookies or adopting some anonymizing technology may be so large for some consumers that it is simply not worth doing.

4.2 Enhanced services

While Maskin and Riley [1984] consider second degree price discrimination when marginal utility may be decreasing, it is interesting in our scenario to consider what happens when incremental purchases of the good have higher utility that earlier purchases due to to enhanced service based on the information previously provided by the customer.

Let v_{H1} represents the value of the first unit of consumption for the high-value consumer, and v_{H2} the value of the second unit of consumption. Define v_{L1} and v_{L2} similarly. Of course, we assume that

$$v_{H1} > v_{L1} \tag{20}$$

Utility for the high-value consumer can take on 3 values, $(0, v_{H1}, v_{H1} + v_{H2})$ and likewise for the low-value consumer. Note that it is, by definition, impossible to receive a utility of v_{H2}. Thus there are 2^3 cases, which are summarized in Table 2.

The analysis mimics that in Section 3.2. The seller can again flat-price at the high or low prices, but the most interesting case is where the seller uses

High type	Low type	Maximum revenue
$v_{H1} + v_{H2}$	$v_{L1} + v_{L2}$	$v_{L1} + v_{L2}$
$v_{H1} + v_{H2}$	v_{L1}	$v_{L1} + \pi v_{L2}$
$v_{H1} + v_{H2}$	0	$\pi(v_{H1} + v_{H2})$
v_{H1}	$v_{L1} + v_{L2}$	Not incentive compatible
v_{H1}	v_{L1}	v_{L1}
v_{H1}	0	πv_{H1}
0	$v_{L1} + v_{L2}$	Not incentive compatible
0	v_{L1}	Not incentive compatible

Table 2: Payoffs and profits for multiple consumption case.

price-conditioning to have the high-value type consume twice and the low-value type consume once. The self-selection constraints for this conditioning solution in this case are

$$
\begin{aligned}
v_{H1} + v_{H2} - p_H &\geq v_{H1} - p_L & (21)\\
v_{H1} + v_{H2} - p_H &\geq 0 & (22)\\
v_{L1} - p_L &\geq v_{L1} + v_{L2} - p_H & (23)\\
v_{L1} - p_L &\geq 0, & (24)
\end{aligned}
$$

which can be transformed to

$$
\begin{aligned}
v_{H2} + p_L &\geq p_H & (25)\\
v_{H1} + v_{H2} &\geq p_H & (26)\\
p_H &\geq v_{L2} + p_L & (27)\\
v_{L1} &\geq p_L. & (28)
\end{aligned}
$$

These inequalities are plotted in Figure 3.

We are interested in the pricing plan $p_H = v_{H2} + v_{L1}$ and $p_L = v_{L1}$ which induces the high-value consumers to purchase twice, and the low-value consumer to purchase only once.

For this to be profitable, we need to verify that the horizontal line determined by $p_H = v_{H1} + v_{H2}$ passes above this optimum. Algebraically, this requires:

$$v_{H1} + v_{H2} > v_{H2} + v_{L1}, \tag{29}$$

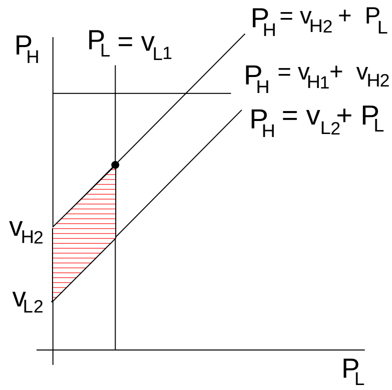

Figure 3: Self-selection constraints.

Making the cancellation we see that this condition reduces to assumption (20).

The revenue from price conditioning exceeds the revenue from flat pricing when

$$\pi v_{H2} + v_{L1} > \pi v_{H1} + \pi v_{H2} \tag{30}$$

$$\pi v_{H2} + v_{L1} > v_{L1} + v_{L2}. \tag{31}$$

Making the obvious cancellations gives us the following result.

Fact 2 (When is conditioning profitable?) *Conditioning prices will be profitable when*

$$v_{L1} > \pi v_{H1}$$

$$\pi v_{H2} > v_{L2}.$$

in which case $p_H = v_{H2} + v_{L1}$ *and* $p_L = v_{L1}$.

