RAND Journal of Economics
Vol. 32, No. 4, Winter 2001
pp. 1–27

Competitive price discrimination

Mark Armstrong*

and

John Vickers**

We model firms as supplying utility directly to consumers. The equilibrium outcome of competition in utility space depends on the relationship $\pi(\mathrm{u})$ *between profit and average utility per consumer. Public policy constraints on the "deals" firms may offer affect equilibrium outcomes via their effect on* $\pi(\mathrm{u})$. *From this perspective we examine the profit, utility, and welfare implications of price discrimination policies in an oligopolistic framework. We also show that an equilibrium outcome of competitive nonlinear pricing when consumers have private information about their tastes is for firms to offer efficient two-part tariffs.*

1. Introduction

■ This article has two main aims. First, we propose a framework that we hope will be useful in analyzing situations where firms compete by offering various "deals" involving tariffs, bundles of outputs, and so on. Models of markets where firms simultaneously compete on a number of dimensions are often hard to work with. In many cases, however, this analysis can be simplified by taking a dual approach and modelling firms as competing directly in *utility* space. Viewed from this perspective, a firm's strategy is just to choose the (scalar) level of utility, or "value for money," offered to customers. A firm's market share naturally increases the more utility it offers relative to its rivals. Central to the analysis of competition in utility space is the function $\pi(u)$, which gives the maximum profit per consumer a firm can extract when it provides utility u to its consumers. This function succinctly captures not just the relevant cost and demand information; it also reflects any public policy (or other) constraints on the deals firms may offer. The second aim of the article is to illustrate the usefulness of this dual approach by examining the profit, utility, and welfare implications of public policy toward various kinds of price discrimination in oligopoly.[1]

* Nuffield College, Oxford; mark.armstrong@nuffield.oxford.ac.uk.

** All Souls College, Oxford; john.vickers@economics.oxford.ac.uk.

We are grateful for comments from Simon Anderson, Frank Fisher, Anke Kessler, Preston McAfee, John Moore, Babu Nahata, Barry Nalebuff, and Dick Schmalensee, and we especially thank Lars Stole and two referees for very detailed advice that improved the clarity and accuracy of this article. All remaining errors are our own.

[1] The modelling of firms as competing in utility space is also a useful way to think about competition in network industries. For instance, see Armstrong (forthcoming) for a survey of recent contributions to this topic, viewed from the perspective of firms offering utility directly to their subscribers.

□ **Competition in utility space.** We suppose that consumers differ in their tastes for the utilities supplied by the firms (although the framework continues to be valid in the extreme case of perfect substitutes). We do this by using a discrete-choice framework, so that if u_i is the utility supplied by firm i, a consumer's net utility is u_i plus an idiosyncratic shock. (Thus we can think of u_i as the "average" utility offered by firm i to its population of consumers.) This framework includes such familiar models of product differentiation as the Hotelling location model and the logit model. As well as being convenient to work with from a theoretical perspective, such a framework accords well with the simplest forms of many econometric models of consumer choice.

A convenient feature of this approach is that the effect of the profit function in equilibrium is completely separable from the effect of the discrete-choice model. For instance, in a symmetric model with a fixed number of firms, we show that equilibrium utility maximizes an additively separable function

$$\phi(u) + \log \pi(u),$$

where ϕ is an increasing function depending only on the discrete-choice framework (see Proposition 1 below). Similarly, in a symmetric circular Hotelling model with free entry, we show that equilibrium utility maximizes

$$\frac{u}{\tau^2} - \frac{1}{\pi(u)},$$

where τ is a parameter measuring the competitiveness of the market. Thus one can perform comparative statics analysis on the profit function and on the discrete-choice framework independently, and in the article we do this repeatedly. This analysis is contained in Section 2.

Perhaps the article closest to our competition-in-utility-space approach is Bliss's (1988) model of retail competition, in which the role of a shop is to assemble a collection of goods in one place and offer a price list to consumers. The value of a price list to a consumer is measured by an indirect utility function. A recurrent theme in Bliss's analysis is that retailers choose Ramsey-like markups. This is because a firm's problem can often be decomposed into (i) how much utility to offer to consumers, and (ii) how to set prices to maximize profit while offering a given consumer utility. Whereas Bliss focused on (ii), which is a Ramsey problem, our main concern is with (i), though (ii) is implicit in our analysis.

While our competition-in-utility-space method is perhaps less familiar, we embed this approach into the well-known discrete-choice framework for consumer choice—see Anderson, de Palma, and Thisse (1992) for a description of discrete-choice models, and Caplin and Nalebuff (1991) for a general model of imperfect competition that includes the discrete-choice model as a special case. Although there are exceptions, most discrete-choice models involve, first, an assumption that product characteristics are exogenously fixed before competition takes place, and, second, that consumers have unit demands. In particular, a firm's only decision is the price at which it offers its single product, while a consumer's only decision is which firm (if any) to buy from. The latter assumption is too strong for many applications of interest—for instance, situations where firms offer multiple products and where consumers may wish to purchase multiple units of a given product—and one aim of the article is to extend the discrete-choice framework to cover multiproduct, multiunit situations.

Caplin and Nalebuff (1991) and Anderson, de Palma, and Nesterov (1995) provide an earlier analysis of elastic consumer demand, though the focus is still on single-product firms. These analyses make the simplifying assumption that consumers make *all* their purchases from a single firm, and we shall follow the same approach here.[2] This "one-stop-shopping" assumption is not innocuous, and in future work it would be useful to analyze the implications of allowing consumers

[2] Rochet and Stole (forthcoming) have consumers choosing *qualities* rather than quantities, although formally the analysis is similar. However, the interpretation of the choice as quality—so that consumers have unit demands—has the benefit that there is no need to make an assumption that consumers purchase from a single firm.

to pick and choose from the various products offered by firms to make up their own bundle. It is certainly true in much of our analysis that in equilibrium, consumers *will* choose to obtain all products from a single firm, since that way they incur the transport cost only once (in the Hotelling context) and pay the fixed charge of a two-part tariff only once (when such tariffs are offered). However, it does not follow from this that the ability of consumers to shop from two firms has no effect on the equilibrium. (This issue is discussed further in Section 4.)

Once we allow for elastic consumer demand, we have to decide how the additive shock that underpins the discrete-choice model interacts with a consumer's volume of demand. For instance, in a Hotelling transport cost model, one could imagine incurring the transport cost per unit of product consumed, or, alternatively, the cost could just be incurred per trip, independently of the volume of consumption. Because it greatly simplifies the analysis, we take the latter approach.[3] However, our model of price discrimination in Section 4 could be interpreted as one in which consumers pay a transport cost *per product* (but not per unit of each product).[4]

□ **Competitive price discrimination.** The second main aim of the article is to use this competition-in-utility-space framework to investigate the profit, consumer surplus, and welfare implications of various price discrimination policies in oligopoly. We define a firm's pricing to be *discriminatory* if it sells units of its output(s) at different prices when marginal costs do not differ correspondingly. We will see that the welfare effects of permitting price discrimination in oligopoly depend in large part on whether discrimination takes the form of what we term *intrapersonal* discrimination or whether it is *interpersonal* discrimination. To illustrate the distinction, consider the case of air travel. Different kinds of tickets, say "business" or "leisure" tickets, can be charged at different rates, depending on the consumer demand characteristics. It may be that travellers are segmented, so that some people always want one kind of ticket while others always want the other. When these tickets have different prices we would call this practice interpersonal price discrimination, since different consumers face different prices. Alternatively, it could be that a given traveller demands a *mix* of business and leisure tickets, depending on the nature of the planned journeys, and differential charging here is an example of intrapersonal price discrimination.[5]

We analyze three information structures:

Homogeneous case. In this case, consumers are assumed to be "homogeneous," in the sense that all consumers generate the same level of profits for a given level of utility, i.e., that the function $\pi(u)$ is the same for all consumers. (This is a slight abuse of terminology, however, since consumers continue to differ in their tastes for the utilities offered by firms.)

Observable heterogeneity case. Here consumers are heterogeneous and differ by an observable characteristic on which firms can base their tariffs. Examples might be where consumers are regionally segmented, or where the age or occupation of the consumer can be determined.

Unobservable heterogeneity case. Consumers are heterogeneous and differ by unobservable characteristics (in addition to their differing tastes for the utilities offered by firms). Thus the profit function $\pi(u \mid \theta)$ here depends on a parameter θ that is private information to consumers.

In Section 3 we discuss in more detail the simple case of homogeneous consumers who make all their purchases from a single firm, and here we can apply our competition-in-utility-space framework directly. Clearly, in the homogeneous consumer case, the only possible kind of

[3] The same assumption is made in Anderson, de Palma, and Nesterov (1995). Similarly, in their model of quality choice, Rochet and Stole (forthcoming) assume that the transport cost does not depend upon a consumer's choice of quality.

[4] See Stole (1995), and the references therein, for an analysis of the more complicated case where transport costs (or brand preferences) do interact with the quantity/quality choice.

[5] It is fair to say that many economists would reserve the term "price discrimination" for what we call interpersonal price discrimination. For instance, Varian (1989) defines (third-degree) discrimination to be where different consumers are charged different prices.

price discrimination is of the intrapersonal variety. The key idea is that a more constrained firm will be able to make less profit for a given utility level than would a less constrained firm. Thus, the function $\pi_D(u)$ if a firm practices price discrimination of some kind will generally differ from the function $\pi_N(u)$ if it cannot: indeed, $\pi_D(u) \geq \pi_N(u)$ with strict inequality almost always. Both nonlinear pricing—which, in the homogeneous case, reduces to two-part tariffs—and third-degree forms of discrimination are discussed, although the analysis is more general than this and covers virtually any constraints placed on firms in the way they can offer deals to their customers.

When there are fixed costs associated with serving consumers, so that marginal-cost pricing of services does not allow firms to break even, greater flexibility allows firms to offer a higher level of utility before losses are made. In such cases, allowing firms in competitive markets greater flexibility causes an increase in consumer utility and overall welfare (Proposition 2). If there is no such fixed cost, then the analysis required is more delicate, and we show that greater flexibility causes profits and overall welfare to increase in competitive markets, although now consumers may be made worse off (Proposition 3). The results for welfare are true both when there is a fixed number of firms and when there is free entry. The key feature of the homogeneous consumer framework, combined with the one-stop-shopping assumption, is that firms, when permitted to do so, choose socially desirable prices—i.e., profits are maximized subject to a consumer utility constraint. Thus, firms will make their prices relate to the underlying costs and to the demand elasticities in the correct Ramsey manner. In particular, (i) when two-part tariffs are used then marginal prices equal marginal costs, and (ii) when linear pricing is used, then prices are, roughly speaking, higher for those products with less elastic demand.

Section 4 covers cases where consumers are heterogeneous, and we analyze both observable and unobservable heterogeneity. The focus here is naturally on interpersonal price discrimination. We cannot *directly* apply the competition-in-utility-space framework here, although various features of that model will be useful ingredients in the analysis. In the first part of Section 4 we discuss the case of interpersonal third-degree discrimination, and in contrast with the homogeneous consumer case, the welfare consequences of competitive price discrimination are ambiguous. For instance, it may be that competition forces firms to set a *lower* price in less elastic markets compared to the no-discrimination case. More generally, it will only be by chance that prices are Ramsey-like. There is thus a fundamental difference between the "intrapersonal" and "interpersonal" forms of discrimination in terms of their welfare effects. But although the total welfare effect is ambiguous, the effect on consumers and profits separately is not: profits increase but aggregate consumer utility decreases if discrimination is permitted in competitive markets (Proposition 4).

