MPRA
Munich Personal RePEc Archive

Discriminating Against Captive Customers

Armstrong, Mark and Vickers, John

University of Oxford

8 February 2019

Online at https://mpra.ub.uni-muenchen.de/92189/
MPRA Paper No. 92189, posted 18 Feb 2019 17:49 UTC

Discriminating Against Captive Customers

Mark Armstrong John Vickers*

February 2019

Abstract

We analyze a market where some consumers only consider buying from a specific seller while other consumers choose the best deal from several sellers. When sellers are able to discriminate against their captive customers, we show that discrimination harms consumers in aggregate relative to the situation with uniform pricing when sellers are approximately symmetric, while the practice tends to benefit consumers in sufficiently asymmetric markets. We also show how the asymmetry of markets may be affected by the information that firms have on consumer captivity.

1 Introduction

In a market where some customers are "captive" to particular sellers while others choose freely among alternative offers, is it good or bad for consumers overall if sellers can discriminate against their captive customers? Such discrimination is clearly bad for the captives because they are monopolized, but competition then prevails for the custom of non-captives. With uniform pricing, on the other hand, captives get some benefit of competition, but competition is weakened by their presence, making the net effect unclear.

In this paper we show by way of a parsimonious duopoly model with homogeneous products that the answer depends on the degree of symmetry between firms. The key step in our analysis, following Armstrong and Vickers (2001), is to think of a consumer's surplus as a function of the profit generated. While a consumer's surplus is always a convex function of the price she pays, under a mild condition it is a *concave* function of the profit she generates. It is as though consumers in aggregate are risk-averse to profit variation. With symmetric firms, discrimination against captive customers harms consumers overall because it does

*Both authors at the Department of Economics and All Souls College, University of Oxford. We are grateful to three referees, and to Yongmin Chen, Martin Obradovits, David Ronayne, David Sappington and Jidong Zhou for helpful comments.

not affect profits but widens the variation of profit across consumers. Given that consumer surplus is a concave function of profit, this mean-preserving spread of profit is harmful to consumers. But if monopoly profit exceeds the associated deadweight loss, the comparison is reversed if there is enough asymmetry between firms. That is because uniform pricing, by softening competition, raises profits by enough to make consumers worse off despite their aversion to the greater profit variation that comes with discrimination.

Our model applies to situations where a seller has information about whether or not a prospective customer is able or willing to consider rival sellers for her purchase. For instance, some consumers might use a comparison website to choose between multiple offers while others shop more randomly, and a seller engages in price discrimination if it chooses different prices on the comparison site and when consumers buy from it directly.[1] A chain store may face varying degrees of local competition across its stores, and can choose higher prices in those outlets where consumers are more captive. An insurance seller (say) might offer a customer a relatively expensive deal, which is then discounted if the customer says she has found a better deal. A consumer's previous behaviour might reveal her likely switching costs, and a supplier might then offer an existing customer with high switching costs a high price. An energy firm might offer a range of different tariffs for its product, where inert customers end up on the most expensive "default" tariff while more active consumers shop around for cheaper (but often short term) offers. Price discrimination in such markets is a live policy issue, as regulators in the energy sector consider whether to require suppliers to put all customers on their cheapest available tariff (or more generally to limit the gap between the cheapest and the default tariffs).

After presenting our modelling framework in the next section, where we show how price discrimination based on whether a consumer is captive cannot increase industry profit, we specialise the market in section 3 to duopoly. There we show how the impact of price discrimination on consumers depends on the degree of asymmetry between sellers and the degree of "risk aversion" to profit by consumers, where the former makes discrimination more likely to benefit consumers and the latter makes it less likely. In section 4 we extend the analysis to more general information structures, where sellers might observe a noisy signal of captivity or a signal about which seller a consumer is captive to. When the overall market is symmetric, such price discrimination can only increase profit and harm

[1] Baye and Morgan (2002) study a model where firms pay to advertise their price on a comparison website, and show that firms set a lower price there than when consumers buy from them directly.

consumers.

The model we analyze involves a market with homogeneous products where different consumers are able or willing to consider different subsets of firms for their purchase. When firms use uniform pricing, the equilibrium in Bertrand competition is typically that firms use mixed strategies for their prices and there is price dispersion in the market. Classic papers in this tradition include Butters (1977), Varian (1980), and Burdett and Judd (1983). We follow the framework in Narasimhan (1988), who studies a duopoly model where firms can be asymmetric. The advantage of studying a duopoly market is that it is easily solved, while asymmetric models with more than two firms are currently little understood when firms use uniform prices.

Whereas most of the literature on price discrimination explores the implications of differences of preferences across markets, our model abstracts from this issue to focus on discrimination on the basis of whether or not a consumer is captive. In Armstrong and Vickers (1993) we studied a dynamic model where an incumbent seller active in two markets with the same demand function faced potential entry in one market while the other was captive. When the incumbent was able to set different prices across its markets, entry might be deterred, with the result that both prices could rise relative to the regime with uniform pricing. More recent papers that examine price discrimination not based on consumer preferences include Chen and Schwartz (2015) on cost-based differential pricing, and Heidues and Köszegi (2017) on discrimination based on indicators of consumer naivety. Chen and Schwartz (2015) provides an interesting contrast with our results. They suppose a monopoly serves two markets, each with same demand function but with different unit costs. If consumer surplus with monopoly pricing is a convex function of cost, as is commonly the case, then consumer surplus is higher with differential pricing than with uniform pricing, even though average price might increase.

