

Personalized Pricing and Competition^{*}

Andrew Rhodes[†] Jidong Zhou[‡]

May 2022

Abstract

We study personalized pricing (or first-degree price discrimination) in a general oligopoly model. In the short-run, when the market structure is fixed, the impact of personalized pricing hinges on the degree of market coverage (i.e., how many consumers buy). If coverage is high (e.g., because the production cost is low, or the number of firms is large), personalized pricing intensifies competition and so harms firms but benefits consumers, whereas the opposite is true if coverage is low. However in the long-run, when the market structure is endogenous, personalized pricing always benefits consumers because it induces the socially optimal level of firm entry. We also study the asymmetric case where some firms can use consumer data to price discriminate while others cannot, and show it can be worse for consumers than when either all or no firms can personalize prices.

Keywords: personalized pricing, competition, price discrimination, consumer data

JEL classification: D43, D82, L13

^{*}We are grateful to Heski Bar-Isaac, Ioana Chioveanu, Alexandre de Cornière, Joyee Deb, Florian Ederer, Christian Hellwig, Johannes Hörner, Bruno Jullien, David Martimort, Leslie Marx, Barry Nalebuff, Ben Polak, Patrick Rey, David Ronayne, Larry Samuelson, Ali Shourideh, Alex Smolin, Guofu Tan, Nikhil Vellodi and various seminar audiences for helpful comments. Rhodes acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d’Avenir) program (grant ANR-17-EURE-0010). Some of the material in this paper was previously circulated under the title “Personalized Pricing and Privacy Choice.”

[†]Toulouse School of Economics. andrew.rhodes@tse-fr.eu

[‡]Yale School of Management. jidong.zhou@yale.edu

1 Introduction

Thanks to advances in information technology, firms are increasingly able to do personalized pricing, i.e., offer different consumers different prices based on their individual preferences.¹ Many firms have access to rich consumer-level data, which they have either collected themselves or acquired from data brokers. Using Artificial Intelligence (AI), firms can then glean information about a consumer’s preferences and target her with personalized offers. According to a recent survey, around 40% of firms that have adopted AI for personalization use it to set real-time prices and promotions (Deloitte, 2018). Indeed, personalized pricing has been documented in a wide range of industries, including retailing, travel, and personal finance (see., e.g., OECD, 2018). As far back as 2012, the Wall Street Journal uncovered evidence that retailers like Staples and Home Depot were personalizing prices on their websites, based on a consumer’s browsing history and distance from a competitor’s store.² In the last decade, firms’ ability to personalize prices has likely grown significantly—but so too has their ability to do it surreptitiously, so as to avoid a potential consumer backlash. As a result, personalized prices are often concealed, for example as personalized discounts which are sent by email or smartphone app.³

What is the welfare impact of personalized pricing? This is an important and timely question, given ongoing debates about big-tech firms and how they use consumer data. Policymakers are often wary of personalized pricing, raising concerns that it “transfers value from consumers to shareholders” (Council of Economic Advisers, 2015). In this paper we evaluate such concerns using a general oligopoly model. To do this, we focus on the limit case of perfect, or first-degree, price discrimination; as firms gain access to richer data and more sophisticated AI, this type of very fine-tuned personalization is likely to become increasingly feasible. We show that the welfare impact of such personalized pricing depends on market characteristics—such as cost conditions and the degree of competition (which both influence the extent of market coverage), the endogeneity of market structure, as well as asymmetries across firms in terms of their access to data.

We begin in Section 2 by reviewing two well-known benchmarks from the literature. The first benchmark is monopoly: here personalized pricing allows the firm to extract

¹See <https://bit.ly/3A4Rk10> and <https://bit.ly/38Ygqz6> for a history of personalized pricing.

²See <https://on.wsj.com/39sHIFf> for further details.

³As an example, see <https://bit.ly/370ftAc> for how Kroger uses its mobile app to offer personalized coupons. There are several other ways that firms can disguise their usage of personalized pricing. Certain products, such as financial products, are already somewhat personalized, so it is easy for a firm to also personalize their prices without consumers realizing (see, e.g., FCA, 2019). Even for more standardized products, firms can use “sticky targeting” whereby prices are fixed for all consumers (including the one being targeted) for a short period (see Shiller, 2021), or can personalize rankings and search results in order to steer a consumer towards products with a certain price point (e.g., Hannak et al., 2014).

all the social surplus, and so is good for the firm but bad for consumers. The second benchmark is the classic linear Hotelling model: Thisse and Vives (1988) show that personalized pricing (i.e., offering each consumer a different price based on their location on the Hotelling line) leads to a reduction in the price paid by *every* consumer. (Intuitively, each firm tries to poach consumers on its rival’s “turf” with low prices, which then forces the rival to charge less to its customer base.) Therefore going from monopoly to duopoly completely reverses the impact of personalized pricing—it now harms firms but benefits consumers. This insight has been very influential: as we discuss further in the literature review, the model of Thisse and Vives (1988) is an important building block for many subsequent papers that touch on issues such as data privacy and data brokers.

The first contribution of this paper is to reconcile the opposing impacts of personalized pricing under monopoly and Hotelling duopoly. In Section 3 we introduce a discrete-choice model which nests both monopoly and Hotelling as special cases. There is an arbitrary number of (single-product) firms, and consumers’ valuations for their products are drawn from a joint distribution. Consumers either buy one of the products or take an outside option. Our model is based on Perloff and Salop (1985), but is more general because it allows for correlated product valuations and partial market coverage (i.e., some consumers may take the outside option). Under uniform pricing firms cannot use information about individual consumers’ preferences, and so offer all consumers the same price. Under personalized pricing firms know each consumer’s valuations for all the products, and make personalized offers accordingly.⁴

Section 4 compares market performance in these two regimes. We first study the short-run case where the market structure is fixed. Contrary to Thisse and Vives (1988), we show under a mild regularity condition that some personalized prices *exceed* the uniform price: although consumers who regard their two best products as close substitutes pay less under personalized pricing, consumers with a strong preference for one product end up paying more. Nevertheless, if the market is fully covered under uniform pricing (i.e., if all consumers buy in equilibrium), competitive personalized pricing does lower industry profit and increase *aggregate* consumer surplus under a log-concavity condition. (The log-concavity condition ensures that there are relatively few consumers with strong preferences compared to those with weak preferences.) We therefore significantly generalize the aggregate welfare results from Thisse and Vives (1988).

We then show, however, that if the market is not fully covered, the impact of personalized pricing can be completely reversed. In particular, competitive personalized pricing can now increase industry profit and lower consumer surplus, as in the monopoly case.

⁴In Section 6 we consider an alternative information structure where firms only observe each consumer’s valuation for their own product.

This always happens—irrespective of the number of firms in the market—when the production cost is sufficiently high, such that market coverage is low under both pricing regimes. Indeed, when product valuations are independent across firms and follow an exponential distribution, it happens whenever marginal cost is such that the market is not fully covered under uniform pricing. Using numerical examples, we also show that the welfare impact of personalized pricing follows a cut-off rule for common distributions such as the Extreme value (which gives the logit model) and the Normal (which gives the probit model). Specifically, for a given number of firms, when the production cost is sufficiently low the impact is similar to Thisse and Vives (1988), when the cost is sufficiently high the impact is similar to monopoly, and when the cost is intermediate personalized pricing benefits both consumers and firms.⁵ Similarly, for a fixed production cost, with relatively few firms the impact is like in monopoly, with many firms the impact is similar to Thisse and Vives (1988), and otherwise industry profit and consumer surplus both increase.⁶

The intuition for why competitive personalized pricing can benefit firms and harm consumers is as follows. First, consider the case where all consumers value each product above marginal cost. Here, partial coverage arises when the uniform price excludes some low-valuation consumers from the market. Personalized pricing brings these consumers into the market, but since they have low valuations the positive effect on their surplus is relatively small. On the other hand, consumers who bought under uniform pricing have a high valuation for at least one product, and so relatively many of these consumers have a strong preference for one product over another—meaning that personalized pricing can raise the average price they pay. When this happens, personalized pricing can make consumers worse off overall, even though it expands demand. Second, consider the case where marginal cost is higher, such that for each product some consumers have a valuation below marginal cost. Now each firm faces a new “monopoly segment” of consumers, who value only its product above marginal cost. This gives firms some monopoly power, which is an additional force for personalized pricing to harm consumers.

We also study the long-run case where the market structure is endogenous. Specifically, we consider a free-entry game where firms choose whether or not to pay a fixed cost to enter the market, and then engage in price competition. We show that if the entry of a new product does not change consumers’ preferences over existing products, then with personalized pricing the new entrant fully extracts the increase in match efficiency caused by its entry. Consequently, in the long run, personalized pricing leads to the socially

⁵Note that if uniform pricing leads to only partial coverage, personalized pricing increases total surplus by expanding the market. This explains why profit and consumer surplus can both increase.

⁶Contrary to Thisse and Vives (1988), in our numerical examples the welfare impact of personalized pricing is often qualitatively the same under both monopoly and duopoly; in particular, it can raise industry profit and harm consumers even when the number of firms is relatively large.

optimal market structure. If we ignore integer constraints, this implies that personalized pricing must benefit consumers in the long run relative to uniform pricing.

In Section 5 we consider the case of asymmetrically informed firms, where only some firms have consumer data and can use it to price discriminate. We show that this “mixed” case can be worse for consumers than the symmetric cases where either all or no firms can personalize prices. Intuitively, when a firm with consumer data competes with other firms that can only do uniform pricing, it is able to “poach” some consumers for whom it is not their favorite product via a low personalized price. This results in match inefficiency compared to the symmetric cases and can make consumers worse off in aggregate. This suggests that it is sometimes desirable (at least from the perspective of consumers) to force a seller with superior information to share its data with its competitors or prevent it from personalizing prices.

We discuss other extensions in Section 6, including the case when firms offer personalized discounts and are constrained in how large they can be. Section 7 concludes.

Related literature. The literature on price discrimination is extensive, but it mainly focuses on imperfect price discrimination. (See the survey papers by Varian, 1989; Armstrong, 2007; Fudenberg and Villas-Boas, 2007; and Stole, 2007).⁷ One exception is the study of spatial price discrimination, where firms can charge customers in different locations different prices. An important contribution to this literature is Thisse and Vives (1988), which can also be reinterpreted as a model of competitive personalized pricing. They consider a two-stage game where firms first choose whether or not to price discriminate and then compete in prices. Using a Hotelling model with uniformly distributed consumers, they show that discriminatory pricing is a dominant strategy for each firm, and so the unique equilibrium features price discrimination. When firms have the same cost, as discussed earlier, they are trapped in a Prisoner’s dilemma because every personalized price is below the uniform price.⁸

The Hotelling setup in Thisse and Vives (1988) has been widely used in the subsequent literature. For example, Shaffer and Zhang (2002) use it to study personalized pricing when one firm has a brand advantage over the other, while Chen and Iyer (2002) use

⁷Note that perfect price discrimination can be regarded as the limit case of third-degree price discrimination when each consumer is treated as a separate market. Our paper is therefore more related to the literature on competitive third-degree price discrimination (see, e.g., Holmes, 1989; Corts, 1998; Armstrong and Vickers, 2001; Chen, Li, and Schwartz, 2021; and Adachi, 2022). However, the approaches in that literature (e.g., the idea of best-response asymmetry in Corts, 1998 or the indirect utility approach in Armstrong and Vickers, 2001) are not directly useful for studying our problem.

⁸When firms have different costs, the low-cost firm can earn more than under uniform pricing, but within the parameter range in Thisse and Vives (1988) industry profit is still lower and consumer surplus is still higher under discriminatory pricing.

it to study personalized pricing when firms first need to advertise to reach consumers. Montes, Sand-Zantman, and Valletti (2019) use it to study whether a monopolistic data intermediary should sell data to one or both competing firms who can use the data to conduct personalized pricing. Chen, Choe, and Matsushima (2020) use it to study consumer identity management which helps consumers avoid being exploited by firms via personalized pricing. In all these studies, an implicit underlying assumption is that competitive personalized pricing in the benchmark case intensifies competition, harms firms and benefits consumers. Our paper shows that this is not necessarily true in a more general model which allows for partial market coverage.⁹

Our paper is closely related to Anderson, Baik, and Larson (2021) (ABL henceforth), who also use a general discrete-choice framework to study competitive personalized pricing.¹⁰ One important difference is that they have full market coverage—whereas our paper allows for partial market coverage, and emphasizes that this can qualitatively change the impact of personalized pricing. Another important difference is that in our paper firms can freely offer personalized prices, leading to a relatively simple pure-strategy pricing equilibrium; ABL, by contrast, assume that it is costly for firms to send targeted discounts, which leads to a mixed strategy equilibrium in both pricing and advertising.¹¹ (Our modeling choice captures the idea that the cost of making personalized offers is mainly a fixed one, due to investments in buying consumer data and developing AI tools.)

