

Recent Developments in the Economics of Price Discrimination*

Mark Armstrong
Department of Economics
University College London

October 2005

Abstract

This paper selectively surveys the recent literature on price discrimination. The focus is on three aspects of pricing decisions: the information about customers available to firms; the instruments firms can use in the design of their tariffs; and the ability of firms to commit to their pricing plans. Developments in marketing technology mean that firms often have access to more information about individual customers than was previously the case. The use of this information might be restricted by public policy towards customer privacy. Where it is not restricted, firms may be unable to commit to the use they make of the information. With monopoly supply, an increased ability to engage in price discrimination will boost profit unless the firm cannot commit to its pricing policy. With competition, the effects of price discrimination on profit, consumer surplus and overall welfare depend on the kinds of information and/or instruments available to firms. The paper investigates the circumstances in which price discrimination causes all prices (and hence profit) to fall.

*Paper presented at the 2005 World Congress of the *Econometric Society*. Much of this paper reflects joint work with John Vickers. In preparing this paper I have greatly benefited from consulting the earlier and more comprehensive survey by Lars Stole (2003), and the reader is referred to that survey for a more complete account of the important contributions to this topic. Thanks are also due to Bhaskar, Yongmin Chen, Drew Fudenberg, Ken Hendricks, Marco Ottaviani, Barry Nalebuff, Pierre Regibeau, Patrick Rey, John Thanassoulis, Nir Vulkan, and Mike Whinston.

1 Introduction

This paper surveys, in a highly selective manner, the recent progress which has been made in the economic understanding of price discrimination. One can say that price discrimination exists when two “similar” products with the same marginal cost are sold by a firm at different prices.¹ There are many forms of price discrimination, including: charging different consumers different prices for the same good (third-degree price discrimination); making the marginal price depend on the number of units purchased (nonlinear pricing); making the marginal price depend on whether other products are also purchased from the same firm (bundling); making the price depend on whether this is the first time a consumer has purchased from the firm (introductory offers; customer “poaching”).

In broad terms, this paper is about what happens to profit and consumer surplus when firms use more ornate tariffs to sell their products. There are two reasons why a firm might be able to tune its tariff more finely: it might obtain higher quality *information* about its potential customers, or it might be able to use additional *instruments* in its tariff design.

A firm can become better informed about its potential customers if it purchases customer data from a marketing company or from another firm. It can use this data to send personalized price offers to new customers (an example of third-degree price discrimination).² Alternatively, a firm might keep records of its customers’ past purchases, and use this information to update its future price offers to those customers. Firms’ access to better information is affected by public policy towards consumer privacy (for instance, whether firms are permitted to pass information about their customers to other firms). It is also affected by a consumer’s ability to become “anonymous” when dealing with firms, and to pretend to be a new customer.

Examples of the use of more instruments include: using two-part tariffs instead of linear prices; charging different identifiable consumer groups different prices instead of a common price; offering a discount if two products are jointly purchased; or making the price for an item depend on whether a customer has previously purchased similar items from the firm.³ Public policy towards price discrimination affects the range of instruments which firms can

¹Stigler (1987) suggests a definition that applies to a wider class of cases: discrimination exists when two similar products are sold at prices that are in different ratios to their marginal costs. (This definition makes more sense when discussing “versioning”, where slightly different versions of a product—such as hardback and paperback books—are offered for sale at very different prices.) Which of these definitions we use makes no difference for the purposes of this paper. An alternative definition might be that price discrimination is present when a similar product is sold to different consumers at different prices. However, this definition rules out cases of “intra-personal” discrimination which are sometimes relevant.

²See Taylor (2004, section 1) for a summary of the market for customer information. For instance, he reports that a good customer mailing list can sell for millions of dollars on its own.

³*Pure* bundling—where two products are made available only as a joint purchase—is not a more ornate tariff compared to separable prices, but rather just a different kind of tariff. However, mixed bundling, where individual products as well as the bundle are offered for sale, is a more ornate tariff compared to either pure bundling or separable prices.

use. Firms are also constrained in the range of instruments they can use by the possibility of arbitrage and resale between consumers.⁴

With monopoly supply, except for issues to do with commitment problems (see sections 2.2 and 2.3 below), the use of more ornate tariffs must lead to higher profits. When the firm has access to more detailed information about its customers or can use a wider range of instruments in its tariff, it can do no worse than before and generally it can do better. With competition, though, the effects of using more ornate tariffs are less clear cut. In particular, in section 3 a Hotelling example is used to argue that the impact of more information on profits and prices depends crucially on the *kind* of information which is available. Some information will cause firms to make higher profits in equilibrium, whereas other kinds of information will cause all prices to fall compared to the situation with uniform pricing.

Similarly, the availability of an additional tariff instrument has ambiguous effects on profit and consumer surplus (see section 4). Competing multiproduct firms often make less profit when they practice mixed bundling than when they sell their products separately. In contrast, when consumers buy all relevant products from one firm (i.e., in a one-stop shopping framework), when consumers are homogenous the effect of more ornate tariffs in competitive markets is to boost profit and to harm consumers. Price discrimination also allows a firm to target price reductions more accurately at market segments where competition is most intense. Doing so can harm rivals and deter entry, and as such can harm consumers compared to the case where discrimination is banned.

A third theme of the paper, in addition to the effects of more information and instruments, is how the ability of firms to commit to their pricing plans affects industry outcomes. In a dynamic context, it is well known that a monopolist's ability to commit to future prices affects its profit. Advances in marketing may mean that the commitment problem has become more severe. The finely-tuned customer data which firms often possess permits the use of personalized prices. Such prices are often "secret" rather than public and, as such, it is unlikely that firms will be able to commit to such prices. In both monopoly and oligopoly settings, we will see that when firms can commit to price plans (or when they cannot engage in behaviour-based price discrimination, which often amounts to the same thing), this will increase industry profit and harm consumers.

The welfare effects of allowing price discrimination (or more ornate tariffs) is ambiguous, both with monopoly and with oligopoly supply. There is no justification for public policies that prohibit price discrimination in general. Price discrimination can lead to efficient pricing, for instance (see section 2.1). Price discrimination can lead to more intense competition which benefits consumers (see sections 3.3 and 4.2). When firms have difficulty committing to prices, they often are forced to charge low prices. In such situations, a policy which forbids discrimination can act to restore commitment power, to the detriment of consumers (sections 2.2 and 2.3). In dynamic settings, the effect of price discrimination on consumer

⁴European competition law forbids dominant firms from preventing arbitrage (or "parallel imports") between countries. See Motta (2004), for instance.

welfare may depend upon whether consumers are sophisticated or naive (see section 2.2). When firms offer different prices to their loyal customers and to their rival's previous customers this can make competition more intense, but it can also induce excessive switching between firms (section 5). By contrast, multiproduct firms might induce excessive loyalty, or one-stop shopping, by means of bundling discounts (section 4.2). Discrimination can also lead firms to leave consumers with less surplus than they would enjoy in its absence (sections 3.2 and 4.1). In addition, the freedom of an incumbent firm to engage in price discrimination will typically have an discouraging effect on entry (section 4.3). Sensible policy towards price discrimination needs to be founded on good economic understanding of the market in question.

2 Monopoly supply

2.1 Information and Instruments

In many cases, the welfare losses caused by firms exploiting market power are caused by the firms having insufficient information about their consumers' preferences, or being constrained in their ability to condition prices on their information about consumers. In some circumstances, allowing firms to engage in price discrimination can implement efficient prices. In these cases, welfare is unambiguously improved.

One familiar example of this is first-degree discrimination, where a monopolist has perfect information about each consumer's valuation for its products and has the ability to set personalized prices. To be concrete, suppose there is a population of consumers, each of whom wishes to consume a unit of the firm's product. A consumer's valuation for this unit is denoted v and this varies among consumers according to the distribution function $F(v)$. Suppose the firm has unit cost c . If price discrimination is not possible (e.g., because the firm does not have the necessary information, cannot prevent arbitrage between consumers, or is not permitted to engage in discrimination), the firm will choose a uniform price p to maximize profit $(p - c)(1 - F(p))$. Clearly, this uniform price will be above cost, and total surplus is not maximized. (It is efficient to serve those consumers with $v \geq c$, but only those with $v \geq p > c$ are served.) If the firm can observe each consumer's v and is permitted to discriminate on that basis, it will charge the type- v consumer $p = v$, provided this price covers its cost of supply. In other words, an efficient outcome is achieved. However, the firm appropriates the entire gains from trade and consumers are left with nothing.

There are situations when price discrimination can lead to *approximately* efficient prices, even when the monopolist has relatively little information about consumer tastes.⁵ A supermarket, say, supplies a large number of products. Consumers have a wide variety of preferences over these products—some people prefer tea to coffee, and so on. Suppose that the supermarket sets the price of each item equal to its marginal cost of supply for that

⁵See Armstrong (1999) and Bakos and Brynjolfsson (1999) for this analysis.

item and also sets an entry fee (or annual fee) to get into the shop. In this case, provided they come in at all, consumers will buy a product whenever it is efficient for them to do so. Suppose that consumer valuations for the various products are independently distributed product by product (so the fact a consumer likes coffee, say, gives no guidance about whether she also likes potatoes). For a given list of prices, the “law of large numbers” implies that each consumer’s total surplus is approximately the same, even though consumers differ significantly in their individual purchases. Therefore, the firm can extract almost all this total surplus without excluding many consumers. In these circumstances the firm can obtain approximately the first-best profit level by setting its marginal prices equal to marginal costs and extracting the resulting consumer surplus by means of a lump-sum fee.⁶ The key insight is that a multiproduct firm can better predict a consumer’s total surplus than it can predict a consumer’s surplus derived from any individual product. As in the previous example of first-degree price discrimination, the efficient outcome is approximately achieved, and in the case of monopoly the firm extracts almost all the gains from trade.

Monopoly first-degree price discrimination is merely an extreme form of a fairly common situation. Lack of information about consumer tastes (or an inability to set suitably finely-tuned prices) in combination with market power leads to welfare losses as the firm faces a trade-off between volume of demand and the profit it makes from each consumer. In many cases, if the firm obtains more detailed information about its consumers (or if it is permitted to price discriminate when it was not previously) this will enable the firm to extract consumer surplus more efficiently, and this will often lead to greater overall welfare. However, it is consumers’ private information that protects them against giving up their surplus to a monopoly. Therefore, there will often be a reduction in consumer surplus when the monopoly firm obtains better information about its consumers.

Finally, in the context of *regulated* monopoly, socially optimal prices almost always exhibit price discrimination.⁷ For instance, Ramsey prices—the set of linear prices which maximize total welfare subject to the regulated firm covering its costs—depend on demand conditions in much the same way as an unregulated monopolist’s prices do. For instance, when the firm serves a number of independent markets, each with the same marginal cost of supply, Ramsey principles suggest that the most efficient way to cover the firm’s costs is to set a higher price in those markets where consumer demand is less elastic, exactly as would be the case with an unregulated profit-maximizing firm. Permitting price discrimination will surely improve welfare in such situations. The same point applies to the use of other tariff instruments, such as two-part tariffs and bundling.

⁶As emphasized by Bakos and Brynjolfsson (1999), if marginal costs are zero (as with electronic distribution of software and other information), this profit-maximizing outcome is implemented by pure bundling.

⁷Edgeworth (1910, page 462) writes “a regulated discrimination of prices, such as might conceivably be practised by a Socialist Directory, but is not possible in a *regime* of competition, tends to increase the sum-total of utility.”

2.2 Dynamic pricing

A topic which has received much recent attention is dynamic price discrimination.⁸ There are many aspects to this phenomenon. A publisher sets a high price for a new (hardback) book, then subsequently the price is reduced. Or a retailer might use information it has obtained from its previous dealings with a customer to offer that customer a special deal (or, as we will see, sometimes a bad deal). This latter form of discrimination, sometimes termed “behaviour-based price discrimination”, could be highly complex. If a supermarket has sufficient information, it could offer those customers who have purchased, say, nappies, a voucher offering discounts to a particular brand of baby food. Or if a customer regularly spends £80 per shopping trip, the supermarket might send the customer a discount voucher if he spends more than £100 next time. Or if the consumer appears to have starting shopping elsewhere recently, the supermarket will send a generous discount voucher to attempt to regain that consumer.⁹

In this section and in section 5 I will focus on simple forms of dynamic price discrimination. Specially, I will assume that consumers have unit demand for a single product per period, and that there is no scope for firms to tailor their products to what they judge to be a consumer’s particular tastes.¹⁰ I will focus on the case where firms are able, where policy permits, to make their price depend on whether or not the consumer has already purchased from the firm.

The reader might ask: what is the difference between dynamic discrimination and static multi-product discrimination such as mixed bundling (which is discussed in section 4.2)? There are two chief differences. First, unlike the static case, consumers might not know their preferences for future consumption at the time of their initial dealings with a firm. Second, firms might not be able to commit to their future price policy at the time of their initial dealings with consumers. It is perhaps especially plausible that firms cannot commit when they offer personalized discounts (such as the supermarket examples just mentioned). In the following discussion I will focus on this second aspect of the dynamic interaction.