A feature of some oligopoly models of price discrimination is that, unlike the monopoly case, discrimination reduces equilibrium profit—see, for example, Thisse and Vives (1988) and Corts (1998) for analyzes with product differentiation and deterministic prices. The same is true in our main model with asymmetric sellers, but with symmetry equilibrium profits are the same with and without discrimination, which is the key to the mean-preserving spread argument central to our analysis. In the more general framework in section 4, other kinds of information about captivity can, for instance, convert a symmetric market into a pair of asymmetric markets, with the result that price discrimination instead causes profit and prices to rise. This more general framework relates to the analysis of price discrimination in Bergemann, Brooks and Morris (2015), who analyze how different ways of partitioning consumers affects profit and consumer surplus, with the differences that we analyze duopoly rather than monopoly and information concerns consideration sets rather than consumer valuations.

2 A framework

There are n sellers which costlessly supply a homogeneous product. Consumers differ according to which sellers they are able or willing to buy from, and an exogenous fraction consider a given subset $S \subset \{1, ..., n\}$ of sellers. Since consumers who do not consider any sellers play no role in the analysis, suppose all consumers consider at least one seller and normalize the measure of consumers to 1. A consumer is *captive* to a seller if she considers only that seller. Suppose seller $i = 1, ..., n$ has γ_i captive customers, and let $\gamma = \Sigma_{i=1}^{n} \gamma_i$ be the total number of captives.

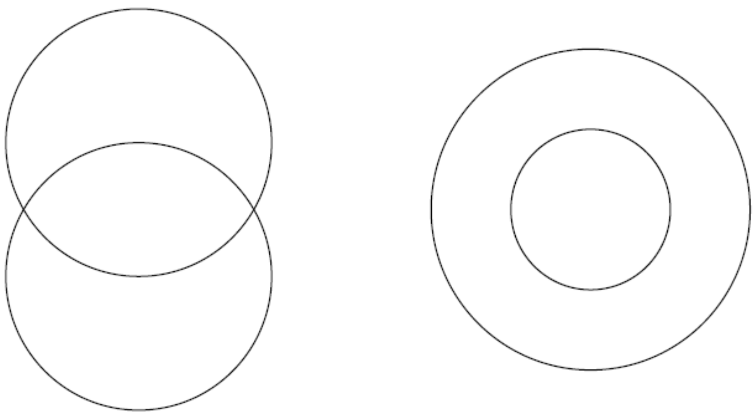

Figure 1: Markets with symmetric and nested reach

Figure 1 illustrates two patterns of consumer consideration in duopoly (where the consumers who consider a seller lie inside that seller's "circle"). Here, the left-hand Venn diagram shows a symmetric pattern of consideration sets (the two sellers have the same number of captive customers), while the right-hand diagram depicts a situation where a smaller seller's potential customers all also consider the larger seller (i.e., the smaller seller has no captive customers). This case of nested reach is relevant when, for instance, the smaller firm is a recent entrant which is considered by only a subset of the incumbent's customer base.

Sellers compete in Bertrand manner, and a consumer buys from the seller she considers with the lowest price. Each consumer demands $q(p)$ units of the product if the price paid is p, where $q(\cdot)$ is a smooth and weakly decreasing function when positive.[2] Thus, if a consumer buys from a seller at price p she generates profit $\pi(p) \equiv pq(p)$ for that seller. Denote the profit-maximizing price by p^* and maximum profit by $\pi^* = \pi(p^*)$. A consumer's net surplus if she pays price p, $v(p)$, is the usual area under the demand curve, so that $v'(p) = -q(p)$.

Suppose all sellers know for sure whether a consumer is captive or not and price accordingly, in which case there is a unique equilibrium and this involves pure strategies. If a consumer is contested, i.e., she considers at least two sellers, then Bertrand competition forces the price to that consumer down to marginal cost, so that $p = \pi = 0$ and the consumer enjoys surplus $v(0)$. When the consumer is captive, her seller will charge the monopoly price p^*, so that $\pi = \pi^*$ and the consumer obtains surplus $v(p^*)$. Thus, aggregate consumer surplus in this scenario is $\gamma v(p^*) + (1-\gamma)v(0)$ while aggregate profit is $\gamma\pi^*$.

When sellers either do not know when a consumer is captive, or are not permitted to discriminate against captive customers, a seller must offer a uniform price to all potential customers. If all consumers are captive ($\gamma = 1$) then all sellers choose the monopoly price, while if no consumer is captive ($\gamma = 0$) all sellers choose the competitive price,and in either of these extremes the outcome is the same with or without price discrimination. When $0 < \gamma < 1$, however, the equilibrium with uniform pricing involves at least some sellers using mixed strategies for their prices. Since aggregate profit is a continuous function of the vector of prices chosen by the n sellers, existence of equilibrium is ensured by Dasgupta and Maskin (1986, Theorem 5). Except in symmetric and other special cases—such as the duopoly market studied in section 3—the form of the equilibrium is not known. However, since seller i can always choose the monopoly price and sell at least to its γ_i captive customers, in any equilibrium its expected profit must be at least $\gamma_i\pi^*$. Therefore, industry profit in any equilibrium with uniform pricing must be at least equal to $\gamma\pi^*$, which was the equilibrium profit with price discrimination.