There is also growing empirical research on personalized pricing. One strand looks for evidence of firms engaging in personalized pricing. As discussed earlier, detecting personalized pricing is usually hard because sellers have incentives to disguise personalized offers, but nevertheless there is some suggestive evidence. For instance, Hannak et al. (2014) find evidence of some form of personalization on 9 out of 16 e-commerce sites in their study, while Aparicio, Metzman, and Rigobon (2021) document evidence that increasing use of algorithmic pricing is associated with increasing price differentiation (for the same product at the same time but across different delivery zipcodes). The other strand of the empirical literature assesses the impact of personalized pricing (see, e.g., Waldfogel, 2015; Dube and Misra, 2019; Shiller, 2020; and Kehoe, Larsen, and Pastorino,

⁹Jullien, Reisinger, and Rey (2022) develop a generalized Hotelling setup, and show for example that personalized pricing can raise profit when a manufacturer competes with a retailer that sells its product.

¹⁰See also Section 4 of Ali, Lewis, and Vasserman (2020) which uses a similar oligopoly discrete-choice model with full market coverage to study endogenous consumer information disclosure. They show that the Thisse and Vives (1988) welfare result continues to hold in a partial revelation equilibrium.

¹¹Such randomized offers cause consumers to sometimes buy the wrong product, which harms match efficiency. As a result, in ABL personalized pricing can make both firms and consumers worse off compared to uniform pricing. In contrast, in our model it is possible that both firms and consumers are better off under personalized pricing as it can expand demand.

2020). For instance, Shiller (2020) shows that if Netflix could use consumer information from rich web-browsing data to implement price discrimination, its profit could increase by about 13%, while the profit improvement would be tiny if it only relied on demographic information.

2 Two Benchmarks

We start by briefly recapping two well-known benchmarks from the existing literature.

Monopoly The impact of personalized pricing (or perfect price discrimination) under monopoly is straightforward. Suppose consumers wish to buy at most one unit of a product, and have heterogeneous valuations for it. Under uniform pricing, the firm sets a standard monopoly price. Consumers who value the product more than the monopoly price buy and obtain positive surplus; all other consumers are excluded from the market. Under personalized pricing, each consumer with a valuation above marginal cost is offered a personalized price exactly equal to their valuation, and they all buy. As a result, total surplus is maximized but it is fully extracted by the monopolist. Personalized pricing therefore increases total welfare and firm profit but reduces consumer surplus.

Hotelling duopoly The other well-known case is the linear Hotelling model studied by Thisse and Vives (1988). Suppose consumers are uniformly distributed along a unit-length Hotelling line. Suppose the two firms have cost normalized to zero, with firm 1 located at the leftmost point on the line, and firm 2 located at the rightmost point. A consumer with location x values firm 1's product at $v_1 = V - x$ and firm 2's product at $v_2 = V - (1 - x)$, where V is large enough that the market is fully covered in equilibrium. Under uniform pricing firms set the standard Hotelling price of 1. Under personalized pricing the firms compete for each consumer individually. Consumers with location $x < 1/2$ prefer product 1, while consumers with location $x > 1/2$ prefer product 2. Firms therefore engage in asymmetric Bertrand competition. Thisse and Vives (1988) derive the following equilibrium price schedules:

$$\begin{aligned} p_1(x) = v_1 - v_2 = 1 - 2x \quad \text{and} \quad p_2(x) = 0 \quad \text{for } x \in [0, \tfrac{1}{2}] \\ p_1(x) = 0 \quad \text{and} \quad p_2(x) = v_2 - v_1 = 2x - 1 \quad \text{for } x \in [\tfrac{1}{2}, 1] \end{aligned} \quad (1)$$

where $p_i(x)$ is the price offered by firm $i = 1, 2$ to the consumer at x , and each consumer buys her preferred product, with those having stronger preferences paying more. Note that each consumer pays (weakly) less under personalized pricing because $p_i(x) \leq 1$. Therefore, unlike with monopoly, personalized pricing harms firms and benefits consumers. (Under

both uniform and personalized pricing the market is fully covered and consumers buy their preferred product, so personalized pricing has no impact on total welfare.)

However, the observation that each personalized price is lower than the uniform price, which is highlighted in Thisse and Vives and many subsequent works, can easily be overturned. To see this, suppose instead that consumers are distributed along the Hotelling line according to a symmetric and strictly log-concave (so single-peaked) density. Discriminatory prices are independent of the distribution and so are the same as in (1), but the uniform price, which equals 1 over the density of consumers at $x = 1/2$, is now strictly below 1. As a result, consumers near the two ends of the line (i.e., those with relatively strong preferences) now pay more under personalized pricing, while consumers near the middle of the line (i.e., those with relatively weak preferences) still pay less.¹² The impact of personalized pricing on industry profit and (aggregate) consumer surplus is then less clear. In the next section, we develop a more general oligopoly model—which includes the Hotelling model as a special case—and investigate to what extent competitive personalized pricing is overall pro-competitive.

3 A General Oligopoly Model

There are n competing firms in a market, each supplying a differentiated product at constant marginal cost c . There is also a unit mass of consumers, each wishing to buy at most one of the products. If a consumer buys nothing she obtains an outside option with zero surplus.¹³ Let $\mathbf{v} = (v_1, \dots, v_n) \in \mathbf{R}^n$ denote a consumer's valuations for the n products. In the population \mathbf{v} is distributed according to an *exchangeable* joint cumulative distribution function (CDF) $\tilde{F}(\mathbf{v})$, with corresponding density function $\tilde{f}(\mathbf{v})$. (The exchangeability means that any permutation of (v_1, \dots, v_n) has the same joint CDF; it implies that there are no systematic quality differences across products.) Let F and f be respectively the common marginal CDF and density function of each v_i , and let $[\underline{v}, \bar{v}]$ be its support. (We need $c < \bar{v}$ to have an active market.) To ease the exposition, we assume that \tilde{F} has full support on $[\underline{v}, \bar{v}]^n$, but this is not crucial for the main results.

Note that although we allow a consumer's valuations for different products to be correlated,¹⁴ sometimes we focus on the IID case where the v_i 's are independent across

¹²Armstrong (2007) makes the same point by considering a specific non-uniform distribution.

¹³To have our model cover both the case of full market coverage and the case of partial market coverage, we have chosen to normalize the outside option and vary the marginal cost c . The same qualitative insights obtain if instead we normalize the production cost and vary the outside option.

¹⁴In the duopoly case, our set-up nests Hotelling with a symmetric location distribution if v_1 and v_2 are large enough (to cover the market) and we treat $v_1 - v_2$ as a consumer's location. For any location distribution (e.g., uniform), there is at least one correlation structure over (v_1, v_2) that generates it.

products (which is the leading case in the literature on random-utility oligopoly models).

We consider two different pricing regimes. Under uniform pricing, firms set the same price for every consumer (either because they have no data on consumer preferences, or are forbidden from using it). Under personalized pricing, firms perfectly observe each consumer's vector of valuations $\mathbf{v} = (v_1, \dots, v_n)$, and offer them a personalized price. (See Section 6 for an alternative information structure where firm i only observes v_i .) In either regime, consumers perfectly know their own valuations for the competing products and make the best choice after seeing their prices. Both firms and consumers are assumed to be risk neutral.

Notation. It will be convenient to introduce the following notation. Let $G(\cdot|v_i)$ and $g(\cdot|v_i)$ be respectively the CDF and density function of $\max_{j \neq i} \{v_j\}$, the valuation for firm i 's best competing product, conditional on v_i . Let $v_{n:n}$ and $v_{n-1:n}$ be the highest and second-highest order statistics of (v_1, \dots, v_n) , and let $F_{(n)}(v)$ and $F_{(n-1)}(v)$ be their respective CDFs. Then

$$F_{(n)}(v) = \tilde{F}(v, \dots, v) = \int_{\underline{v}}^v G(v|v_i) dF(v_i) , \quad (2)$$

and

$$F_{(n-1)}(v) = F_{(n)}(v) + n \int_v^{\bar{v}} G(v|v_i) dF(v_i) . \quad (3)$$

To understand $F_{(n-1)}(v)$, notice that for the second-highest valuation to be below v , either all the v_i 's must be less than v , or exactly one of them must be above v and the others be below v . Let $f_{(n)}(v)$ and $f_{(n-1)}(v)$ be respectively their density functions. In the IID case we have $\tilde{F}(\mathbf{v}) \stackrel{\text{IID}}{=} \prod_{i=1}^n F(v_i)$, $G(v|v_i) \stackrel{\text{IID}}{=} F(v)^{n-1}$, $F_{(n)}(v) \stackrel{\text{IID}}{=} F(v)^n$, and

$$F_{(n-1)}(v) \stackrel{\text{IID}}{=} F(v)^n + n(1 - F(v))F(v)^{n-1} .$$

In order to solve the uniform pricing game, it is useful to define the random variable

$$x_z \equiv v_i - \max_{j \neq i} \{z, v_j\} , \quad (4)$$

where z is a constant. Since $x_z = v_i - z - \max_{j \neq i} \{0, v_j - z\}$, one can interpret it as a consumer's preference for product i relative to the best alternative (including the outside option) when all products are sold at price z . Let $H_z(x)$ and $h_z(x)$ be respectively the CDF and density function of x_z . More explicitly,

$$1 - H_z(x) = \Pr[v_i - x > \max_{j \neq i} \{z, v_j\}] = \int_{z+x}^{\bar{v}} G(v_i - x|v_i) dF(v_i) , \quad (5)$$

and

$$h_z(x) = G(z|z+x)f(z+x) + \int_{z+x}^{\bar{v}} g(v_i - x|v_i) dF(v_i) . \quad (6)$$

When z is irrelevant (i.e., when $z \leq \bar{v}$), let $H(x)$ and $h(x)$ be respectively the CDF and density function of $x \equiv v_i - \max_{j \neq i} \{v_j\}$; we use them for the case of full market coverage.¹⁵

3.1 Uniform pricing

We first study the regime of uniform pricing, where firms are unable to price discriminate. We focus on a symmetric pure-strategy pricing equilibrium, and let p denote the equilibrium uniform price.¹⁶ Recalling the definition of x_z and $H_z(x)$ in equations (4) and (5), when firm i unilaterally deviates to a price p_i its deviation demand is

$$\Pr[v_i - p_i > \max_{j \neq i} \{0, v_j - p\}] = \Pr[v_i - \max_{j \neq i} \{p, v_j\} > p_i - p] = 1 - H_p(p_i - p) ,$$

and its deviation profit is $(p_i - c)[1 - H_p(p_i - p)]$. It is clear that a firm will never set a price below marginal cost c or above the maximum valuation \bar{v} .

To ensure that the uniform pricing equilibrium is uniquely determined by the first-order condition, we make the following assumption:

Assumption 1. $1 - H_z(x)$ is log-concave in x and $\frac{1 - H_z(0)}{h_z(0)}$ is non-increasing in z .

In the Appendix we report some primitive conditions under which this assumption holds. For example, the first condition holds if the joint density \tilde{f} is log-concave (Caplin and Nalebuff, 1991), and both conditions hold in the IID case with a log-concave f . (The second condition must hold if, for $z < z'$, x_z is greater than $x_{z'}$ in the sense of hazard rate dominance.) Assumption 1 also holds in the Hotelling case (see footnote 14) provided that $v_1 - v_2$ has a log-concave density.

Given the first condition in Assumption 1, firm i 's deviation profit is log-concave in p_i , and so the equilibrium price p must solve the first-order condition

$$p - c = \phi(p) , \tag{7}$$

where

$$\phi(p) \equiv \frac{1 - H_p(0)}{h_p(0)} = \frac{\int_p^{\bar{v}} G(v|v) dF(v)}{G(p|p)f(p) + \int_p^{\bar{v}} g(v|v) dF(v)} . \tag{8}$$

To interpret this, note that $1 - H_p(0)$ is each firm's equilibrium demand,¹⁷ while $h_p(0)$ is the absolute value of the equilibrium demand slope and it measures how many consumers

¹⁵Anderson, Baik, and Larson (2021) also use such notation to simplify demand expressions when the market is assumed to be fully covered.

¹⁶If the joint density \tilde{f} is log-concave, the pricing equilibrium is unique and symmetric in the duopoly case (Caplin and Nalebuff, 1991) and in the IID case (Quint, 2014).