In the second part of Section 4 we discuss a model with unobserved heterogeneity, where consumers have private, product-specific information about their tastes (in addition to their private information about their tastes for the deals offered by firms). We make two important assumptions. First, we assume that the market is competitive, in the sense that all consumers choose to participate in the market in equilibrium. Second, we assume that the "vertical" taste parameter does not interact with the "horizontal" taste parameter. (This assumption has two parts: the two taste parameters are statistically independently distributed, and, as discussed above, the horizontal taste parameter is unaffected by the quantities chosen by consumers.) Under these conditions it is an equilibrium for firms to offer cost-based two-part tariffs (Proposition 5). In particular, there is no screening in this equilibrium, and all consumers face the same marginal prices. In effect, the need to compete dominates any benefit derived from screening consumers. Thus, just as in the homogeneous consumer case, allowing firms to have maximum flexibility in tariff choice can induce an efficient outcome.

□ **Contrasts with monopoly price discrimination.** The above results contrast with familiar results for monopoly; see Varian (1989) for an analysis of price discrimination with monopoly. A general disadvantage of price discrimination—both in monopoly and in oligopoly—is that output is suboptimally distributed to consumers because their marginal utilities will be unequal. This drawback may be overcome if price discrimination leads to a sufficient increase in output. We can deduce that if output falls when discrimination is introduced, then welfare must also fall. (When we discuss interpersonal third-degree discrimination in duopoly, we shall use this fact to find conditions under which discrimination is socially undesirable.) For instance, in the case of third-degree price discrimination, with independent linear demands (and all markets served), total output is exactly the same with discrimination as without, so discrimination reduces welfare. More generally, the sign of the total output effect in the monopoly case can be related to the convexity/concavity properties of demands—see Varian (1989). The monopoly model of price discrimination does not distinguish between the intrapersonal and the interpersonal kinds of discrimination, since all that matters for the analysis is aggregate demand. As discussed above, though, we argue that this distinction makes a crucial difference to the desirability of discrimination in oligopoly, and in the homogeneous consumer case we argue that price discrimination is unambiguously good for welfare in competitive markets. In particular, we can simply *infer* without explicit analysis that the output effect is positive if discrimination is permitted.

Turning next to price discrimination when there is unobservable heterogeneity, Armstrong (1996) and Rochet and Choné (1998) analyze multiproduct nonlinear pricing by a monopolist facing consumers with private information about their tastes. It is argued that the problem of designing the profit-maximizing nonlinear tariff is complex, and explicit solutions seem obtainable in only a few cases. In particular, Armstrong shows that a monopolist will generally choose to exclude some consumers from the market at the optimum, while Rochet and Choné show that for those consumers who do participate there will be "bunching," i.e., that some consumers with different tastes will choose to purchase the same bundle of products. In practice, these two features of the optimum make it hard to find closed-form solutions. By contrast, in Section 4 we shall show that these same models often have a very straightforward equilibrium—without exclusion or bunching—when there is a reasonable degree of competition.

□ **Other work on price discrimination in oligopoly.** The literature on price discrimination in competitive settings is not as extensive as the analysis for monopoly, though see Varian (1989) and Wilson (1993) for some references. One very relevant article, however, is Holmes (1989), who examines the output and profit effects of (interpersonal) third-degree price discrimination in a symmetric duopoly model with product differentiation. (We examine a very similar model in Section 4.) He shows that the sign of the total output effect depends not only on the convexity/concavity properties of industry demands—as in the monopoly case—but also upon a term involving interfirm cross-price elasticities, which reflect competitive interactions. Holmes also shows that it is possible for equilibrium profit to be higher when price discrimination is banned. However, few clear-cut results are derived, and little is said about the welfare effects of allowing discrimination. In Section 4, however, we shall argue that as markets become very competitive the Holmes model becomes easier to interpret and stronger results emerge; we show that the effect on profits and on aggregate consumer welfare is unambiguous, and we provide a simple criterion for predicting when the total output effect is negative (which implies that discrimination is undesirable).[6]

Turning to nonlinear pricing in oligopoly, Stole (1995) considers—as we shall—nonlinear pricing in oligopoly when firms are spatially differentiated. Consumers differ both in their brand preferences (horizontal differentiation) and their quality preferences (vertical differentiation). He characterizes optimal pricing schedules and shows how they depend on consumers' private information about brand and quality preferences, and on the intensity of competition. A major difference between Stole's model and ours is that he allows the two kinds of private information to interact, in that transport costs are assumed to depend on the quantities consumed (and hence on a consumer's taste for the output), whereas we make the simplifying assumption that transport costs are a lump-sum cost. A model very similar to ours is Rochet and Stole (forthcoming), which was prepared independently of this work. They assume the same lump-sum transport cost technology as we do, and find conditions implying that efficient two-part tariffs emerge as an equilibrium.

[6] See also Katz (1984) and Borenstein (1985) for earlier analyses in a monopolistic competition framework.

□ **Customer poaching.** Finally, we mention one important kind of price discrimination that our framework cannot easily handle, when a firm offers a deal to those customers who "belong" (in some sense) to rival firms that differs from the deal it offers its "own" customers. In most cases we expect the firm to offer a *better* deal to rival firms' customers in order to build market share. A feature of such models is that, compared to the situation where deals are uniform, price discrimination can act to *intensify* competition, in the sense that *all* prices come down and equilibrium profits decrease.[7] (This is in direct contrast with, say, our model in the first part of Section 4, where allowing price discrimination causes some prices to rise and some to fall, while total profits increase.) The basic reason why this kind of discrimination intensifies competition across the board is that it makes it cheaper for a firm to target its rivals' customers without damaging the profits it can extract from its own customer base: when all firms target each others' customers, though, all prices come down.[8]

In our model there is no sense in which customers "belong" in any observable way to one or another firm. However, the kinds of situations where customers are associated with particular firms include the following: First, in a spatial model Thisse and Vives (1988) suppose that a consumer's location is observable, so that consumers are known to be more loyal to one firm than the other. In this case, when firms can base their prices on consumer location, all prices might fall compared to the case where prices are constrained to be uniform. Second, Corts (1998) examines a market where one firm supplies a higher-quality product than the other. There are two identifiably different consumer groups, one that cares about quality and one that cares only about price. Corts shows that allowing discrimination could cause both prices to fall compared to the uniform price regime. The reason is that the two firms behave differently in response to the ability to discriminate: the high-quality firm will choose to set the higher price in the "choosy" market, whereas the low-quality firm sets its low price in that market, and so competition is stronger in both markets. Finally, there is the interesting recent literature on customer poaching, including Villas-Boas (1999) and Fudenberg and Tirole (2000). Here, the focus is on the dynamics of market share. Firms know whether a consumer was a previous customer and so can target other consumers accordingly. Again, a result is that all prices might fall with this kind of discrimination.[9]

2. Competition in utility space

■ **A framework.** There are two symmetrically placed firms, A and B, facing differentiated consumers. (This analysis applies equally to any fixed number of firms—see below for the extension to endogenous numbers of firms.[10]) Consumers buy all products from one or the other firm, or they consume an outside option. Suppose that the "deal" offered by firm i gives each consumer a utility u_i, and that the outside option gives utility u_0. We use a discrete-choice framework and assume that a consumer of type $\zeta = (\zeta_0, \zeta_A, \zeta_B)$ has net utility $u_i + \zeta_i$ if he buys from firm i, and net utility $u_0 + \zeta_0$ if no purchase is made. Therefore, a consumer's maximum utility is $\max\{u_A + \zeta_A, u_B + \zeta_B, u_0 + \zeta_0\}$. The triplet $\zeta = (\zeta_0, \zeta_A, \zeta_B)$ is distributed across the population according to a known distribution. There is a continuum of potential consumers, with measure normalized to one.

Aggregate consumer utility is given by

$$V(u_A, u_B) \equiv E\left[\max\{u_A + \zeta_A, u_B + \zeta_B, u_0 + \zeta_0\}\right],$$

[7] In fact, Holmes (1989) shows that the same effect can occur when consumers do not "belong" to particular firms. Moreover, Nahata, Ostaszewski, and Sahoo (1990) show that when profit functions are not single-peaked, the same effect can occur even in monopoly.

[8] Armstrong and Vickers (1993) present a model where price discrimination can cause all prices to *rise*. An incumbent multimarket firm faces a threat of entry in one market. Permitting discrimination makes the incumbent react aggressively to entry—since it does not have to sacrifice profits in the captive market—which can deter entry.

[9] For instance, Villas-Boas has infinitely-lived firms facing an overlapping sequence of two-period consumers, and he shows that in the steady state, prices can be lower when there is "customer recognition" than without.

[10] The only result that depends on duopoly seems to be the fact that we can rule out asymmetric equilibria in Proposition 1 below. See Perloff and Salop (1985) for a discussion of the issue.

where E denotes expectation with respect to the distribution for ζ. By the envelope theorem the function V has partial derivative[11]

$$V_1(u_A, u_B) \equiv s(u_A, u_B),$$

where $s(u_A, u_B)$ is the number of consumers who choose to buy from firm A given the pair of gross utilities u_A and u_B. (Differentiating V with respect to u_A has two effects: (i) the direct effect on those consumers who buy from A anyway, which is just $s(u_A, u_B)$, and (ii) a shift in the set of consumers who choose to buy from A. However, those consumers on the boundary in $(\zeta_0, \zeta_A, \zeta_B)$–space between buying from A and other options are indifferent between A and the other options, so there is no aggregate consumer welfare effect from (ii) at the margin.)

We focus throughout on the case where consumers' tastes for the two firms' products are symmetrically distributed, so that $V(u_A, u_B) \equiv V(u_B, u_A)$. This implies that the number of B's consumers is just $s(u_B, u_A)$. Clearly it is necessary that $s(u_A, u_B) + s(u_B, u_A) \leq 1$, with equality if all consumers participate. It is also necessary that s is increasing in the first argument and decreasing in the second.[12]

Suppose that if a firm offers consumers a "deal"—i.e., package of products for sale using a certain tariff—that gives them utility u, then its maximum profit per consumer is $\pi(u)$. For this function to be well defined, we assume the firms have constant-returns-to-scale technology in serving consumers, so that profit per consumer does not depend on the number of consumers served. Until we reach Section 4, we assume consumers are homogeneous in the sense that the function $\pi(u)$ is the same for all consumers. The form the relationship $\pi(u)$ takes depends on the cost structure of the firms, on the consumer demand functions they face, and also on the types of tariff they are permitted to offer—see the next subsection for examples. Using this notation, firm $i = A, B$ wishes to maximize $s(u_i, u_j)\pi(u_i)$ given the rival's choice of utility u_j.

We need to make some regularity assumptions in order to guarantee existence of equilibrium. First, we impose the following condition on s:[13]

$$\frac{s_1(u_A, u_B)}{s(u_A, u_B)} \text{ is weakly increasing in } u_B. \tag{1}$$

Since s is decreasing in the second argument, this condition requires that s_1 not be decreasing too fast in the second argument. In economic terms, this implies that firms' reaction functions in utility space are upward sloping. Second, we need to ensure that very low utility levels ultimately drive all consumers away. To be precise, we assume that the "collusive" utility level u^*, which is defined to be the utility that maximizes (symmetric) joint profits, exists. More specifically, we assume

$$\text{there exists } u^* > -\infty \text{ that maximizes } s(u, u)\pi(u). \tag{2}$$

Finally, it is economically natural to assume that $\pi(u)$ is ultimately negative for large values of u, so let $\bar{u}$ be the largest utility level that allows a firm to break even. (For instance, supplying very high utility ultimately involves offering products at prices below costs.) Therefore, assume there exists $\bar{u}$ defined by

$$\pi(\bar{u}) = 0, \pi(u) < 0 \text{ if } u > \bar{u}. \tag{3}$$

[11] We assume that this derivative exists, which requires that there is only a measure-zero set of potential consumers who can be indifferent between any pair of options. A sufficient condition for this is for ζ to be distributed according to a bounded density function.

[12] To see this algebraically, one can show, for example, that the set of consumers who buy from A contracts if u_B is increased.

[13] For s to be differentiable, we can assume that the density function for ζ is continuous. In fact, the Hotelling example we often use has kinked market-share functions. However, when all consumers participate in the market—as is the case below—this problem does not arise.