[2]The analysis in this paper applies equally if consumers have heterogeneous demand functions, provided that their demand is independent of their consideration set.

Stating this conclusion formally:

Proposition 1 *Industry profit when sellers can discriminate against captive customers is no higher than industry profit with uniform pricing.*

Consider the special case of unit demand, i.e., where $q(p) = 1$ if $p \leq 1$ and $q(p) = 0$ for $p > 1$, in which case $p^* = \pi^* = 1$. Then total welfare (profit plus consumer surplus) does not depend on price and is identically equal to 1 regardless of the pricing strategies followed by sellers. Since profit is weakly greater with uniform pricing, we have the following corollary to Proposition 1:

Corollary 1 *If consumers have unit demand then aggregate consumer surplus with price discrimination is no lower than consumer surplus in any equilibrium with uniform pricing.*

In the next section we put more structure on the model to gain further insight into when price discrimination of this form is harmful or beneficial for consumers and for overall welfare.

3 A duopoly market

In broad terms, when sellers engage in price discrimination the result is that the average profit generated from consumers falls but the variability of profit across consumers rises, relative to the regime with uniform pricing. In this section we consider consumer surplus as a function of the profit a consumer generates. In regular cases, this consumer surplus is a concave function of profit, in which case consumers are "risk averse" towards variation in profit, and whether they prefer the regime with price discrimination depends on how much industry profit falls.

In more detail, if $\eta(p) \equiv -pq(p)/q'(p)$ denotes elasticity of demand, $\pi'(p)$ has the sign of $1 - \eta(p)$, and so $\pi(p)$ is strictly single-peaked in p if

$$\eta(p) \text{ strictly increases with } p \ , \tag{1}$$

which is assumed henceforth. As before, denote the profit-maximizing price by p^*, in which case only prices in the interval $[0, p^*]$ will be chosen by sellers.[3] Since profit $\pi(p)$ is strictly

[3] Since unit cost has been normalized to zero, price p is net of cost. With positive cost, condition (1) is met with constant elasticity of demand. Condition (1) implies that profit $\pi(p)$ is concave in p in the range $[0, p^*]$. Profit being concave in quantity q is sufficient, but not necessary, for condition (1). Our method yields welfare results without needing to determine the effect of discrimination on total quantity.

increasing in $[0, p^*]$, and since $v(p)$ is strictly decreasing in p, we can construct a decreasing function $V(\pi)$ such that if the consumer generates profit π she enjoys net surplus $V(\pi)$, so that

$$v(p) \equiv V(\pi(p)) \ . \tag{2}$$

Differentiating (2) shows that $-q(p) = V'(\pi(p))\pi'(p)$, or

$$-V'(\pi(p)) = \frac{1}{1-\eta(p)} \ .$$

In particular, condition (1) implies $V(\pi)$ is strictly concave on $[0, \pi^*]$. Since profit $\pi(p)$ is strictly increasing over the relevant range $[0, p^*]$, as in Armstrong and Vickers (2001) we can view sellers as choosing the per-consumer profit π rather than the price p they ask from their customers, and a consumer buys from the seller with the smallest π from the set of sellers she considers.

In the remainder of the paper we consider a duopoly market, where seller $i = 1, 2$ has γ_i captive consumers (and remaining consumers consider both sellers). Thus seller i reaches (i.e., is considered by) $\sigma_i \equiv 1 - \gamma_j$ consumers, and the proportion of seller i's reach which is captive is denoted $\rho_i = \gamma_i/\sigma_i$, or

$$\rho_i = \frac{\gamma_i}{1-\gamma_j} \ . \tag{3}$$

Throughout the following analysis we label firms so that $\rho_1 \geq \rho_2$ (in which case $\gamma_1 \geq \gamma_2$ and $\sigma_1 \geq \sigma_2$). Suppose that $0 < \rho_1 < 1$, i.e., there are some captive and some contested consumers, in which case the equilibrium with uniform pricing involves mixed strategies, as described in the following standard result:

Lemma 1 *The unique equilibrium with uniform pricing involves the two sellers choosing profit in the same interval* $[\pi_0, \pi^*]$*, where the minimum profit is* $\pi_0 = \rho_1\pi^*$*, seller 1 has an atom at* $\pi = \pi^*$ *with probability* $(\sigma_1 - \sigma_2)/\sigma_1$ *(while seller 2 has no such atom), and seller* $i = 1, 2$ *obtains profit* $\sigma_i\pi_0$*.*