¹⁷Due to firm symmetry we can also write that $1 - H_p(0) = \frac{1}{n}[1 - F_{(n)}(p)]$.

are marginal for each firm. (The first term in $h_p(0)$ captures the extensive margin, and the second term captures the intensive margin due to competition.)

Given the second condition in Assumption 1, $\phi(p)$ is non-increasing. Notice also that $\phi(p)$ is constant for $p \leq \underline{v}$. Then we can show the following result. (All omitted proofs are available in the Appendix.)

Lemma 1. *Suppose Assumption 1 holds.*

(i) *If $c \leq \underline{v} - \phi(\underline{v})$, the equilibrium uniform price satisfies*

$$p - c = \phi(\underline{v}) = \frac{1/n}{\int_{\underline{v}}^{\bar{v}} g(v|v) dF(v)} \quad (9)$$

and $p \leq \underline{v}$, such that the market is fully covered in equilibrium.

(ii) *Otherwise, the equilibrium uniform price uniquely solves (7) and $p > \underline{v}$, such that the market is not fully covered in equilibrium.*

Intuitively, when cost is relatively low ($c \leq \underline{v} - \phi(\underline{v})$), marginal consumers are sufficiently valuable that firms choose to cover the whole market; when cost is relatively high ($c > \underline{v} - \phi(\underline{v})$), firms optimally exclude some low-valuation consumers. Note that a sufficient (but by no means necessary) condition for partial coverage is that $\underline{v} \leq c$, i.e., some consumers value a product less than marginal cost.

The literature on random-utility oligopoly models usually studies the IID case, such that $G(v|v) \stackrel{\text{IID}}{=} F(v)^{n-1}$ and $g(v|v) \stackrel{\text{IID}}{=} (n-1)F(v)^{n-2}f(v)$, and so

$$\phi(p) \stackrel{\text{IID}}{=} \frac{[1 - F(p)^n]/n}{F(p)^{n-1}f(p) + \int_p^{\bar{v}} f(v)dF(v)^{n-1}}. \quad (10)$$

Most papers further assume that the market is covered (e.g., Perloff and Salop, 1985; Anderson, Baik, and Larson, 2021), in which case $\phi(p)$ simplifies to $1/[n \int_{\underline{v}}^{\bar{v}} f(v)dF(v)^{n-1}]$.¹⁸

Example: uniform distribution. Suppose the v_i 's are independent and uniformly distributed on $[\underline{v}, \underline{v}+1]$. Here, if $p \leq \underline{v}$ then $\phi(p) = \frac{1}{n}$, and if $p > \underline{v}$ then $\phi(p) = \frac{1}{n}[1 - (p - \underline{v})^n]$. Therefore if $c + \frac{1}{n} \leq \underline{v}$ the market is fully covered and the equilibrium price is $p = c + \frac{1}{n}$; otherwise the market is not fully covered and $p > \underline{v}$ uniquely solves

$$p - c = \frac{1 - (p - \underline{v})^n}{n}. \quad (11)$$

Example: exponential distribution. Suppose the v_i 's are independent and exponentially distributed with $F(v) = 1 - e^{-(v-\underline{v})}$ on $[\underline{v}, \infty)$. Here we have $\phi(p) = 1$, and so the

¹⁸An exception is Section 4.2 of Zhou (2017), which shows that in the IID case $\phi(p)$ in (10) is decreasing and the equilibrium price decreases in n if f is log-concave.

equilibrium price is $p = c + 1$ regardless of whether or not the market is covered (and irrespective of the number of firms).¹⁹ Therefore the market is fully covered in equilibrium if and only if $c + 1 \leq \underline{v}$.

Industry profit under uniform pricing is

$$\Pi_U \equiv n(p - c)[1 - H_p(0)] = n \frac{[1 - H_p(0)]^2}{h_p(0)}, \quad (12)$$

where we have used the equilibrium price condition (7). Since all consumers buy their favorite product as long as it has a positive surplus, (aggregate) consumer surplus is

$$V_U \equiv \mathbb{E}[\max\{0, v_{n:n} - p\}] = \int_p^{\bar{v}} (v - p) dF_{(n)}(v) = \int_p^{\bar{v}} [1 - F_{(n)}(v)] dv, \quad (13)$$

where the last equality is from integration by parts. Notice that these expressions are valid regardless of whether or not the market is fully covered.

3.2 Personalized Pricing

Now consider the regime where firms perfectly observe each consumer's vector of valuations $\mathbf{v} = (v_1, \dots, v_n)$ and set personalized prices accordingly. In this case, firms engage in a game of asymmetric Bertrand competition for each consumer. To rule out uninteresting equilibria, we assume that firms do not play dominated strategies, and that when consumers are indifferent between several offers they buy the product with the highest valuation (so that total welfare is maximized). Consider a consumer who values, say, firm 1's product the highest and firm 2's product the second highest. Competition is then essentially between these two firms. Suppose $v_1 \geq c$ (otherwise the consumer takes the outside option). Competition forces firm 2 to price at marginal cost. Firm 1 prices at $c + v_1 - v_2$ if $v_2 \geq c$, and otherwise acts as a monopolist and charges v_1 . In both cases the other $n - 2$ firms charge weakly more than c , and firm 1 sells to the consumer. (Since the prices of these $n - 2$ firms can be anything above c , there are multiple equilibria, but they are all outcome-equivalent.) To ease the exposition, we henceforth focus on the equilibrium in which these other $n - 2$ firms charge c . Firm i 's equilibrium pricing schedule can then be written as:

$$p(v_i, \mathbf{v}_{-i}) = \begin{cases} c + v_i - \max_{j \neq i} \{c, v_j\} & \text{if } v_i \geq \max_{j \neq i} \{c, v_j\} \\ c & \text{otherwise} \end{cases} \quad (14)$$

¹⁹Using integration by parts, the denominator in (10) can be rewritten as $f(\bar{v}) - \int_p^{\bar{v}} F(v)^{n-1} f'(v) dv$. For the exponential distribution $f(\bar{v}) = 0$ and $f(v) = -f'(v)$, so this equals the numerator of (10).

where \mathbf{v}_{-i} denotes a consumer's valuations for all products other than i . Intuitively, if a firm's product is a consumer's favorite and has a valuation above cost, the firm charges the consumer a price equal to the difference between her valuation for its product and that of the best alternative (which is either the outside option, or the best rival product sold at marginal cost).²⁰ Note that if $\underline{v} \geq c$, the “max” constraint from the outside option in (14) is irrelevant and so all consumers buy their favorite product in equilibrium (i.e., the market is fully covered under personalized pricing).

To calculate profit, notice that when firm i wins a consumer, its profit margin is $p(v_i, \mathbf{v}_{-i}) - c = v_i - \max_{j \neq i} \{c, v_j\} = x_c$, where x_c has a CDF $H_c(x)$. That is, firm i earns x_c whenever $x_c \geq 0$, so its equilibrium profit is $\int_0^\infty x dH_c(x)$, and industry profit under personalized pricing is

$$\Pi_D = n \int_0^\infty x dH_c(x) = n \int_0^\infty [1 - H_c(x)] dx . \quad (15)$$

(There are alternative ways to calculate profit as we will show later.)

Consumers always buy their favorite product (as long as it is valued above c). Given the equilibrium pricing schedule in (14), it is clear that a consumer's favorite firm sets its price such that the consumer is indifferent between its product and the next best option (which is either the outside option or the second-best product sold at marginal cost). Therefore, consumer surplus under personalized pricing is

$$V_D \equiv \mathbb{E}[\max\{0, v_{n-1:n} - c\}] = \int_c^{\bar{v}} (v - c) dF_{(n-1)}(v) = \int_c^{\bar{v}} [1 - F_{(n-1)}(v)] dv . \quad (16)$$

Notice that expressions (15) and (16) are both valid regardless of whether or not $c < \underline{v}$.

4 The Impact of Personalized Pricing

We now examine how a shift from uniform to personalized pricing affects market performance. We first study the short-run impact, when the number of firms n is taken as given. In the case of full market coverage, we significantly generalize the insight from Thisse and Vives (1988), showing that competitive personalized pricing harms firms and benefits consumers (on average). However, we also show that when the market is only partially covered, their insight can be completely overturned. We then study the long-run impact, when n is determined by firms' free-entry decisions.

²⁰Note that a consumer-specific pricing schedule includes uniform pricing as a special case. Therefore, in an extended game where firms simultaneously choose whether to adopt discriminatory pricing and what prices to offer, the only equilibrium is that all firms adopt personalized pricing.

4.1 The short-run impact with a fixed market structure

Suppose the number of firms is fixed. Our first result shows that, under a mild regularity condition, the highest personalized price exceeds the uniform price. As a result, some consumers benefit from personalized pricing while others suffer. Recall that $h(x)$ is the density of $v_i - \max_{j \neq i} \{v_j\}$.

Lemma 2. *Suppose Assumption 1 holds, and that $h(x) < h(0)$ for $x > 0$. Then the highest personalized price exceeds the uniform price.*

Proof. Using equation (14) the highest personalized price is $p_{\max} = c + \bar{v} - \max\{c, \underline{v}\}$. If $\underline{v} \leq c$, $p_{\max} = \bar{v}$ and so it must exceed the uniform price. If $\underline{v} > c$, $p_{\max} = c + \bar{v} - \underline{v}$ and so $p < p_{\max}$ if and only if $p - c < \bar{v} - \underline{v}$. Under Assumption 1, $\phi(p)$ is decreasing and so the uniform price must satisfy $p - c = \phi(p) \leq \phi(\underline{v}) = \frac{1}{nh(0)}$. At the same time,

$$\frac{1}{n} = \int_0^{\bar{v}-\underline{v}} h(x) dx < h(0)(\bar{v} - \underline{v}) ,$$

where the equality is from the fact that the probability of $v_i \geq \max_{j \neq i} \{v_j\}$ is $\frac{1}{n}$, and the inequality is from the assumption that $h(x) < h(0)$ for $x > 0$. Therefore we have $p - c < \bar{v} - \underline{v}$. \square

Note that the hypotheses of the lemma hold, for example, in the IID case with a log-concave f . However they fail in the linear Hotelling model studied earlier in Section 2, because there $h(x)$ is constant in $x \geq 0$ (which explains why, in that case, the highest personalized price exactly equals the uniform price as shown in Thisse and Vives, 1988).

The remainder of this subsection addresses the subtler question of how personalized pricing affects profit and aggregate consumer surplus.

4.1.1 The case of full market coverage

We first study the case where the market is fully covered under uniform pricing, i.e., where $p \leq \underline{v}$. From Lemma 1, this happens when $c \leq \underline{v} - \phi(\underline{v})$. This condition in turn implies $c < \underline{v}$, which means that the market is also fully covered under personalized pricing. Total welfare is therefore the same under uniform and personalized pricing, because in both cases all consumers buy their preferred product. The following result reports the impact of personalized pricing on profit and consumer surplus.

Proposition 1. *Suppose Assumption 1 holds and $c \leq \underline{v} - \phi(\underline{v})$ (in which case the market is fully covered under both pricing regimes). Then for any $n \geq 2$, relative to uniform pricing, personalized pricing harms firms and benefits consumers.*

Proof. Under the stated full-coverage condition, $x_z = v_i - \max_{j \neq i} \{z, v_j\}$ simplifies to $x = v_i - \max_{j \neq i} \{v_j\}$ for both $z = p$ and $z = c$ as $c < p \leq \underline{v}$. Recall that H and h are respectively the CDF and density function of x . Then industry profit under uniform pricing is

$$\Pi_U = p - c = \frac{1}{nh(0)} ,$$

while under personalized pricing it is

$$\Pi_D = n \int_0^\infty [1 - H(x)] dx = n \int_0^\infty \frac{1 - H(x)}{h(x)} dH(x) \leq n \frac{[1 - H(0)]^2}{h(0)} = \frac{1}{nh(0)} .$$

The inequality follows because, under Assumption 1, $1 - H$ is log-concave and therefore $\frac{1-H}{h}$ is decreasing. The final equality follows because firm symmetry implies $1 - H(0) = \frac{1}{n}$. Therefore, firms suffer from personalized pricing. Since total welfare is unchanged, consumers benefit from personalized pricing. \square

The intuition for this result is as follows. Notice that consumers with a relatively small gap between their top two valuations pay less under personalized pricing, while the reverse is true for consumers with a relatively large gap between their top two valuations. Under log-concavity (in Assumption 1) there are relatively more of the former consumers, and so personalized pricing harms firms but benefits consumers in aggregate. Note that since our set-up includes Hotelling as a special case (see footnote 14), Proposition 1 significantly generalizes the result in Thisse and Vives (1988).²¹

4.1.2 The case of partial market coverage

We now turn to the (perhaps more realistic) case where the market is not fully covered under uniform pricing. From Lemma 1, we know this happens when $c > \underline{v} - \phi(\underline{v})$. One simple impact of personalized pricing is that it now expands total demand: under uniform pricing, a consumer buys if the best match is above the uniform price $p > c$; under personalized pricing, a consumer buys if the best match is above c . Personalized pricing therefore strictly improves total welfare.