In more detail, suppose a firm can sell its product over two periods. Suppose there is a diverse population of consumers, each of whom potentially wishes to buy a single unit of the product in each of the two periods. A consumer’s valuation of the unit, v , is uniformly distributed on $[0, 1]$, and this valuation is the same in the two periods. Suppose the firm has

⁸See Fudenberg and Villas-Boas (2005) for a more detailed survey of this material, covering both monopoly and oligopoly supply.

⁹Casinos also use dynamic pricing techniques. For instance, they monitor activity on their machines and might offer a someone a cash voucher if he has been “unlucky”, or they might set high room rates to customers who have historically been unprofitable gamblers. See Christina Binkley, “Harrah’s uses science to lure gamblers”, *Wall Street Journal*, 28 November 2004.

¹⁰For instance, someone might return to the same hairdresser since that hairdresser knows how best to cut his/her hair. Or *Amazon* suggests new books that you might like, given your past purchases. See Acquisti and Varian (2004) for a fuller account of the various ways in which a monopoly might use a customer’s history to tailor its current deal.

zero cost of production. Suppose the firm and consumers share the discount factor $\delta \leq 1$. The firm chooses three prices: p_1 , the price for a unit in the first period; p_2 , the price in the second period if the consumer did not purchase in the first period; and $\hat{p}_2$, the price for a unit in the second period if the consumer also purchased in the first period. In this framework there are two forms of price discrimination possible: (i) the firm can base its second-period price on whether the consumer purchased in the initial period (i.e., $p_2 \neq \hat{p}_2$), or (ii) the firm sets the same price the second period, but this common price is different from the initial price ($p_2 \neq p_1$). Case (i) might be termed “behaviour-based” discrimination, while case (ii) could be termed “inter-temporal” discrimination.

Three distinct settings are discussed in this section: the case where consumers are sophisticated and the firm can commit to its pricing plans; the case where consumers are “naive” and do not foresee the firm may react strategically to their initial choices; and the case where consumers are sophisticated but the firm cannot commit to its pricing plans.

Sophisticated consumers and commitment

Suppose the firm announces its three prices $\{p_1, p_2, \hat{p}_2\}$ at the start of period 1 and the firm cannot alter these prices in the light of a consumer’s first-period decision. In this case, the profit-maximizing policy is to reproduce the optimal static policy over the two periods, so that¹¹

$$p_1 = p_2 = \hat{p}_2 = \frac{1}{2}. \quad (1)$$

When the firm can commit to its second-period prices, it is optimal not to price discriminate (in either sense).¹² Clearly, this is also the outcome when the firm is unable to practice price discrimination.

Naive consumers

Suppose next the consumers do not take into account the effect that their initial purchasing decision has on the second-period prices they will face.¹³ This framework might be relevant for a new market, for instance, where consumers have not yet learned the firm’s

¹¹See Hart and Tirole (1988) and Acquisti and Varian (2004). More generally, it is a standard result in principal-agent theory that when the agent’s private information does not change over time, the optimal dynamic incentive scheme simply repeats the optimal static incentive scheme. (For instance, see section 8.1 of Laffont and Martimort (2002) and the references therein.)

¹²If the model is modified so that either consumers have different discount factors or consumers as a whole have a different discount factor to the firm, then it can be optimal to engage in price discrimination. In addition, Crémer (1984) presents a model with commitment where the price declines over time. He supposes that consumers, at the time of their first purchase, are uncertain about their valuation for the product. In this setting it is optimal to set the second-period price equal to marginal cost, since that maximizes a consumer’s “option value” from future consumption. All the monopolist’s profits are therefore extracted in the first period.

¹³Alternatively, we could think of consumers as simply being myopic. However, the stated “behavioural” assumption seems more plausible.

pricing strategy. If the firm chooses the first-period price p_1 , naive consumers will purchase in the first period whenever $v \geq p_1$, in which case the firm makes profit in the first period equal to $p_1(1 - p_1)$. In the second period, the firm knows whether a consumer's valuation lies in the interval $[0, p_1)$, if the consumer did not purchase in period 1, or $[p_1, 1]$, if the consumer did purchase in period 1, and it will choose its second-period prices accordingly. In the lower-value segment it will set the price $p_2 = \frac{1}{2}p_1$, and in the high-value segment it sets the higher price $\hat{p}_2 = \max\{\frac{1}{2}, p_1\}$. In particular, the firm views its previous customers as its strong market in the second period, while new customers constitute the weak market.¹⁴ (This will continue to hold in oligopolistic settings in section 5.)

The static profit-maximizing price in the first period is $p_1 = \frac{1}{2}$. However, the firm will wish to raise its price above this level, since this strategy renders its customer information in the second period more valuable.¹⁵ If it chooses the initial price $p_1 \geq \frac{1}{2}$, its discounted profit is

$$p_1(1 - p_1) + \delta \left[p_1(1 - p_1) + \frac{1}{4}p_1^2 \right],$$

and so the profit-maximizing initial and subsequent prices are

$$p_1 = \hat{p}_2 = \frac{1 + \delta}{2 + \frac{3}{2}\delta} > \frac{1}{2}; \quad p_2 = \frac{1}{2}p_1. \quad (2)$$

Thus, consumers face an unchanged price in the second period if they purchase in the first period, while the offered price is halved if they did not initially purchase. Here, discounted consumer surplus is lower compared to when the firm cannot practice price discrimination.¹⁶

Sophisticated consumers and no commitment

When consumers are sophisticated, the price plan in (1) is not feasible when the firm cannot commit to its future prices. Once those consumers with $v \geq \frac{1}{2}$ have purchased in the first period, the firm is left with an unidentifiable pool of low-value consumers. Given this, in the second period the profit-maximizing policy is to set $p_2 = \frac{1}{4}$ to those consumers who have not already purchased. Sophisticated consumers foresee the firm will behave in this opportunistic manner, and some consumers (with v slightly greater than $\frac{1}{2}$) will strategically delay their purchase to receive the discounted price in the next period.

We calculate the non-commitment policy as follows. Suppose that in the first period the firm chooses price p_1 . What must the time-consistent second-period prices be? Given p_1 , suppose those consumers with value $v \geq v^*$ choose to buy in the first period, where the

¹⁴In the price discrimination literature, a market is said to be “strong” (“weak”) if a firm wishes to raise (lower) its price there compared to the situation where it must charge a uniform price in the two markets.

¹⁵The period-1 price which generates the highest period-2 profit is $p_1 = \frac{2}{3}$.

¹⁶Discounted consumer surplus with the prices in (2) is equal to $\frac{1}{2}(5\delta + 2\delta^2 + 4)(\delta + 1)/(3\delta + 4)^2$. This is greater than the case without price discrimination, when prices are given by (1) and discounted consumer surplus is $\frac{1}{8}(1 + \delta)$.

threshold v^* is to be determined. The firm will optimally choose the second-period price for consumers who did not purchase in the first period to be

$$p_2 = \frac{1}{2}v^* . \quad (3)$$

The second-period price for those who did buy in the first period depends on v^* : if $v^* < \frac{1}{2}$ then $\hat{p}_2 = \frac{1}{2}$ whereas if $v^* > \frac{1}{2}$ then $\hat{p}_2 = v^*$. In either case, the consumer who is indifferent between buying in the first period and buying only in the second period, v^* , satisfies

$$v^* - p_1 = \delta(v^* - p_2) ,$$

and so from expression (3) $v^* = 2p_1/(2 - \delta)$. If $v^* \geq \frac{1}{2}$ (as will happen in equilibrium) the firm's overall profit is

$$p_1(1 - v^*) + \delta \left\{ v^*(1 - v^*) + \left(\frac{1}{2}v^*\right)^2 \right\} .$$

After substituting $v^* = 2p_1/(2 - \delta)$ and maximizing this with respect to p_1 , we see that the equilibrium prices are

$$p_1 = \frac{4 - \delta^2}{2(4 + \delta)} ; p_2 = \frac{2 + \delta}{2(4 + \delta)} ; \hat{p}_2 = v^* = \frac{2 + \delta}{4 + \delta} . \quad (4)$$

Notice that $p_2 \leq p_1 \leq \frac{1}{2} \leq \hat{p}_2$. Therefore, when the firm cannot commit to its future prices, it will set a relatively low first-period price ($p_1 \leq \frac{1}{2}$), followed by a high second-period price for those consumers who purchased in the first period ($\hat{p}_2 \geq \frac{1}{2}$) and an even lower second-period price aimed at those who did not purchase in the first period ($p_2 \leq p_1$). In this example, the second-period price for those who previously purchased from the firm is exactly twice the price for those consumers who did not already purchase.¹⁷ All consumers are better off when the firm cannot commit, compared to the price plan (1), while the firm is obviously worse off.¹⁸ Total welfare is also higher, despite the restricted consumption in the first period and the high price for repeated sales compared to the commitment regime.

The effect of a ban on price discrimination in this case is to restore the firm's commitment power, to the detriment of all consumers. (By contrast, when the consumers are naive, a ban on price discrimination will make consumers in aggregate better off.) All that is needed to restore commitment power is to forbid behaviour-based discrimination, i.e., to require $p_2 = \hat{p}_2$. When second-period prices are constrained to be equal, a consumer has no incentive to behave strategically in the first period (i.e., he will buy if $v \geq p_1$), and the firm has no incentive to lower the second-period price below the commitment level. If the firm could

¹⁷Of course, this form of discrimination is not feasible if past consumers can pretend to be new customers, for instance by deleting their computer "cookies" or using another credit card when they deal with an online retailer.

¹⁸With commitment prices (1), the total discounted payment for two units is $(1 + \delta)/2$. With the no-commitment prices (4), if a consumer buys two units the total discounted payment is less than $(1 + \delta)/2$, and so all consumers must be better off when the firm cannot commit.

commit not to practice behaviour-based discrimination, for instance by being seen not to invest in consumer tracking technology, its profits would rise.¹⁹

Taylor (2004) analyzes a related model where one firm sells a product in period 1 and a separate firm sells a related product in period 2.²⁰ The first firm is able to sell its information about which consumers purchased from it in the first period to the second firm. The second firm is willing to pay for this information, since it provides the basis for behaviour-based price discrimination towards its consumers. Since the first firm can fully extract the second firm's benefit from the information, the scenario is essentially the same as when an integrated firm supplies in both periods and cannot commit to its second-period price. Taylor also distinguishes between sophisticated and naive consumers. If consumers are naive, in the sense that they do not foresee their decisions with one firm might affect their offers from the subsequent firm, the first firm has an incentive to raise its price above the monopoly level in order to boost the value of information to the second firm (just as in expression (2) above). Public policy towards privacy might prohibit the passing of consumer information between firms, and this would make naive consumers better off and reduce the level of industry profit. On the other hand, when consumers are sophisticated, a ban on information transfer will surely increase the industry profits in this setting.

2.3 The “secret deals” problem

Suppose a vertically separated upstream monopolist sells an essential input, which has marginal cost c , to two competing downstream firms, A and B . For simplicity, suppose that A and B supply an identical product and compete in a Cournot fashion, given the input prices set by the upstream firm.²¹ Suppose that the contract between the monopolist and downstream firm i takes the form of a two-part tariff, with fixed charge f^i and per unit of input price p^i . First, note that by an appropriate choice of a non-discriminatory two-part tariff, the monopolist can ensure (i) that the industry profit is maximized and (ii) that it appropriates the entire industry profit. (This is achieved by setting a high input price p , a

¹⁹Villas-Boas (2004) analyzes a related model with a infinitely-lived firm facing a sequence of (sophisticated) two-period consumers arranged in overlapping generations. The firm can make its prices depend on whether a consumer has previously purchased, but it cannot determine whether a new consumer is “young” or whether she is “old” but chose not to consume in her first period. An interesting result in this richer framework is that there are cycles in the prices offered to new consumers. In addition, as in the two-period framework presented in the text, the firm is better off if it is unable to practice behaviour-based price discrimination.

²⁰In Taylor's model consumer valuations are binary and are not perfectly correlated for the two products. See Calzolari and Pavan (2005) for a related and more general model.

²¹The key papers relevant to this “secret deals” problem are Hart and Tirole (1990), O'Brien and Shaffer (1992) and McAfee and Schwartz (1994). See Rey and Tirole (2005) for a fuller account of this literature. In particular, the results reported are also valid when downstream firms supply differentiated products and compete in prices, and when the monopolist offers arbitrary nonlinear tariffs for its input.

price above cost c , which generates high retail prices downstream, and then extracting the downstream profit via the fixed charge f .) In many relevant cases, however, the fact that the downstream firms are “big” customers implies that these contracts for input prices will often be negotiated bilaterally rather than simply announced by the monopolist. In addition, contracts for inputs are often much more detailed and specific than the simple prices in retail markets. In such cases, it is plausible that each downstream firm’s contract with the monopolist will not be observed by the rival downstream firm. In this situation of “secret deals” the monopolist finds it hard to avoid offering each downstream firm a cost-based two-part tariff, with $p^i = c$. Of course, if both A and B pay for the input at cost, then the industry outcome will correspond to the (moderately competitive) Cournot outcome rather than the industry profit-maximizing outcome.