In this framework a symmetric Nash equilibrium in utility space exists, the equilibrium is often unique, and, most importantly for our purposes, there is a simple formula for the equilibrium utility:

Proposition 1. Assume (1)–(3) hold. Let $\phi(u)$ be a function satisfying

$$\phi'(u) = \frac{s_1(u,u)}{s(u,u)} \geq 0. \tag{4}$$

(This is unique up to a constant.) Then a symmetric equilibrium is given by each firm choosing utility $\hat{u}$ that maximizes

$$\phi(u) + \log \pi(u). \tag{5}$$

This equilibrium utility $\hat{u}$ is strictly greater than the collusive utility level u^*. There are no asymmetric equilibria. This is the unique (pure-strategy) equilibrium if (5) is concave.

Proof. All proofs are in the Appendix.

Thus equilibrium in this market involves the maximization of the sum of log profits and an increasing function ϕ of consumer utility. In this result, the term s_1/s in (4) gives the proportional increase in a firm's customer base caused by unilaterally raising utility when both firms offer the same utility level. This is a natural measure of the degree of product substitutability in the market. The larger the term, the more weight is placed on consumer utility compared to profits in the equilibrium.[14]

□ **Examples of profit functions $\pi(u)$.** In this section we list a few examples of $\pi(u)$ derived from primitives in the market. In all cases except the last these will be necessary ingredients in the subsequent analysis of price discrimination. In all cases we assume consumer preferences are quasi-linear in income.

Single-product linear pricing. Suppose firms offer a single product and each consumer has the demand function $Q(P)$. Let $V(P)$ be the consumer surplus function associated with $Q(P)$, so that $V'(P) \equiv -Q(P)$. (Recall that preferences are assumed to be quasi-linear.) Therefore, a consumer's utility is just $u = V(P)$. Each firm has a constant marginal cost c for providing the service, together with a fixed cost per consumer for providing service to a consumer of $k \geq 0$.[15] A firm's profit per consumer if it offers utility $u = V(P)$ is therefore $Q(P)(P-c) - k$, so that $\pi(u)$ is defined implicitly by

$$\pi(V(P)) \equiv Q(P)(P-c) - k. \tag{6}$$

If $k = 0$, then the maximum break-even level of utility, which defines $\bar{u}$ in (3), corresponds to marginal-cost pricing, i.e., $\bar{u} = V(c)$. On the other hand, if $k > 0$, then the maximum break-even level of utility is the one corresponding to average-cost pricing, i.e., if $P^{av} > c$ is the lowest price that allows a firm to break even with linear pricing then $\bar{u} = V(P^{av})$. Differentiating (6) yields

$$\pi'(u) = -(1 - \sigma(u)), \tag{7}$$

where

$$\sigma(V(P)) \equiv (P-c)\frac{-Q'(P)}{Q(P)}.$$

[14] See Bergstrom and Varian (1985) and Slade (1994) for an analysis of when an oligopoly in equilibrium can be considered as maximizing a single objective function (as here). More generally, the duopoly game described in Proposition 1 is a "potential game"—see Monderer and Shapley (1996) for details.

[15] These per-consumer fixed costs are often relevant, as with the costs of handsets in mobile telephony, the costs of laying cables into homes in fixed telephony, the cost of digital decoders in pay-TV markets, or even the costs of parking spaces and checkout staff in shops.

(The function σ represents the elasticity of demand at a given utility level expressed in terms of the *markup* $(P - c)$ instead of price P.) If the weak condition $\sigma'(u) \leq 0$ holds when $u \leq V(c)$, then (7) implies that $\pi(u)$ is concave in u when $\pi(u) \geq 0$.[16] It will be useful for later to note that when $k = 0$, so that $\bar{u} = V(c)$, (7) implies that

$$\pi'(\bar{u}) = -1. \tag{8}$$

Multiproduct linear pricing. Suppose, generalizing the above, that firms sell n products. Consumer demands may be interdependent across products, and if $\mathbf{p} = (p_1, \ldots, p_n)$ is the vector of prices, then each consumer's demand for product i is $q_i(\mathbf{p})$. A consumer's surplus is $v(\mathbf{p})$, where $\partial v/\partial p_i \equiv -q_i$. Firms incur a fixed cost $k \geq 0$ of supplying a customer, and a vector of constant marginal costs $\mathbf{c} = (c_1, \ldots, c_n)$ for supplying the n products.

In this framework, $\pi(u)$ is given by the solution to the Ramsey problem:

$$\pi(u) = \max_{\mathbf{p} \geq 0} : \left\{ \sum_{i=1}^{n} q_i(\mathbf{p})(p_i - c_i) - k \mid v(\mathbf{p}) = u \right\}. \tag{9}$$

If $k = 0$, then $\bar{u}$, the maximum utility that allows a firm to break even, is again induced by marginal-cost pricing, i.e., $\bar{u} = v(\mathbf{c})$. On the other hand, if $k > 0$, then pricing at marginal cost will not cover costs, and $\bar{u}$ will be lower. Indeed, unless at least one product has perfectly inelastic demand (which we rule out from now on), there will be deadweight losses and $\bar{u} < v(\mathbf{c}) - k$. In sum,

$$\bar{u} \begin{cases} = v(\mathbf{c}) & \text{if } k = 0 \\ < v(\mathbf{c}) - k & \text{if } k > 0. \end{cases} \tag{10}$$

For a given u let $\lambda(u)$ be the Lagrange multiplier used to solve problem (9).[17] For any $\lambda \leq 1$ let $\hat{\mathbf{p}}(\lambda)$ be the price vector that maximizes the function

$$\sum_{i=1}^{n} q_i(\mathbf{p})(p_i - c_i) + \lambda v(\mathbf{p}). \tag{11}$$

In particular, $\hat{p}_i(1) = c_i$. The price vector that solves problem (9) is therefore $\hat{\mathbf{p}}(\lambda(u))$, where $\lambda(u)$ is defined implicitly by

$$v(\hat{\mathbf{p}}(\lambda(u))) \equiv u. \tag{12}$$

The envelope theorem implies that

$$\pi'(u) = -\lambda(u). \tag{13}$$

Note that $\lambda(u)$ is increasing, so that $\pi(u)$ is concave.[18] Suppose that $k = 0$, so that $\bar{u} = v(\mathbf{c})$. Then marginal-cost pricing (which corresponds to $\lambda = 1$) is the policy that generates utility $\bar{u}$, i.e.,

$$\pi'(\bar{u}) = -\lambda(\bar{u}) = -1. \tag{14}$$

Two-part tariffs. Consider the same framework as above but allow firms to use two-part tariffs instead of linear prices. Suppose a firm chooses a vector of marginal prices $\mathbf{p}$ and a fixed charge

[16] This discussion is very closely related to Section 3 of Anderson, de Palma, and Nesterov (1995), where the case of single-product firms competing for consumers with elastic demand is analyzed. In particular, the function η in their assumption A4 is our function σ in the text, i.e., $\sigma(V(P)) \equiv \eta(P)$.

[17] We assume demand functions are well behaved enough so that the Lagrangian method is valid—see Diamond and Mirrlees (1971) for a full discussion of when the Lagrange method works for the Ramsey-unique solution.

[18] To see this, note that the optimized value of (11), viewed as a function of λ, is convex in λ and has derivative equal to $v(\mathbf{p}(\lambda))$. Therefore, $v(\mathbf{p}(\lambda))$ is increasing in λ, which, in combination with (12), implies that $\lambda(u)$ is increasing.

f. Then consumer utility is $u = v(\mathbf{p}) - f$. The profit-maximizing method of generating u is to set prices equal to marginal cost, in which case utility is $u = v(\mathbf{c}) - f$. Profit as a function of u is therefore given by the linear relationship

$$\pi(u) = v(\mathbf{c}) - k - u, \tag{15}$$

and the maximum utility that allows a firm to break even is

$$\bar{u} = v(\mathbf{c}) - k. \tag{16}$$

Comparing this expression with (10), we see that when there are per-consumer fixed costs, the use of two-part tariffs enables a higher level of utility to be achieved compared with linear pricing.

Advertising in media markets. Although the remainder of the article is focused on pricing issues, here we briefly consider how the competition-in-utility-space framework can be applied to the equilibrium balance between advertising and retail charges in media markets such as newspapers and pay-TV. Suppose that consumers obtain utility from consuming the media product equal to

$$u = v - D(a) - p,$$

where v is the gross surplus from consuming the service, a is the volume of advertising, $D(\cdot)$ is the function giving how much consumers dislike advertising, and p is the charge to the consumer for the product. Advertising revenues for a media firm are closely related to their subscriber base, so suppose that if a firm chooses advertising volume a it obtains advertising revenue $R(a)$ per subscriber. Suppose that a firm incurs a cost k per subscriber in providing the service. Therefore, its profit per subscriber is $p + R(a) - k$, which in terms of u is $v - u - k + R(a) - D(a)$, and so its maximum profit per subscriber—if the firm is free to choose the balance between consumer prices and advertising—is given by the function

$$\pi(u) = v - u - k + \max_a \{R(a) - D(a)\},$$

which is of the same linear form as the two-part tariff case above.

□ **Examples of discrete-choice models.** Here we provide two examples of preferences that generate particular instances of the market-share function $s(u_A, u_B)$. (See Anderson, de Palma, and Thisse (1992) for further examples.)

Hotelling preferences. The simplest Hotelling model involves customers being uniformly located on the unit interval [0, 1]. A customer located at $0 \leq x \leq 1$ receives net utility $u_A - tx$ if she buys from A and net utility $u_B - t(1 - x)$ if she buys from B. Thus firm A is located at zero and firm B is located at one. The parameter $t > 0$ is a consumer's "transport cost" per unit distance of buying from one firm or the other. If an agent buys from neither firm, she receives zero net utility.[19] Therefore, a consumer at x will buy from firm A if (i) $u_A - tx \geq u_B - t(1 - x)$ and (ii) $u_A - tx \geq 0$. These inequalities can be rearranged to give

$$x \leq \min : \left\{ \frac{1}{2}\left(1 + \frac{u_A - u_B}{t}\right), \frac{u_A}{t} \right\}.$$

Therefore, s is given by

$$s(u_A, u_B) = \min : \left\{ \frac{1}{2}\left(1 + \frac{u_A - u_B}{t}\right), \frac{u_A}{t} \right\}. \tag{17}$$

[19] Thus, in terms of the notation developed above, $u_0 = 0$, $\zeta_0 \equiv 0$, $\zeta_A = -tx$, and $\zeta_B = -t(1 - x)$, where x is a random variable uniformly distributed on [0, 1].

(There are obvious modifications when u_A and u_B differ by so much that s above lies outside the interval [0, 1], which never occurs in equilibrium.) Unfortunately, the Hotelling model has the well-known technical drawback that market shares are kinked at a point where there exists a consumer who is indifferent between the three options of buying from A, buying from B, or buying nothing. This means that we cannot directly apply Proposition 1. However, the following similar result is easily obtained:

Lemma 1. Suppose that $\log \pi(u)$ is concave and that

$$t \leq 2\beta, \tag{18}$$

where $\beta > 0$ is the value of u that maximizes $(1/2)\log u + \log \pi(u)$. Then equilibrium utility, denoted $\hat{u}(t)$, maximizes

$$\frac{u}{t} + \log \pi(u). \tag{19}$$

Condition (18) implies that if $\hat{u}(t)$ maximizes (19), then $\hat{u}(t) \geq \beta$ and so all consumers are willing to pay the transport cost to obtain this utility, i.e., that the market is covered.