Before investigating the impact on firms and consumers, we offer an alternative formula to calculate industry profit under personalized pricing, which is more convenient to use in some of the subsequent analysis. Conditional on firm i winning a consumer and its

²¹Although not highlighted in Anderson, Baik, and Larson (2021), this generalization of Thisse-Vives is also implied by their Proposition 6 which does comparative statics with respect to the advertising cost in their model. Our proof is similar to that of Proposition 7 in ABL which shows the opposite result when $1 - H$ is log-convex. If we were to instead work directly with the primitive valuation distribution \tilde{F} , this result would be considerably harder to prove, even in the IID case.

product being valued at v_i , its expected profit margin is

$$m(v_i) \equiv v_i - \int_{\underline{v}}^{v_i} \max\{c, x\} d \frac{G(x|v_i)}{G(v_i|v_i)} = \frac{\int_c^{v_i} G(x|v_i) dx}{G(v_i|v_i)}, \quad (17)$$

where we have used (14) and integration by parts. Then industry profit under personalized pricing is

$$\Pi_D = n \int_c^{\bar{v}} m(v_i) G(v_i|v_i) dF(v_i) = n \int_c^{\bar{v}} [G(v|\bar{v}) - G(v|v)F(v)] dv, \quad (18)$$

where the second equality is from using (17) and integration by parts. In the IID case, $G(v|\bar{v}) = G(v|v) = F(v)^{n-1}$, so it simplifies to

$$\Pi_D \stackrel{\text{IID}}{=} \int_c^{\bar{v}} \frac{1 - F(v)}{f(v)} dF(v)^n. \quad (19)$$

We will now show that when the market is only partially covered, competitive personalized pricing can raise profit and lower aggregate consumer surplus. To understand why, it is useful to first investigate why the simple proof in Proposition 1 breaks down with partial coverage. Under Assumption 1, we still have that

$$\Pi_D = n \int_0^\infty [1 - H_c(x)] dx \leq n \frac{[1 - H_c(0)]^2}{h_c(0)}, \quad (20)$$

but now the last term is greater than

$$\Pi_U = n \frac{[1 - H_p(0)]^2}{h_p(0)},$$

because $p > c$ and both $1 - H_z(0)$ and $\frac{1 - H_z(0)}{h_z(0)}$ decrease in z . (In the full-coverage case, $c < p \leq \underline{v}$ and so $H_c = H_p = H$.) This observation also suggests that if $1 - H_z(x)$ is log-linear in x , then the inequality in (20) binds and so we have $\Pi_D > \Pi_U$ whenever the market is not fully covered. That is indeed what we show in the following example.

An exponential distribution example. Before pursuing some general analytic results, it is illuminating to first consider an IID example with an exponential distribution $F(v) = 1 - e^{-(v-\underline{v})}$. As shown on page 11, in this exponential example the uniform price is $c + 1$ irrespective of whether or not the market is fully covered. With full market coverage (which requires $c \leq \underline{v} - 1$), personalized pricing has *no* impact on profit or consumer surplus in this exponential example. This can be seen from noticing that $1 - H(x)$ is log-linear.²²

²²Alternatively, notice that with a covered market $\Pi_U = 1$. To show $\Pi_D = 1$, use equation (15) and

$$1 - H(x) = \int_{\underline{v}+x}^\infty F(v-x)^{n-1} dF(v) = \int_{\underline{v}}^\infty F(t)^{n-1} f(t+x) dt = e^{-x} \int_0^1 (1-u)^{n-1} du = \frac{e^{-x}}{n},$$

where the third equality is from changing the integral variable from t to $u = e^{-(t-\underline{v})}$.

We now show that with partial coverage (meaning that $c > \underline{v} - 1$), personalized pricing always *benefits* firms but *harms* consumers. This example suggests that, at least for some distributions close to the exponential one, moving away from full market coverage can qualitatively change the impact of personalized pricing.

Under uniform pricing with price $1 + c$, a fraction $F(1 + c)^n$ of consumers are excluded from the market, so industry profit is $\Pi_U = 1 - F(1 + c)^n$. Under personalized pricing, using (19) and the fact that $1 - F(v) = f(v)$ in this exponential example, we immediately have $\Pi_D = 1 - F(c)^n$. Therefore, under the condition for partial coverage ($1 + c > \underline{v}$), $\Pi_D > \Pi_U$, i.e., personalized pricing *boosts* profit.

The impact on consumer surplus is still unclear, because as noted earlier personalized pricing also increases total welfare. Consumer surplus under uniform pricing and personalized pricing are respectively

$$V_U = \int_{1+c}^{\infty} (v - c) dF(v)^n - \Pi_U \text{ and } V_D = \int_c^{\infty} (v - c) dF(v)^n - \Pi_D ,$$

where the integral term in each expression is the total welfare in each regime. The former is greater than the latter if and only if

$$F(1 + c)^n - F(c)^n > \int_c^{1+c} (v - c) dF(v)^n ,$$

which is true as $v - c < 1$ for $v \in (c, 1 + c)$. Therefore, personalized pricing boosts profit so much that consumers always *suffer* from it.

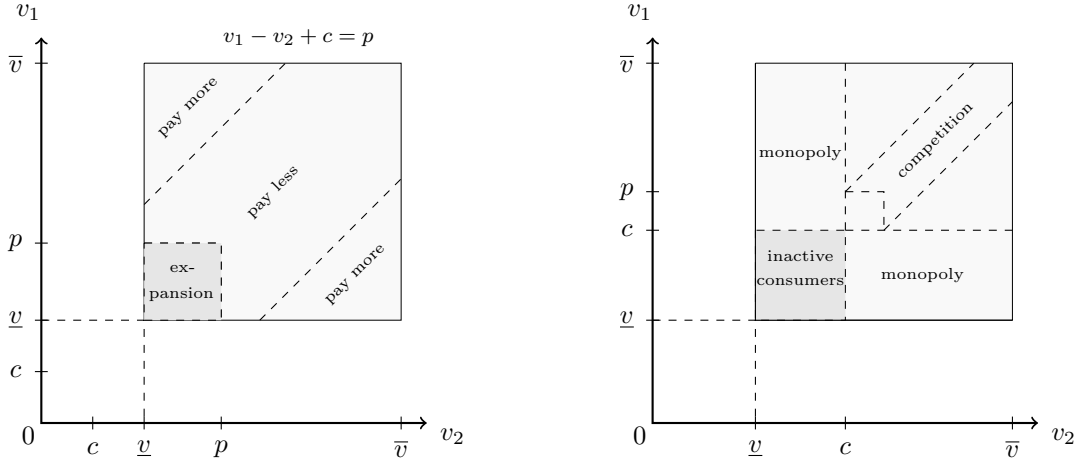
One way to see the intuition is as follows. Notice that under personalized pricing total demand is $1 - F(c)^n$, and so the *average* price that consumers pay is $1 + c$, which is exactly equal to the uniform price. Personalized pricing therefore raises profit, because it expands the size of the market. At the same time, this market expansion is from consumers whose highest valuation is between c and $1 + c$ —and since this is below the average price, personalized pricing lowers aggregate consumer surplus. This insight can be generalized: if personalized pricing weakly raises the average price paid by consumers, it must harm consumers overall even though it also expands the demand.

The production cost and market coverage. Given the full-coverage result in Proposition 1, it is clear that for a more general (regular) distribution, the impact of personalized pricing can only be reversed when the market is sufficiently far away from being fully covered. As we saw earlier, by changing the marginal cost c we can change the degree of market coverage. In particular, when c is sufficiently close to the valuation upper bound \bar{v} , most consumers are excluded from the market. In that case we can show that the

impact of personalized pricing is completely different from the full coverage case.²³

Proposition 2. *Suppose $\bar{v} < \infty$ and $f(\bar{v}) > 0$. For any given $n \geq 2$, there exists $\hat{c} < \bar{v}$ such that when $c > \hat{c}$ personalized pricing benefits firms and harms consumers compared to uniform pricing.*

To understand this result—and more generally how varying c affects the impact of personalized pricing—we refer to the following graphs which illustrate the duopoly case. (The graphs also work for the $n > 2$ case if we interpret v_2 as $\max_{j \geq 2} \{v_j\}$.)



(a) The case of $c < \underline{v}$. As c increases, consumers in the expansion region benefit, but others suffer.

(b) The case of $c > \underline{v}$. As c increases, the monopoly region becomes more important relative to the competition region.

Figure 1: The impact of personalized pricing with partial market coverage

Consider first the case depicted in Figure 1a, where $c < \underline{v}$ (so the market is fully covered under personalized pricing) but $p > \underline{v}$ (so the market is only partially covered under uniform pricing). Under personalized pricing, the consumers with $v_1 > v_2$ buy from firm 1 and pay $v_1 - v_2 + c$, and the other half buy from firm 2 and pay $v_2 - v_1 + c$. Compared to the regime of uniform pricing with price p , those consumers in the northwest and southeast corners pay more, those with weaker preferences in the middle region pay less, and those in the “expansion” region, who were excluded from the market under uniform pricing, now buy. As c increases, the “expansion” triangle grows, which is a positive effect for consumers; at the same time the two corner regions expand as the lines

²³One may wonder whether this result is immediate from (13) and (16). From those expressions, it is clear that if c and p were sufficiently close to each other and meanwhile both were bounded away from \bar{v} , then $V_U > V_D$ given $v_{n:n}$ is greater than $v_{n-1:n}$ in the sense of first-order stochastic dominance. This ranking, however, is no longer obvious when both c and p approach \bar{v} .

of $|v_1 - v_2| + c = p$ move toward the center, which is a negative effect for consumers. (The latter effect happens because the uniform price p increases in c slower than the personalized prices, which can be seen from equation (7). This reflects the fact that under personalized pricing firms fully pass cost increases through to consumers, whereas under uniform pricing firms share some of the burden.) Since the consumers in the expansion region have low valuations, the positive effect is relatively small, and so it becomes possible for personalized pricing to hurt consumers.

When c exceeds \underline{v} , as depicted in Figure 1b, a new effect emerges. Now some consumers value only one product above cost. These consumers lie in the “monopoly” regions in the figure. Each firm acts as a true monopolist over these “captured” consumers and extracts all their surplus under personalized pricing. For the consumers in the “competition” region who value both products above the cost, the situation is the same as in Figure 1a. As c increases, both the monopoly and competition regions shrink, but the monopoly region becomes proportionally more important. When c is close to \bar{v} , the monopoly region dominates, so the impact of personalized pricing becomes qualitatively the same as in the monopoly case, as proved in Proposition 2.²⁴ (Notice also that since both regions shrink as c increases, the size of the impact goes to zero in the limit.)

Following this discussion, we have an even stronger result when c is sufficiently large. Using the results in the proof of Proposition 2, we can show the following:

Corollary 1. *Suppose $\bar{v} < \infty$ and $f(\bar{v}) > 0$. For any given $n \geq 2$, there exists $\hat{c} < \bar{v}$ such that for $c > \hat{c}$ each firm is better off and consumers are worse off under competitive personalized pricing than under monopoly uniform pricing.*

Proposition 2 and the above discussion suggest a possible cutoff result, whereby personalized pricing benefits firms if and only if c exceeds a threshold c' , and harms consumers if and only if c exceeds another threshold c'' . (Where $c'' > c'$ because personalized pricing raises total welfare.) Although it appears hard to formally prove such a cutoff result, numerical simulations suggest it is true. In particular, Figure 2 plots the impacts of personalized pricing on industry profit ($\Pi_D - \Pi_U$) and on consumer surplus ($V_D - V_U$) for four common distributions (all in the IID case) and for different values of c .²⁵

Figure 2a considers the exponential case with $F(v) = 1 - e^{-(v-1)}$ on $[1, \infty)$. At $c = 0$ the market is (just) covered under uniform pricing and so the impact is zero, but for higher values of c the market is only partly covered, so as explained before personalized

²⁴Alternatively, when c is close to \bar{v} , conditional on a consumer valuing one firm’s product more than c , it is very unlikely that the consumer values any other product more than c . Therefore when c is high, each firm is essentially a monopolist competing only against the outside option.

²⁵As indicated by some of these examples, the assumption $\bar{v} < \infty$ in Proposition 2 does not appear to be crucial, but we have not been able to extend that result to the case with an unbounded support.