Why is the monopolist forced to set $p^i = c$ in this framework? Since it is secret, the choice of contract between the monopolist and firm A , say, cannot affect the expected output from firm B (or its input choice). Therefore, the contract with A will maximize the combined profits of A and the monopolist, taking the contract with B as given. But the joint profit-maximizing contract will involve $p^A = c$, since in that way the downstream firm’s incentives are in line with the monopolist’s. A similar contract will be secretly agreed with firm B . The monopolist’s market power is eroded by its inability not to negotiate efficient bilateral deals secretly with the downstream firms.

This effect is closely related to the discussion in section 2.2 about a monopolist selling to final consumers which cannot commit to its future prices. In the dynamic context, the monopolist cannot commit not to offer a good deal to the remaining (low value) consumers in the second period, and this acts to undermine its market power. Just as in that setting, a ban on price discrimination in input prices will act to restore the monopolist’s market power. If the monopolist were not allowed to offer different terms to different downstream firms, then it cannot secretly negotiate efficient bilateral contracts, and it can then implement the monopoly outcome. As in the dynamic pricing problem, then, a policy which prohibits price discrimination will end up raising all prices and being detrimental to (final) consumers.

3 The effects of more information in oligopoly

The discussion to this point has focussed on monopoly, and argued that consumers are protected from having their surplus extracted when (i) they possess private information about their tastes and (ii) the firm is unable to commit to its pricing plans. Competition between suppliers provides a third means by which consumers are protected against surplus extraction. Even if firms know everything about a consumer’s tastes, competition ensures that the consumer will still be left with positive surplus. In competitive environments, whether consumers are better or worse off when firms obtain more detailed information about them is a subtle question, as we will see in section 3. In general, when a fixed set of competing firms know everything about consumer tastes, (one) equilibrium necessarily

involves the efficient outcome.²²

Much of the recent literature has focussed more on price discrimination in competitive settings. Many of the extra effects that can appear with competition can be illustrated using a simple Hotelling duopoly example.²³ Suppose two firms, A and B , compete in prices in a symmetric Hotelling fashion. A consumer wishes to buy a single unit from either firm A or firm B , and if he buys from firm $i = A, B$ his net surplus is

$$u^i = v - p^i - td^i ,$$

where v is the consumer's valuation for the unit (which is the same at either firm), p^i is firm i 's price, d^i is the distance the consumer travels to firm i , and t is the transport cost per unit of distance faced by the consumer. The two firms are situated at each end of the unit interval $[0, 1]$ and consumers are uniformly located along this interval. A consumer located at $x \in [0, 1]$ is a distance $d^A = x$ from firm A and $d^B = 1 - x$ from firm B . A consumer's preferences in this framework are defined by three parameters: v is the consumer's valuation for the product, x represents his relative preference for firm A over firm B , and t represents his "choosiness", i.e., how much he dislikes buying his less preferred brand. (The parameter t might also represent a consumer's "laziness", and so a high t is closely related to having high search costs.)

The parameters (v, x, t) are distributed among the consumers in some way. In the various examples which follow, I will assume for simplicity that all consumers choose to buy from one firm or the other. This assumption largely, but not entirely, eliminates the ability to compare welfare with or without price discrimination, but it does hugely simplify the calculations. In addition, I will assume that the parameters (v, x, t) are independently distributed. (I will in the following sections suppose that each of the three parameters in turn is observable to the firms, and I do not wish to consider whether observing v , say, allows firms to obtain a signal about (x, t) .) Suppose that production costs are normalised to zero. The consumer with preferences (v, x, t) will buy from A rather than B if his surplus u is higher there. Since his valuation of the product v is the same from either firm, he will choose to buy from the firm with the lower total cost of purchase, and he will go to firm A if $p^A + xt \leq p^B + (1 - x)t$.

²²See Spulber (1979) for this analysis. Bernheim and Whinston (1986) provide related analysis in a more general context. Also, notice that not all equilibria need be efficient. Consider two firms which each supply one of two perfectly complementary products. Suppose that consuming one unit of each product yields utility 2 to the consumer, while consuming two units of each product yields utility 3. Suppose that costs of production are zero for both firms. If both firms offer the nonlinear tariff such that $p_1 = 1$ and $p_2 = 3$ (where p_i is the firm's price for buying i units), the consumer will choose to buy one unit of each product, which is inefficient. Nevertheless, neither firm has a profitable deviation. Bhaskar and To (2004) show in a free entry model that, even when firms possess full information about the market and can use the full range of tariff instruments, there is a socially excessive number of firms in the market.

²³A similar exercise was performed in the pioneering paper by Borenstein (1985). He analyzes a free entry model rather than a duopoly, and so also considers the effect of permitting price discrimination on the equilibrium number of firms. He does not consider the effect of discriminating on the basis of brand preference.

3.1 Discriminating on valuation: profit neutrality

Suppose that firms each observe a consumer's valuation v and can target a personalized price to that consumer. Does this information affect prices and profits in equilibrium? Clearly, with monopoly supply the ability to observe valuations would be valuable—see section 2.1—but in this particular competitive environment it actually has no effect. For simplicity, suppose that all consumers have the same choosiness parameter t .²⁴ Given v , if the two firms' prices are p^A and p^B , a type- x consumer will buy from A when $p^A + tx \leq p^B + t(1 - x)$, i.e., when $x \leq \frac{1}{2} - \frac{p^A - p^B}{2t}$. Therefore, firm i 's profit from the type- v segment is

$$\pi^i = \left(\frac{1}{2} - \frac{p^i - p^j}{2t} \right) p^i$$

and it is straightforward to show that the equilibrium price in this segment is $p^A = p^B = t$. Since the equilibrium price does not depend on v , whether or not v is observable has no effect on outcomes. Notice that this is true even in asymmetric cases: if firm B has no information about v while A does, A has no incentive to use its superior information in its pricing decisions. Therefore, a marketing firm with customer data about v would, in this setting, be unable to sell this data to one or both duopolists.

This example extends to situations where consumers buy multiple units and multiple products, and where consumers' preferences over these various units is private information.²⁵ Suppose that a type- θ consumer obtains gross utility $u(\theta, q)$ if she buys the vector of quantities q from a firm, excluding his transport cost and the price he must pay. The net surplus obtained by a type- θ consumer located at x if he purchases quantities q_A from firm A in return for total payment P_A is

$$u(\theta, q_A) - P_A - tx ,$$

while his net surplus if he purchases q_B from firm B in return for payment P_B is

$$u(\theta, q_B) - P_B - t(1 - x) .$$

Implicit in this formulation is the assumption that brand preferences (or transport costs) do not depend on the quantities purchased. Suppose first that firms can observe the type θ of each consumer (but not the location x), and that θ is distributed independently from the brand preference parameter x . Then the most profitable way for a firm to attract a

²⁴The argument works just as well when there is unobserved heterogeneity in t , provided that all consumers participate.

²⁵For this analysis see Armstrong and Vickers (2001, section 4) and Rochet and Stole (2002). This analysis assumes that consumers purchase all products from a single firm and do not “mix-and match” across firms. Rochet and Stole show that this “no discrimination” result is not robust to a number of changes to the model, including: firms having different costs; brand preference x being correlated with the vertical taste parameter θ , or the total market size being affected by the contracts offered.

consumer is to set its marginal prices equal to marginal costs and to extract profit by means of a lump sum charge. If each firm's marginal cost for supplying product i is c_i , each firm will set marginal price $p_i \equiv c_i$ and the fixed charge t . This cost-based two-part tariff does not depend on θ . Therefore, this two-part tariff remains an equilibrium even when firms cannot observe the taste parameter θ .²⁶ If firm B offers the cost-based two-part tariff $T(q) = t + \sum_i c_i q_i$, then the same tariff is firm A 's best response if A does not observe θ (since it is also A 's best response when it can observe θ).²⁷

In sum, firms' ability to observe "vertical" taste parameters (v or θ in the previous discussions) has no effect on outcomes in these models of competitive price discrimination.

3.2 Discriminating on choosiness: best-response symmetry

Now return to the unit demand Hotelling framework and suppose consumers differ in their choosiness parameter t . Suppose that t is distributed on the interval $[t_L, t_H]$ according to some probability distribution, and that location x is independently and uniformly distributed on the unit interval $[0, 1]$. Assume that $t_L > 0$, so that no consumers view the firms' services as perfect substitutes. If firms were able to observe t but not x , the equilibrium price to the type- t consumers is just given by $p(t) = t$, as in section 3.1. This reveals a major difference between discrimination based on choosiness t and based on valuation v : in the latter case firms could not extract anything extra from high-value consumers due to competitive pressure, but when a consumer is known to be choosy firms can extract high profit. Industry profit when firms price discriminate in this way is

$$\Pi_T = E(t) , \tag{5}$$

the expectation of t .

If firms cannot price discriminate, they will set uniform prices p_A and p_B . A type- (x, t) consumer will buy from A at these prices if $x \leq \frac{1}{2} - \frac{p_A - p_B}{2t}$. When the two prices are not too different, the number of such consumers is²⁸

$$\frac{1}{2} - \frac{p_A - p_B}{2\hat{t}} ,$$

where $\hat{t}$ is the harmonic mean of t (so $\hat{t} = (E\{1/t\})^{-1}$). Therefore, the equilibrium non-

²⁶However, in general it may not be the only equilibrium.

²⁷Miravete and Röller (2004) fit a model of duopoly competition in nonlinear tariffs to data from cellular telephone markets. (The model assumes that consumers buy services from both firms, in contrast to the one-stop shopping framework used in the text.) They estimate that if firms were restricted to offer two-part tariffs rather than fully nonlinear tariffs, in equilibrium they would obtain 94% of the equilibrium profits with unrestricted tariffs.

²⁸The condition for this to be the correct formula for A 's demand is that $|p_A - p_B| < t_L$, so that no firm has a monopoly over even the most price-sensitive of consumers.

discriminatory price (and equilibrium industry profit) is²⁹

$$\Pi_{ND} = p_{ND} = \hat{t}. \quad (6)$$

Since the harmonic mean is necessarily lower than the (arithmetic) mean, it follows that industry profits when firms can discriminate according to choosiness (5) are always higher than those generated without such discrimination (6). Since total welfare is not affected by price discrimination in this full-participation framework, aggregate consumer surplus decreases with this form of discrimination, although clearly the more price-sensitive consumers are better-off with discrimination.³⁰

In this example a firm will raise some prices and lower others if it is permitted to engage in price discrimination. That is to say, the non-discriminatory price is an “average” of the discriminatory prices (in this case the harmonic mean). In the remainder of this section we investigate in more detail when this phenomenon occurs.

First, consider the straightforward case of monopoly supply. Suppose a firm serves two markets, 1 and 2, which have independent consumer demands. The firm’s profit in market i when it sets the price p_i in that market is denoted $\pi_i(p_i)$. Then the profit-maximizing discriminatory prices are characterized by $\pi'_i(p_i) = 0$, while the profit-maximizing uniform price $\bar{p}$ satisfies $\pi'_1(\bar{p}) + \pi'_2(\bar{p}) = 0$. Except in the fluke case where there is no gain from discrimination, it follows that in market 1, say, we have $\pi'_1(\bar{p}) > 0$ and in market 2 $\pi'_2(\bar{p}) < 0$. Assuming profit functions are single-peaked it follows that if the firm can price discriminate it will raise its price in market 1 (the strong market) and lower its price in market 2.³¹

²⁹One important issue not discussed here is when a pure strategy equilibrium exists. If some consumers view the two firms as very close substitutes then no pure strategy equilibrium exists, as in Varian (1980). If t_L is close to zero the candidate equilibrium price $\hat{t}$ is also close to zero, and it will be worthwhile for a firm to deviate from the candidate equilibrium, and instead to set a high price which targets the choosier consumers. For instance, take the special case where there are just two values of t , t_L and t_H , which are equally likely. One can show that the candidate equilibrium in expression (6) is an equilibrium when $p_L/p_H > 0.093$.

³⁰This example is closely related Proposition 4 (part (i)) in Armstrong and Vickers (2001), specialized to the case of inelastic demand. Similar effects are seen, in a more extreme way, if the model of “tourists and locals” in Varian (1980) is extended to allow for price discrimination. Local consumers are assumed to know the full range of prices offered by the competing firms and to buy from the lowest price firm. Tourists are assumed to know nothing, and randomly choose a firm. If firms cannot price discriminate between these two groups, Varian shows that firms choose prices according to a mixed strategy. However, if firms can distinguish between the two groups, they would charge price equal to marginal cost to the local consumers and price equal to the reservation value to tourists. Therefore, as in the example in the text, the choosy group (here, the ignorant group) is treated badly by price discrimination. In Varian’s model, however, industry profit is unchanged when price discrimination is practiced.

³¹This discussion has only been about third-degree discrimination. At least in the case of monopoly, the typical case with other forms of discrimination is also that some prices rise while others fall. For instance, when a multiproduct firm practises mixed bundling, in many cases it will choose to raise its prices for individual purchase, and lower its price for joint purchase, compared to the case where the firm sets a separable price for each product. Similarly, if a firm offers a two-part tariff instead of a linear price, the typical case is that the overall tariff rises for consumers who buy little but falls for high-volume consumers.

Matters are more complicated when there are competing firms, as discussed in Corts (1998). The chief aspect which differs from monopoly is that a market might be strong for one firm but weak for its rival. For now, though, suppose firms do not differ in their judgement of which markets are strong. Corts uses the term “best response symmetry” for such cases.