For an illustration of Lemma 1, consider the case where two-part tariffs are used, so that the profit function is given by (15). Then Lemma 1 shows that, provided parameters are such that the market is covered in equilibrium, which in this case reduces to requiring[20]

$$t \leq \frac{2}{3}\left[v(\mathbf{c}) - k\right], \tag{20}$$

in equilibrium we have

$$\hat{\pi}(t) \equiv t, \tag{21}$$

so that a firm's profit per customer is equal to the differentiation parameter t. Although tariffs are not the focus of our approach in this section, it is useful for later to note that the tariff that generates this level of profits is

$$T^*(\mathbf{q}) = t + k + \sum_{i=1}^{n} c_i q_i. \tag{22}$$

This is a two-part tariff with prices equal to marginal costs, and the margin of the fixed charge over the fixed cost is the product-differentiation parameter t. In particular, this equilibrium tariff does not depend on any consumer-demand characteristics, and this feature will play a major role when we discuss tariff design under asymmetric information at the end of the article.

It will be useful to understand the asymptotic properties of the equilibrium as the market becomes highly competitive. (Note that Lemma 1 certainly applies when t is close to zero.) As described above, $\hat{u}(t)$ denotes equilibrium consumer utility with transport cost t. Similarly, let $\hat{\pi}(t) \equiv \pi(\hat{u}(t))$ be equilibrium profits, and let $\hat{w}(t) \equiv \hat{u}(t) + \hat{\pi}(t)$ be unweighted welfare (not including consumers' transport costs which, except in the next subsection, are identical in all equilibria, and hence irrelevant for welfare). Then we have the following result:

Lemma 2. Suppose $\pi(u)$ is at least twice continuously differentiable. In the Hotelling duopoly model for small t, the equilibrium utility, profits, and welfare are given by the following second-order Taylor approximations:

$$\hat{u}(t) \approx \bar{u} - t + \frac{t^2}{2}\left[\frac{\pi''(\bar{u})}{\pi'(\bar{u})}\right]$$

$$\hat{\pi}(t) \approx -t\pi'(\bar{u}) + t^2\pi''(\bar{u})$$

[20] In this case, $\beta = (1/3)[v(\mathbf{c}) - k]$ in Lemma 1.

$$\hat{w}(t) \approx \bar{u} - t(1 + \pi'(\bar{u})) + \frac{t^2}{2}\left[2\pi''(\bar{u}) + \frac{\pi''(\bar{u})}{\pi'(\bar{u})}\right].$$

When we come to discuss the impact in competitive markets of various price discrimination policies, we will use these Taylor expansions—which relate outcomes to the slope and curvature of π at $\bar{u}$—to evaluate a change in the tariff regime in terms of its effect on consumers, firms, and overall welfare.

Logit preferences. With the normalization $u_0 = 0$ the logit model with product-differentiation parameter μ implies that the market-share function takes the form

$$s(u_A, u_B) = \frac{e^{u_A/\mu}}{e^{u_A/\mu} + e^{u_B/\mu} + 1}.$$

(A significant difference between the logit and the Hotelling models is that in the former there are always some consumers who choose not to participate.) Then

$$\frac{s_1(u, u)}{s(u, u)} = \frac{1}{\mu}\frac{e^{u/\mu} + 1}{2e^{u/\mu} + 1},$$

and so Proposition 1 implies that equilibrium utility $\hat{u}$ maximizes

$$\log \pi(u) + \frac{u}{\mu} - \frac{1}{2}\log\left(2e^{u/\mu} + 1\right).$$

(One can check that (1) is satisfied.) When μ is small, this maximization problem is closely approximate to maximizing $\log \pi(u) + u/(2\mu)$, which is the same problem as (19) when $t = 2\mu$. In particular, all the asymptotic results for the Hotelling model contained in Lemma 2 above are valid for the logit model as well if we write $t = 2\mu$.

□ **Free entry.** Here we extend the model to allow for free entry of firms. We need to be more specific about consumer preferences here, and so we use the familiar Salop (1979) extension of the Hotelling model where consumers of mass one are uniformly distributed around a circle of circumference one.[21] Transport cost is t and there is an overall fixed cost per firm K. Let $\pi(u)$ be the same per-consumer profit function as previously. We assume that market equilibrium implies that firms are equally spaced around the circle, and that integer concerns are not significant.

Analogously to Lemma 1 above, equilibrium u for a given number of firms n satisfies

$$\frac{\pi'(u)}{n} + \frac{\pi(u)}{t} = 0. \tag{23}$$

Given equilibrium utility u, the number of firms n is given by the zero-profit condition

$$\pi(u) = nK.$$

We assume that transport cost t is large relative to the fixed cost K so that free-entry equilibrium allows a reasonable number of firms to enter the market. Define τ by $\tau^2 = Kt$. Then the above pair of conditions implies that

$$(\pi(u))^2 + \tau^2\pi'(u) = 0 \tag{24}$$

[21] An alternative approach would be to follow Perloff and Salop (1985) and Anderson, de Palma, and Nesterov (1995) and model endogenous entry by assuming that a consumer's utility shocks ζ_i are independently and identically distributed across all potential firms.

defines free-entry equilibrium level of gross utility as a function of τ, denoted $\hat{u}(\tau)$. Assuming that $\log \pi(u)$ is concave, this implies that $-1/\pi$ is also concave, and so (24) implies that equilibrium $\hat{u}(\tau)$ maximizes the function

$$\frac{u}{\tau^2} - \frac{1}{\pi(u)},$$

which is the free-entry version of (19) above. Clearly $\hat{u}(\tau)$ is decreasing in τ and tends to $\bar{u}$, the maximum break-even level of utility, as $\tau \to 0$. Notice that, like (21) in the duopoly model, when firms offer two-part tariffs the expression (24) implies that equilibrium profits, excluding K, are $\pi(\hat{u}(\tau)) \equiv \tau$.

Since free entry drives profits to zero, total welfare with parameter τ, denoted $\hat{w}(\tau)$, is just net consumer utility, which is utility $\hat{u}(\tau)$ minus average transport costs, and is given by

$$\hat{w}(\tau) = \hat{u}(\tau) - \frac{t}{4n} = \hat{u}(\tau) + \frac{\pi(\hat{u}(\tau))}{4\pi'(\hat{u}(\tau))}, \tag{25}$$

where the second equality comes from (23).

The following lemma, which is analogous to Lemma 2 derived in the duopoly case, describes the asymptotic properties of $\hat{w}$:

Lemma 3. The second-order Taylor expansion of $\hat{w}(\tau)$ about $\tau = 0$ is

$$\hat{w}(\tau) \approx \bar{u} - \frac{5\tau}{4}\frac{1}{\sqrt{-\pi'(\bar{u})}} + \frac{\tau^2}{8}\frac{\pi''(\bar{u})}{(\pi'(\bar{u}))^2}.$$

3. Discrimination with homogeneous consumers

■ For the remainder of the article we examine the welfare effects of allowing firms more discretion over their choice of tariffs. In this section we focus on the case of homogeneous consumers, so that, using the terminology of the Introduction, price discrimination must be of the intrapersonal variety.

Let $\pi_D(u)$ be the profit function if firms can engage in price discrimination in some way, and let $\pi_N(u)$ be the corresponding function if no discrimination is possible. Note that the cause of a change in tariff regime need not be a change in public policy, but could result from technological developments that facilitate more complex tariffs. To illustrate the following discussion, consider two natural kinds of price discrimination that use the profit functions derived in Section 2.

□ **Two-part tariffs versus linear prices.** Here, π_N is the profit function corresponding to multiproduct linear pricing, as in (9), and the break-even level of utility, $\bar{u}_N$, is given by (10). When two-part tariffs are used, the profit function, π_D, is given by (15), and the break-even level of utility, $\bar{u}_D$, is given by (16). We will see that the relative performance of these tariffs in competitive markets depends on whether or not there is a fixed cost k associated with serving a customer. Comparing (16) with (10), we see that when $k > 0$ we have $\bar{u}_D > \bar{u}_N$. In contrast, when $k = 0$, we have $\bar{u}_D = \bar{u}_N = v(\mathbf{c})$, while (14) and (15) give $\pi'_D(\bar{u}) = \pi'_N(\bar{u}) = -1$.

□ **Intrapersonal third-degree price discrimination.** Suppose the multiproduct firms have the same marginal cost c for each product. A ban on price discrimination requires a firm to set a uniform (linear) price $p_i \equiv P$ for all products, in which case the demand function aggregated over all products is $Q(P) \equiv \sum_{i=1}^{n} q_i(P, \ldots, P)$. Let $V(P) \equiv v(P, \ldots, P)$ be consumer surplus at the uniform price P. Then the single-product linear pricing analysis in Section 2 applies, so π_N, the profit function with the uniform pricing constraint, is given by (6). If $k = 0$, then $\bar{u}_N$, the maximum utility that allows break-even without price discrimination, is just $V(c)$, and (8) implies that $\pi'_N(\bar{u}_N) = -1$. If $k > 0$, let P^{av} denote the minimum *uniform* price that allows a firm to break even, in which case $\bar{u}_N = V(P^{av})$.

If firms can offer different prices for different services, then $\pi_D(u)$ is given by (9) above (with $c_j \equiv c$). As above, if $k = 0$, then $\bar{u}_D = V(c)$, and from (14) we see that $\pi'_D(\bar{u}_D) = -1$. If $k > 0$, then—except for the knife-edge case where uniform pricing happens to be the utility-maximizing way to cover the fixed cost k—we have $\bar{u}_D > \bar{u}_N$, since the firm *could* choose to set price P^{av} in each market, but there is generically a more profitable way to generate utility $\bar{u}_N$.

We wish to emphasize, however, that the following argument applies more generally than to this pair of choices, and virtually any pair of profit functions such that π_D corresponds to a situation where firms have more flexibility over their choice of tariff than is the case with π_N will work. For instance, another case might be where π_D corresponds to unrestricted multiproduct linear pricing, and π_N is the function induced when firms are not permitted to set any prices below the associated marginal cost (i.e., where "loss-leading"is banned).[22]

When firms can engage in price discrimination, they have more ways to generate a given level of utility, and so it follows that $\pi_D(u) \geq \pi_N(u)$, where the inequality will often be strict. In particular, if $\bar{u}_D$ and $\bar{u}_N$ are the respective maximum utilities that allow firms to break even, then $\bar{u}_D \geq \bar{u}_N$. The above examples—of linear versus two-part pricing and uniform versus nonuniform pricing—suggest that the inequality $\bar{u}_D \geq \bar{u}_N$ typically is *strict* when there are positive fixed costs of supplying services to a consumer ($k > 0$), for in that case discrimination provides a more efficient means with which to cover fixed costs. In sum,

$$k > 0 \Longrightarrow \bar{u}_D > \bar{u}_N. \tag{26}$$

The fixed per-consumer cost case is illustrated in Figure 1, which applies to the duopoly Hotelling model. From (19), the equilibrium utilities, denoted u_N and u_D in the figure, satisfy $t\pi'_i(u_i) + \pi_i(u_i) = 0$, the tangent to $\pi_i(\cdot)$ at the point u_i intersects the horizontal axis a distance t to the right of equilibrium utility. In fact, because $\pi_D(u)$ is a straight line with $\pi'_D(u) \equiv -1$, this diagram illustrates the case of two-part tariffs versus linear pricing—see (15) above. Therefore equilibrium profits with discrimination in this case are t, as in (21) above. The figure is intended to show what may happen when the market is not particularly competitive (i.e., t is quite large), in which case equilibrium utility is *lower* when discrimination is allowed.