Proof. This result is taken from Narasimhan (1988). For completeness we construct the equilibrium as follows. Let seller i choose its per-consumer profit π according to the CDF $F_i(\pi)$. Then for $i \neq j$ in equilibrium these CDFs need to satisfy

$$\pi \times \sigma_i \left[\rho_i + (1-\rho_i)(1-F_j(\pi))\right] \equiv \sigma_i\pi_0$$

for any π in seller i's support. (Here, seller i will always sell to its $\rho_i\sigma_i$ captive customers, and when it chooses profit π it will also sell to the $(1-\rho_i)\sigma_i$ contested customers if its rival offers a higher profit, which occurs with probability $1-F_j(\pi)$.) This defines two functions, F_1 and F_2, which are increasing on the interval $[\pi_0,\pi^*]$, are both zero at $\pi=\pi_0$, and where $F_2(\pi^*)=1$ (so seller 2 has no atom at $\pi=\pi^*$) and $1-F_1(\pi^*)=(\sigma_1-\sigma_2)/\sigma_1$. Each seller is indifferent over any profit in the interval $[\pi_0,\pi^*]$, and neither seller has an incentive to choose profit outside this interval. ■

We next present our main result, which is that consumers in aggregate prefer uniform pricing if sellers are sufficiently symmetric (as with the left-hand diagram in Figure 1) while they usually prefer price discrimination if sellers are sufficiently asymmetric (as with the right-hand diagram). Here, the precise bounds for when parts (i) and (ii) of this result hold are contained within the proof.

Proposition 2
(i) Consumer surplus is higher with uniform pricing than with price discrimination when ρ_2 is sufficiently close to ρ_1.
(ii) If the deadweight loss associated with monopoly is no greater than monopoly profit, which is the case if demand $q(p)$ is log-concave, then consumer surplus is higher with price discrimination than with uniform pricing when ρ_2 is sufficiently small.

Proof. As in section 2, with price discrimination consumer surplus is

$$(1-\gamma)V(0)+\gamma V(\pi^*)\ , \tag{4}$$

while industry profit is $\gamma\pi^*$, where $\gamma=\gamma_1+\gamma_2$ is the fraction of captive customers in the market. The proof for part (i) finds a lower bound on consumer surplus with uniform pricing and shows when this lower bound is greater than (4), while part (ii) finds an upper bound on consumer surplus with uniform pricing and shows when this upper bound is below (4). In the following analysis we parameterize the market in terms of (ρ_1,ρ_2), in which case the numbers of captive customers and reach can be expressed as

$$\gamma_i=\frac{\rho_i(1-\rho_j)}{1-\rho_1\rho_2}\ ;\ \gamma=1-\frac{(1-\rho_1)(1-\rho_2)}{1-\rho_1\rho_2}\ ;\ \sigma_i=\frac{1-\rho_j}{1-\rho_1\rho_2}\ . \tag{5}$$

(i) Lemma 1 shows that industry profit with uniform pricing is $(\sigma_1+\sigma_2)\pi_0$, where $\pi_0\equiv\rho_1\pi^*$ is the minimum profit with uniform pricing. This industry profit is unchanged if the distribution of profit across consumers is altered so that σ_2 consumers generate profit π_0 and the remainder generate profit π^*, i.e., $(\sigma_1+\sigma_2)\pi_0 = (1-\sigma_2)\pi^*+\sigma_2\pi_0$. This hypothetical profit distribution is therefore a mean-preserving spread of the true distribution under uniform pricing, in the sense of Rothschild and Stiglitz (1970). Since $V(\cdot)$ is a concave function, aggregate consumer surplus with this hypothetical profit distribution, which is

$$\sigma_2 V(\pi_0) + (1-\sigma_2)V(\pi^*) \ , \tag{6}$$

cannot be greater than the equilibrium consumer surplus with uniform pricing. Since consumer surplus with price discrimination is (4), a sufficient condition for consumers to prefer uniform pricing is that (6) be no lower than (4), which entails

$$V(\pi_0) \geq V(\pi^*) + \frac{1-\gamma}{\sigma_2}(V(0) - V(\pi^*)) = \rho_2 V(\pi^*) + (1-\rho_2)V(0) \ , \tag{7}$$

where the equality follows from (5). Condition (7) requires that consumers prefer profit π_0 ($= \rho_1\pi^*$) for sure to the simple lottery consisting of $\pi = \pi^*$ with probability ρ_2 and $\pi = 0$ otherwise. This condition holds strictly when $\rho_2 = \rho_1$ due to the strict concavity of $V(\cdot)$ and assumption that $0 < \rho_1 < 1$, and hence it also holds for ρ_2 close to ρ_1.

(ii) Lemma 1 shows that industry profit with uniform pricing is $\Pi \equiv (\sigma_1 + \sigma_2)\pi_0$ and that firm 1 chooses the monopoly profit π^* with probability $(\sigma_1 - \sigma_2)/\sigma_1$. Therefore, a consumer will pay π^* if she is captive to firm 1 and that firm chooses π^*, and so the fraction of consumers who pay the monopoly price is $a \equiv \rho_1(\sigma_1 - \sigma_2)$.