Note that the second inequality is more likely to hold when the seller can offer an enhanced service that is worth relatively more to the high-value type than to the low-value type. For example, one-click shopping may be more valuable to those who consume more frequently, or to those who have a higher value of time. Similarly, personalized coupons for baby food might be more valuable to consumers who have previously purchased diapers.

Case	Consumer Surplus	Producer Surplus	Total Surplus
Sell only to high-value	0	$\pi[v_{H1} + v_{H2}]$	$\pi[v_{H1} + v_{H2}]$
Condition prices	$\pi[v_{H1} - v_{L1}]$	$\pi v_{H2} + v_{L1}$	$\pi[v_{H1} + v_{H2}] + (1 - \pi)v_{L1}$
Sell to both	$\pi[v_{H1} + v_{H2} - v_{L1} - v_{L2}]$	$v_{L2} + v_{L1}$	$\pi[v_{H1} + v_{H2}] + (1 - \pi)[v_{L1} + v_{L2}]$

Table 3: Surplus calculations.

Indeed, if both types have the same value for the enhanced service, the necessary inequalities cannot both be satisfied. To see this, assume the contrary:

$$v_{H2} - v_{H1} = v_{L2} - v_{L1} = e > 0. \tag{32}$$

Now subtract the second inequality from the first in Fact 2, to find

$$\pi(v_{H2} - v_{H1}) > v_{L2} - v_{L1}. \tag{33}$$

Substituting, and recalling that $\pi < 1$, we have the contradiction

$$\pi e > e. \tag{34}$$

It is also worth noting that we have never needed to assume that $v_{H2} > v_{H1}$ or $v_{L2} > v_{H1}$. Hence Fact 2 applies in the classic case of "diminishing marginal utility" (as in Maskin and Riley [1984]) just as well as it does in the "enhanced service" case in our application.

4.3 Welfare effect of conditioning

How does price conditioning affect social welfare? The appropriate surplus calculations are shown in Table 3. Note that in terms of total welfare, conditioning fits between the two other cases. Conditioning dominates flat pricing when the alternative is selling only to the high-value type, but not when the alternative is selling to both types.

More specifically, if $v_{L1} + v_{L2} < \pi(v_{H1} + v_{H2})$, and the inequalities in Fact 2 are satisfied, allowing firms to use cookies makes the society as a whole better off. The welfare ordering of the outcomes is the same as the ordering of total quantity sold, which is consistent with the welfare analysis in Varian [1985].

5 Timing

We have seen that the seller will condition prices on purchase history when it is able to provide an enhanced service that has differential value to the consumers. In this case, the present value of the payments will be

$$p_H = v_{H2} + v_{L1} \tag{35}$$
$$p_L = v_{L1}. \tag{36}$$

Since we are assuming that the seller can commit to price plans, it appears that this present value can be divided between the two periods an arbitrary ways. However, whether or not that is the case depends on the tools that the high-value buyer has to defend himself against the price discrimination.

Let us return to the setup described in the introduction and consider an overlapping generations model where consumers visit an online store at most twice. If they have no cookie indicating a prior visit, they are charged p_0. If they have a cookie indicating that they bought on a prior visit, they are charged p_b. If they have a cookie indicating that they did not buy on their earlier visit, they are charged p_n.

One way to implement the pricing system described in equations (35-36) is to charge

$$p_0 = v_{H1} \tag{37}$$
$$p_b = v_{L1} + v_{H2} - v_{H1} \tag{38}$$
$$p_n = v_{L1}. \tag{39}$$

In the second visit the high type pays v_{L1} plus a premium equal to the incremental value of the enhanced service. This strategy is commonly implemented by special offers to "new customers only," with repeat customers paying the "standard" price.

But another way to implement the same present value is to charge

$$p_0 = v_{L1} \tag{40}$$
$$p_b = v_{H2} \tag{41}$$
$$p_n = v_{L1}. \tag{42}$$

Here everyone is charged a low price on first visit and a high price on the second visit. Essentially, the seller is collecting information on the first visit which is then used to provide the enhanced service that only the high-value people are willing to pay for on the second visit.