On the other hand, when there are no fixed costs to cover, so that $k = 0$, then we expect that $\bar{u}_N = \bar{u}_D = \bar{u}$, say. The reason for this is that marginal-cost pricing—which is *non*discriminatory—is the most efficient way to supply utility to consumers. Moreover, when $k = 0$, we expect that $\pi'_D(\bar{u}) = \pi'_N(\bar{u}) = -1$. Again, this makes sense from an economic standpoint: when $k = 0$, we know that $\bar{u}$ corresponds to marginal-cost pricing, and since total welfare per consumer, $u + \pi(u)$, is maximized at marginal-cost pricing, it follows that $\pi'(\bar{u}) = -1$. Finally, note that since $\pi_D(u) \geq \pi_N(u)$ and the first two terms in a Taylor expansion about $\bar{u}$ are equal for the two profit functions, i.e., $\pi_D(\bar{u}) = \pi_N(\bar{u}) = 0$ and $\pi'_D(\bar{u}) = \pi'_N(\bar{u}) = -1$, it is necessarily the case that $\pi''_D(\bar{u}) \geq \pi''_N(\bar{u})$. Except in knife-edge situations, this last inequality will be strict.[23] In sum,

$$k = 0 \Longrightarrow \bar{u}_D = \bar{u}_N = \bar{u},\ \pi'_D(\bar{u}) = \pi'_N(\bar{u}) = -1,\ \pi''_D(\bar{u}) > \pi''_N(\bar{u}). \tag{27}$$

Unfortunately, we have mostly been able to obtain general results only in the competitive limit, e.g., as the product-differentiation parameter t or τ is close to zero in the Hotelling models. Results for less competitive markets have been possible to obtain only by specifying the market and the kinds of tariffs in more detail. For instance, by specifying a functional form for consumer demand (such as linear independent demand functions) one can explicitly work out the welfare comparison between nonuniform versus uniform multiproduct linear pricing over the whole range

[22] Alternatively, in the media-markets model above, π_D could be the profit function if firms are unrestricted, and π_N could be the one corresponding to when firms are restricted in some way, for instance in the amount of advertising they can supply.

[23] For instance, in the case of intrapersonal third-degree price discrimination, one can show that $\pi''_D(\bar{u}) = \pi''_N(\bar{u})$ only if all the appropriate elasticities of demand for the n products are *equal* (when evaluated at marginal-cost pricing).

FIGURE 1
DUOPOLY MODEL WITH $k > 0$

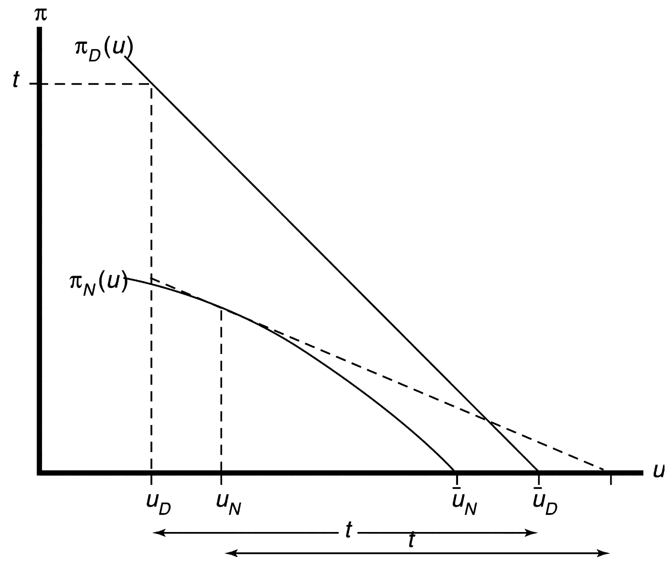

of transport costs.[24] However, one result that holds outside the competitive limit is the following application of monotone statics:

Lemma 4. In the duopoly model, suppose that $\bar{u}_N = \bar{u}_D = \bar{u}$ and $\log \pi_D(u) - \log \pi_N(u)$ is decreasing in u for $u < \bar{u}$. Then equilibrium utility is lower when price discrimination is permitted than when it is not.

Thus we see that if π_D/π_N decreases in u, so that the proportional benefit of engaging in price discrimination is higher when less utility is offered to consumers, in equilibrium consumer utility will fall with discrimination. (Of course, this result says nothing about the effect of discrimination on overall welfare; we will see below that welfare often rises despite a fall in consumer surplus when discrimination is introduced.) One application of this result is to the comparison between linear and two-part pricing.

Corollary 1. Suppose $k = 0$. In the duopoly model, equilibrium utility is lower with two-part tariffs than with linear pricing.

For the remainder of this section we focus on the competitive limit. We use the Hotelling discrete-choice model, in both the duopoly and free-entry frameworks. This is because it is convenient to use a parameterized family of preferences that can allow services to become closely substitutable. But since the logit model closely approximates the Hotelling model in competitive markets—see above—all the following results for competitive markets apply equally with that discrete-choice example as well. A simple corollary of Lemmas 2 and 3 is that in the fixed per-consumer cost case of (26), we have consumer utility and total welfare being higher when discrimination is allowed, provided the market is sufficiently competitive:

Proposition 2 (fixed per-consumer cost case). Suppose $\bar{u}_D > \bar{u}_N$. Then

(i) in the duopoly Hotelling model if there is sufficiently strong competition—t close to zero—both consumer utility and welfare are higher if price discrimination is permitted;

(ii) in the free-entry Hotelling model if there is sufficiently strong competition—τ close to zero—welfare (which equals consumer utility) is higher if price discrimination is permitted.

[24] This was done in an earlier version of the article, and we showed that allowing discrimination was better for overall welfare for a large range of transport costs, not just for those near zero.

Notice that Figure 1 applies to the case where $\bar{u}_D > \bar{u}_N$ and where the market is *not* especially competitive, when consumers are worse off when discrimination is permitted. However, one can see from the diagram that as t becomes small we have equilibrium utility u_D being greater than the maximum possible utility without discrimination, $\bar{u}_N$, and hence greater than the equilibrium utility without discrimination.

The intuition for this result is straightforward: in competitive markets firms can attract consumers only by delivering utility that is close to the maximum possible; restricting the ways in which firms can deliver utility reduces this maximum utility. In the duopoly model, profits are approximately zero in the competitive case, so welfare then is also higher with discrimination.[25] In the free-entry model, profits are always zero, but in the competitive case, transport costs are close to zero (either because the per-unit transport cost t is small or because there are many firms and so the distance travelled is small), and therefore welfare here is also increased by allowing discrimination. This simple result does not depend on the particular Hotelling specification of tastes. Any parameterized model in which equilibrium utility tends to the maximum possible break-even utility as competition becomes more intense will have the same property.[26]

The analysis required is more delicate in the case where there are no per-consumer costs, as in (27). However, we have the following corollary of Lemmas 2 and 3:

Proposition 3 (no fixed per-consumer cost case). Suppose $\bar{u}_N = \bar{u}_D = \bar{u}$, $\pi_N'(\bar{u}) = \pi_D'(\bar{u}) = -1$ and $\pi_D''(\bar{u}) > \pi_N''(\bar{u})$. Then

(i) in the duopoly Hotelling model if there is sufficiently strong competition—t close to zero—consumer utility decreases and profits and welfare increase if price discrimination is permitted;

(ii) in the free entry Hotelling model if there is sufficiently strong competition—τ close to zero—welfare (which equals consumer utility) is higher if price discrimination is permitted.

See Figure 2—which is the same as Figure 1 except that $\bar{u}_N = \bar{u}_D = \bar{u}$—for an illustration of this result in the Hotelling duopoly case. In fact, this figure precisely illustrates the case covered by Corollary 1 comparing linear pricing with two-part tariffs.

Comparing Propositions 2 and 3, notice that, with the duopoly framework, changing one apparently minor feature of the model specification—whether or not there exist per-consumer fixed costs—reverses the ranking of the two tariffs as far as consumers are concerned (though not for overall welfare). A technical reason for this can be found by looking at Lemma 2. When $\bar{u}_D > \bar{u}_N$, the impact of allowing discrimination in competitive markets is immediately seen simply by the limiting value of the functions $\hat{u}$ and $\hat{w}$ as t tends to zero, and no information about the derivatives of the functions π_N and π_D is needed for the comparison. Thus, only "0th-order" information is needed for the comparison. When $\bar{u}_D = \bar{u}_N$, however, not only are the limiting values of π_N and π_D equal, but so are their first derivatives. (Compare (27) with (26).) This means that second-order information on these profit functions is required to be able to understand the impact of price discrimination. Roughly speaking, whenever 0th-order effects are sufficient to compare the two regimes, everyone is better off with discrimination in the competitive limit; when we have to move to second-order comparisons, the ranking is reversed for consumers.

Why, when there are no fixed per-consumer costs, are consumers worse off with discrimination in competitive markets? Because second-order effects are involved, it seems hard to get much direct economic intuition. However, one explanation can be found by looking at Lemma 4. Taking a second-order expansion of $\pi_i(u)$ about $\bar{u}$, one obtains (for $i = N, D$)

$$\pi_i(u) \approx \bar{u} - u + \frac{1}{2}(\bar{u} - u)^2 \pi_i''(\bar{u}).$$

Since we know that $\pi_D''(\bar{u}) > \pi_N''(\bar{u})$, it follows that *locally* around $\bar{u}$ we have that $\log \pi_D - \log \pi_N$

[25] In fact, Lemma 2 shows that in competitive markets, equilibrium profits are higher with discrimination whenever $\pi_D'(\bar{u}_D) < \pi_N'(\bar{u}_N)$, which is certainly true in the two-part tariff versus linear-pricing comparison.

[26] See Perloff and Salop (1985) for a fairly general analysis of when this happens with unit demands.

FIGURE 2
DUOPOLY MODEL WITH $k = 0$

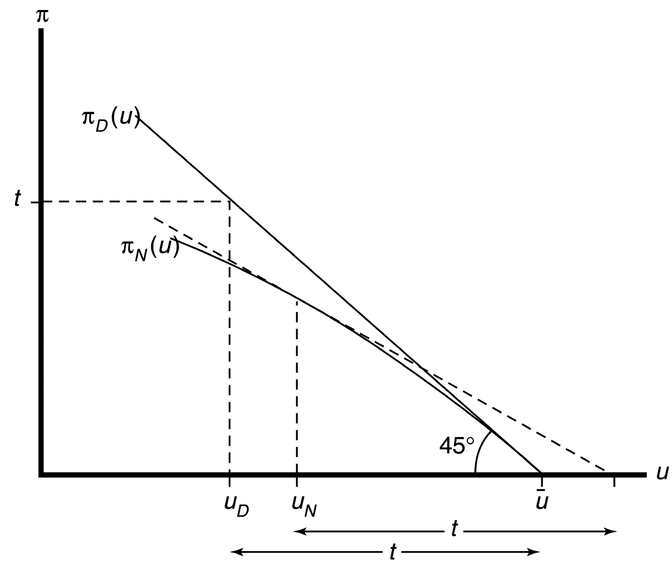

is decreasing in u, and we can apply Lemma 4. (In the competitive limit we know that we can restrict attention to $u \approx \bar{u}$.)

Finally, we mentioned in the Introduction that one concern of the previous literature has been to determine the effect of price discrimination on total output. In all the many cases where we have seen that discrimination raises welfare, we can simply *infer* that total output rises when discrimination is allowed. Indeed, it must rise by enough to offset the welfare cost of unequal marginal utilities.

4. Discrimination with heterogeneous consumers

■ **Interpersonal third-degree price discrimination.** In this section we discuss a model first analyzed by Holmes (1989), in which two firms operate in two distinct markets, denoted 1 and 2, and each supplies a single product in each market. For simplicity, we analyze the duopoly Hotelling model as before. Firms have the same constant marginal cost c in each market, and there are no per-consumer fixed costs. Firms offer linear tariffs, and let p_i denote a firm's price in market i. Suppose the fraction of customers in market 1 is α (and the fraction in market 2 is $1-\alpha$). Each customer in market i has the demand function $q_i(p_i)$. The two markets are separate, and a customer in one market is not able to purchase the product in the other market. Let $\bar{q}(p) \equiv \alpha q_1(p) + (1-\alpha)q_2(p)$ denote the average demand across the two markets with common price p. Let $v_i(p_i)$ denote the consumer-surplus function associated with demand function q_i, and let $\bar{v}(p)$ be that associated with $\bar{q}$. Let $\pi_i(u)$ be the per-customer profit function in market i, so that $\pi_i(v_i(p_i)) = q_i(p_i)(p_i - c)$ as in (6). Let $\eta_i = -cq_i'(c)/q_i(c)$ be the elasticity of demand in market i with marginal-cost pricing. Without loss of generality, suppose that markets are labelled so that market 1 is the more elastic:

$$\eta_1 \geq \eta_2. \tag{28}$$

We allow the product-differentiation parameter (or transport cost) to differ in the two markets, so let t be the parameter in market 1 and γt be that in market 2. In sum, if firms A and B set the prices p_1^A and p_1^B respectively in market 1, then firm i obtains total profits of

$$\alpha q_1(p_1^i)(p_1^i - c)\left(\frac{1}{2} + \frac{v_1(p_1^i) - v_1(p_1^j)}{2t}\right) \tag{29}$$

in that market (and similarly for market 2).