Since industry profit consists of the profit from those consumers paying $\pi = \pi^*$ and those paying $\pi < \pi^*$, we have

$$\Pi = a\pi^* + (1-a)\mathbb{E}[\pi \mid \pi < \pi^*]$$

so that

$$\mathbb{E}[\pi \mid \pi < \pi^*] = \frac{\Pi - a\pi^*}{1-a} = \frac{2\sigma_2\pi_0}{1-a} = \frac{2\pi_0}{1+\rho_1} \ ,$$

where the final equality follows using (5). It follows that expected consumer surplus with uniform pricing satisfies

$$\begin{aligned} \mathbb{E}(V(\pi)) &= aV(\pi^*) + (1-a)\mathbb{E}[V(\pi) \mid \pi < \pi^*] \\ &\leq aV(\pi^*) + (1-a)V\left(\frac{2\pi_0}{1+\rho_1}\right) \ , \end{aligned} \tag{8}$$

where the inequality follows from the concavity of $V(\cdot)$. Therefore, consumer surplus is higher with price discrimination if (8) if no higher than (4), i.e., if

$$V\left(\frac{2\pi_0}{1+\rho_1}\right) \le V(\pi^*) + \frac{1-\gamma}{1-a}(V(0) - V(\pi^*)) = \frac{\rho_1+\rho_2}{1+\rho_1}V(\pi^*) + \frac{1-\rho_2}{1+\rho_1}V(0) \ , \quad (9)$$

where the equality follows using (5).

We claim that (9) holds strictly when $\rho_2 = 0$ (and $\rho_1 > 0$), and hence when ρ_2 is small enough, provided that the deadweight loss from monopoly pricing is less than monopoly profit, i.e., if

$$V(0) - V(\pi^*) - \pi^* \le \pi^* \ . \quad (10)$$

Since total surplus $V(\pi) + \pi$ is strictly decreasing in π over the range $[0, \pi^*]$, we have

$$V\left(\frac{2\pi_0}{1+\rho_1}\right) < V(0) - \frac{2\pi_0}{1+\rho_1} = V(0) - \frac{\rho_1}{1+\rho_1}2\pi^* \le \frac{\rho_1 V(\pi^*) + V(0)}{1+\rho_1} \ ,$$

where the second inequality follows from (10), and this demonstrates the claim.

Finally, we show that (10) holds when $q(p)$ is log-concave.[4] Log-concavity implies

$$\log q(p) \le \log q(p^*) + (p - p^*)\frac{q'(p^*)}{q(p^*)} = \log q(p^*) + \frac{p^* - p}{p^*} \ ,$$

where the equality follows from the first-order condition for p^* to maximize profit. It follows that $q(p) \le q(p^*)e^{1-p/p^*}$, in which case

$$V(0) - V(\pi^*) - \pi^* = \int_0^{p^*} [q(p) - q(p^*)]dp \le q(p^*) \int_0^{p^*} [e^{1-p/p^*} - 1]dp = (e-2)\pi^*$$

which is indeed smaller than π^*. ■

Intuitively, part (i) of this result is true since in near-symmetric markets industry profit is similar when sellers engage in price discrimination and when they cannot. (In either case, industry profit is approximately equal to the number of captive customers times π^*.) However, the distribution of profit across consumers is riskier with price discrimination—it is either 0 or π^*—and since consumers are "risk averse" towards variation in profit they are worse off with price discrimination. When sellers are very asymmetric, though, profit is considerably lower with price discrimination. With uniform pricing the seller with

[4] Note that log-concavity of demand also implies (1). Log-concavity is stronger than required to show (10). A weaker, though less familiar, condition which ensures this is that $1/\sqrt{q(p)}$ be convex (or, in the terminology of ρ-concavity, q is a $(-1/2)$-concave function).

many captive customers is unwilling to compete aggressively, and this enables the smaller firm to achieve profit well in excess of its "captive profit" (which is all it can get with price discrimination). Condition (9) describes when this reduction in profit is enough to outweigh the greater variability of profit with price discrimination. Provided that demand is not too convex (e.g., if $q(p)$ is log-concave), then price discrimination benefits consumers with nested reach, when only the larger seller has any captive customers.

In the limit case of unit demand, where $\pi^* = 1$ and $V(\pi) = 1 - \pi$, part (ii) of the result applies in all situations (condition (9) then holds always), as is consistent with Corollary 1. This case corresponds to "risk neutral" preferences over profit, when consumers care only about average profit and not its variation.

Total welfare—industry profit plus consumer surplus—is $V(\pi) + \pi$, which is also a concave function of π given condition (1). Therefore, the same method of analysis used for Proposition 2 can be used to obtain the following result.

Proposition 3
(i) Total welfare is higher with uniform pricing than with price discrimination when ρ_2 is close to ρ_1.
(ii) Total welfare is higher with price discrimination than with uniform pricing when ρ_1 is close to 1.

Proof. (i) The analysis we used for part (i) of Proposition 2 is valid here, if we replace $V(\pi)$ with $V(\pi) + \pi$. Therefore, from (7) total welfare is higher with uniform pricing if

$$V(\pi_0) + \pi_0 \geq \rho_2(V(\pi^*) + \pi^*) + (1 - \rho_2)V(0) \ . \tag{11}$$

Again, this is satisfied for ρ_2 close to ρ_1.