Suppose there are two markets, 1 and 2, and two firms, A and B . Suppose there are no cross-price effects across the two markets, and that firm A 's profit in market i is $\pi_i^A(p_i^A, p_i^B)$ if it sets the price p_i^A there while its rival sets the price p_i^B . If discrimination is allowed, write $p_i^A = R_i^A(p_i^B)$ for firm A 's profit-maximizing price in market i if its rival sets the price p_i^B . (Similar notation is used for firm B 's reaction functions.) In reasonable situations these reaction functions are upward sloping. Suppose that both firms view market 1, say, as the strong market, in the sense that

$$R_1^A(p^B) > R_2^A(p^B) ; R_1^B(p^A) > R_2^B(p^A) .$$

See Figure 1 for a depiction of this situation.




Figure 1: Duopoly Reaction Functions

When firms can price discriminate, the equilibrium prices in market 1 are at γ on the figure, and the prices in market 2 are at α . Now suppose that firms cannot price discriminate.

As in the monopoly case just described, if the profit functions are single-peaked, firm A 's best response to a uniform price p^B from its rival will lie between its pair of reaction functions on the figure. Similarly, firm B 's response function if it cannot discriminate lies between its two reaction functions. We can deduce that the equilibrium prices when discrimination is not possible lie inside the diamond $a\beta\gamma\delta$. In particular, the comparison between discriminatory and non-discriminatory prices is clear: permitting discrimination increases the prices in market 1 (the strong market) and decreases the prices in market 2 (the weak market).

3.3 Discriminating on brand preference: best-response asymmetry

Now return to the specific Hotelling example, and suppose for simplicity that all consumers have the same choosiness parameter t .³² Suppose that firms can observe a consumer's location x and price accordingly. Consider firm A 's best response to its rival's price p_x^B when a consumer has brand preference parameter x . Firm A can profitably serve this consumer segment provided that

$$tx < p_x^B + t(1 - x) ,$$

in which case it will offer the limit price that just prevents the consumer from being tempted by the rival offer p_x^B , so that $p_x^A + tx = p_x^B + t(1 - x)$, or

$$p_x^A = p_x^B + t(1 - 2x) .$$

Thus, firm A prices high to those consumers who prefer its product but prices low to consumers who prefer its rival's product, so that small x consumers are the "strong market" for firm A . Similarly, those consumers who prefer firm B 's product are that firm's strong market. We deduce that one firm's strong market is the other's weak market. Corts terms this situation "best-response asymmetry".

With this form of price discrimination, the equilibrium price paid by a consumer located at x is

$$p_x = \begin{cases} (1 - 2x)t & \text{if } x \leq \frac{1}{2} \\ (2x - 1)t & \text{if } x \geq \frac{1}{2} \end{cases} \quad (7)$$

Consumers will obtain the product from the closer (preferred) firm, which is efficient, and those consumers closer to the middle will obtain the best deal (even when account is taken of their greater transport costs). A firm will set a high price to consumers with a strong brand preference for its product to exploit those consumers' distaste for the rival's product.

Suppose next that firms must set a uniform price to consumers. If consumers are uniformly distributed along the interval then the equilibrium uniform price is $p = t$. This

³²This analysis is due to Thisse and Vives (1988).

uniform price is above all the discriminatory prices in expression (7). Thus, this is an example where *all* prices decrease with price discrimination.³³ Price discrimination has no impact on total welfare since all consumers just wish to buy a single unit, and they buy this unit from the closer firm with either pricing regime. All consumers clearly benefit from price discrimination. Firms, however, make lower profits—in fact, in this example they make precisely *half* the profits—when they engage in this form of price discrimination compared to when they must offer a uniform price.

The fact that firms might be worse off when they practice price discrimination is one of the key differences between monopoly and competition. Ignoring issues of commitment, a monopolist is better off when it can price discriminate: the firm is free to choose a uniform price when discrimination is permitted but in general it is better off setting different prices. In the same way, an oligopolistic firm is always better off if it can price discriminate compared to when it cannot, for *given* prices offered by its rivals. However, once account is taken of what rivals too will do, firms in equilibrium can be worse off when discrimination is used. Firms then find themselves in a classic prisoner’s dilemma.

A closely related model, which is also useful for understanding models of dynamic price discrimination in section 5, is by Bester and Petrakis (1996).³⁴ Instead of being able to condition prices on a consumer’s precise location, here a firm merely observes whether a consumer has a brand preference for its product or its rival’s product. That is to say, firms observe whether a consumer has location $x \leq \frac{1}{2}$ or location $x \geq \frac{1}{2}$. The authors suppose that firms can target different prices to consumers in different regions by placing targeted coupons, which promise a discount if the consumer brings the coupon to the store, in different regional newspapers. When consumers are uniformly located along the interval one can show that the equilibrium discriminatory prices are

$$\hat{p} = \frac{2}{3}t ; p = \frac{1}{3}t . \quad (8)$$

Here, $\hat{p}$ is a firm’s price to a consumer on that firm’s “turf” (i.e., when the consumer is known to prefer that firm) and p is a firm’s price to a consumer on the rival’s turf. Thus, each consumer is offered two prices: a low price from the less preferred firm and a high price from the preferred firm. However, as in the Thisse-Vives model, these prices are below the equilibrium uniform price (which is $p = t$). Those consumers close to the middle of the interval, who have little preference for one firm over the other, will clearly choose the low price from the (slightly) more distant firm. This is inefficient, since consumers should buy from their preferred firm regardless of prices. However, consumers with a strong preference for a firm will buy from that firm despite the high price. Again, a prisoner’s dilemma

³³In this framework it is perfectly possible that some prices increase with discrimination. Suppose that instead of following a uniform distribution, the density of x on $[0, 1]$ is $f(x) = 6x(1 - x)$, a density which puts more consumers located close to the mid-point. Then one can show that the equilibrium uniform price is $p = \frac{2}{3}t$, which is lower than discriminatory prices for those consumers close to the ends of the interval.

³⁴See also Shaffer and Zhang (1995).

emerges, and both firms are better off without this form of price discrimination even though each would individually like to discriminate.³⁵

Thus, there are competitive situations where price discrimination causes all prices to fall.³⁶ In such cases, discrimination acts to intensify competition. The analysis in section 3.2 indicates that, at least in the case of third-degree discrimination, this situation can only occur when firms differ in their view of which markets are strong and which are weak. In the early literature on competitive price discrimination it was not always made clear that the crucial feature that can cause discrimination to intensify competition is best response asymmetry. For instance, Thisse and Vives (1988, page 134) wrote that “denying a firm the right to meet the price of a competitor on a discriminatory basis provides the latter with some protection against price attacks. The effect is then to weaken competition [...]” And Anderson and Leruth (1993, page 56) argue that price discrimination reduces profits since “firms compete on more fronts”. The example of discrimination based on choosiness in section 3.2 above, illustrates that it is not the number of fronts on which firms compete which is relevant. (In that example, discrimination raises equilibrium profit.) Rather, in the case of third-degree price discrimination what matters is whether firms have divergent views about which markets are strong and which are weak.³⁷

3.4 Private information

Next, consider a variant of the Bester and Petrakis (1996) model where firms have *private* information about consumer brand preferences.³⁸ Perhaps each firm purchases customer data from a different marketing company, for example. Suppose that if a consumer prefers firm A (i.e., she has location parameter $x \leq \frac{1}{2}$), firm i receives the signal $s^i = s_L$ with probability $\alpha \geq \frac{1}{2}$ and the signal $s^i = s_R$ with probability $1 - \alpha$. Similarly, if a consumer prefers firm B then the firm receives a signal $s^i = s_R$ with probability α and the signal

³⁵Notice that profit with discrimination in the Thisse-Vives model is always lower than that in the Bester-Petrakis model. Liu and Serfes (2004) propose a model that encompasses these two as extremes. They show that equilibrium profits are U-shaped in the precision of information about brand preferences: profit is lowest when the firms have information which is less precise than the perfect information in the Thisse-Vives framework.

³⁶Price discrimination might also cause all prices to fall when there is monopoly supply. Nahata, Ostaszewski, and Sahoo (1990) show that if the profit functions are not single-peaked then all prices might decrease, or all might increase, when a monopolist engages in third-degree price discrimination. In addition, as discussed in section 2.2, when a monopolist sells a durable good over time and cannot commit to future prices, all prices might fall compared to the case where the firm commits to set the same price over time.

³⁷Nevo and Wolfram (2002) present evidence consistent with the hypothesis that price discrimination via coupons in the breakfast cereal market exhibits best response asymmetry, and that the introduction of coupons leads to a fall in all prices. They also document how firms allegedly colluded to stop the use of coupons. Odlyzko (2003) discusses how competing railway companies welcomed tariff regulation in order to avoid profit-destroying price discrimination.

³⁸The discussion of the framework where firms possess a symmetric private signal is a simplified version of the model in chapter 2 of Esteves (2004).

$s^i = s_L$ with probability $1 - \alpha$. Conditional on a consumer's location, the signals s^A and s^B are independently distributed. A symmetric equilibrium will consist of a firm choosing the price $\hat{p}$ for those consumers they believe prefer their product and choosing price p when they think the consumer is likely to prefer their rival's product. (Specifically, if firm A observes the signal $s^A = s_L$ it will set the price $\hat{p}$, while if it sees the other signal it will set the price p . Firm B will follow the reverse strategy.)

Then Appendix A shows that equilibrium prices are

$$\hat{p} = \frac{t}{\alpha + \frac{1}{2}} ; p = \frac{t}{\alpha + 2\alpha^2} . \quad (9)$$

When $\alpha > \frac{1}{2}$ it follows that $\hat{p} > p$ and firms charge more to those consumers they consider likely to have a brand preference for them. When $\alpha = \frac{1}{2}$ (i.e., when the signal has no informational content) it follows that $\hat{p} = p = t$, just as in the standard Hotelling model without information. When the signal is perfectly accurate, i.e., when $\alpha = 1$, the prices are as given in expression (8), just as in the Bester-Petrakis model. More generally, the availability of the private signal causes both prices to fall compared to the case when the signal is not available. Finally, one can show that equilibrium industry profit falls monotonically as the accuracy of the private signal rises.

One can perform the same exercise when signals instead give information about a consumer's choosiness. In this case, industry profit is increasing in the precision of the signal. More interesting, though, is to present an asymmetric variant of this model which is suited to discussing whether a firm has an incentive to share information with its rival. Suppose there are two consumer segments: a consumer has choosiness parameter $t = t_L$ with probability $\frac{1}{2}$ and choosiness parameter $t = t_H$ with probability $\frac{1}{2}$. Suppose that firm A knows each consumer's choosiness precisely, but firm B knows nothing. Does A have an incentive to share its information with its rival? Without information sharing, firm A sets the two prices, p_L^A and p_H^A , respectively to the type- t_L and the type- t_H consumers, but firm B is constrained to choose the uniform price, p^B . One can show that the equilibrium prices are given by

$$p^B = 2 \frac{t_L t_H}{t_L + t_H} , p_L^A = \frac{3t_L t_H + t_L^2}{2t_L + 2t_H} , p_H^A = \frac{3t_L t_H + t_H^2}{2t_L + 2t_H} .$$

One can show that $p_L^A \leq p^B \leq p_H^A$, and so the better-informed firm has the greater market share in the price-sensitive market, but the smaller market share in the choosy market.³⁹ With these equilibrium prices, one can show that firm A 's profit is

$$\pi_{PRIV}^A = \frac{14t_L t_H + t_L^2 + t_H^2}{16(t_L + t_H)} .$$

³⁹One can also show neither market is cornered by one firm with these prices. Interestingly, firm B 's price is the same as when firm A is not informed. When neither firm is informed, expression (6) reveals that each firm sets the uniform price $2 \frac{t_L t_H}{t_L + t_H}$.

Firm B 's profit in this case is equal to $\frac{t_L t_H}{t_L + t_H}$. This profit is the same as if firm A were not informed. Therefore, in this case, an uninformed firm is indifferent about whether or not its rival has the ability to practice price discrimination. Firm A obtains higher profit compared to the case in which it was not informed.

Suppose now that firm A shares its information with its rival. (Assume that the firm has no scope to distort the information it passes to B .) In this case, both firms' equilibrium prices are $p_L = t_L$ in the price-sensitive segment and $p_H = t_H$ in the choosy segment. Firms share each segment equally, and so firm A obtains profit

$$\pi_{PUB}^A = \frac{t_L + t_H}{4} .$$

One can verify that $\pi_{PUB}^A > \pi_{PRIV}^A$ except in the extreme case when $t_L = t_H$ (in which case the information has no value). Thus, the well-informed firm has an incentive to provide its rival with information (and the rival is willing to accept this information). In this example, when a firm is uninformed about consumer choosiness it will price low, and this low price disadvantages the well-informed firm.