The regimes of no discrimination and discrimination are described below.

Price discrimination. In this case we can consider the two markets separately. From (8) we know that $\pi_i'(v_i(c)) = -1$, and so Lemma 2 implies that equilibrium utilities in the two markets are, to first order in t, approximately given by

$$u_1 \approx v_1(c) - t; u_2 \approx v_2(c) - \gamma t. \tag{30}$$

Therefore, aggregate consumer welfare is approximately given by

$$u_D(t) \approx \bar{v}(c) - t\left[\alpha + \gamma(1-\alpha)\right]. \tag{31}$$

From (30) the prices in the two markets are approximately given by

$$p_1 \approx c + \frac{t}{q_1(c)}; p_2 \approx c + \frac{\gamma t}{q_2(c)},$$

and so market 2 is the more "competitive," in the sense that equilibrium price is lower there if price discrimination is permitted, provided that

$$q_2(c) > \gamma q_1(c). \tag{32}$$

Notice that this condition governing which market has the lower price has nothing to do with the industry-level demand elasticities that are relevant for welfare: a market could be highly inelastic at the industry level and yet competition could force firms to choose the lower price there. In more detail, from (29) a firm's own-price elasticity in market 1 when the two firms choose the same price p_1 is

$$\eta_1(p_1) + \frac{p_1 q_1(p_1)}{t},$$

where $\eta_1(p_1) = -p_1 q_1'(p_1)/q_1(p_1)$ is the industry-level elasticity. Thus a firm's own-price elasticity in a market, which is inversely related to the equilibrium price, does depend on the industry-level elasticity η_1, but in highly competitive markets (t close to zero) this factor is swamped by the second competitive interaction term, which yields the relation (32).

From Lemma 2, aggregate profits are approximately given by

$$\pi_D(t) \approx t\left[\alpha + \gamma(1-\alpha)\right]. \tag{33}$$

Finally, total output under discrimination is approximately given by

$$Q_D(t) \approx \bar{q} - \frac{t}{c}\left[\alpha\eta_1 + \gamma(1-\alpha)\eta_2\right], \tag{34}$$

where $\bar{q} = \bar{q}(c)$.

No price discrimination. A ban on price discrimination requires a firm to set $p_1 = p_2 = p$. Using the expression (29), and the corresponding expression for market 2, the first-order condition for the equilibrium uniform price p is

$$\alpha q_1(p)\pi_1(v_1(p)) + \frac{1}{\gamma}(1-\alpha)q_2(p)\pi_2(v_2(p)) = t\left[\alpha\frac{d}{dp}\pi_1(v_1(p)) + (1-\alpha)\frac{d}{dp}\pi_2(v_2(p))\right]$$

or

$$\phi(p)(p-c) = t\left[\bar{q}(p) + \bar{q}'(p)(p-c)\right],$$

where

$$\phi(p) \equiv \alpha\,(q_1(p))^2 + \frac{1-\alpha}{\gamma}\,(q_2(p))^2\,.$$

Writing $p = p(t)$, differentiating this first-order condition gives

$$\left\{\frac{d}{dp}\left[\phi(p)(p-c)\right] - t\frac{d}{dp}\left[\bar{q}(p) + \bar{q}'(p)(p-c)\right]\right\} p'(t) = \bar{q}(p) + \bar{q}'(p)(p-c),$$

which implies that $p'(0) = \bar{q}/\phi$, where these quantities are evaluated at marginal-cost prices. Therefore, aggregate consumer surplus for small t is approximately

$$u_N(t) = \bar{v}(p(t)) \approx \bar{v}(c) - t\frac{\bar{q}^2}{\phi}. \tag{35}$$

Similarly, aggregate profits are approximately given by

$$\pi_N(t) \approx t\frac{\bar{q}^2}{\phi}. \tag{36}$$

Finally, total output without discrimination is approximately

$$Q_N(t) \approx \bar{q} + t\frac{\bar{q}}{\phi}\bar{q}'. \tag{37}$$

Proposition 4. For sufficiently small t,

(i) except in the knife-edge case where $q_2(c) = \gamma q_1(c)$, aggregate consumer surplus is lower and profits are higher if price discrimination is permitted;

(ii) if (32) holds, then welfare is lower when price discrimination is permitted.

Notice in particular that profits always increase if discrimination is permitted in competitive markets, which shows that Holmes's (1989) result that profits may fall with discrimination requires markets to be reasonably uncompetitive. Also, it is *aggregate* consumer surplus that falls with discrimination: clearly consumers in the "weak" market benefit from the lower prices caused by allowing discrimination.[27]

Figure 3 illustrates this result. The point (p, p) is the equilibrium uniform price vector, and iso-welfare, iso-total output, iso-profit, and iso-consumer surplus contours are drawn through this point. The iso-welfare and iso-output contours are tangent at this point, since welfare is maximized for a *given* level of total output when prices are chosen to be uniform. (One can also check that the former lies inside the latter.) Since we assume that market 1 is more elastic, the iso-profit and iso-consumer surplus contours are less steep than the iso-welfare contour through this point, with the former being less steep than the latter. Improvements to consumer surplus and welfare that maintain the same level of profitability all lie in the region $p_2 > p_1$. However, assumption (32) implies that firms choose a lower price in market 2 if they are allowed to discriminate, something

[27] It is straightforward to show that the equilibrium uniform price lies between the two discriminatory prices. As discussed in the Introduction, this feature contrasts with models of "customer poaching" in which *all* prices might fall with discrimination.

FIGURE 3
INTERPERSONAL THIRD-DEGREE PRICE DISCRIMINATION

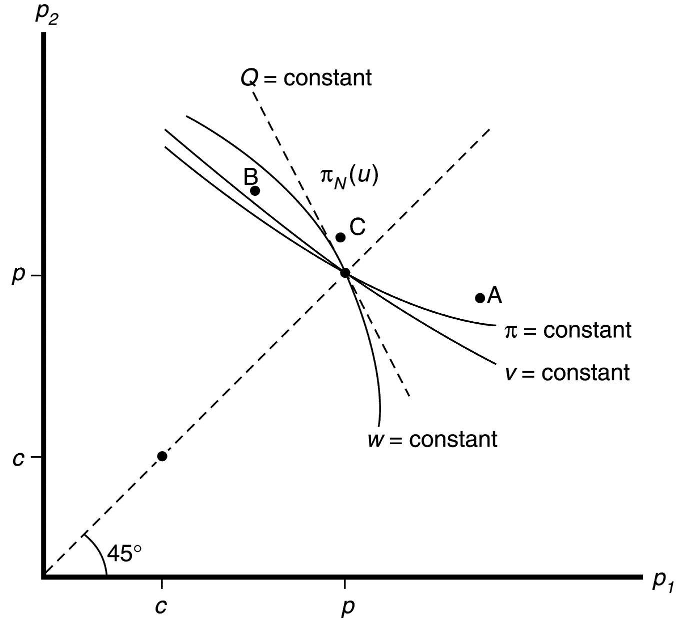

that increases profits but reduces consumer surplus and overall welfare—this is a point like A on the figure.

Intuitively, the fact that firms reduce the price in the inelastic market and raise the price in the elastic market causes total output to fall—this is shown rigorously in the proof—and this must reduce welfare compared to the no-discrimination case. In sum, and in contrast to our earlier analysis with homogeneous consumers, a firm's incentives when choosing its pattern of relative prices may diverge from those of the economy as a whole, and Ramsey-like prices are offered only by chance.

If the inequality (32) is reversed but not too much, then firms choose to raise price in the "correct" market, but this effect is not sufficient to cause total output to rise, and the welfare effect is still negative—this is a point like C on the figure. On the other hand, if (32) is reversed sufficiently strongly, then the welfare effect is reversed, as in point B on the figure.

Inter- or intrapersonal price discrimination? In this section we have chosen to interpret the industry as one where there are two disjoint groups of consumers.[28] However, to bring out connections with the analysis of intrapersonal price discrimination in Section 3, we can also interpret the model as one where there is a single group of "homogeneous" consumers, each of whom buys two products.

To see this, suppose that a consumer incurs the "transport cost" of going to one firm or the other on a *per-product* basis (rather than the per-trip basis of Section 3). Specifically, suppose that the transport cost is t for product 1 and γt for product 2 (per unit of distance). Suppose a consumer lies at location $x_1 \in [0, 1]$ with respect to product 1, and lies at location $x_2 \in [0, 1]$ with respect to product 2, and that the marginal distribution of each x_i is uniform on [0, 1]. A consumer's demand for product i with the price p_i is $(1/2)q_i(p_i)$, and there are no cross-price effects across the two markets. Then for any pair of price vectors offered by the two firms, it is clear that the outcomes of this homogeneous consumer model coincide precisely with those of the segmented market model (setting $\alpha = 1/2$) discussed earlier in this section: because transport

[28] We are grateful to a referee for encouraging us to provide the following discussion.

costs are incurred per product and there are no cross-price effects, the prices offered in one market do not affect a consumer's preferences over the choice of firm in the other market. For instance, in this homogeneous-consumer framework, the profits generated in market 1 for firm i do not depend at all on firm behavior in market 2 and are given by the expression (29) above (with $\alpha = 1/2$). The extent of correlation between a consumer's location parameters x_1 and x_2, i.e., whether or not consumers systematically favor one firm or the other for the two products, plays no role in this analysis.

In particular, when transport costs are levied per product rather than per trip, there is no reason to suppose that a firm will offer Ramsey-like prices for its two products when price discrimination is permitted, by contrast with the analysis in Section 3. This discussion suggests that the welfare effects of discrimination are not determined so much by the "homogeneous-" versus "heterogeneous-"consumer dichotomy that we have so far emphasized; rather, it is determined by whether the "one-stop-shopping" assumption we imposed in Section 3 is appropriate. If consumers choose products on a mix-and-match basis from the two firms, then these consumers behave much as if there were two segmented consumer groups, each of which buys only one product. In particular, while price discrimination is desirable in competitive markets where one-stop shopping is the relevant model, when mix-and-match shopping is the more appropriate framework the welfare effects of price discrimination are less clear-cut.

Finally, it should be emphasized that this close parallel between multiproduct consumers and segmented, single-product consumers only applies to this setting with linear pricing. For instance, even with per-product transport costs it will most likely be in the interests of firms to *bundle* the two products together when they face multiproduct consumers, so that joint consumption is encouraged (or perhaps discouraged) *endogenously*.[29] However, this bundling policy makes no sense when the industry is made up of two segmented groups of single-product consumers.

□ **Price discrimination with unobserved consumer heterogeneity.** Here, suppose consumers have private information about their tastes for the services (and not just about their location). As in Section 2, each firm offers n products, and if it supplies quantities $\mathbf{q} = (q_1, \ldots, q_n)$ to a consumer a firm incurs the cost $k + \sum_{i=1}^{n} c_i q_i$.