For example, an online merchant learns billing information and shipping address on the first visit. On the second visit, the merchant can offer, for example, "one-click shopping," a service that frequent purchasers, or those with high time value, might find particularly valuable.[2]

Which of these two pricing patterns might we expect? The answer depends on the nature of the technology at the buyer's disposal. If the only way that the high-value consumer can imitate the low-value consumer is to delay purchase when faced with a high-price during the first visit to a store, then these two price profiles are equivalent.

But if the high-value consumer can delete his cookies, and return to the seller appearing to be a consumer who never bought before, the profile that involves charging $p_0 = v_{L1}$ cannot be an equilibrium. For if this profile were offered, the high-value consumer would buy on its first visit (taking the "low introductory offer for new consumers"), delete his cookie, and then return to buy again at the same price. True, he would not get the enhanced service, but his payoff would be $2v_{H1} - 2v_{L1}$, which is larger than $v_{H1} - v_{L1}$, the payoff from pricing plan (35-36).

Hence the only equilibrium price plan when "delete cookies" (or, more generally, anonymous shopping) is possible is to charge the high price first.

6 No commitment

What happens when the seller cannot commit to its second period behavior?

Let us return to the baseline model described in Section 1. If the most profitable strategy is to sell at the low price, the inability to commit doesn't affect the outcome.

However, when the most profitable strategy is to charge the high price to all, it can be shown that inability to commit induces the high-value consumer to pursue a mixed strategy. At least some fraction of the time the high-value consumer will emulate the behavior of the low-value type, hoping that the seller will then cut the price in the second period (see Taylor [2002] and

[2]Below we describe some empirical results from Goolsbee and Chevalier [2002] that show that Amazon customers are much less price sensitive that Barnes and Noble customers, possibly because of the enhanced services that Amazon offers. See also Economist [2001].

Fudenberg and Tirole [1991], pp. 402-405). In equilibrium, the seller will charge the same prices as in the case of full commitment, but will make less profit due to the randomized strategy of the high-value type.[3]

Turning now to the case of enhanced services, we ask: "Can price conditioning be an equilibrium when sellers cannot commit to future pricing?" The answer is "yes," but there is a subtlety. When commitment is not possible, we have to worry about the sequencing of price offers.

Any first period price in which the high-value and low-value types behave differently will allow the seller to enforce a separating equilibrium second period. Hence the equilibrium in which the seller conditions must be interpreted as one in which the seller offers to the same price to everyone first period, v_{L1}, and offers a price of v_{H2} second period.

However, we saw in section 5 that offering the low price in the first period is an equilibrium only when the high-value buyer cannot delete cookies; that is, the only way the high-value consumer has to imitate the low-value consumer is to delay purchase.

In practice, this is a case in which a seller is able to offer a coupon to new users only, in the hope of converting them into second-period customers. Of course, if anyone can pretend to be a new customer, this strategy is not effective in enforcing price discrimination. (We examine a competitive equilibrium of this sort in a later section.) If the high-value consumer can "delete" rather than just "delay," being unable to commit imposes a cost on the seller, in that it will not be able to implement a price conditioning solution.

In addition, if the value of the enhanced service is such that price discrimination is not optimal, and flat-pricing at the high price is better than flat pricing at the low price, the lack of a commitment device might force the seller to adopt a mixed strategy in the first period just as in the baseline case.

Which is the more realistic model, commitment or no commitment? The answer will depend on the circumstances. One way for the seller to commit to flat pricing is to publicly post prices. This is common in both supermarkets and online shopping, where most price discounting takes place via coupons of one sort or another. Posted prices, with coupons, gives the seller the best

[3]Because profit is reduced due to randomization, there may be conditioning solutions that yield more profit than the flat price outcome when the discount rate is greater than one. We thank Curtis Taylor for this observation.

of both worlds, allowing it to "commit" to prices, yet still offer discounts.

This analysis thus far depends the fact that the model terminates after two periods. It would be desirable to examine a no-commitment model with several periods, but such an extension brings up several additional issues outside the scope of this paper.[4]

7 Competition and conditioning

Up until now we have been considering a monopoly seller. In this section we examine what happens when identical sellers compete for customers. We assume that these sellers cannot commit to future prices, and cannot tell whether customers have bought before from other firms.

As before, we assume each seller sets prices of p_0 if the customer has no cookie, p_b if a cookie shows customer bought at p_0, and p_n if the cookie shows that the consumer did not buy at p_0.