(ii) Likewise, (9) implies that total welfare is higher with price discrimination if

$$V\left(\frac{2\pi_0}{1+\rho_1}\right) + \frac{2\pi_0}{1+\rho_1} \leq \frac{\rho_1 + \rho_2}{1+\rho_1}(V(\pi^*) + \pi^*) + \frac{1-\rho_2}{1+\rho_1}V(0) \ . \tag{12}$$

When $\rho_1 \approx 1$, this reduces to the requirement $V(\pi^*) + \pi^* \leq V(0)$, which is true since total welfare $V(\pi) + \pi$ is maximized at $\pi = 0$. ■

For given $\rho_2 < 1$, the condition that ρ_1 be close to 1 requires that σ_2 be close to zero—see (5)—so that the smaller firm is very small. Therefore, part (ii) of this result applies to a narrow set of cases. To illustrate Propositions 2 and 3, consider the example with linear demand $q(p) = 2 - p$, in which case $p^* = \pi^* = 1$ and $V(\pi) = 1 + \sqrt{1-\pi} - \frac{1}{2}\pi$. Figure 2 depicts the impact of price discrimination in terms of (ρ_1, ρ_2), where recall that $\rho_2 \leq \rho_1$. Expression (7) shows that a sufficient condition for uniform pricing to be preferred by consumers overall is that (ρ_1, ρ_2) lies above the upper solid curve, while expression (9) shows that a sufficient condition for price discrimination to be preferred is that (ρ_1, ρ_2) lies below the lower solid curve. Expression (11) shows that total welfare is greater with uniform pricing when (ρ_1, ρ_2) lies above the upper dashed curve, while (12) shows that discrimination raises total welfare if (ρ_1, ρ_2) lies to the right of the lower dashed curve.

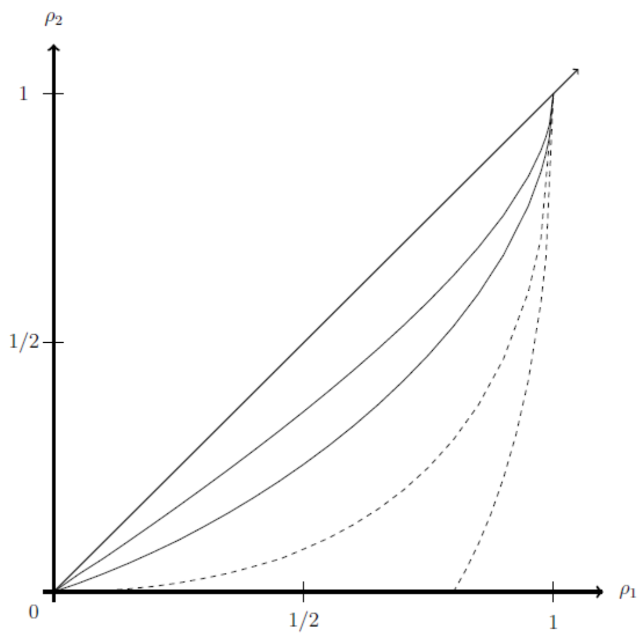

Figure 2: Impact of price discrimination with linear demand

4 More general information structures

Our analysis so far has assumed that sellers observe perfectly whether a consumer is captive or contested, in which case there is perfect competition and zero profit when a consumer is contested. Moreover, when the market is symmetric, expected profits are the same whether or not sellers can engage in this stark form of price discrimination. A natural question is how these results are altered when sellers have other kinds of information about consumer consideration sets. For instance, sellers might see a noisy rather than a perfect signal about whether the consumer is captive (but not to which seller she is captive), or sellers might see a signal about which seller the consumer might be captive to (but not about whether she is captive).[5]

To investigate these issues, suppose that the overall duopoly market (where seller $i = 1, 2$ has γ_i captive customers and there are $1 - \gamma_1 - \gamma_2$ contested customers) is segmented into $N \geq 2$ sub-markets labelled $j = 1, ..., N$. (These sub-markets could be geographic regions, for instance.) The fraction of consumers in segment j is α_j, and within segment j the fraction of consumers who are captive to seller i is γ_i^j and the fraction who are contested is $1 - \gamma_1^j - \gamma_2^j$. Here, $\sum_{j=1}^N \alpha_j \gamma_i^j = \gamma_i$ for $i = 1, 2$. We wish to compare the outcome when sellers can observe and condition their prices on the particular sub-market with the outcome when pricing is uniform. In each sub-market or in the market as a whole, equilibrium strategies are as described in Lemma 1.

The impact of price discrimination on profit is easily understood in the two configurations depicted in Figure 1. If the underlying market has nested reach (say, $\gamma_2 = 0$) then any form of discrimination by market segment can only reduce equilibrium profits. Since $\gamma_2^j = 0$ in all sub-markets, seller 1 makes exactly its captive profit in each segment, which adds up to the same profit it obtains without price discrimination. Seller 2's profit in sub-market j is its reach there, $1 - \gamma_1^j$, multiplied by seller 1's captive-to-reach ratio there, which is γ_1^j. Since $\gamma_1^j(1 - \gamma_1^j)$ is a concave function, its overall profit, $\sum_{j=1}^N \alpha_j \gamma_1^j (1 - \gamma_1^j)$, is strictly below its profit without discrimination, $\gamma_1(1 - \gamma_1)$, except in the trivial case where all segments are the same. Thus, with nested reach uniform pricing is the most profitable information structure.