Finally, we can perform the same exercise for the case where one firm has information about consumer brand preference, while its rival does not.⁴⁰ Specifically, suppose all consumers have the same choosiness parameter t , and firm A knows whether $x < \frac{1}{2}$ or $x > \frac{1}{2}$ while firm B knows nothing. Suppose firm A sets the price $\hat{p}^A$ to those consumers it knows prefer its product and the price p^A to those consumers who prefer B 's product, while firm B sets the uniform price p^B . Then one can show the equilibrium prices are

$$p^B = \frac{1}{2}t ; \hat{p}^A = \frac{3}{4}t ; p^A = \frac{1}{4}t .$$

The profits of the two firms are

$$\pi_{PRIV}^A = \frac{5}{16}t ; \pi_{PRIV}^B = \frac{1}{4}t .$$

Suppose instead that neither firm has information. In this case, each firm sets the uniform price $p = t$ and makes profit of $\frac{1}{2}t$. Therefore, firm A is actually made worse off by its private information about consumer tastes.⁴¹ Clearly, if firm A could *secretly* obtain this information (so that firm B continues to set the price $p^B = t$) then it is made better off by its information. However, so long as it is common knowledge that firm A has this information (and will surely use it) then firm A is worse off once B 's aggressive response is considered. If firm A has this information, then it would like to *commit* not to price discriminate. Another

⁴⁰See also section III of Thisse and Vives (1988) and section 3 of Chen (2004). These authors assume that when one firm chooses to discriminate while the other does not, the non-discriminating firm acts as a Stackleberg price leader.

⁴¹If firm A gives firm B this information, then we are in the Bester-Petrakis situation, and each firm will make profit of $\frac{5}{18}t$. This means that firm A 's profits decrease, and so it has no incentive to release its information to B .

way to think about this is to suppose that the Bester-Petrakis model is extended to a two-stage interaction, where firms first (simultaneously) decide whether or not to “invest” in the technology or information gathering procedures needed to be able to price discriminate, and in the second stage they choose their price(s). In this two-stage game, one equilibrium is for neither firm to choose to be able to price discriminate in the first period. Thus, the prisoner’s dilemma aspect of the Bester-Petrakis model falls away in this dynamic version.^{42,43}

4 The effects of more instruments in oligopoly

4.1 Homogeneous consumers

The model presented in Armstrong and Vickers (2001, section 3) provides one framework in which to discuss the effects of increasing the range of tariff instruments available to firms. Suppose there are two firms, A and B . Suppose that firm i ’s maximum profit per consumer is $\pi_i(u)$ when the firm offers each of its consumers the surplus u . For this function to be well behaved, assume that firms have constant-returns-to-scale in serving consumers, so profit per consumer does not depend on the number of consumers served. Assume consumers are homogeneous in the sense that the relationship $\pi_i(u)$ is the same for all consumers. Assume further that consumers choose to purchase all relevant products from one firm. The shape of $\pi_i(u)$ embodies the firm’s cost function, the demand function of consumers, and—most important for the current purpose—the kinds of tariffs the firm can employ.

If consumers choose their supplier according to a Hotelling specification with homogeneous transport cost parameter t , firm i ’s profit is⁴⁴

$$\Pi_i = \left(\frac{1}{2} + \frac{u_i - u_j}{2t} \right) \pi_i(u_i) , \quad (10)$$

where u_i is its chosen consumer surplus and u_j is its rival’s consumer surplus. Let $\bar{u}_i$ denote the maximum level of consumer surplus that allows firm i to break even. In most natural cases, this is the level of consumer surplus associated with marginal-cost pricing, a pricing

⁴²There is also another, lower profit, equilibrium where both firms choose to be able to discriminate in the first period. Thus, there is a complementarity in the firms’ decisions about whether to price discriminate, and a firm’s incentive to discriminate is higher if its rival also discriminates. A similar feature is seen also in Ellison (2005).

⁴³A related issue is whether firms choose to compete in (i) “posted prices” or by (ii) “secret deals”. For instance, one might envisage a two-stage interaction in which each firm chooses its pricing strategy in the first stage and then chooses its actual prices in the second stage. With (i) a firm knows nothing about its individual consumers and posts a uniform price. With (ii) a firm might learn about the consumers from the negotiation process and price accordingly, and moreover the firm might take the rival’s posted price as given if the rival chooses strategy (i). Thus, when one firm chooses (i) and the other chooses (ii), the former acts as a Stackleberg price leader (as in section III of Thisse and Vives (1988)).

⁴⁴Note that transport cost is assumed to be a lump-sum cost which does not depend on the quantities chosen, i.e., it is a “shopping cost” rather than a “shipping cost”.

policy that is assumed to be feasible for either firm. Assume firms are symmetric except possibly for the range of tariff instruments they employ. This implies that each firm delivers the same consumer surplus with marginal-cost pricing, and so $\bar{u}_A = \bar{u}_B = \bar{u}$, say. Moreover, since pricing at marginal cost maximizes total per-consumer surplus ($u + \pi_i(u)$), it follows that $\pi'_i(\bar{u}) = -1$ for each firm. In sum, it is natural to assume

$$\pi_A(\bar{u}) = \pi_B(\bar{u}) = 0 ; \pi'_A(\bar{u}) = \pi'_B(\bar{u}) = -1 . \quad (11)$$

The analysis in Appendix B shows that in competitive markets (t small), firm i 's equilibrium consumer surplus (u_i), firm i 's total profit (Π_i), and total welfare (W) are approximately:

$$u_i \approx \bar{u} - t - \frac{t^2}{6} (2\pi''_i(\bar{u}) + \pi''_j(\bar{u})) \quad (12)$$

$$\Pi_i \approx \frac{t}{2} + \frac{t^2}{6} (2\pi''_i(\bar{u}) + \pi''_j(\bar{u})) \quad (13)$$

$$W \approx \bar{u} - \frac{t}{4} + \frac{t^2}{4} (\pi''_A(\bar{u}) + \pi''_B(\bar{u})) . \quad (14)$$

Next, suppose there are two possible profit functions, $\pi(\cdot)$ and $\hat{\pi}(\cdot)$, where $\pi(\cdot) \leq \hat{\pi}(\cdot)$. Therefore, $\hat{\pi}$ is the profit function when a firm has access to a wider range of instruments when designing its tariff compared to the situation associated with profit function $\pi(\cdot)$. (For instance, π might be the profit function when a firm uses linear prices, while $\hat{\pi}$ is the corresponding profit function when a firm uses two-part tariffs.) Since the two profit functions satisfy (11) and $\hat{\pi} \geq \pi$, it follows that $\hat{\pi}''(\bar{u}) \geq \pi''(\bar{u})$, with strict inequality except in knife-edge cases. Using expressions (12)–(14) we can draw the following conclusions for competitive markets.

Suppose the market environment changes so that both firms have access to more instruments in their tariff design. This can be modeled by supposing that both firms use the profit function $\hat{\pi}(\cdot)$ instead of $\pi(\cdot)$. In this case (i) consumer surplus falls, (ii) profit rises, and (iii) total welfare rises.⁴⁵ Thus, this model exhibits the same qualitative features as monopoly first-degree price discrimination discussed in section 2.1: the use of more instruments enables the industry to better extract consumer surplus, which causes profits and welfare to rise but consumer surplus to decline.

Notice also that there is no prisoner's dilemma in this setting, and a firm wishes to use the more ornate tariff even if its rival does not.⁴⁶ In sum, in this framework we predict that

⁴⁵This is part (i) of Proposition 3 in Armstrong and Vickers (2001). For instance, from expression (12) $u_i \approx \bar{u} - t - \frac{1}{2}t^2\pi''(\bar{u})$ if $\pi(\cdot)$ is used, but $u_i \approx \bar{u} - t - \frac{1}{2}t^2\hat{\pi}''(\bar{u})$ if $\hat{\pi}(\cdot)$ is used, and so consumers are worse off in the latter case. Expression (13) shows the firms' benefit from using more ornate tariffs is approximately twice the loss which consumers then suffer.

⁴⁶From expression (13), since $\hat{\pi}''(\bar{u}) > \pi''(\bar{u})$ it is a dominant strategy for a firm to use the profit function $\hat{\pi}$ instead of π . Also, in this case (12) shows that the firm which uses the more ornate tariff will obtain a smaller market share in equilibrium.

firms will use whichever tariff instruments they can, and the result is that profit and welfare rise (compared to a situation where less ornate tariffs are employed), while consumers suffer.

This result seems sometimes to be at odds with casual observation. Firms in many industries do not use two-part tariffs, for instance, even though such tariffs could be offered. (Supermarkets do not generally offer two-part tariffs, for example.) That is to say, there is sometimes *too little* price discrimination observed compared to what simple theory suggests. Therefore, this model is failing to capture an important aspect of many real-world markets. One possibility is that consumers have an (“irrational”) aversion to fixed charges, say, for psychological reasons, which is not captured in the model. Another possibility is that arbitrage between consumers renders some instruments infeasible. Alternatively, the presence of substantial consumer heterogeneity might overturn the result, and a firm might sometimes be made worse off by unilaterally choosing to compete with two-part tariffs given the response this induces from its rivals. A fourth aspect of some real-world markets which this model ignores is that consumers often make purchases from more than one supplier. This extension is pursued in the next section, where it is seen that the use of more ornate tariffs can depress equilibrium profit.

4.2 Mixed bundling

In addition to the case of discrimination based on brand preference (section 3.3), a second important example where price discrimination can intensify competition is competitive bundling. Consider for instance the two-dimensional Hotelling model presented by Matutes and Regibeau (1992).⁴⁷ Two firms, A and B , each offer their own brand of two products, 1 and 2. Consumers are uniformly distributed on the unit square, and the location (or brand preferences) of a consumer is denoted $(x_1, x_2) \in [0, 1]^2$. Let x_1 represent a consumer’s distance from firm A ’s brand of product 1 and x_2 represent a consumer’s distance from the same firm’s brand of product 2. Transport cost (or choosiness) is t for both products.

In general, each firm sets three prices. Let p_1^i denote firm i ’s price for its product 1, let p_2^i be its price for its product 2, and let p_{12}^i be its total price when a consumer buys both of its products. For simplicity, suppose that conditions in the market are such that all consumers buy both products.⁴⁸ Then a type- (x_1, x_2) consumer’s total cost if he buys both products from firm A is $p_{12}^A + t(x_1 + x_2)$, his total cost if he buys both products from B is $p_{12}^B + t([1 - x_1] + [1 - x_2])$, and his total cost if he buys product i from A and product $j \neq i$ from B is $p_i^A + p_j^B + t(x_i + [1 - x_j])$. The consumer will choose the option from among these four possibilities which involves the smallest outlay. Whenever $p_{12}^A + p_{12}^B \leq p_1^A + p_2^A + p_1^B + p_2^B$ (which corresponds to the case where firms do not charge a premium for joint consumption), the pattern of demand is shown in Figure 2.

⁴⁷See Anderson and Leruth (1993) for related analysis using a Logit model of consumer demand.

⁴⁸This assumption implies that the two products could equally well be taken to be perfect complements, so that each consumer needs to buy both products if he is to gain any utility, as is assumed in the more general model with partial consumer participation in Matutes and Regibeau (1992).

Suppose that production is costless. Consider first the case where firms must choose from the restrictive set of tariffs which do not involve discounts for joint consumption. That is to say, firms must choose separable tariffs, so that $p_{12}^i \equiv p_1^i + p_2^i$. In this case firms compete product-by-product, and the equilibrium price for each product from each firm is $p = t$. Here, each firm makes a total profit of $\pi = t$.




Figure 2: Pattern of Consumer Demand with Duopoly Bundling

Next suppose firms can offer tariffs involving the additional instrument of discounts for joint consumption (i.e., firms can set $p_{12}^i \neq p_1^i + p_2^i$). By calculating the areas depicted in Figure 2, one can show the symmetric equilibrium prices are

$$p_1^A = p_2^A = p_1^B = p_2^B = \frac{11}{12}t ; p_{12}^A = p_{12}^B = \frac{4}{3}t . \quad (15)$$

These prices are uniformly below the equilibrium prices without bundling. A firm's equilibrium profit with the prices in (15) is $\pi \approx 0.7 \times t$, compared to $\pi = t$ when the bundling instrument is not used. Similarly to the model in section 3.3, firms play a prisoner's dilemma: given a rival's prices, a firm is always better off if it has the flexibility to engage in mixed bundling, but in equilibrium profit is reduced when this extra instrument is used by both firms.

As discussed in section 3.4, in some circumstances it is natural to model the firms' interaction as a two-stage game, where firms first decide whether to compete using separable

prices or using mixed bundling, and then in the second stage they choose their prices. Suppose in the second stage of this interaction that firm A can practice mixed bundling whereas firm B cannot. In the case where consumer locations are uniformly distributed one can show equilibrium prices are

$$p_1^A = p_2^A = \frac{13}{12}t ; p_{12}^A = \frac{3}{2}t ; p_1^B = p_2^B = \frac{3}{4}t .$$

Firm A 's equilibrium profit in this case is approximately $0.82 \times t$, while firm B 's profit is $\frac{3}{4}t$. Recall that if neither firm chooses to compete with mixed bundling each will make profit t . Therefore, in the first stage a firm has no incentive unilaterally to price discriminate in this manner if its rival does not. Moreover, even if its rival chooses to bundle, it is still in a firm's interest to choose not to bundle. Therefore, in this extended two-stage game it is no longer an equilibrium to practice bundling, and it is a dominant strategy for a firm to commit to price its products separably.⁴⁹

When this two-stage interaction is not the natural model for the market, however, equilibrium prices are given in expression (15). Consumers with strong tastes for one firm's product (in the top quartile) and strong tastes for the rival firm's other product (in the top quartile) will buy one product from each firm. Clearly, all consumers are better off as a result of the price reductions caused by price discrimination. However, there is excessive joint consumption (or one-stop shopping): too many consumers buy both products from the same firm than is efficient, and welfare falls with this form of discrimination. These are exactly the opposite comparative statics to those obtained in the exclusively one-stop shopping model described in section 4.1.