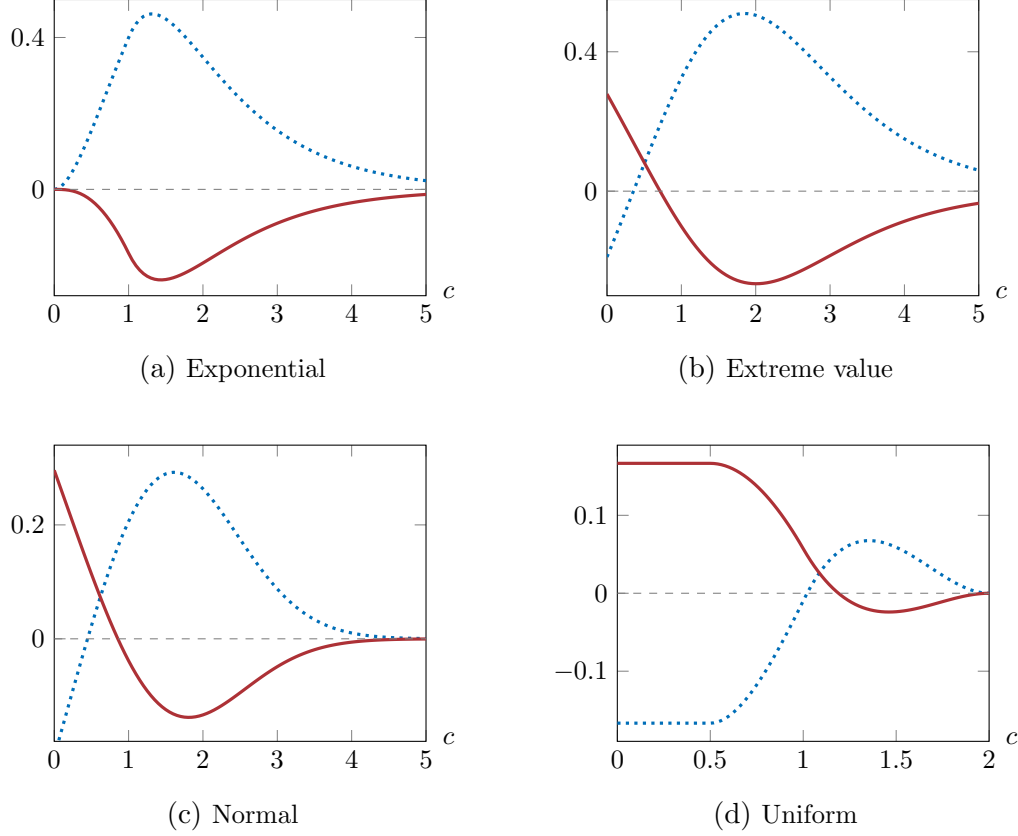


Figure 2: The impact of personalized pricing when $n = 2$, for different values of c
(The dotted and the solid lines represent, respectively, industry profit and consumer surplus.)

pricing benefits firms and harms consumers. (This example also demonstrates that to reverse the impact of personalized pricing, we do not necessarily need the cost to be high enough so that the “monopoly” effect in Figure 1b arises.) Figures 2b and Figure 2c consider, respectively, the Extreme value distribution with $F(v) = e^{-e^{-(v-2)}}$ (which leads to the logit model), and the normal distribution with mean 2 and variance 1 (which leads to the probit model). In both cases, for low values of c (when coverage is high) personalized pricing benefits consumers and harms firms as in the full-coverage case, for high values of c (when coverage is low) personalized pricing has the opposite impact, while for intermediate c both consumers and firms benefit from personalized pricing. Finally, the same pattern is also observed in Figure 2d, which considers the case where valuations are uniformly distributed on $[1, 2]$. In this example, when $c \leq 1/2$ the market is covered, so we know from earlier that personalized pricing harms firms and benefits consumers. When $c > 1/2$ the market is only partially covered, and personalized pricing benefits firms whenever c is above about 1.02, and harms consumers whenever c is above about 1.19.

The number of firms and market coverage. Another parameter which influences market coverage is the number of firms. When n is small, the impact of personalized pricing should be similar to under monopoly; when n is large, the best match should be relatively high, and so the impact of personalized pricing should be similar to the full-coverage case.

We first report an analytical result for the case where n is large, using an approximation argument similar to the one for the case of large c . However, the approximation of the uniform equilibrium price when n is large is technically more difficult. We rely on the approximation results developed in Gabaix et al. (2016) which use extreme value theory, but that technique works only in the IID case. Let

$$\gamma = \lim_{v \rightarrow \bar{v}} \frac{d}{dv} \left(\frac{1 - F(v)}{f(v)} \right) \quad (21)$$

denote the tail index of the valuation distribution of each product. When f is log-concave, we must have $\gamma \in [-1, 0]$.²⁶

Proposition 3. *Consider the IID case without full market coverage for any n . Suppose f is strictly positive and finite everywhere on its support and has a tail index $\gamma \in (0, 1)$. Then there exists \hat{n} such that when $n > \hat{n}$ personalized pricing harms firms and benefits consumers.*

Given the opposite is true in the monopoly case, this again suggests the possibility of a cutoff result in terms of n . Since an analytic result seems hard to obtain, we instead report some numerical examples in Figure 3 below (the IID case with $c = 0$). Figure 3a is for the exponential distribution, and confirms our earlier analytic result that personalized pricing always benefits firms and harms consumers. Figure 3b shows that for the Extreme value distribution, personalized pricing benefits firms if and only if $n < 10$, and harms consumers if and only if $n < 7$. Hence the predictions from Thisse and Vives (1988) fail for a relatively large range of n . A qualitatively similar pattern emerges in Figure 3c for the Normal distribution. Figure 3d considers the uniform distribution with support $[0, 1]$, where the impact of personalized pricing is reversed once we move beyond monopoly. (However, simulations for higher values of c show that the impact of personalized pricing can be similar to in monopoly for some values of $n > 1$.)

²⁶Given f is log-concave, $1 - F$ is log-concave, so $\gamma \leq 0$. To see $\gamma \geq -1$, notice that

$$\frac{d}{dv} \left(\frac{1 - F(v)}{f(v)} \right) = -1 - \frac{1 - F(v)}{f(v)} \frac{f'(v)}{f(v)}.$$

If $\lim_{v \rightarrow \bar{v}} f'(v) \leq 0$, the claim is obvious. If $\lim_{v \rightarrow \bar{v}} f'(v) > 0$, then we must have $\bar{v} < \infty$ and $f(\bar{v}) > 0$, in which case $\frac{1 - F(\bar{v})}{f(\bar{v})} = 0$ and given the log-concavity of f we also have $\frac{f'(\bar{v})}{f(\bar{v})} < \infty$. Then $\gamma = -1$. (Without the log-concavity of f , $\frac{f'(\bar{v})}{f(\bar{v})}$ can be ∞ , in which case γ can be less than -1 .)

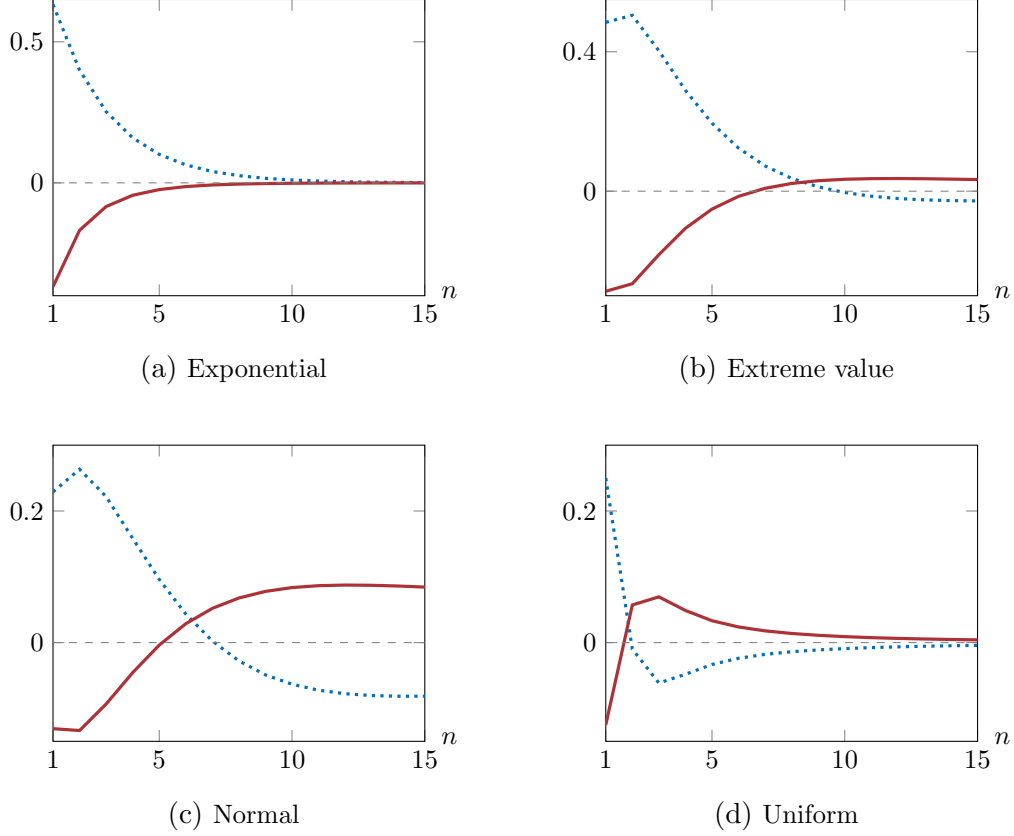


Figure 3: The impact of personalized pricing when $c = 0$, for different values of n
(The dotted and the solid lines represent, respectively, industry profit and consumer surplus.)

4.2 The long-run impact in a free-entry market

In the long-run, the ability to do personalized pricing may also influence firms' incentives to enter the market. To investigate this we now consider a free-entry game, where firms first decide whether or not to pay a fixed entry cost, and then after entering they compete in prices. The free-entry equilibrium is determined by the usual zero-profit condition. (As in the literature, we implicitly assume a sequential entry game to avoid coordination problems, and we ignore integer constraints on n .)

Let us first study the case with personalized pricing. Due to Bertrand competition, the profit on each consumer is simply the difference between her best and second-best product valuations, adjusted for the marginal cost. Hence, with n firms in the market, each firm's profit can be expressed as

$$\frac{1}{n}\Pi_D = \frac{1}{n}\mathbb{E}[\max\{c, v_{n:n}\} - \max\{c, v_{n-1:n}\}] . \quad (22)$$

On the other hand, the increase in match efficiency when the number of firms goes from $n - 1$ to n is

$$\mathbb{E}[\max\{c, v_{n:n}\}] - \mathbb{E}[\max\{c, \hat{v}_{n-1:n-1}\}] , \quad (23)$$

where $\hat{v}_{n-1:n-1}$ denotes the best match among the original $n - 1$ products. (We use \hat{v}_i to denote the valuation for product $i \leq n - 1$ when there are only $n - 1$ firms in the market.) To determine whether there is too much or too little entry relative to the social optimum, it then suffices to compare (22) and (23).

Assumption 2. *Entry of a new firm does not affect consumers' valuations for existing products. That is, the distribution of $(\hat{v}_1, \dots, \hat{v}_{n-1})$ when there are $n - 1$ firms in the market is the marginal distribution of (v_1, \dots, v_{n-1}) when there are n firms in the market.*

Under the above assumption (which is clearly true, e.g., in the IID case),²⁷ it turns out that (22) and (23) are actually equal to each other. Hence one can show that:

Lemma 3. *Under Assumption 2, the free-entry equilibrium under personalized pricing is unique and it is also socially optimal.*

The intuition for this result is straightforward. Suppose $n - 1$ firms are already in the market, and consider the entry of an n^{th} firm. Amongst consumers for whom $v_n \leq \max_{j \leq n-1} \{c, v_j\}$, this additional firm creates no social surplus and earns zero profit. However, amongst consumers with $v_n > \max_{j \leq n-1} \{c, v_j\}$, this new firm raises total surplus by $v_n - \max_{j \leq n-1} \{c, v_j\}$, and fully extracts it via Bertrand competition. As a result, the incentives of the social planner and this new firm are perfectly aligned.²⁸

Now consider the case with uniform pricing. Let n^* denote the socially optimal number of firms. A simple corollary of Lemma 3 is the following:

Corollary 2. *Suppose Assumption 2 holds and each firm's profit under uniform pricing decreases in n .*

²⁷Assumption 2 fails if the entry of a new product induces existing firms to reposition their products, or if consumers have consideration-set dependent preferences. For example, Assumption 2 fails in the Salop circle model because entry of a new firm causes existing ones to relocate, and this changes consumer valuations for their products. (Contrary to Lemma 3 below, Section 3.5 of Stole, 2007 shows that entry is socially excessive in the Salop circle model under perfect price discrimination.) However Assumption 2 can hold in other spatial models, such as in Chen and Riordan (2007) where entry of a new firm does not lead to repositioning by existing firms.