By contrast, when the overall market is symmetric ($\gamma_1 = \gamma_2$) any form of discrimination by market segment can only boost profit: when sellers cannot discriminate they each obtain exactly their captive profit, while with discrimination they can obtain at least this profit by setting $\pi = \pi^*$ in all sub-markets and selling to their captive customers. If any sub-markets are asymmetric (so $\gamma_1^j \neq \gamma_2^j$ for some j) then profit in that segment will exceed the captive profit, in which case expected profits will strictly increase with discrimination. A more detailed analysis shows that in a symmetric market risk-averse consumers are harmed by any form of price discrimination by market segment, as reported in the next result.

Proposition 4 *In a symmetric market consumer surplus is weakly higher, while each*

[5] We assume that sellers possess the same information about consumers, so that this is a situation of third-degree price discrimination. A richer specification would allow sellers to observe private and potentially different signals about consumers.

firm's profit is weakly lower, with uniform pricing than with any form of price discrimination by market segment.

Proof. We proceed in two steps. First, we show that in an asymmetric sub-market the distribution of profit is a mean-preserving spread of that in the corresponding symmetric market with the same profit. Consider an arbitrary duopoly market, where seller $i = 1, 2$ has reach σ_i (so that seller i's captives number $1-\sigma_j$ and the number of contested customers is $\sigma_1+\sigma_2-1 \geq 0$) and sellers are labelled as $\sigma_1 \geq \sigma_2$. Lemma 1 shows that industry profit is

$$\Pi = (\sigma_1 + \sigma_2)\pi_0 \ , \tag{13}$$

where $\pi_0 = \frac{1-\sigma_2}{\sigma_1}\pi^*$ is the minimum profit in this market, and the two CDFs satisfy

$$\sigma_i F_i(\pi) = \frac{\sigma_1\sigma_2}{\sigma_1+\sigma_2-1}\left(1 - \frac{\pi_0}{\pi}\right) \ . \tag{14}$$

Compare this outcome with the symmetric market constructed to yield the same industry profit. If each seller in the symmetric market has reach σ, then to achieve profit (13) in the symmetric market requires

$$2\sigma - 1 = \frac{\sigma_2}{\sigma_1}(\sigma_1+\sigma_2-1) \ ,$$

so that the contestable portion is a fraction $\sigma_2/\sigma_1 \leq 1$ of its size in the asymmetric market.

Let $G(\pi)$ be the probability that a consumer is offered (minimum) profit no greater than π in the asymmetric market, so that

$$\begin{aligned} G(\pi) &= \sigma_1 F_1(\pi) + \sigma_2 F_2(\pi) - (\sigma_1+\sigma_2-1)F_1(\pi)F_2(\pi) \\ &= \frac{\sigma_1\sigma_2}{\sigma_1+\sigma_2-1}\left(1 - \left(\frac{\pi_0}{\pi}\right)^2\right) \ , \end{aligned} \tag{15}$$

where the equality follows from (14). If $\bar{G}(\pi)$ denotes the corresponding probability in the symmetric market, routine calculations show that for π in the supports of both CDFs we have

$$\frac{\bar{G}'(\pi)}{G'(\pi)} = \left(\frac{1}{2}\left(1 + \frac{\sigma_1}{\sigma_2}\right)\right)^2 \geq 1 \ . \tag{16}$$

Therefore, if $\sigma_1 > \sigma_2$ then $\bar{G}$ crosses G only once and from below, and it follows that the distribution of profit in the asymmetric market is a mean-preserving spread of that in the corresponding symmetric market.[6]

The second step demonstrates that the distribution of profit across several symmetric sub-markets is a mean-preserving spread of that in the single symmetric market with the same profit. For simplicity, suppose there are two symmetric sub-markets, one where each seller has reach $\sigma_L \geq \frac{1}{2}$ and the other where each seller has reach σ_H, where $\sigma_H \geq \sigma_L$ and where a fraction α of consumers are in the latter, more competitive, sub-market. (The result for more than two symmetric markets follows by induction.) The symmetric market with the same overall profit as these two symmetric sub-markets has sellers with reach $\sigma = \alpha\sigma_H + (1-\alpha)\sigma_L$. Minimum profit in the single market, $(1-\sigma)/\sigma$, lies between the minimum profits in the two sub-markets. If $G(\pi)$ denotes the CDF for the distribution of profit in the two sub-markets, and $\bar{G}(\pi)$ denotes the corresponding distribution in the single market, then calculations similar to (16) show that $\bar{G}$ crosses G from below at some profit below $(1-\sigma_L)/\sigma_L$ and does not cross it again thereafter until both $\bar{G}$ and G end at 1 at $\pi = \pi^*$. Therefore, since they have the same means, G corresponds to a mean-preserving spread of $\bar{G}$.

These two steps establish that when any set of sub-markets are replaced with a single symmetric market with the same overall profit, the distribution of profit in the former is a mean-preserving spread of that in the latter. To complete the proof, suppose the overall market is symmetric and initially there is discrimination across N sub-markets. If we replace these sub-markets with a single symmetric market with the same overall profit, the distribution of profit becomes less risky and consumers are made better off. Since profit is weakly higher with discrimination than without, sellers in this hypothetical symmetric market have weakly lower reach than they do in the overall market. One can verify from (15) that in a symmetric market with lower reach, consumers are offered higher profit, in the sense of first-order stochastic dominance, than in a symmetric market with higher reach, and hence they are worse off in the former case. This completes the proof. ■

Thus in a symmetric market any additional information about consumer consideration sets, not just the perfect information about captivity studied in the previous section,

[6] In general, suppose G and $\bar{G}$ are two CDFs corresponding to distributions for π in the range $[0, \pi^*]$, possibly with an atom at $\pi = \pi^*$ but otherwise continuous. If the distributions have the same mean, then the respective areas under G and $\bar{G}$ over $[0, \pi^*]$ are equal and hence the two continuous functions must cross. A sufficient condition for G to correspond to a mean-preserving spread of $\bar{G}$ is for the latter to cross the former only once and from below.

benefits firms and harms consumers relative to the regime with uniform pricing.