The economic reason why bundling depresses profit is not easy to come by. (It cannot be anything to do with "best response asymmetry" since different firms do not view the various kinds of consumer as being strong or weak in different ways.) Some intuition is available if one restricts attention to a choice between separable pricing and *pure* bundling.⁵⁰ Suppose consumer preferences are as described above. As before, if firms compete with separable prices then they compete product-by-product and each firm makes profit t . Alternatively, if firms bundle their products together then a type- (x_1, x_2) consumer will buy from firm A if

$$p_{12}^A + t(x_1 + x_2) \leq p_{12}^B + t([1 - x_1] + [1 - x_2]) ,$$

where p_{12}^i is the price for firm i 's bundle. One can then show that the equilibrium bundled price is $p_{12}^A = p_{12}^B = t$, and firms each make profits of just $\frac{1}{2}t$. Thus, with pure bundling the

⁴⁹Gans and King (2005) consider a variant of this model, where there are four separate firms each offering one of the four product variants. They show that it can be profitable for two firms to agree to set a discount for joint purchase of their two products. (The authors assume that the two firms commit to the size of the discount, which they agree to fund equally, and subsequently set their prices non-cooperatively.) If both pairs of firms do this, then there is no effect on profit compared to the case where neither pair of firms offers bundling discounts (although there is again socially excessive bundling in this case).

⁵⁰See Matutes and Regibeau (1988), Economides (1989), and Nalebuff (2000) for this analysis.

industry profit *halves*. Why is this? Equilibrium prices are determined as a balance between building market share and exploiting infra-marginal consumers. If there are many marginal consumers compared to consumers overall (as is the case with homogeneous consumers), then equilibrium prices are close to cost. The crucial point here is that there is a greater number of marginal consumers when firms compete in bundles. (Geometrically, with bundling the marginal consumers lie along the diagonal of the square, with length $\sqrt{2}$, whereas with separable prices the marginal consumers lie along a line of length 1.) This effect—that bundling has a “homogenizing” effect on consumer valuations—is essentially the same as discussed in the monopoly context of section 2.1. With monopoly, when consumers become more homogeneous because of bundling, the firm can extract more consumer surplus. With competition, though, when bundling makes consumers more homogeneous this leads to more aggressive competition and low prices.^{51,52}

4.3 Third-degree price discrimination

This section discusses the asymmetric situation in which a firm which operates in several markets faces a rival which is (potentially) active in just one market.⁵³ Price discrimination can enable a multi-market firm to target price cuts selectively in those markets where competition is present. The ability to make selective price cuts is therefore likely to have an adverse effect on a single-market firm’s profits.

To see this formally, suppose firm A serves two independent markets, 1 and 2. Market 1 is monopolized by firm A , whereas in market 2 there is a rival firm B . Firm A ’s (linear) prices in the two markets are p_1^A and p_2^A , while firm B ’s price in market 2 is p_2^B . In market 1, firm A ’s profit function is $\pi_1(p_1^A)$ whereas in market 2 the firm’s profit function given its rival’s price is $\pi_2^A(p_2^A, p_2^B)$. Similarly, firm B ’s profit function is $\pi_2^B(p_2^A, p_2^B)$. If firm A can price discriminate, in the sense that it may set different prices in the two markets, it will set the monopoly price p_1^M in market 1 (where this price maximizes $\pi_1(\cdot)$) and in market 2 prices

⁵¹Nalebuff (2000) shows that as the number of products which can be bundled becomes large, the profit obtained with pure bundling become an ever smaller proportion of the profit obtained with separable pricing.

⁵²The fact that (pure) bundling can make firms price aggressively appears in a different context in Whinston (1990). There, a multiproduct firm faces competition from a single-product rival. In certain conditions, when the multiproduct firm commits to bundle its products together before prices are chosen, it will choose a lower effective price in the rival’s market compared to when it prices its products separably. Therefore, in some circumstances a commitment to bundle can act to deter single-product entry. (See section 4.3 for a related model in a more conventional price discrimination framework.) In other circumstances, however, when the multiproduct firm commits to bundle its products, this can act to *relax* subsequent competition with a single-product rival. See Carbajo, De Meza, and Seidman (1990), Whinston (1990), Chen (1997a), and section III.E of Nalebuff (2004) for further details. In such cases, an incumbent might wish to commit to prices its products separably to convince potential entrants that it will be aggressive in the event of entry.

⁵³This discussion is loosely based on Armstrong and Vickers (1993). That paper also argues that the effect of allowing price discrimination on the aggressiveness of competition is exacerbated when the multi-product firm is regulated and operates under an average-price constraint. (If it reduces its price in the competitive market it is then allowed to raise its price in the captive market.)

will be determined by the intersection of reaction functions. In market 2, let $p_2^A = R_2^A(p_2^B)$ be firm A 's best response to firm B 's price when price discrimination is practiced, and let $p_2^B = R_2^B(p_2^A)$ be firm B 's best response to A 's price. These two reaction functions are depicted on Figure 3, and the equilibrium prices in market 2 are at the point β on the figure.




Figure 3: The Effect of Banning Price Discrimination

Next suppose that firm A cannot set different prices in the two markets. Suppose further that for all relevant prices set by B , firm A prefers to set a higher price in the monopoly market than in the competitive market, so that $p_1^M > R_2^A(p_2^B)$. In other words, the captive market is firm A 's strong market, which seems plausible in a wide variety of contexts. Then, with single-peak assumptions on firm A 's profit functions, it follows that when it must set a common price in the two markets, firm A 's best-response to its rival's price, $R_{ND}^A(p_2^B)$ say, is shifted upwards, so that $R_{ND}^A(p_2^B) > R_2^A(p_2^B)$. (See Figure 3.) It follows that the market 2 prices when price discrimination is permitted are all lower than the corresponding prices when A must set a uniform price (denoted by γ on the figure). Therefore, the effect of a ban on price discrimination is to reduce the price in the captive market, and to raise both prices in the competitive market. The profit of the single-market rival clearly increases with such a ban, while the effect on firm A 's overall profit is not clear-cut in general.⁵⁴

⁵⁴Dobson and Waterson (2005) present a related model where a national retailer operates in a number

Of course, this has implications for the rival's decision about whether to enter the market. If the entrant has a fixed cost of entry, it will enter only if it expects its post-entry profit to cover its entry cost. There are then three cases to consider. If the entry cost is large, entry will not take place regardless of whether the incumbent can price discriminate. In this case, the social desirability of price discrimination is exactly as in the standard monopoly case (which is ambiguous in general). Similarly, if the entry cost is very small, then entry will take place regardless of policy towards price discrimination. The interesting case is when the cost of entry lies in the intermediate range where entry is profitable only if the incumbent is not permitted to discriminate. In such cases, a ban on price discrimination acts to induce entry.⁵⁵ In such cases it is plausible that a ban on price discrimination will cause prices in *both* markets to fall: if discrimination is possible, there will be no entry and the incumbent will charge monopoly prices in each market; if the incumbent must charge a common price in the two markets, then this will bring in the entrant, and this will bring both prices down from monopoly levels.

The general principle, as in the Thisse-Vives quote in section 3.3, is that denying an incumbent the right to meet the price of a competitor on a discriminatory basis provides the latter with some protection against price attacks. While the effect of a ban on price discrimination is indeed to weaken competition if the entrant is already in the market, once the *ex ante* incentives to enter are considered, the effect of a ban on price discrimination might actually be pro-competitive. However, the welfare effects of a ban on price discrimination in this context are not clear cut. For instance, since the incumbent is reluctant to cut its profits in the captive market by meeting its rival's price in the competitive market, even a highly inefficient entrant might prosper. While preventing an incumbent from engaging in selective price cuts is likely to be a powerful means with which to assist entry, as with many forms of entry assistance the danger of inefficient entry is never far away.⁵⁶

5 Dynamic pricing and revealed preference in oligopoly

This section extends the discussion of behaviour-based price discrimination in section 2.2 to competitive situations. Two models are presented, and in each model firms learn about a consumer's relative preference for the firms' future offerings from a consumer's initial choice of firm. In the first model, consumers have a stable brand preference for one of the two firms. If a consumer buys from firm *A* in the first period, she will most likely prefer to buy from firm *A* in the second period as well, all else equal, and both firms will price accordingly. In

of markets, in some of which it is the sole supplier and in the remainder of which it faces a single local competitor. They show that it is possible for the chain store to benefit if it commits to a national pricing policy (i.e., it does not price differently depending on competitive markets in each local market).

⁵⁵Armstrong and Vickers (1993) present a model where the entrant can choose its scale of entry, and we show that a ban on discrimination will cause the entrant to increase its scale of entry.

⁵⁶See section 2.2 of Vickers (2005) for an account of the recent theory and policy towards selective price cuts.

the second model, consumers initially view the two firms as perfect substitutes, but in the second period they incur a switching cost if they wish to change supplier. In each model, the second period closely resembles the static model of Bester and Petrakis (1996) discussed in section 3.3: when price discrimination is permitted firms will price low to “poach” their rival’s previous customers and price high to their own previous customers. Each firm regards its previous customers as its strong market, and there is “best response asymmetry” in the terminology of Corts (1998). Therefore, the models will share the feature the second-period prices are all lower than they would be if behaviour-based discrimination were not feasible.

Stable brand preferences

First, consider Fudenberg and Tirole (2000)’s model with stable brand preferences. Suppose there are two periods, and consumers wish to buy a single unit from one of two firms (A and B) in each of the two periods. Consumer preferences are as specified in the Hotelling framework, and a consumer located at $x \in [0, 1]$ incurs total cost $p^A + tx$ if she buys from firm A at the price p^A and she incurs total cost $p^B + t(1 - x)$ if she buys the unit from B at the price p^B . Assume that a consumer’s brand preference parameter x is the same in the two periods.

Suppose that firms cannot commit to future prices, in which case the model is analyzed by working back from period 2. In the second period we have, in general, an asymmetric version of the Bester and Petrakis (1996) model. Specifically, suppose firm A managed to attract a fraction $\frac{1}{2} + \gamma$ of consumers in the first period. That is to say, firm A ’s turf is the interval $[0, \frac{1}{2} + \gamma]$, while firm B ’s turf is the remaining $[\frac{1}{2} + \gamma, 1]$. For simplicity, suppose production is costless. In firm A ’s turf one can show the equilibrium second-period prices are⁵⁷

$$\hat{p}_2^A = \frac{2}{3}t(1 + \gamma) ; p_2^B = \frac{1}{3}t(1 + 4\gamma) . \quad (16)$$

Similarly, the prices aimed at firm B ’s turf are

$$\hat{p}_2^B = \frac{2}{3}t(1 - \gamma) ; p_2^A = \frac{1}{3}t(1 - 4\gamma) . \quad (17)$$

These prices generalize expression (8) above.⁵⁸ Thus, the larger “turf” has the higher prices, which implies that, all else equal, a sophisticated consumer would prefer to purchase from the firm with fewer consumers in the first period. Perhaps surprisingly, with these prices firms make the same profit in the second period, and this common profit is

$$\pi_2^A = \pi_2^B = \left(\frac{5}{18} + \frac{10}{9}\gamma^2 \right) t .$$

⁵⁷For these prices to be valid the initial market shares cannot be too asymmetric, and we require that $-\frac{1}{4} \leq \gamma \leq \frac{1}{4}$. For more asymmetric market shares, the firm with the larger share will set a price of zero to its distant market.

⁵⁸Notice that in very asymmetric situations, where one firm has more than two-thirds of the first-period consumers, the smaller firm sets a *higher* price on its rival’s turf than on its own. Related issues are explored in Shaffer and Zhang (2000).