²⁸Levin and Smith (1994) show a related result in the context of IPV auctions (see their Proposition 6). The difference is that in their setup bidders simultaneously decide whether to participate in an auction and end up playing a symmetric mixed-strategy equilibrium. The corresponding social planner's problem is to choose the (symmetric) probability that each bidder participates. Although the outcome is not the first best due to the randomness of the number of bidders, the free-entry outcome coincides with the social planner's solution and the underlying intuition is similar to that in our model. In particular, they also show the equivalence of (22) and (23) in their IID case when c is irrelevant. See also Chapter 6 in Milgrom (2004) for some related discussions.

- (i) Entry under uniform pricing is excessive if $\Pi_U > \Pi_D$ at $n = n^*$, but insufficient if $\Pi_U < \Pi_D$ at $n = n^*$.
- (ii) Uniform pricing therefore leads to excessive entry if the market is fully covered and Assumption 1 holds.

To understand part (i) of the corollary, note that we have just shown that the increase in match efficiency due to an extra firm entering the market is exactly equal to that firm's profit under personalized pricing. Hence, with uniform pricing, entry is either excessive or insufficient depending upon whether Π_U is respectively above or below Π_D .²⁹ To understand part (ii), recall that Proposition 1 showed that with full coverage and Assumption 1 it is always the case that $\Pi_U > \Pi_D$. Anderson, de Palma, and Nesterov (1995) and Tan and Zhou (2021) also prove this excessive entry result in the IID case; our proof is much simpler than theirs, and our result is more general since it also potentially allows for correlated valuations.

Finally, another simple but important consequence of Lemma 3 is the following:

Proposition 4. *Compared to uniform pricing, personalized pricing (weakly) benefits consumers in the long run.*

In the long-run firms earn zero profit (after accounting for the fixed entry cost) irrespective of the pricing regime. Therefore since total welfare is maximized under personalized pricing, so is aggregate consumer surplus.

5 Asymmetrically Informed Firms

So far we have assumed that either all firms do uniform pricing or all firms do personalized pricing. Implicitly, therefore, we have assumed that all firms have access to the same data and technology. However, in certain markets, some firms may have more data and better technology than others. For example, Amazon possesses lots of information about customer shopping behavior, and in principle it can use such information to offer personalized prices for its own products—whereas third-party sellers of similar products on Amazon are often smaller retailers who lack such information. Similarly, in music

²⁹The assumption that each firm's profit under uniform pricing decreases in n ensures the uniqueness of the free-entry equilibrium. (It must hold if the uniform price decreases in n , which, as shown in Zhou, 2017, holds at least in the IID case with a log-concave f .) To see this, let $\pi_n(p_i, p)$ be firm i 's profit when it offers a uniform price p_i and the other $n - 1$ firms offer a price p . Let p_n be the equilibrium uniform price. Then $\pi_n(p_n, p_n) < \pi_{n-1}(p_n, p_{n-1}) \leq \pi_{n-1}(p_{n-1}, p_{n-1})$, where the first inequality is because a firm's profit must increase when the number of competitors drops and they further set higher prices $p_{n-1} > p_n$, and the second inequality is simply from the no-deviation equilibrium condition when there are $n - 1$ firms in the market.

streaming, large players like Amazon and Apple may have more data, and also the ability to make more refined personalized offers, compared to smaller players like Spotify and Pandora. The goal of this section is to investigate market performance in such a “mixed” case, where some firms do uniform pricing and others do personalized pricing. By comparing with the two symmetric regimes studied earlier, we can then evaluate the impact of policies that prevent large retailers such as Amazon from price discriminating or which force them to share their data.

In order to capture this “mixed” case, suppose that firms 1 to k have consumer data and can price discriminate, while firms $k + 1$ to n have no consumer data and therefore have to offer a uniform price. When $k = 0$ all firms do uniform pricing as in Section 3.1, whereas when $k = n$ all firms do personalized pricing as in Section 3.2. When $0 < k < n$ a subtle technical issue arises: if all firms set prices simultaneously, there is no pure-strategy pricing equilibrium,³⁰ and the mixed-strategy equilibrium is rather complicated to characterize. To avoid this problem, we assume that the $n - k$ firms with no consumer data simultaneously set their uniform prices first, and after seeing their prices the other k firms use their data to simultaneously offer personalized prices. This timing seeks to capture the idea that firms with lots of data, such as Amazon, often also have more advanced pricing technologies and so can adjust their prices more frequently.³¹

5.1 Equilibrium analysis

Consider $0 < k < n$ and let p denote the equilibrium price of each firm that has no data. Given that uniform price, let us first derive the equilibrium personalized prices charged by firms 1 to k . Let

$$\hat{x}_{p,c} \equiv v_1 - c - \max\{0, v_2 - c, \dots, v_k - c, v_{k+1} - p, \dots, v_n - p\}$$

be the advantage of firm 1 relative to a consumer’s other alternatives (including the outside option) when each firm prices at its lowest equilibrium price, i.e., when firms 1 to k price at c , and firms $k + 1$ to n price at p . Let $\hat{H}_{p,c}$ be the CDF of $\hat{x}_{p,c}$. Firms 1 to k compete exactly as in Section 3.2, except that now consumers have an outside option

³⁰In any hypothetical pure-strategy equilibrium, each uniform-pricing firm has a positive measure of consumers who are (almost) indifferent between it and some personalized-pricing firm but buy from the latter because of its precisely targeted prices. If the uniform-pricing firm slightly reduces its price, it will win these consumers and so its demand will increase discontinuously. This discontinuity in demand leads to non-existence of pure-strategy pricing equilibrium.

³¹Such a timing is also often assumed in the literature. See, e.g., Thisse and Vives (1988) and Jullien, Reisinger, and Rey (2022) for a similar treatment in their duopoly models when one firm price discriminates but the other does not.

$\max\{0, v_{k+1} - p, \dots, v_n - p\}$. Therefore in equilibrium firm 1, say, offers personalized prices

$$p_1(v_1, \dots, v_n) = \begin{cases} c + \hat{x}_{p,c} & \text{if } \hat{x}_{p,c} > 0 \\ c & \text{otherwise} . \end{cases}$$

Since firm 1 wins a consumer if and only if $x_{p,c} > 0$, its equilibrium profit is

$$\hat{\pi}_D = \int_0^\infty x d\hat{H}_{p,c}(x) = \int_0^\infty [1 - \hat{H}_{p,c}(x)] dx . \quad (24)$$

Note that provided $p > c$, $\hat{x}_{p,c}$ exceeds x_c in the sense of first order stochastic dominance. Therefore comparing with equation (15) from earlier, each personalized-pricing firm here in the mixed regime earns more than when all firms price discriminate (i.e., $\hat{\pi}_D > \frac{1}{n}\Pi_D$).

We now derive the equilibrium uniform price p charged by the $n - k$ firms that have no data. Suppose firm n unilaterally deviates to price p_n . To calculate firm n 's demand we should set the prices of the firms that do personalized pricing to c . (To understand this, note that if, say, $v_n - p_n < v_1 - c$, firm n cannot win the consumer because firm 1 can offer her more surplus with a personalized price close to marginal cost.) Therefore, a consumer buys from firm n if and only if

$$v_n - p_n \geq \max\{0, v_1 - c, \dots, v_k - c, v_{k+1} - p, \dots, v_{n-1} - p\} , \quad (25)$$

i.e., firm n offers more surplus than the outside option, each other firm that sets a uniform price, and all firms that set personalized prices when they charge c . This suggests that under regularity conditions, (i) the equilibrium uniform price p in this mixed regime is lower than that in the regime of uniform pricing, and so (ii) for any given deviation price p_n , firm n 's deviation demand in this mixed regime is smaller than in the regime of uniform pricing, which further implies that each uniform-pricing firm earns less in this mixed regime than in the regime of uniform pricing.

Let $\tilde{x}_{p,c}$ be the difference between the left and righthand sides of (25) at $p_n = p$, i.e.,

$$\tilde{x}_{p,c} \equiv v_n - p - \max\{0, v_1 - c, \dots, v_k - c, v_{k+1} - p, \dots, v_{n-1} - p\} ,$$

and let $\tilde{H}_{p,c}$ be its CDF. Firm n 's deviation demand can then be expressed as $1 - \tilde{H}_{p,c}(p_n - p)$. Suppose this demand behaves well (e.g., it is log-concave in the IID case when each v_i has a log-concave density). Then the equilibrium uniform price solves

$$p - c = \frac{1 - \tilde{H}_{p,c}(0)}{\tilde{h}_{p,c}(0)} . \quad (26)$$

(If $c = p$, this equation degenerates to the first-order condition in the regime of uniform pricing. Therefore, whenever the right-hand side increases in c , p is lower than the

equilibrium price in the regime of uniform pricing.) Firm n 's profit is

$$\tilde{\pi}_U = (p - c)[1 - \tilde{H}_{p,c}(0)] = \frac{[1 - \tilde{H}_{p,c}(0)]^2}{\tilde{h}_{p,c}(0)}. \quad (27)$$

As argued before, whenever p is lower than the equilibrium price in the regime of uniform pricing, we have $\tilde{\pi}_U < \frac{1}{n}\Pi_U$.

Industry profit in the mixed regime is $\Pi_M = k\hat{\pi}_D + (n - k)\tilde{\pi}_U$. The expected consumer surplus is

$$V_M = \mathbb{E}[\max\{0, v_{k-1:k} - c, v_{n-k:n-k} - p\}] , \quad (28)$$

where $v_{k-1:k}$ denotes the second best among $\{v_1, \dots, v_k\}$ and $v_{n-k:n-k}$ denotes the best among $\{v_{k+1}, \dots, v_n\}$. (If $k = 1$ the $v_{k-1:k} - c$ term vanishes.) To see this, let $v_0 = \max\{0, v_{n-k:n-k} - p\}$ be a consumer's outside option when the first k firms compete for her by offering personalized prices. If $\max\{v_1 - c, \dots, v_k - c\} \geq v_0$, the consumer buys from one of the first k firms, in which case her surplus is $\max\{v_0, v_{k-1:k} - c\}$. (If the second best among the k firms is worse than v_0 , the consumer's surplus is v_0 ; otherwise it is $v_{k-1:k} - c$.) Otherwise, she takes the outside option v_0 . Therefore, the expected consumer surplus is $\mathbb{E}[\max\{v_0, v_{k-1:k} - c\}]$, which leads to the expression for V_M in equation (28).

5.2 Comparison

We now compare the mixed case with the uniform and personalized pricing regimes that we studied earlier. For brevity we focus on consumer surplus—and show that under certain conditions the mixed case is the worst for consumers.

When the production cost c is sufficiently high we can compare the three regimes analytically.³² In particular, similar to Proposition 2, we can show that if $\bar{v} < \infty$ and $f(\bar{v}) > 0$, there exists a \hat{c} such that when $c > \hat{c}$ we have $V_D < V_M < V_U$, i.e., consumers rank the mixed case between the other two regimes.³³ Intuitively, recall that when c is sufficiently large, each firm approximately acts like a monopolist. As a result, as more firms are able to personalize prices, consumers must become worse off.

When valuations are IID exponential and one firm can personalize prices (i.e., $k = 1$), we can also provide an analytical comparison for any level of c . In particular, we can

³²In general, the comparison between V_M and V_U and V_D is not trivial. To see this, recall that $V_U = \mathbb{E}[\max\{0, v_{n:n} - p_U\}]$ where p_U is the equilibrium price in the uniform pricing regime, while $V_D = \mathbb{E}[\max\{0, v_{n-1:n} - c\}]$. The comparison with V_M is non-trivial because $v_{n:n} > \max\{v_{k-1:k}, v_{n-k:n-k}\} > v_{n-1:n}$, while under regularity conditions we should have $p_U > p > c$.