To illustrate, suppose a symmetric market is partitioned into a mirror pair of nested sub-markets as shown on Figure 3, one where firm 1 has no captives and the other where firm 2 has no captives (where contested customers are divided equally across the sub-markets). This information structure could arise if erstwhile regional energy monopolies with an existing customer base are permitted to serve each other's markets, where some consumers have low and others have high switching costs. The policy issue is whether or not a seller should be permitted to set distinct prices to its own customer base and to customers attached to the rival. Proposition 4 shows sellers are better off, while consumers in aggregate are worse off, with this form of price discrimination. It is clear that within a sub-market the seller with captive customers sets higher prices (in the sense of first-order stochastic dominance) than its weaker rival. Less obviously, even the weaker seller sets higher prices than it would in the uniform pricing regime. (The strong seller has the same number of captives in the nested sub-market as in the overall market but fewer contested customers, and so the weaker seller must set higher prices for the strong seller to be indifferent over its prices.)

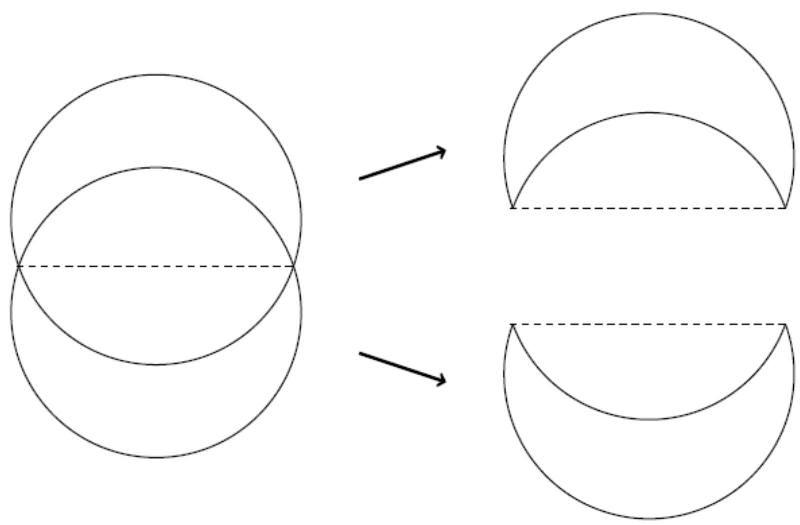

Figure 3: Partitioning a symmetric market into nested sub-markets

Therefore, this form of discrimination induces both sellers to raise their prices relative to the uniform pricing regime, and all consumers are harmed. This contrasts with our symmetric model in section 3, where price discrimination benefitted the contested consumers and average profit was unaffected. This example illustrates how information about consumer captivity might affect not only the variability of profits but also the degree of asymmetry and hence the competitive intensity of the market.

References

Armstrong, Mark, and John Vickers. 1993. "Price Discrimination, Competition and Regulation." *Journal of Industrial Economics* 41 (4): 335–357.

Armstrong, Mark, and John Vickers. 2001. "Competitive Price Discrimination." *Rand Journal of Economics* 32 (4): 579–605.

Baye, Michael, and John Morgan. 2002. "Information Gatekeepers and Price Discrimination on the Internet." *Economics Letters* 76 (1): 47–51.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The Limits of Price Discrimination." *American Economic Review* 105 (3): 921-957.

Burdett, Kenneth and Kenneth Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–969.

Butters, Gerard. 1977. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies* 44 (3): 465–491.

Chen, Yongmin, and Marius Schwartz. 2015. "Differential Pricing When Costs Differ: A Welfare Analysis." *Rand Journal of Economics* 46 (2): 442–460.

Corts, Kenneth. 1998. "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment." *Rand Journal of Economics* 29 (2): 306–323.

Dasgupta, Partha, and Eric Maskin. 1986. "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory." *Review of Economic Studies* 53 (1): 1–26.

Heidhues, Paul, and Botond Köszegi. 2017. "Naïveté-Based Discrimination." *Quarterly Journal of Economics* 132 (2): 1019–1054.

Narasimhan, Chakravarthi. 1988. "Competitive Promotional Strategies." *Journal of Business* 61 (4): 427–449.

Rothschild, Michael, and Joseph Stiglitz. 1970. "Increasing Risk I: A Definition." *Journal of Economic Theory* 2 (3): 225–243.

Thisse, Jacques-François, and Xavier Vives. 1988. "On the Strategic Choice of Spatial Price Policy." *American Economic Review* 78 (1): 122–137.

Varian, Hal. 1980. "A Model of Sales." *American Economic Review* 70 (4): 651–659.