We employ a similar utility specification as in the monopoly model of Armstrong (1996), but extended to a duopoly Hotelling model with transport cost t. Specifically, a consumer's type is an $m + 1$ vector $(\boldsymbol{\theta}, x)$, where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$ is a vector of preference parameters for the n products and, as earlier, $0 \leq x \leq 1$ denotes the location of the consumer between the firms. As before, suppose that consumer location x is uniformly distributed on the interval [0, 1], although this is not important for the argument. A consumer of type $\boldsymbol{\theta}$ has utility, excluding transport costs, given by

$$U(\boldsymbol{\theta}, \mathbf{q}) - T$$

if she consumes quantities $\mathbf{q}$ in return for the payment T. Unlike Armstrong (1996), we need make no assumptions about the form of the utility function U (e.g., concerning single crossing), nor about the distribution of $\boldsymbol{\theta}$ (such as whether $\boldsymbol{\theta}$ is continuously distributed or about independence of $\theta_1, \ldots, \theta_m$). Finally, let

$$v(\boldsymbol{\theta}) \equiv \max_{\mathbf{q}} : \left\{ U(\boldsymbol{\theta}, \mathbf{q}) - \sum_{i=1}^{n} c_i q_i \right\}$$

be the type-$\boldsymbol{\theta}$ consumer's utility, excluding transport costs, when there is marginal-cost pricing.

Suppose firm $j = A, B$ offers the multiproduct nonlinear tariff $T_j(\mathbf{q})$. Let

$$u_j(\boldsymbol{\theta}) = \max_{\mathbf{q}} : \{U(\boldsymbol{\theta}, \mathbf{q}) - T_j(\mathbf{q})\}$$

[29] See Matutes and Regibeau (1992) for a model along these lines when the two products are perfect complements.

be the type-$\boldsymbol{\theta}$ consumer's maximum utility if she buys from firm j, excluding transport costs. Then the type-$(\boldsymbol{\theta}, x)$ consumer obtains net utility $u_A(\boldsymbol{\theta}) - tx$ if she buys from A, $u_B(\boldsymbol{\theta}) - t(1-x)$ if she buys from B, and zero utility if she buys from neither.

We then have the following result:

Proposition 5. Suppose that $\boldsymbol{\theta}$ and x are independently distributed, and x is uniformly distributed on [0, 1]. Provided

$$t \leq \frac{2}{3}\,[v(\theta) - k] \text{ for all } \boldsymbol{\theta}, \tag{38}$$

each firm offering the cost-based two-part tariff (22) is an equilibrium.

Condition (38) is just the condition that ensures all consumers participate with this choice of tariff. (It is the generalization of (20) to the case of heterogeneous consumers.) We do not make any claims about the uniqueness of this equilibrium. Indeed, in the general multiproduct framework used here it does not seem feasible even to calculate A's best response given a tariff chosen by B unless B's tariff is a cost-based tariff of the form (22).

One intuitive way to see why the tariff (22) is an equilibrium when $\boldsymbol{\theta}$ is private information—which is the basis of the proof of the result—is that this tariff would be the equilibrium if $\boldsymbol{\theta}$ was *observable* by firms, provided that the distribution of $\boldsymbol{\theta}$ and x is independent. If a given consumer is known to have taste parameters $\boldsymbol{\theta}$ but unknown location x, then, by independence, this location is uniformly distributed on [0, 1] as in the duopoly Hotelling model developed in Section 2, in which we showed that for these type-$\boldsymbol{\theta}$ consumers, firms would offer the two-part tariff (22), provided the type-$\boldsymbol{\theta}$ market is covered (which is ensured by (38)). Since this tariff does not depend on $\boldsymbol{\theta}$, it is also an equilibrium if $\boldsymbol{\theta}$ is private information to consumers.

In sum, provided that the market is sufficiently competitive so that (38) holds, one equilibrium is given by each firm setting the simple cost-based tariff (22). Since this tariff achieves the first-best outcome, allowing firms to offer general multiproduct nonlinear tariffs is clearly superior in welfare terms to placing any form of restriction on the tariffs permitted (provided that this equilibrium is chosen). This is a strong result, and it shows how competition can have the effect of dramatically simplifying the tariffs offered by firms with market power. The simplicity of this result contrasts strongly with the complex monopoly analysis in Armstrong (1996) and Rochet and Choné (1998). In particular, as discussed in the Introduction, the monopoly model had the feature that the firm would choose to exclude some consumers from its products, a feature that made computation of the optimal nonlinear tariff difficult. With competition, a firm no longer has the ability to exclude consumers, and this, among other factors, greatly simplifies the analysis.

A similar result has also been obtained by Rochet and Stole (forthcoming).[30] They study the problem under the assumptions that firms offer a single product and consumers have single-dimensional private information θ that enters into utility functions so as to satisfy the single-crossing property, and they obtain the parallel result. Importantly, they show that it is not important that consumers be symmetrically distributed between the two firms. Obviously, in that case the two firms offer different tariffs in equilibrium, but these tariffs are still cost-based two-part tariffs. (Only the fixed charge differs across the firms.) The advantage of their more restrictive specification is that they are able to find equilibria when the strong assumptions needed for Proposition 5 are relaxed. In particular, they show that the result is highly sensitive to the assumptions that θ and x are distributed independently and that all consumer types are willing to participate with the candidate tariff. When the market is not fully covered, i.e., when (38) is violated, they show that the market is segmented into two classes: one where consumers have strong tastes and where competition is intense enough to drive marginal prices down to marginal costs, and one with low-demand consumers over whom firms have a local monopoly.

[30] For related analyses of screening models in duopoly, see also Desai (2001), who has a binary distribution for the "vertical" taste parameter and who obtains a similar efficiency result, and Schmidt-Mohr and Villas-Boas (1997).

5. Conclusions

■ This article has had two main themes. The first was to show some advantages of analyzing firms as competitors in utility space, i.e., as simply supplying utility to consumers. Using the basic relationship $\pi(u)$ between profit and utility per consumer, we obtained directly—without needing to examine costs, prices, and outputs—various results about the utility, profit, and welfare effects of price discrimination. The second main theme was that freedom to engage in price discrimination tends to be desirable in sufficiently competitive conditions: constraining the ways that competitive firms supply utility to consumers usually reduces total welfare. (However, we also saw that consumers often suffered when firms were granted extra freedom to extract surplus.) The main exception to this rule was when consumers were segmented into identifiably separate markets—or, alternatively, when consumers buy some products from one firm and some from another, so that "one-stop shopping" is not the relevant framework—and when elasticities go the "wrong" way.

Appendix

■ The proofs for Propositions 1, 4, and 5, Lemmas 1–4, and Corollary 1 follow.

Proof of Proposition 1.

Claim. The function (5) does have a maximum.

Proof. From (3), the function tends to $-\infty$ as $u \to \bar{u}$. Next, notice that the function

$$\{\log s(u, u) + \log \pi(u)\} - \{\phi(u) + \log \pi(u)\}$$

is decreasing, as its derivative is $s_2/s \leq 0$. From (2) we know that the first term $\{\cdot\}$ is greater at u^* than at any $u < u^*$. Since the above function is decreasing, we deduce that the same is true for the second term $\{\cdot\}$, i.e., that $\phi(u^*) + \log \pi(u^*) \geq \phi(u) + \log \pi(u)$ whenever $u < u^*$. Therefore, we know that there are points in the compact set $[u^*, \bar{u} - \varepsilon]$, for small $\varepsilon > 0$, that generate a higher value for (5) than any choice for u outside the set. Since (5) is continuous in this set, we deduce that a maximum exists. Moveover, since the derivative of (5) is positive at u^*, this maximizer is strictly greater than the collusive utility level u^*. *Q.E.D.*

Therefore, let $\hat{u}$ be a utility level that maximizes (5). Suppose B chooses the utility $\hat{u}$. If A chooses utility u_A, the log of its total profit is

$$\log \pi(u_A) + \log s(u_A, \hat{u}) = \{\log \pi(u_A) + \phi(u_A)\} + \{\log s(u_A, \hat{u}) - \phi(u_A)\}\,.$$

The first term $\{\cdot\}$ is maximized by assumption at $u_A = \hat{u}$ and the second term $\{\cdot\}$ has derivative

$$\frac{s_1(u_A, \hat{u})}{s(u_A, \hat{u})} - \frac{s_1(u_A, u_A)}{s(u_A, u_A)}.$$

Clearly this derivative is zero when $\hat{u} = u_A$, and from (1) it is positive when $u_A < \hat{u}$ and negative when $u_A > \hat{u}$. In other words, the second term $\{\cdot\}$ is single-peaked in u_A and maximized at $u_A = \hat{u}$. We deduce that $\hat{u}$ is a best response for firm A, and hence that $\hat{u}$ is a symmetric equilibrium.

Because the game is symmetric and involves upward-sloping reaction functions, there are no asymmetric equilibria. Finally, suppose u' is another symmetric equilibrium. The fact that u' is A's best response when B plays u' implies that it satisfies the first-order condition

$$\phi'(u') + \frac{\pi'(u')}{\pi(u')} = 0.$$

Therefore, if the function (5) is strictly concave, $\hat{u}$ is the unique maximizer of (5). *Q.E.D.*

Proof of Lemma 1. This proof follows the method used for Proposition 1. Suppose that firm B offers utility $\hat{u}$ that maximizes (19). If A chooses u_A, the log of its profit is

$$\log \pi(u_A) + \log s(u_A, \hat{u}) = \left\{\log \pi(u_A) + \frac{u_A}{t}\right\} + \left\{\log s(u_A, \hat{u}) - \frac{u_A}{t}\right\},$$

where s is given by (17). By construction, $u_A = \hat{u}$ maximizes the first term $\{\cdot\}$ above, while the second term $\{\cdot\}$ has derivative given by

$$\frac{1}{t + u_A - \hat{u}} - \frac{1}{t} \tag{A1}$$

if $t < u_A + \hat{u}$, and otherwise is equal to

$$\frac{1}{u_A} - \frac{1}{t}. \tag{A2}$$

We claim that the derivative of the second term $\{\cdot\}$ is positive if $u_A < \hat{u}$ and negative if $u_A > \hat{u}$, which implies that $u_A = \hat{u}$ is the best response for firm A. When $t < u_A + \hat{u}$, (39) is positive if $u_A < \hat{u}$ and otherwise is negative. When $t > u_A + \hat{u}$, it follows that $u_A < t$, and so from (40) the derivative of the second term $\{\cdot\}$ is positive. All that remains to show is that it is not possible that $u_A > \hat{u}$, that $t > u_A + \hat{u}$, and for t to satisfy (18). But since $\log \pi$ is concave, (18) implies that the maximizer of $\log \pi(u) + u/t$ is no less than the maximizer of $\log \pi(u) + u/(2\beta)$, the latter by construction being β. We deduce that $\hat{u} \geq \beta$. But (18) together with $t > u_A + \hat{u}$ and $\hat{u} \geq \beta$ implies that $u_A < \hat{u}$ as required. *Q.E.D.*

Proof of Lemma 2. From Lemma 1 we see that for sufficiently small t, $\hat{u}(t)$ satisfies the first-order condition

$$t\pi'(\hat{u}(t)) + \pi(\hat{u}(t)) = 0,$$

which can be differentiated to give

$$\pi' + \left[t\pi'' + \pi'\right]\hat{u}'(t) = 0$$

(where the dependence of π on $\hat{u}(t)$ has been suppressed). This implies that $\hat{u}'(0) = -1$. Differentiating again gives

$$\pi''\hat{u}' + \left[t\pi'' + \pi'\right]\hat{u}'' + \left[\pi'' + t\pi'''\hat{u}' + \pi''\hat{u}'\right]\hat{u}' = 0,$$

which implies $\hat{u}''(0) = \pi''(\bar{u})/\pi'(\bar{u})$. This gives the second-order Taylor expansion for $\hat{u}$.