We also assume that the good can be provided at constant marginal cost of $c \geq 0$. To avoid trivial cases we also assume $v_{H1} > c$. Since we normalize the population size to 1, the cost of selling one unit to a fraction π of the population is πc. We also assume that the enhanced service can be provided at zero marginal cost; this makes no difference as long as the consumers' valuations of the enhanced service exceed its marginal cost.

There are several conceivable equilibria. Consumers could make their first purchase from a firm and then stay with it in order to receive enhanced services on the next purchase. Some consumers could switch to a competitor or delete their cookies in order to receive the "introductory" price of p_0. Or, possibly, everyone could switch sellers every period.

We use the notation introduced earlier for the incremental value of the enhanced service:

$$\begin{aligned} e_H &= v_{H2} - v_{H1} \\ e_L &= v_{L2} - v_{L1}. \end{aligned}$$

We will spell out the analysis for the case where the Spence-Mirrlees condition holds,

$$e_L < e_H,$$

[4]See Villas-Boas [1999], Villas-Boas [2001], Fudenberg and Tirole [1998] and Fudenberg and Tirole [2000] for models of this type.

and simply state the results for the reverse inequality, since the analysis is completely parallel.

There are three equilibrium conditions that must be satisfied:

1. Consumers must make optimal choices, which will impose a set of inequalities.
2. Profits are driven to zero, which is an equality.
3. Firms are profit maximizing, which requires comparing their price choices to alternative choices they might make.

Case 1. We first show that it is not an equilibrium for all firms to charge a flat price p at which all consumers purchase. The zero profit condition requires $p = c$. Consider a single firm that raises its price by any amount less than $\min\{e_H, e_L\}$ and provides the (free) enhanced service. This will be an attractive option for some or all consumers, thereby increasing profit, showing that charging flat prices is not an equilibrium.

Case 2. All customers shop at the same store twice rather than switch. Consumer optimization requires

$$v_{H2} - p_b \geq v_{H1} - p_0 \tag{43}$$
$$v_{L2} - p_b \geq v_{L1} - p_0, \tag{44}$$

or

$$p_b \leq p_0 + e_H \tag{45}$$
$$p_b \leq p_0 + e_L. \tag{46}$$

Profit maximization will drive p_b to satisfy

$$p_b = p_0 + e_L,$$

and competition ensures profits are driven to zero, which means

$$p_0 + p_b = 2c.$$

Solving these two equations in two unknowns we have

$$p_0 = c - \frac{e_L}{2} \tag{47}$$
$$p_b = c + \frac{e_L}{2}. \tag{48}$$

In order to show that this is an equilibrium, we need to show that no single firm can increase its profit by changing its behavior.

Clearly no firm will want to lower its price. Will a single firm want to raise its price? By raising p_b to $p_0 + e_H$ the deviating firm will induce its low-value customers to switch to the competition, or delete their cookies, in order to purchase at price p_0. On the other hand the high-value customers will choose to pay the higher price. The profit from this pricing deviation will be less than the profit from the presumed equilibrium when

$$p_0 + \pi(p_0 + e_H) + (1 - \pi)p_0 < p_0 + p_b = 2p_0 + e_L,$$

which reduces to

$$\pi e_H < e_L.$$

Note that this is a "lock-in" equilibrium: consumers face a cost of switching in the second period, because they would lose the enhanced service. While in Villas-Boas [1999] firms in a duopoly lower prices to attract the competitor's previous customers, here firms in a perfectly competitive market respond by charging low prices in the first period, then high prices in the second period, as in typical lock-in equilibria behavior described in Klemperer [1989, 1995]. In our case, the switching costs the consumers face are due to the fact that their current seller provides them with personalized enhanced services and switching to a new seller would eliminate these services.

Case 3. The low-value type switches to another seller or deletes its cookie, the high-value type remains.