³³More precisely, for c close to \bar{v} , following a similar logic as in the proof of Proposition 2, we can show $V_M \approx (n - k)f(\bar{v})\frac{\varepsilon^2}{8}$. (We can also prove that $\Pi_M \approx kf(\bar{v})\frac{\varepsilon^2}{2} + (n - k)f(\bar{v})\frac{\varepsilon^2}{4}$, and hence $\Pi_U < \Pi_M < \Pi_D$.) The details are available upon request.

prove that there exists a \tilde{c} such that when $c < \tilde{c}$ we have $V_M < V_D \leq V_U$ i.e., consumers are worst off in the mixed case. Intuitively, this happens because the firm that is able to personalize can “poach” some consumers for whom it is not their favorite product via a low personalized price—but this harms match efficiency, and can reduce total welfare by so much that consumers end up worse off in aggregate compared to the uniform and discriminatory pricing regimes. On the other hand, we can also prove that when $c > \tilde{c}$ we have $V_D < V_M < V_U$, i.e., the same ranking as in the previous paragraph for large c . (The proof of these consumer surplus rankings is available in the Appendix.)

Numerical simulations with other distributions confirm the patterns outlined above. Figure 4 depicts consumer surplus in the three regimes as a function of c , for respectively the Extreme value and Normal distributions (for $n = 2$, and in the mixed case also $k = 1$). At low values of c coverage is high, and the mismatch effect highlighted above causes the mixed regime to be worst for consumers. On the other hand, at higher values of c firms are more like local monopolists, and hence the discriminatory regime when all firms do personalized pricing is worst for consumers.

Overall, our results suggest that when coverage is relatively high, policies which force a large firm to share its data or prevent it from personalizing prices can benefit consumers.

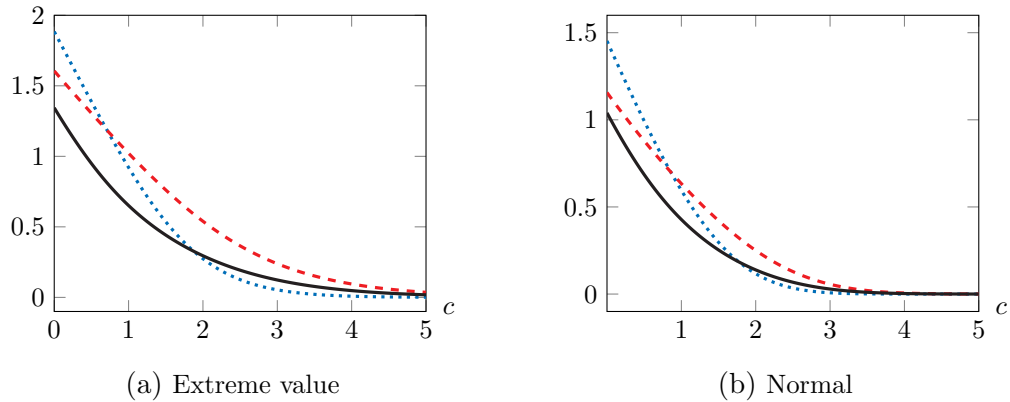


Figure 4: Consumer surplus when $n = 2$, for different values of c
(The solid, dashed, and dotted lines are respectively the mixed, uniform, and discriminatory cases.)

6 Discussion and Extensions

6.1 Valuation correlation

Although our model allows for valuations to be correlated across products, so far we have not discussed how the impact of personalized pricing might vary with the degree of this correlation. Although it appears hard to obtain general results, it is clear that in

the limit case with perfectly (positively) correlated valuations the impact of personalized pricing disappears—because firms will price at marginal cost under both the uniform and discriminatory regimes. The following example suggests that, as one might expect, a greater degree of correlation weakens the impact of personalized pricing.

Example: bivariate Normal distribution. Suppose $n = 2$ and valuations are drawn from a bivariate Normal distribution with mean μ , variance σ , and correlation coefficient $\rho \in (-1, 1)$. Suppose the market is fully covered.³⁴ Then one can show that:³⁵

$$\Pi_D = 2\sigma\sqrt{\frac{1-\rho}{\pi}} \quad \text{and} \quad \Pi_U = \sigma\sqrt{\pi(1-\rho)}, \quad \text{and hence} \quad \Pi_D - \Pi_U \propto -\sqrt{1-\rho}.$$

In this example, as product valuations become more correlated (i.e., as ρ increases), profit falls under both uniform and discriminatory pricing, and the impact of personalized pricing on profit also becomes smaller. (Given the assumption of a covered market, the impact on consumer surplus also gets smaller.) Intuitively, higher correlation means that products are less differentiated, and so firms compete more fiercely. In the limit case of perfect positive correlation (i.e., as $\rho \rightarrow 1$), personalized pricing has no impact on profit or consumer surplus because firms' products become homogeneous.

6.2 An alternative information structure

So far we have assumed that under personalized pricing, firms observe a consumer's valuation for *each* product. Here we consider a natural alternative, where firms observe only a consumer's valuation for *their own* product, i.e., each firm i observes only v_i . For convenience, we refer to these two cases as “full” and “partial” discrimination respectively.³⁶

Under partial discrimination a firm offers a price $p(v)$ to a consumer who has valuation v for its product. It turns out that $p(v)$ can be derived in the same way as a bid in a standard first-price auction. Recall that $G(\cdot|v_i)$ and $g(\cdot|v_i)$ are respectively the CDF and density function of $\max_{j \neq i} \{v_j\}$ conditional on v_i , and note that in a standard first-price auction with interdependent values (see, e.g., Milgrom and Weber, 1982) a bidder with

³⁴Strictly speaking, full market coverage is impossible because the lower support of the valuation distribution is unbounded. However the market is almost covered if μ is large enough for a given c .

³⁵Since the joint density of the bivariate Normal is log-concave, the first part of Assumption 1 is satisfied (the second part is not needed since the market is covered). Also notice that $v_1 - v_2$ is normally distributed with mean 0, variance $2\sigma^2(1-\rho) \equiv \tau^2$, and density function $h(x) = \frac{1}{\sqrt{2\pi}\tau} e^{-\frac{x^2}{2\tau^2}}$. Hence

$$\Pi_U = \frac{1}{2h(0)}, \quad \text{and} \quad \Pi_D = 2 \int_0^\infty xh(x)dx = -2\tau^2 \int_0^\infty h'(x)dx = 2\tau^2 h(0),$$

where we used $xh(x) = -\tau^2 h'(x)$. Substituting in for $h(0)$ and τ gives the stated profit expressions.

³⁶Note that one could loosely interpret full discrimination as when firms have access to comprehensive third-party data, and partial discrimination as when firms only have access to first-party data.

value z bids $b(z)$, where

$$b(z) \equiv \int_c^z [1 - L(x|z)]dx \quad \text{with} \quad L(x|z) = \exp\left(-\int_x^z \frac{g(t|t)}{G(t|t)}dt\right). \quad (29)$$

We now impose the following regularity condition:

Assumption 3. *For all $z \in [\underline{v}, \bar{v}]$, $[v - b(z) - c] \frac{g(z|v)}{G(z|v)}$ increases in v whenever it is positive.*

Assumption 3 definitely holds if $\frac{g(z|v)}{G(z|v)}$ increases in v , which is satisfied when product valuations are either IID or (positively) affiliated in the sense of Milgrom and Weber (1982). However the assumption can also hold even when $\frac{g(z|v)}{G(z|v)}$ decreases in v , provided it does not decrease too quickly.

Lemma 4. *Suppose Assumption 3 holds. Under partial discrimination there exists a symmetric equilibrium in which each firm uses the price schedule $p(v) = v - b(v)$, where $b(v)$ is defined in equation (29) and is strictly increasing in v .*

Competition under partial discrimination is similar to that in a first-price auction. Formally, we can interpret a firm as “bidding” surplus of $b(v) = v - p(v)$ to a consumer who has valuation v for its product, and then the consumer picking the best (non-negative) “bid.” The lemma implies that a consumer buys provided at least one product valuation exceeds cost, because the surplus $v - p(v)$ offered by a firm is positive if and only if $v > c$. The lemma also implies that, conditional on buying, a consumer buys the highest-valuation product because each firm’s offered surplus $v - p(v)$ strictly increases in v .³⁷

Proposition 5. *Suppose Assumption 3 holds. Comparing partial and full discrimination:*

- (i) *Total welfare is the same under both regimes.*
- (ii) *When valuations are IID, profit and consumer surplus are the same in both regimes.*
- (iii) *Otherwise firms earn more (less) under partial discrimination if $\frac{g(z|v)}{G(z|v)}$ increases (decreases) in v , and the opposite is true for consumer surplus.*

The intuition for these results is as follows. For part (i), total welfare is the same under both partial and full discrimination because in both regimes consumers buy the best product conditional on its valuation exceeding marginal cost. Parts (ii) and (iii) then exploit the connection with auction theory. In particular, notice that under *full* discrimination competition is the same as in a second-price auction—because the winning firm earns a profit equal to the difference between the highest and second-highest valuations

³⁷Interestingly, note that Assumption 3 does *not* guarantee that a firm charges a higher price to a consumer with a higher valuation for its product. Nevertheless, a sufficient condition for $p'(v) > 0$ is that $\frac{g(v|v)}{G(v|v)} \leq \frac{g(x|x)}{G(x|x)}$ for $x \leq v$; this holds, for example, in the IID case if each v_i has a log-concave density.

(including the outside option). The well-known revenue equivalence theorem then implies that, when valuations are IID, firm and consumer payoffs are the same under both full and partial discrimination. Meanwhile the theory of auctions with interdependent valuations (e.g., Milgrom and Weber, 1982) implies that bidders (i.e., firms) are better off and the auctioneer (i.e., consumers) are worse off with partial information if valuations are positively affiliated, while the reverse is true if they are negatively affiliated.³⁸

Proposition 5 implies that with IID valuations, all our earlier results about full discrimination (for both the short- and long-run) carry over to partial discrimination. Outside the IID case, this equivalence remains approximately true if the production cost c is sufficiently high and if $\bar{v} < \infty$ and $f(\bar{v}) > 0$. This is simply because when each firm acts approximately like a monopolist, it does not matter whether or not they know consumers' valuations for other firms.³⁹ More generally, though, Proposition 5 implies that correlation in valuations matters for whether consumers prefer firms to have more or less information about their tastes.⁴⁰

Although beyond the scope of the current paper, it would of course be interesting to also consider other information structures and investigate how the welfare impact of price discrimination changes in the amount of information that firms have about consumer preferences. This is, however, a challenging question in general because the space of information structures is large and the pricing equilibrium under some information structures is complicated to characterize. One way to proceed is to consider a special class of information structures. For example, in the linear Hotelling model, some papers (e.g., Liu and Serfes, 2004; Bounie, Dubus, and Waelbroeck, 2021) have studied the class of interval information structures and shown that providing firms with finer information has a non-monotonic impact on profit and consumer surplus. Another way to proceed is to consider an information design approach and explore the welfare limits when arbitrary information structures are feasible. For example, Bergemann, Brooks, and Morris (2015) and Elliott, Galeotti, Koh, and Li (2021) have respectively studied this issue in the monopoly case and the competition case.⁴¹ Our approach, by focusing on simpler in-

³⁸We note that while the auctions literature focuses on the case of positive affiliation, negative affiliation may be reasonable in our context, e.g., if products differ in characteristics space, and consumers have different preferences over different characteristics.

³⁹More formally, industry profit under partial discrimination is

$$\Pi_{PD} = n \int_c^{\bar{v}} [p(v) - c] G(v|v) dF(v) = n \int_c^{\bar{v}} [v - b(v) - c] G(v|v) dF(v) = n \int_c^{\bar{v}} \int_c^v L(x|v) dx G(v|v) dF(v).$$

Following Proposition 2, for high c we find that $\Pi_{PD} \approx n f(\bar{v}) \frac{\varepsilon^2}{2}$ which equals the approximation of Π_D .

⁴⁰One may wonder, in the case of full market coverage, whether partial discrimination also reduces profit and improves consumer surplus relative to uniform pricing. According to Proposition 5, this must be true if $\frac{g(z|v)}{G(z|v)}$ decreases in v , but we have not been able to show it is true more generally.

⁴¹The information structures available to firms, however, might be constrained by communication

formation structures, is arguably more suitable for evaluating policies which simply either allow for or ban the use of consumer data.

6.3 Constrained personalized pricing

Up until now, in the regime of personalized pricing we have assumed that firms have complete flexibility over the prices they charge. However in practice a firm may be constrained—if its personalized prices differ by too much, it may face a consumer backlash or scrutiny from policymakers. Therefore here we consider the case where the maximum price difference across consumers is restricted. We show that the main insights from our earlier analysis carry over to such a “constrained” setting—in particular, the impact of personalized pricing on profit and consumer surplus depends on market coverage.