The other Taylor expansions follow from the definitions for $\hat{\pi}$ and $\hat{w}$. *Q.E.D.*

Proof of Lemma 3. Total differentiation of (24) yields

$$[2\pi\pi' + \tau^2\pi'']\hat{u}' + 2\tau\pi' = 0,$$

and differentiating again gives

$$[2\pi\pi' + \tau^2\pi'']\hat{u}'' + [2\pi\pi'' + 2\pi'\pi' + \tau^2\pi'''](\hat{u}')^2 + 4\tau\pi''\hat{u}' + 2\pi' = 0. \tag{A3}$$

Expression (24) then implies that

$$\frac{\pi(\hat{u}(\tau))}{\tau} \longrightarrow \sqrt{-\pi'(\bar{u})} \text{ as } \tau \longrightarrow 0.$$

Letting $\tau \longrightarrow 0$ in (A3) implies that

$$\hat{u}'(0) = -\frac{1}{\sqrt{-\pi'(\bar{u})}}.$$

Using these facts, dividing both sides of (A3) by 2τ, and letting $\tau \longrightarrow 0$, we obtain that

$$\left(\pi'(\bar{u})\right)^{3/2}\hat{u}'' + \pi''(\bar{u})\sqrt{-\pi'(\bar{u})}\left(\hat{u}'(0)\right)^2 + 2\pi''(\bar{u})\hat{u}'(0) + \lim_{\tau \to 0}\frac{\left(\pi'(\hat{u}(\tau))\right)^2\left(\hat{u}'(\tau)\right)^2 + \pi'(\hat{u}(\tau))}{\tau} = 0.$$

Using L'Hôpital's rule to evaluate the final term on the left-hand side of this expression, we have

$$\left(\pi'\right)^{3/2}\hat{u}'' + \pi''\sqrt{-\pi'}\left(\hat{u}'\right)^2 + 2\pi''\hat{u}' - \pi''\hat{u}' + 2(\pi')^2\hat{u}'\hat{u}'' = 0,$$

which, after substituting for the value for $\hat{u}'$, implies

$$\hat{u}''(0) = 0.$$

From (25) we see that $\hat{w}(0) = \bar{u}$. Differentiating (25) we get

$$\hat{w}'(\tau) = \left[1 + \frac{1}{4}\left(1 - \frac{\pi\pi''}{(\pi')^2}\right)\right]\hat{u}'(\tau)$$

and so

$$\hat{w}'(0) = -\frac{5}{4}\frac{1}{\sqrt{-\pi'(\bar{u})}},$$

which give the first two terms in the Taylor expansion. Differentiating again, and using the fact that $\hat{u}''(0) = 0$, gives

$$\hat{w}''(0) = \frac{\pi''(\bar{u})}{4(\pi'(\bar{u}))^2},$$

which completes the proof. *Q.E.D.*

Proof of Lemma 4. From (5), let u_i maximize of $\log \pi_i(u) + \phi(u)$. Then by revealed preference

$$\log \pi_D(u_D) + \phi(u_D) \geq \log \pi_D(u_N) + \phi(u_N)$$
$$\log \pi_N(u_N) + \phi(u_N) \geq \log \pi_N(u_D) + \phi(u_D),$$

which implies that

$$\log \pi_D(u_D) - \log \pi_N(u_D) \geq \log \pi_D(u_N) - \log \pi_N(u_N).$$

Hence from the assumption we have $u_D \leq u_N$ as required. *Q.E.D.*

Proof of Corollary 1. From (15) we have

$$\frac{\pi_D'(u)}{\pi_D(u)} = -\frac{1}{v(\mathbf{c}) - u}.$$

Since π_N in (9) is concave,

$$0 = \pi_N(v(\mathbf{c})) \leq \pi_N(u) + (v(\mathbf{c}) - u)\pi_N'(u),$$

and so

$$\frac{\pi_N'(u)}{\pi_N(u)} \geq -\frac{1}{v(\mathbf{c}) - u} = \frac{\pi_D'(u)}{\pi_D(u)}.$$

This implies that $\log \pi_D - \log \pi_N$ is decreasing in u, and we can apply Lemma 4 to this pair of profit functions and deduce that consumer utility falls if two-part tariffs are used. *Q.E.D.*

Proof of Proposition 4. (i) Comparing (31) with (35), and (33) with (36), we see that this is true provided that

$$\frac{\bar{q}^2}{\phi} < \alpha + \gamma(1 - \alpha), \tag{A4}$$

which from the definition of $\bar{q}$ and ϕ is true if and only if

$$\gamma q_1^2 + \frac{1}{\gamma}q_2^2 - 2q_1q_2 = \left(\sqrt{\gamma}q_1 - \frac{q_2}{\sqrt{\gamma}}\right)^2$$

is strictly positive, which is true unless $q_2(c) = \gamma q_1(c)$.

(ii) Using the argument developed by Varian (1989)—see the Introduction to this article—we know that welfare is lower with discrimination if total output is strictly lower with discrimination. Comparing (34) with (37), we see that this is true in competitive markets provided that

$$\alpha\eta_1 + \gamma(1 - \alpha)\eta_2 > -c\frac{\bar{q}}{\phi}\bar{q}' = \frac{\bar{q}^2}{\phi}\frac{\alpha q_1\eta_1 + (1 - \alpha)q_2\eta_2}{\alpha q_1 + (1 - \alpha)q_2}.$$

But from (A4) we see this is true provided that

$$\frac{\alpha\eta_1 + \gamma(1 - \alpha)\eta_2}{\alpha + \gamma(1 - \alpha)} \geq \frac{\alpha q_1\eta_1 + (1 - \alpha)q_2\eta_2}{\alpha q_1 + (1 - \alpha)q_2},$$

which, given (32), holds if and only if $\eta_2 \leq \eta_1$ as claimed. *Q.E.D.*

Proof of Proposition 5. Suppose B offers the tariff (22), so that $u_B(\boldsymbol{\theta}) = v(\boldsymbol{\theta}) - k - t$. Given this tariff, an upper bound on the available profit for firm A is obtained by supposing that it *can* observe a consumer's parameter $\boldsymbol{\theta}$ (but not her location x). Therefore, suppose this is so, and for a consumer of type $\boldsymbol{\theta}$ suppose A offers the gross utility $u(\boldsymbol{\theta})$. The independence assumption implies that the distribution of x is uniform on [0, 1] given $\boldsymbol{\theta}$, so from (17) the number of type-$\boldsymbol{\theta}$ consumers buying from A is

$$\min\left\{\frac{1}{2}\left(1+\frac{u(\boldsymbol{\theta})-\left[v(\boldsymbol{\theta})-k-t\right]}{t}\right), \frac{u(\boldsymbol{\theta})}{t}\right\}.$$

The firm will choose the method of generating utility $u(\boldsymbol{\theta})$ that maximizes its profits, which is done by setting prices equal to costs, and setting the fixed fee $f_A(\boldsymbol{\theta}) = v(\boldsymbol{\theta}) - u(\boldsymbol{\theta})$. Therefore, A's maximum profit available from the type-$\boldsymbol{\theta}$ consumers is

$$\max_{f_A} : (f_A - k) \times \min\left\{\frac{1}{2}\left(1+\frac{k+t-f_A}{t}\right), \frac{v(\boldsymbol{\theta})-f_A}{t}\right\}.$$

When t is large the solution to this problem depends on $\boldsymbol{\theta}$ via $v(\boldsymbol{\theta})$. However, the problem has the solution $f_A(\boldsymbol{\theta}) = k + t$ provided that $t \leq (2/3)[v(\boldsymbol{\theta}) - k]$. In particular, provided that t is small enough to satisfy (38), A's choice of tariff, given that it can observe a consumer's type $\boldsymbol{\theta}$, takes the form (22) for all $\boldsymbol{\theta}$. Since this tariff does not depend on $\boldsymbol{\theta}$, it must be the best response for A when $\boldsymbol{\theta}$ is private information. *Q.E.D.*

References

ANDERSON, S.P., DE PALMA, A., AND THISSE, J.-F. *Discrete Choice Theory of Product Differentiation*. Cambridge, Mass.: MIT Press, 1992.

———, ———, AND NESTEROV, Y. "Oligopolistic Competition and the Optimal Provision of Products." *Econometrica*, Vol. 63 (1995), pp. 1281–1301.

ARMSTRONG, M. "Multiproduct Nonlinear Pricing." *Econometrica*, Vol. 64 (1996), pp. 51–75.

———. "The Theory of Access Pricing and Interconnection." In M. Cave, S. Majumdar, and I. Vogelsang, eds., *The Handbook of Telecommunications Economics*. Amsterdam: North-Holland, forthcoming.

——— AND VICKERS, J. "Price Discrimination, Competition and Regulation." *Journal of Industrial Economics*, Vol. 41 (1993), pp. 335–360.

BERGSTROM, T. AND VARIAN, H. "Two Remarks on Cournot Equilibria." *Economics Letters*, Vol. 19 (1985), pp. 5–8.

BLISS, C. "A Theory of Retail Pricing." *Journal of Industrial Economics*, Vol. 36 (1988), pp. 375–391.

BORENSTEIN, S. "Price Discrimination in Free-Entry Markets." *RAND Journal of Economics*, Vol. 16 (1985), pp. 380–397.

CAPLIN, A. AND NALEBUFF, B. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." *Econometrica*, Vol. 59 (1991), pp. 25–59.

CORTS, K.S. "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment." *RAND Journal of Economics*, Vol. 29 (1998), pp. 306–323.

DESAI, P. "Quality Segmentation in Spatial Markets: When Does Cannibalization Affect Product Line Design?" *Marketing Science*, Vol. 20 (2001), pp. 265-283.

DIAMOND, P.A. AND MIRRLEES, J.A. "Optimal Taxation and Public Production II: Tax Rules." *American Economic Review*, Vol. 61 (1971), pp. 261–278.

FUDENBERG, D. AND TIROLE, J. "Customer Poaching and Brand Switching." *RAND Journal of Economics*, Vol. 31 (2000), pp. 634–657.

HOLMES, T.J. "The Effects of Third-Degree Price Discrimination in Oligopoly." *American Economic Review*, Vol. 79 (1989), pp. 244–250.

KATZ, M.L. "Price Discrimination and Monopolistic Competition." *Econometrica*, Vol. 52 (1984), pp. 1453–1471.

MATUTES, C. AND REGIBEAU, P. "Compatibility and Bundling of Complementary Goods in a Duopoly." *Journal of Industrial Economics*, Vol. 40 (1992), pp. 37–54.

MONDERER, D. AND SHAPLEY, L.S. "Potential Games." *Games and Economic Behavior*, Vol. 14 (1996), pp. 124–143.

NAHATA, B., OSTASZEWSKI, K., AND SAHOO, P.K. "Direction of Price Changes in Third-Degree Price Discrimination." *American Economic Review*, Vol. 80 (1990), pp. 1254–1258.

PERLOFF, J.M. AND SALOP, S.C. "Equilibrium with Product Differentiation." *Review of Economic Studies*, Vol. 52 (1985), pp. 107–120.

ROCHET, J.-C. AND CHONÉ, P. "Ironing, Sweeping, and Multidimensional Screening." *Econometrica*, Vol. 66 (1998), pp. 783–826.

——— AND STOLE, L. "Nonlinear Pricing with Random Participation." *Review of Economic Studies*, forthcoming.

SALOP, S.C. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics*, Vol. 10 (1979), pp. 141–156.

SCHMIDT-MOHR, U. AND VILLAS-BOAS, J.M. "Competitive Product Lines." Mimeo, Haas School of Business, University of California, Berkeley, 1997.

SLADE, M.E. "What Does an Oligopoly Maximize?" *Journal of Industrial Economics*, Vol. 42 (1994), pp. 45–61.

STOLE, L.A. "Nonlinear Pricing and Oligopoly." *Journal of Economics and Management Strategy*, Vol. 4 (1995), pp. 529–562.

THISSE, J.-F. AND VIVES, X. "On the Strategic Choice of Spatial Price Policy." *American Economic Review*, Vol. 78 (1988), pp. 122–137.

VARIAN, H.R. "Price Discrimination." In R. Schmalensee and R.D. Willig, eds., *The Handbook of Industrial Organization*. Amsterdam: North-Holland, 1989.

VILLAS-BOAS, J.M. "Dynamic Competition with Customer Recognition." *RAND Journal of Economics*, Vol. 30 (1999), pp. 604–631.

WILSON, R. *Nonlinear Pricing*. New York: Oxford University Press, 1993.