Thus, each firm's second-period profit is minimized when firms share the first-period market equally. Intuitively, an equal initial market share generates the most informative outcome in the second period, and, in this setting with best-response asymmetry, more information destroys profit. When initial market shares are very asymmetric, on the other hand, little is learned about (most) consumers' brand preferences.⁵⁹

Consider next the choice of first-period prices. Here, the outcome depends on the sophistication of consumers. Suppose first that consumers are sophisticated, and anticipate the effect of initial market share on future prices in expressions (16)–(17). When both firms choose the same initial price p_1 say, they will share the market equally in the first period (so $\gamma = 0$). Suppose instead that firm A slightly undercuts its rival in the first period, and chooses price $p_1 - \varepsilon$. How many more consumers will it attract in the first period? Let us say that the marginal consumer in period 1 is located at $x = \frac{1}{2} + \gamma$. When a consumer lies near the midpoint between the two firms (γ small), she will surely switch firms in the second period in order to take advantage of the poaching discount. Therefore, the marginal consumer is indifferent between buying from A -then- B or from B -then- A . Given the second-period prices in (16)–(17), γ satisfies

$$p_1 - \varepsilon + \left(\frac{1}{2} + \gamma\right)t + \delta \left[\frac{1}{3}t(1 + 4\gamma) + \left(\frac{1}{2} - \gamma\right)t\right] = p_1 + \left(\frac{1}{2} - \gamma\right)t + \delta \left[\frac{1}{3}t(1 - 4\gamma) + \left(\frac{1}{2} + \gamma\right)t\right] .$$

Therefore,

$$\gamma = \frac{\varepsilon}{2t(1 + \frac{1}{3}\delta)} .$$

The key insight is that sophisticated consumers react less sensitively to price reductions in the first period than they would in a static model of this kind (when we would have $\gamma = \varepsilon/(2t)$). The reason for this is that if a firm cuts its price in the first period, that brings a direct benefit to a consumer in the first period, but a disadvantage to the marginal consumer at that firm in the second period, since her “poaching” price is raised. It follows that the equilibrium first-period price is⁶⁰

$$p_1 = t(1 + \frac{1}{3}\delta) \tag{18}$$

while the second-period prices are as in (8) above:

$$p_2 = \frac{1}{3}t ; \hat{p}_2 = \frac{2}{3}t . \tag{19}$$

⁵⁹This is especially clear in the variant of the Fudenberg-Tirole model presented in chapter 3 of Esteves (2004). Esteves assumes that consumer preferences are determined by a binary distribution, and half of consumers prefer A by a fixed amount and half of consumers prefer B by the same fixed amount. If firms set similar prices in the first period, they will share the market and consumer tastes are fully revealed in the second period. If firms set significantly different first-period prices then one firm attracts all consumers, and so nothing is learned in the second period and subsequent profit is high.

⁶⁰In fact, if a firm undercuts its rival in the first period this will affect the firm's equilibrium profits in the second period (which will rise). However, this effect is second-order for small deviations, at least for this particular example, and so it plays no role in the calculation of the equilibrium first-period price.

These second-period prices imply that the middle third of consumers switch firms in the second period.

The sophisticated reasoning of the consumers in this model when there is behaviour-based discrimination might sometimes be thought implausible, especially in new markets when consumers have not yet grasped firms' pricing incentives. Suppose instead consumers are naive and do not foresee that when they buy from the larger firm in the first period they will face higher prices in the second period. Then, competition in the first period is just as in a standard static model, and the price-period price is $p_1 = t$. Second-period prices are unaffected. Therefore, in this model, if consumers are naive (or myopic) this improves their position compared to when they are fully strategic. By contrast, in the monopoly situation of section 2.2, when consumers are naive they are treated less favourably.

If the instrument of behaviour-based price discrimination is not available, a firm's price in each period is $p_1 = p_2 = \hat{p}_2 = t$. This implies that the first-period price is reduced (or unchanged with naive consumers) while second-period prices are raised compared to the situation with discrimination. Just as with the Bester-Petrakis model, behaviour-based price discrimination is socially inefficient, since in the second period a third of consumers buy from the less-preferred firm.⁶¹ When consumers are sophisticated, equilibrium profit is (slightly) lower with discrimination.⁶² (When consumers are naive, firms are even worse off with discrimination.) Therefore, when firms cannot commit to future prices, the ability to engage in behaviour-based price discrimination reduces their profit. Consumers in aggregate are better off with discrimination.⁶³ (If consumers were naive, though, they would be even better off with discrimination.) In sum, a ban on this form of price discrimination would make consumers worse off, regardless of their presumed sophistication.⁶⁴ (By contrast, in the monopoly analysis of section 2.2, a ban on discrimination would make naive consumers worse off but sophisticated consumers better off.)

What happens instead if firms can commit fully to future prices? That is to say, at the start of period 1, each firm i announces its three prices p_1^i, p_2^i and $\hat{p}_2^i$, where p_1^i is its price for a unit in the first period, p_2^i is the (poaching) price for a unit in the second period if

⁶¹As ever, one should be wary of reaching policy conclusions on the basis of these unit demand models since prices have little role to play in welfare terms. If consumers had elastic multi-unit demands in each period, the price reductions in the second period (as well as the high prices in the first period) would have a welfare impact.

⁶²Industry profit with discrimination is $(1 + \frac{8}{9}\delta)t$ and without discrimination it is $(1 + \delta)t$.

⁶³With no discrimination, consumers pay a total discounted charge $(1 + \delta)t$ for the two units. With discrimination, a consumer can buy a unit from the same firm in each period, in which case expressions (18)–(19) imply the total discounted charge is less than $(1 + \delta)t$, and so all consumers are better off with price discrimination.

⁶⁴Villas-Boas (1999) presents related analysis when long-lived firms face overlapping generations of short-lived consumers. One important difference with the Fudenberg-Tirole model is that a firm only knows whether or not a consumer has previously purchased from it, and cannot distinguish between its rival's previous consumer and consumers who are new to the market. He finds that, at least for patient firms and consumers, behaviour-based discrimination causes all prices to fall.

the consumer did not purchase from i in the first period, and $\hat{p}_2^i$ is the firm's second-period price to repeat buyers.⁶⁵ One equilibrium of this game is for firms to offer pure bundling contracts. Specifically, in return for an up-front fee $(1+\delta)t$, any consumer can obtain one unit in period 1 and one unit in period 2 from the firm. (This implements the same outcome as the static Hotelling equilibrium, where price is equal to t , duplicated for two periods.) If one firm offers this pure bundling contract, its rival's best response is to offer the same contract. Thus, if firms can commit to their prices, the result is the same as when behaviour-based price discrimination is not possible. (The same result is true for monopoly.)

Changing brand preferences

The discussion so far has not considered the possibility that brand preferences change over time. If preferences change over time, then even with the ability to commit to future prices, firms will generally have an incentive to practice behaviour-based price discrimination. If consumers know their future preferences at the time of their initial choice of firm, and firms can commit to future prices, the model of Matutes and Regibeau (1992), discussed in section 4.2, applies immediately. If brand preferences for the two firms are independently distributed over time (and follow the uniform distribution), behaviour-based price discrimination causes all prices to fall compared to the case when firms cannot make their second-period prices depend on whether the consumer purchased from them in the first period.⁶⁶ This is clearly beneficial for consumers, although social welfare is reduced since there is excessive joint consumption. This provides an interesting contrast to the Fudenberg-Tirole model—in which brand preferences are stable but where firms cannot commit to future prices—where there is excessive switching between firms. The Matutes-Regibeau framework, where firms commit to rewards for consumer loyalty, seems broadly applicable to frequent flyer programmes and the like.⁶⁷

Switching costs

Next consider Chen (1997b)'s model of switching costs. Again, there are two firms each selling a product over two periods and consumers wish to purchase a single unit in each of the two periods. If a consumer wishes to change supplier in the second period, suppose

⁶⁵Section 5 of Fudenberg and Tirole (2000) analyzes the use of long-term contracts but without full commitment. (Firms have the opportunity to poach their rival's customers in the second period at prices that are determined only in the second period.) They show in the uniform example that, compared to when only short-term contracts are employed, the use of long-term contracts reduces profit, reduces switching (which improves efficiency) and boosts consumer surplus.

⁶⁶Prices are as given in expression (15), and so a repeat buyer receives a discount of $\frac{1}{2}t$ on her second unit.

⁶⁷For related analysis in the case where consumers do not know their future preferences at the time of the initial choice, see Caminal and Matutes (1990) and section 6 of Fudenberg and Tirole (2000). See Kim, Shi, and Srinivasan (2001) for a recent analysis of reward programs as a price discrimination device to segment the high-volume and low-volume markets. This paper also shows that firms might commit to use inefficient reward schemes (e.g., free gifts instead of price reductions), since that can sometimes act to soften subsequent price competition.

she must incur a switching cost s . This cost varies across consumers, and suppose that s is uniformly distributed on the interval $[0, t]$. Thus, if firm A sets the second-period price $\hat{p}_2^A$ to its first-period customers, and if firm B sets the poaching price p_2^B to the same group of customers, a customer will switch to B whenever $\hat{p}_2^A > p_2^B + s$. If firms can discriminate in the second period, it is straightforward to show that the equilibrium prices are just as in the brand loyalty model in (19). (These are the prices regardless of market shares in the first period.) Just as in the Fudenberg-Tirole model, a third of consumers switch suppliers in the second period, which is inefficient. If the total number of consumers is normalized to 1, and if the number of firm A 's first-period customers is denoted n^A (and the number of firm B 's customers is $n^B = 1 - n^A$), then firm A 's second-period profit is⁶⁸

$$\frac{2}{3}n^A\hat{p}_2 + \frac{1}{3}n^B p_2 = \frac{1}{3}t(\frac{1}{3} + n^A) . \quad (20)$$

Next, consider the equilibrium first-period price. Since second-period prices are same for the two firms, a consumer will just choose her initial supplier on the basis of the lowest price. Since each first-period consumer brings with her a second-period profit of $\frac{1}{3}t$ (see expression (20)) which is discounted by δ , the equilibrium first-period price is $p_1 = -\frac{1}{3}\delta t$, which is below cost. This feature is common in models of switching costs, where firms compete hard to attract new consumers anticipating that consumers will generate high profit subsequently. Therefore, we see a major difference between the two models of Fudenberg-Tirole and Chen: in the former, prices start high and then decrease, whereas here the reverse is seen.

Next, consider the situation where firms must charge a uniform price in the second period. Now, in contrast to the case of price discrimination, the equilibrium second-period prices will depend on initial market shares. If firms divide the market equally in the first period, the equilibrium price in the second period is $p_2 = t$. If firms have different market shares in the first period, the larger firm will choose the higher price in the second period since it has a greater number of "captive" consumers. More precisely, as shown by Chen, uniform prices in the second period given firm A 's initial market share $n^A \geq \frac{1}{2}$ are

$$p^A = \frac{n^A + 1}{3n^A}t ; p^B = \frac{2 - n^A}{3n^A}t \leq p^A .$$

Both prices here are lower compared to when firms share the market equally ($n^A = \frac{1}{2}$). Thus, second-period competition is intensified when firms have asymmetric initial market shares.⁶⁹ Therefore, in the first period, sophisticated consumers will not react sensitively to price reductions since they know that if they buy from the larger firm they will face a higher price in the second period. (Somewhat confusingly, this is the same mechanism as occurs

⁶⁸Thus, in contrast to the Fudenberg-Tirole model, here the firm with the larger initial market share makes a higher profit in the second period compared to its rival.

⁶⁹Second-period profits of the two firms are $\pi_2^A = t\frac{(1+n^A)^2}{9n^A}$ and $\pi_2^B = t\frac{(2-n^A)^2}{9n^A}$. Thus, the firm with the larger market share obtains higher second-period profit. However, the profit of both firms is decreasing in the larger firm's market share and both firms are better off if they share the first-period market equally.

in the Fudenberg-Tirole model *with* price discrimination.) As a result, relatively high prices can be sustained in equilibrium in the first period.⁷⁰ Chen shows that firms are always better off when behaviour-based price discrimination is not possible, and it is ambiguous whether consumers are better or worse off.^{71,72}

6 Concluding comments

This paper has surveyed the recent literature on price discrimination, with a focus on effects on industry outcomes when firms: (i) have access to more information about their potential customers; (ii) can use more instruments when choosing their tariffs; and (iii) cannot commit to their pricing policy. The paper argued that the importance of each of these three factors has been increased due to developments in marketing and e-commerce. The analysis reported here is more suggestive than definitive, and was largely presented through a series of worked examples. (Many of the papers from this analysis was taken share this feature.) There is plenty more work to be done in this exciting area.

⁷⁰There are technical complexities in the model, in that there are multiple symmetric first-period prices that can be equilibria. Some of these prices can be below cost, just like in the situation with price discrimination. Take for instance the situation where consumers are naive (which is not considered by Chen), and in the first period consumers simply buy from the firm with the lower price (and the initial market is equally divided if the firms choose the same initial price). Then if a firm undercuts its rival in the first period it will attract all consumers and so make less profit in the second period (see the previous footnote). Similarly, if a firm sets a negative price in the first period, its rival may be unwilling to choose a higher price, because if it does so it will lose all its initial consumers and so make small profits in the second period. One can show that any initial price p_1 in the range $-\frac{7}{9}\delta t \leq p_1 \leq \frac{1}{9}\delta t$ can be a symmetric equilibrium.

⁷¹See Taylor (2003) for an extension to Chen's analysis in a number of important directions, for instance to more than two periods and more than two firms. When there are more than two periods, the fact that a consumer switched supplier in the second period indicates that the consumer has low switching costs, and this could generate more intense competition for him in future periods.

⁷²Asplund, Eriksson, and Strand (2004) show that newspaper subscriptions in local Swedish markets are more frequently sold at (perhaps introductory) discounts in duopoly regions than in monopoly regions, and they argue that this is suggestive of customer poaching.