This requires

$$p_b \leq p_0 + e_H \tag{49}$$
$$p_b \geq p_0 + e_L. \tag{50}$$
$$\tag{51}$$

Profit maximization now implies $p_b = p_0 + e_H$. Profits come from everyone buying at p_0 during the first visit, and high-value types buying at p_b and low-value types buying at p_0 during their second visit. Competition ensures that profits are driven to zero, implying

$$p_0 + \pi(p_0 + e_H) + (1 - \pi)p_0 = 2c,$$

or

$$2p_0 + \pi e_H = 2c.$$

Solving for equilibrium we have

$$\begin{aligned} p_0 &= c - \frac{\pi e_H}{2} && (52) \\ p_b &= c + \frac{(2-\pi)e_H}{2}. && (53) \\ & && (54) \end{aligned}$$

For this to be an equilibrium no single firm can deviate from these prices and make a profit. If a single firm lowers p_b to $p_0 + e_L$, it will keep its low-value customers but make less revenue on the high-value customers. This will not be profitable when

$$p_0 + p_b = p_0 + (p_0 + e_L) < p_0 + (p_0 + \pi e_H),$$

which is to say when

$$e_L < \pi e_H.$$

This is a "partial lock-in" equilibrium, as only the low-value types find it attractive to switch. The firms find it more profitable to let them go than to keep them, since keeping them would require cutting the price to the high-value types.

Case 4. The high-value type switches sellers, and the low-value type remains. This requires

$$\begin{aligned} p_b &\geq p_0 + e_H && (55) \\ p_b &\leq p_0 + e_L, && (56) \\ & && (57) \end{aligned}$$

which implies

$$p_0 + e_L \geq p_b \geq p_0 + e_H$$

Hence this case cannot occur when $e_H > e_L$.

Here is a summary of the results.

Fact 3 (Equilibria with competition.) *With identical competing firms we have*

- *It is never an equilibrium for all firms to charge a flat price.*
- *If $e_L < e_H$, then in equilibrium*

  - *No consumers will switch when* $e_L > \pi e_H$.
  - *Low-value consumers will switch when* $e_L < \pi e_H$.
  - *High-value consumers will never switch.*

- *If* $e_L > e_H$*, then in equilibrium*
  - *No consumers will switch when* $e_H > (1 - \pi)e_L$.
  - *High-value consumers will switch when* $e_H < (1 - \pi)e_L$.
  - *Low-value consumers will never switch.*

These outcomes exhibit a form of "customer poaching," a term introduced by Fudenberg and Tirole [2000]. They analyze a duopoly in which some consumers remain loyal and others defect to the competitor. In their model switching costs are zero, firms offer partial substitutes a la Hotelling, services are not personalized, and firms can tell which firms consumers bought from previously. Their baseline case is long distance telecommunications service, which is quite different from our situation, due to the undifferentiated nature of the good being sold.

In our situation, the seller is able to provide a *personalized* service that is valuable to at least some of the consumers. This creates switching costs for the consumers, since they would then have to rebuild the relationship with the seller.

In some cases these switching costs may be relatively small—e.g., entering credit card information—but even relatively small switching costs can matter. Goolsbee and Chevalier [2002] estimate demand elasticities facing Amazon and barnesandnoble.com. They find that the demand curve facing Amazon is much more inelastic than that facing barnesandnoble.com, an effect that may be due to the more personalized environment offered by Amazon. If this hypothesis is correct, it may be that Amazon's investment in "enhanced services" may be a significant contribution to its competitive advantage. Note that offering such services often requires large expenditures in fixed cost to implement the system, but very small marginal costs to service each customer, consistent with the zero marginal cost assumption in this paper.

8 Summary

The reduction of the cost of storing and using information has made it possible for firms to adapt their offers, content, or formats for each particular individual. In particular, thanks to information technology sellers can use past purchase data to make current prices dependent on previous behavior. In this paper we have investigated various conditions under which price conditioning will be optimal for a seller when the history of interactions with individual consumers can be tracked.

In our baseline case, with constant values for the product being sold, the seller will not find it profitable to condition prices on past customer behavior. If enough customers are myopic, or the costs of using anonymous technologies are too high, the seller will want to condition pricing on purchase history. In addition, if the seller can offer enhanced services to returning customers due to the information it has received from them during previous transactions, it will often be profitable to condition. We have extended this analysis to consider scenarios where the seller cannot commit to prices and where the seller faces competition. In these cases there may be mixed strategy equilibria and lock-in equilibria.

References

Joseph P. Bailey. Internet price discrimination: Self-regulation, public policy, and global electronic commerce. Technical report, The Robert H. Smith School of Business, University of Maryland, 1998.

Robert C. Blattberg and John Deighton. Interactive marketing: Exploiting the age of addressability. *Sloan Management Review*, 33(1):5–14, 1991.

Giacomo Calzolari and Alessandro Pavan. Optimal design of privacy policies. Technical report, Gremaq, University of Toulouse, 2001.

Yuxin Chen and Ganesh Iyer. Consumer addressability and customized pricing. *Marketing Science*, 21(2):197–208, 2002.

Yuxin Chen, Chakravarthi Narasimhan, and Z. J. Zhang. Individual marketing with imperfect targetability. *Marketing Science*, forthcoming, 2003.

Yuxin Chen and Z. John Zhang. Competitive targeted pricing with strategic customers. Technical report, Leonard N. Stern School of Business, New York University, and Columbia Business School, 2001.

Sandeep Dayal, Helene Landesberg, and Michael Zeisser. Building trust online. *McKinsey Quarterly*, 4, 2001.

The Economist. Amazon, the software company. *Tuesday, December 18th*, 2001.

Drew Fudenberg and Jean Tirole. *Game Theory.* MIT Press, Cambridge, MA, 1991.

Drew Fudenberg and Jean Tirole. Upgrades, trade-ins, and buy-backs. *Rand Journal of Economics*, 29:238–258, 1998.

Drew Fudenberg and Jean Tirole. Customer poaching and brand switching. *Rand Journal of Economics*, 31:634–657, 2000.

Austan Goolsbee and Judith Chevalier. Measuring prices and price competition online: Amazon and Barnes and Nobel. Technical report, University of Chicago, 2002. `http://gsbadg.uchicago.edu/vitae.htm`.

Oliver D. Hart and Jean Tirole. Contract renegotiation and coasian dynamics. *Review of Economic Studies*, 55:509–540, 1988.

Paul Klemperer. Price wars caused by switching costs. *Review of Economic Studies*, 56(3):405–420, 1989.

Paul Klemperer. Competition when consumers have switching costs:an overview with applications to industrial organization, macroeconomics and international trade. *Review of Economic Studies*, 62:515–539, 1995. `http://www.paulklemperer.org/index.htm`.

Eric Maskin and John Riley. Monopoly with incomplete information. *Rand Journal of Economics*, 40:171–196, 1984.

R. McCulloch, Peter E. Rossi, and G. M. Allenby. The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340, 1996.

Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic Theory*, 43:301–317, 1978.

John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *Quarterly Journal of Economics*, 98(2):267–289, 1983.

Peter E. Rossi and G. M. Allenby. Marketing models of consumer heterogeneity. *Journal of Econometrics*, 89:57–78, 1999.

Stephen W. Salant. When is inducing self-selection suboptimal for a monopolist? *Quarterly Journal of Economics*, 104(2):391–397, 1989.

John Schartz. Giving the Web a memory cost its users privacy. *New York Times*, September 4 2001. `http://www.nytimes.com/2001/09/04/technology/04COOK.html`.

Nancy Stokey. Intertemporal price discrimination. *Quarterly Journal of Economics*, 93:355–71, 1979.

David Streifield. On the web price tags blur: What you pay could depend on who you are. *The Washington Post*, 2001.

Curtis R. Taylor. Private demands and demands for privacy: Dynamic pricing and the market for customer information. Technical report, Department of Economics, Duke University, 2002.

Jean Tirole. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA, 1988.

David Ulph and Nir Vulkan. Electronic commerce and competitive first-degree price discrimination. Technical report, University College, London, 2000. `http://www.ecn.bris.ac.uk/www/ecnv/welcome.htm`.

David Ulph and Nir Vulkan. E-commerce, mass customization and price discrimination. Technical report, University College, London, 2001. `http://www.ecn.bris.ac.uk/www/ecnv/welcome.htm`.

Hal R. Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–875, 1985.

Hal R. Varian. *Microeconomic Analysis*. W. W. Norton, New York, 3 edition, 1992.

J. Miguel Villas-Boas. Dynamic competition with customer recognition. *RAND Journal of Economics*, 30(4):604–631, 1999.

J. Miguel Villas-Boas. Price cycles in markets with customer recognition. Technical report, Hass School of Business, UC Berkeley, 2001.