TECHNICAL APPENDIX

A Private information

Here, the equilibrium prices in expression (9) are derived. Suppose that firm B sets the price $\hat{p}$ when it sees the signal s_R and the price p when it sees the signal s_L . First consider the case when firm A sees the signal $s^A = s_L$ for a consumer. What price should firm A offer this consumer? There are four relevant events: (i) the consumer is near to firm A and the rival firm also receives the signal $s^B = s_L$ and so sets the price p ; (ii) the consumer is near to firm A but the rival firm receives the ‘wrong’ signal $s^B = s_R$ and so sets the high price $\hat{p}$; (iii) the consumer is far from A and the rival firm also receives the signal $s^B = s_L$ and so sets the price p , and (iv) the consumer is far from A and the rival firm receives the ‘correct’ signal $s^B = s_R$ and so sets the price $\hat{p}$. Conditional on firm A seeing the signal $s^A = s_L$, the respective probabilities of the four events are $\alpha^2, \alpha(1 - \alpha), (1 - \alpha)^2, \alpha(1 - \alpha)$. If firm A sets a price $\hat{q}$ to this consumer (which lies in the range $p \leq \hat{q} \leq \hat{p}$) then it will make the sale with probability

$$\alpha^2(1 + \frac{p - \hat{q}}{t}) + \alpha(1 - \alpha) + \alpha(1 - \alpha)\frac{\hat{p} - \hat{q}}{t} .$$

(With case (ii) firm A will surely make the sale whenever it sets a price $\hat{q} \leq \hat{p}$; with case (iii) it will surely not make the sale when $\hat{q} \geq p$.)

Suppose next that firm A observes signal $s^A = s_R$. In the same way, if it sets the price q (which lies in the range $p \leq q \leq \hat{p}$) its probability of making the sale is

$$\alpha^2\frac{\hat{p} - q}{t} + \alpha(1 - \alpha)(1 + \frac{p - q}{t}) + (1 - \alpha)^2 .$$

One can then show that symmetric equilibrium prices are given by expression (9).

B Effect of more instruments

Here, the results reported in section 4.1 are derived. Let $\pi_A(u)$ and $\pi_B(u)$ be two profit-as-a-function-of-utility functions which satisfy (11). Firm i will choose u_i to maximize expression (10), and so equilibrium utilities $\{\hat{u}_A(t), \hat{u}_B(t)\}$ satisfy the first-order conditions

$$\pi_i(\hat{u}_i(t)) + (t + \hat{u}_i(t) - \hat{u}_j(t))\pi'_i(\hat{u}_i(t)) , \quad (21)$$

where $\hat{u}_i(t)$ denotes the equilibrium utility offered by firm i when transport cost is t . We wish to investigate the equilibrium in competitive markets, i.e., when $t \approx 0$. Clearly, $\hat{u}_i(0) = \bar{u}$. Differentiating (21) yields

$$\pi'_i\hat{u}'_i + (1 + \hat{u}'_i - \hat{u}'_j)\pi'_i + (t + \hat{u}_i - \hat{u}_j)\pi''_i\hat{u}'_i \equiv 0 .$$

Setting $t = 0$ here yields $\hat{u}'_A(0) = \hat{u}'_B(0) = -1$. Differentiating (21) once more yields

$$\pi''_i(\hat{u}'_i)^2 + \pi'_i\hat{u}''_i + (\hat{u}''_i - \hat{u}''_j)\pi'_i + 2(1 + \hat{u}'_i - \hat{u}'_j)\pi''_i\hat{u}'_i + (t + \hat{u}_i - \hat{u}_j)(\pi'''_i(\hat{u}'_i)^2 + \pi''_i\hat{u}''_i) = 0 .$$

When $t = 0$ this simplifies to

$$\hat{u}''_j(0) - 2\hat{u}''_i(0) = \pi''_i(\bar{u}) .$$

Solving this pair of simultaneous equations yields

$$\hat{u}''_i(0) = -\frac{2\pi''_i(\bar{u}) + \pi''_j(\bar{u})}{3} ,$$

which leads to expression (12).

The equilibrium market share of firm i is approximately

$$n_i = \frac{1}{2} + \frac{u_i - u_j}{2t} \approx \frac{1}{2} + \frac{1}{12}(\pi''_j(\bar{u}) - \pi''_i(\bar{u}))t .$$

Aggregate consumer surplus in equilibrium is approximately equal to

$$U(t) \approx \bar{u} - \frac{5t}{4} - \frac{t^2}{4}(\pi''_A(\bar{u}) + \pi''_B(\bar{u})) . \quad (22)$$

The per-consumer profit for firm i is

$$\pi_i(\hat{u}_i(t)) \approx t + \left(\frac{5}{6}\pi''_i(\bar{u}) + \frac{1}{6}\pi''_j(\bar{u}) \right) t^2 .$$

Therefore, the total profit of firm i (i.e., $n_i \times \pi_i$) is given by expression (13). Industry profit is therefore

$$\Pi(t) \approx t + \frac{t^2}{2}(\pi''_A(\bar{u}) + \pi''_B(\bar{u})) .$$

Finally, from (22), total welfare ($W = \Pi + U$) is given by expression (14).

References

- ACQUISTI, A., AND H. VARIAN (2004): “Conditioning Prices on Purchase History,” *Marketing Science*, forthcoming.
- ANDERSON, S., AND L. LERUTH (1993): “Why Firms May Prefer Not to Price Discriminate via Mixed Bundling,” *International Journal of Industrial Organization*, 11(1), 49–61.
- ARMSTRONG, M. (1999): “Price Discrimination by a Many-Product Firm,” *Review of Economic Studies*, 66(1), 151–168.
- ARMSTRONG, M., AND J. VICKERS (1993): “Price Discrimination, Competition and Regulation,” *Journal of Industrial Economics*, 41(4), 335–360.
- (2001): “Competitive Price Discrimination,” *Rand Journal of Economics*, 32(4), 579–605.
- ASPLUND, M., R. ERIKSSON, AND N. STRAND (2004): “Price Discrimination in Oligopoly: Evidence from Swedish Newspapers,” mimeo.
- BAKOS, Y., AND E. BRYNJOLFSSON (1999): “Bundling Information Goods: Pricing, Profits, and Efficiency,” *Management Science*, 45(12), 1613–1630.
- BERNHEIM, D., AND M. WHINSTON (1986): “Menu Auctions, Resource Allocation, and Economic Influence,” *Quarterly Journal of Economics*, 101(1), 1–31.
- BESTER, H., AND E. PETRAKIS (1996): “Coupons and Oligopolistic Price Discrimination,” *International Journal of Industrial Organization*, 14(2), 227–242.
- BHASKAR, V., AND T. TO (2004): “Is Perfect Price Discrimination Really Efficient? An Analysis of Free Entry,” *Rand Journal of Economics*, 35(4), 762–776.
- BORENSTEIN, S. (1985): “Price Discrimination in Free-Entry Markets,” *Rand Journal of Economics*, 16(3), 380–397.
- CALZOLARI, G., AND A. PAVAN (2005): “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, forthcoming.
- CAMINAL, R., AND C. MATUTES (1990): “Endogenous Switching Costs in a Duopoly Model,” *International Journal of Industrial Organization*, 8(3), 353–373.
- CARBAJO, J., D. DE MEZA, AND D. SEIDMAN (1990): “A Strategic Motivation for Commodity Bundling,” *Journal of Industrial Economics*, 38(3), 283–298.
- CHEN, Y. (1997a): “Equilibrium Product Bundling,” *Journal of Business*, 70(1), 85–103.

- (1997b): “Paying Customers to Switch,” *Journal of Economics and Management Strategy*, 6(4), 877–897.
- (2004): “Marketing Innovation,” *Journal of Economics and Management Strategy*, forthcoming.
- CORTS, K. (1998): “Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment,” *Rand Journal of Economics*, 29(2), 306–323.
- CRÉMER, J. (1984): “On the Economics of Repeat Buying,” *Rand Journal of Economics*, 15(3), 396–403.
- DOBSON, P., AND M. WATERSON (2005): “Chain-Store Pricing Across Local Markets,” *Journal of Economics and Management Strategy*, 14(1), 93–119.
- ECONOMIDES, N. (1989): “The Desirability of Compatibility in the Absence of Network Externalities,” *American Economic Review*, 71(5), 1165–1181.
- EDGEWORTH, F. Y. (1910): “Applications of Probabilities to Economics II,” *Economic Journal*, 20(79), 441–465.
- ELLISON, G. (2005): “A Model of Add-on Pricing,” *Quarterly Journal of Economics*, 120(2), 585–637.
- ESTEVEZ, R. B. (2004): “Competitive Behaviour-Based Price Discrimination,” PhD thesis, Oxford University.
- FUDENBERG, D., AND J. TIROLE (2000): “Customer Poaching and Brand Switching,” *Rand Journal of Economics*, 31(4), 634–657.
- FUDENBERG, D., AND M. VILLAS-BOAS (2005): “Behaviour-Based Price Discrimination and Customer Recognition,” in *Handbook on Economics and Information Systems*. North-Holland, Amsterdam, forthcoming.
- GANS, J., AND S. KING (2005): “Paying for Loyalty: Product Bundling in Oligopoly,” *Journal of Industrial Economics*, forthcoming.
- HART, O., AND J. TIROLE (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55(4), 509–540.
- (1990): “Vertical Integration and Market Foreclosure,” *Brookings Papers on Economic Activity: Microeconomics*, pp. 205–276.
- KIM, B.-D., M. SHI, AND K. SRINIVASAN (2001): “Reward Programs and Tacit Collusion,” *Marketing Science*, 20(2), 99–120.

- LAFFONT, J.-J., AND D. MARTIMORT (2002): *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton, NJ.
- LIU, Q., AND K. SERFES (2004): “Quality of Information and Oligopolistic Price Discrimination,” *Journal of Economics and Management Strategy*, 13(4), 671–702.
- MATUTES, C., AND P. REGIBEAU (1988): “Mix and Match: Product Compatibility without Network Externalities,” *Rand Journal of Economics*, 19(2), 221–234.
- (1992): “Compatibility and Bundling of Complementary Goods in a Duopoly,” *Journal of Industrial Economics*, 40(1), 37–54.
- MCAFEE, R. P., AND M. SCHWARTZ (1994): “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity and Uniformity,” *American Economic Review*, 84(1), 210–230.
- MIRAVETE, E., AND L.-H. RÖLLER (2004): “Competitive Nonlinear Pricing in Duopoly Equilibrium: the Early U.S. Cellular Telephone industry,” mimeo.
- MOTTA, M. (2004): *Competition Policy: Theory and Practice*. Cambridge University Press, Cambridge.
- NAHATA, B., K. OSTASZEWSKI, AND P. SAHOO (1990): “Direction of Price Changes in Third-Degree Price Discrimination,” *American Economic Review*, 80(5), 1254–1258.
- NALEBUFF, B. (2000): “Competing Against Bundles,” in *Incentives, Organization, and Public Economics*, ed. by G. Myles, and P. Hammond. Oxford University Press.
- (2004): “Bundling as an Entry Barrier,” *Quarterly Journal of Economics*, 119(1), 159–188.
- NEVO, A., AND C. WOLFRAM (2002): “Why do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals,” *Rand Journal of Economics*, 33(2), 319–339.
- O’BRIEN, D., AND G. SHAFFER (1992): “Vertical Control with Bilateral Contracts,” *Rand Journal of Economics*, 23(3), 299–308.
- ODLYZKO, A. (2003): “Privacy, Economics, and Price Discrimination on the Internet,” *ICEC2003: Fifth International Conference on Electronic Commerce*, pp. 355–366.
- REY, P., AND J. TIROLE (2005): “A Primer on Foreclosure,” in *Handbook of Industrial Organization: Volume III*, ed. by M. Armstrong, and R. Porter. North-Holland, Amsterdam, forthcoming.
- ROCHET, J.-C., AND L. STOLE (2002): “Nonlinear Pricing with Random Participation,” *Review of Economic Studies*, 69(1), 277–311.

- SHAFFER, G., AND J. ZHANG (1995): “Competitive Coupon Targeting,” *Marketing Science*, 14(4), 395–415.
- (2000): “Pay to Switch or Pay to Stay: Preference-Based Price Discrimination in Markets with Switching Costs,” *Journal of Economics and Management Strategy*, 9(3), 397–424.
- SPULBER, D. (1979): “Noncooperative Equilibrium with Price Discriminating Firms,” *Economics Letters*, 4, 221–227.
- STIGLER, G. (1987): *Theory of Price*. Macmillan, New York.
- STOLE, L. (2003): “Price Discrimination and Imperfect Competition,” in *Handbook of Industrial Organization: Volume III*, ed. by M. Armstrong, and R. Porter. North-Holland, Amsterdam, forthcoming.
- TAYLOR, C. (2003): “Supplier Surfing: Competition and Consumer Behavior in Subscription Markets,” *Rand Journal of Economics*, 34(2), 223–246.
- (2004): “Consumer Privacy and the Market for Customer Information,” *Rand Journal of Economics*, 35(4), 631–650.
- THISSE, J.-F., AND X. VIVES (1988): “On the Strategic Choice of Spatial Price Policy,” *American Economic Review*, 78(1), 122–137.
- VARIAN, H. (1980): “A Model of Sales,” *American Economic Review*, 70(4), 651–659.
- VICKERS, J. (2005): “Abuse of Market Power,” *Economic Journal*, 115, F244–F261.
- VILLAS-BOAS, J. M. (1999): “Dynamic Competition with Customer Recognition,” *Rand Journal of Economics*, 30(4), 604–631.
- (2004): “Price Cycles in Markets with Customer Recognition,” *Rand Journal of Economics*, 35(3), 486–501.
- WHINSTON, M. (1990): “Tying, Foreclosure and Exclusion,” *American Economic Review*, 80(4), 837–859.