Consider the following two-stage model: firms first simultaneously choose a *public list price* which caps their personalized prices, and then offer *personalized discounts* of up to Δ off their list price, where $0 < \Delta < \bar{v} - \max\{c, \underline{v}\}$.⁴² Note that $\Delta = 0$ gives uniform pricing, while equation (14) implies that for $\Delta \geq \bar{v} - \max\{c, \underline{v}\}$ firms charge the “unconstrained” personalized prices studied earlier.⁴³

We look for a symmetric equilibrium where each firm sets the same list price p^* . Suppose that firm i unilaterally deviates and chooses a general list price $p_i \geq c + \Delta$ in the first stage.⁴⁴ Then in the equilibrium of the second stage, firm i offers personalized prices

$$p(v_i, \mathbf{v}_{-i}; p_i) = \begin{cases} \min\{p_i, p^* - \Delta + x_{p^*-\Delta}\} & \text{if } x_{p^*-\Delta} \geq p_i - p^* \\ p_i - \Delta & \text{otherwise} \end{cases}$$

where $x_{p^*-\Delta} = v_i - \max_{j \neq i} \{p^* - \Delta, v_j\}$ is the advantage of firm i relative to the best alternative (including the outside option) when all firms offer the minimum possible equilibrium personalized price $p^* - \Delta$. These prices take a similar form as the “unconstrained” personalized prices from equation (14). When firm i wins a consumer (in which case other firms must be offering the lowest possible price $p^* - \Delta$), it can charge a personalized

incentives between firms and consumers. See, e.g., Ali, Lewis, and Vasserman (2020), and Ichihashi and Smolin (2022) for some recent research in this direction.

⁴²If instead firms simultaneously choose personalized prices (which differ across consumers by no more than Δ), one can show there is no pure-strategy pricing equilibrium. We circumvent this technical issue by considering a two-stage game which arguably captures better the practice of personalized discounts.

⁴³Varying Δ between 0 and $\bar{v} - \max\{c, \underline{v}\}$ therefore traces out a path from uniform to “unconstrained” personalized pricing. This is in the same spirit as some works in the literature on third-degree price discrimination (e.g., Schmalensee, 1981; Aguirre, Cowan, and Vickers, 2010) that use comparative statics in the price difference across two market segments to perform welfare analysis.

⁴⁴Note that a list price satisfying $p_i < c + \Delta$ is weakly dominated since a firm never prices below cost.

price which is $x_{p^*-\Delta}$ higher than its competitors', but this price is constrained by its own list price p_i and the lower bound $p_i - \Delta$.

Using the above, and recalling that H_z denotes the CDF of x_z , firm i 's profit is

$$(p_i - c)[1 - H_{p^*-\Delta}(p_i - p^* + \Delta)] + \int_{p_i - p^*}^{p_i - p^* + \Delta} (p^* - \Delta + x - c) dH_{p^*-\Delta}(x) . \quad (30)$$

The first term is profit from consumers for whom the list price binds because they have a very strong preference for firm i , while the second term is profit from consumers who have a less strong preference for firm i and so receive personalized discounts. Taking the derivative of profit with respect to p_i and imposing $p_i = p^*$, a symmetric equilibrium price p^* must satisfy the following equation

$$p^* - \Delta - c = \frac{1 - H_{p^*-\Delta}(\Delta)}{h_{p^*-\Delta}(0)} . \quad (31)$$

Intuitively, when a firm slightly increases its list price it obtains more revenue from consumers for whom the list price binds (and they have mass $1 - H_{p^*-\Delta}(\Delta)$), but since its minimum price also increases the firm stops selling to consumers who were just indifferent about buying from it (hence the $h_{p^*-\Delta}(0)$ term).

To ensure that the solution to equation (31) is unique and constitutes an equilibrium, we make the following assumption:

Assumption 4. $1 - H_z(x)$ is log-concave in x and logsubmodular in (x, z) .

Note that the logsubmodularity condition requires that $\frac{1-H_z(x)}{h_z(x)}$ be non-increasing in z for all x , and so is stronger than the second condition in Assumption 1 from earlier. It is straightforward to show that Assumption 4 is satisfied if, for example, valuations are IID with a log-concave distribution.

We then obtain the following result about the equilibrium list price (recall that H and h are respectively the CDF and density of $x = v_i - \max_{j \neq i} \{v_j\}$):

Lemma 5. Suppose Assumption 4 holds.

- (i) The unique (symmetric) equilibrium list price p^* satisfies equation (31).
- (ii) The market is fully covered if and only if $c \leq \underline{v} - \frac{1-H(\Delta)}{h(0)}$.
- (iii) Suppose also that $h(x) < h(0)$ for all $x > 0$. Then as Δ increases the equilibrium list price p^* strictly increases and a firm's minimum price $p^* - \Delta$ strictly decreases.

To interpret the lemma, notice that as $\Delta \rightarrow 0$ the equilibrium list price that solves equation (31) coincides with the uniform price that we solved for earlier, and the condition for coverage is also the same. Under the same mild regularity condition as in Lemma 2, an increase in Δ induces firms to raise their list price (to better price discriminate

consumers with a strong preference for their product) but also lower their minimum price (to sell to more consumers). When $\Delta \rightarrow \bar{v} - \max\{c, \underline{v}\}$ the list price equals the highest “unconstrained” personalized price, while the lowest price charged by firms equals c .

We now turn to how the flexibility of personalized pricing impacts market performance:

Proposition 6. (i) Suppose $c \leq \underline{v} - \phi(\underline{v})$ so that the market is fully covered as $\Delta \rightarrow 0$. A larger Δ (i.e., more flexible personalized pricing) harms firms and benefits consumers. (ii) Suppose $c > \underline{v} - \phi(\underline{v})$ so that the market is not fully covered as $\Delta \rightarrow 0$. Then allowing a small extent of personalized pricing (i.e., allowing a small $\Delta > 0$) benefits firms, and it harms consumers if $nG(v|v)f(v) \geq f_{(n)}(v)$ (which is true, e.g., when valuations are IID).

Consistent with our earlier results, part (i) of the proposition shows that in a fully covered market a higher discount harms firms and benefits consumers. Intuitively, as Δ increases, consumers with a strong preference for one product lose out since they pay more as the list price increases, but consumers with lower valuations gain since more of them are able to purchase as the minimum price decreases, and the latter dominates the former due to logconcavity. Also consistent with our earlier results, part (ii) of the proposition shows that when the market is not covered as $\Delta \rightarrow 0$, a small discount benefits firms but harms consumers. Figure 5 depicts the impact of larger Δ on welfare outcomes when $n = 2$, $c = 1.5$, and valuations are IID uniform on $[1, 2]$. Industry profit and consumer surplus are non-monotone in Δ ; since $\Delta = 0$ corresponds to uniform pricing while $\Delta = 0.5$ corresponds to “unconstrained” personalized pricing, consumers prefer uniform pricing while firms prefer an intermediate degree of personalized pricing.

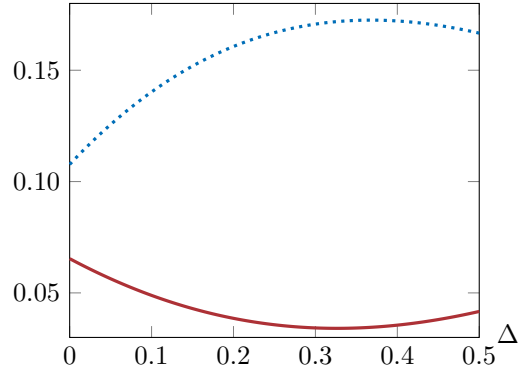


Figure 5: The impact of Δ on profit (dotted line) and consumer surplus (solid line)

In summary, when firms are constrained in how much personalized prices can vary across consumers, our main insights from earlier carry over. Specifically, market coverage matters—in covered markets more flexible personalized pricing harms firms but benefits consumers, whilst in uncovered markets the opposite can be true.

7 Conclusion

This paper has investigated the impact of personalized pricing, a form of price discrimination which is becoming increasingly relevant in the digital economy. The paper delivers three main insights: (i) In a general oligopoly model, competitive personalized pricing intensifies competition, harms firms and benefits consumers under a log-concavity condition if the market is fully covered; however, (ii) the impact can be completely reversed in the (arguably more realistic) case without full market coverage. In this case, personalized pricing tends to amplify the degree of existing competition. Specifically, it intensifies competition and benefits consumers in markets that are already competitive enough (e.g., when production cost is low or the number of competitors is big), whereas in markets where competition is relatively weak (e.g., when product cost is high or the number of competitors is small), personalized pricing further reduces it, harming consumers. Moreover, (iii) in a long-run equilibrium with endogenous market structure, personalized pricing induces the socially optimal level of entry and so favors consumers. We also considered the case where some firms can use consumer data to price discriminate while others cannot. We showed that when market coverage is relatively high, this “mixed” case can be the worst for consumers—and hence policies which prevent data-rich firms from price discriminating, or which force them to share their data, can improve consumer surplus.

One important issue that this paper has not addressed is a consumer’s incentive to allow her data to be collected and then used for personalized pricing (e.g., privacy policies like GDPR in the EU and CCPA in California give consumers some control over what data is harvested and how it is used). This issue is studied in our companion paper Rhodes and Zhou (2022). There we identify a novel externality across consumers: when more consumers share their data, firms may (adversely) revise their offers to consumers who choose to remain anonymous. Due to this externality, we demonstrate that (i) although privacy policies such as GDPR benefit consumers, there are still too many consumers who choose to share their data relative to the consumer optimum, and (ii) more competition can harm consumers in aggregate by inducing more of them to share their data.

Appendix: Omitted Proofs and Details

Primitive conditions for Assumption 1. We report some primitive conditions for Assumption 1. Define a piece of notation

$$G_2(x|y) \equiv \frac{\partial G(x|y)}{\partial y} . \quad (32)$$

Lemma 6. (i) If the joint density \tilde{f} is log-concave, then $1 - H_z(x)$ is log-concave in x .
(ii) $\phi(z) = \frac{1-H_z(0)}{h_z(0)}$ is non-increasing in z if (a) $G_2(v|v) \geq 0$ and $f'(v) \geq 0$, or (b) if \tilde{f} is log-concave and $\frac{G_2(v|v)}{G(v|v)}$ is non-increasing in v . (In particular, condition (b) holds in the IID case with a log-concave f .)

Proof. (i) Note that

$$1 - H_z(x) = \int_{A_x} \tilde{f}(\mathbf{v}) d\mathbf{v} ,$$

where $A_x = \{\mathbf{v} : v_i - \max_{j \neq i} \{z, v_j\} > x\}$. To prove $1 - H_z(x)$ is log-concave in x , according to the Prékopa-Borell Theorem (see, e.g., Caplin and Nalebuff, 1991), it suffices to show that, for any $\lambda \in [0, 1]$, we have

$$\lambda A_{x_0} + (1 - \lambda) A_{x_1} \subset A_{\lambda x_0 + (1 - \lambda) x_1} , \quad (33)$$

where the former is the Minkowski average of A_{x_0} and A_{x_1} . Let $\mathbf{v}^0 \in A_{x_0}$ and $\mathbf{v}^1 \in A_{x_1}$, i.e.,

$$v_i^0 > z + x_0 \quad \text{and} \quad v_i^0 > v_j^0 + x_0 \quad \text{for any } j \neq i ,$$

and

$$v_i^1 > z + x_1 \quad \text{and} \quad v_i^1 > v_j^1 + x_1 \quad \text{for any } j \neq i .$$

These immediately imply that

$$v_i^\lambda > z + \lambda x_0 + (1 - \lambda) x_1 \quad \text{and} \quad v_i^\lambda > v_j^\lambda + \lambda x_0 + (1 - \lambda) x_1 \quad \text{for any } j \neq i ,$$

where $v_i^\lambda = \lambda v_i^0 + (1 - \lambda) v_i^1$. Hence, we have $\mathbf{v}^\lambda \in A_{\lambda x_0 + (1 - \lambda) x_1}$, and so (33) holds.

(ii) Recall that

$$\phi(z) = \frac{\int_z^{\bar{v}} G(v|v) dF(v)}{G(z|z)f(z) + \int_z^{\bar{v}} g(v|v) dF(v)} .$$

For $z \leq \underline{v}$, $\phi(z)$ is a constant and so is non-increasing. In the following, we focus on $z > \underline{v}$.

Using $\frac{dG(v|v)}{dv} = g(v|v) + G_2(v|v)$, one can check that $\phi'(z) \leq 0$ if and only if

$$G(z|z)f(z) + \int_z^{\bar{v}} g(v|v) dF(v) + \left(\frac{f'(z)}{f(z)} + \frac{G_2(z|z)}{G(z|z)} \right) \int_z^{\bar{v}} G(v|v) dF(v) \geq 0 . \quad (34)$$

This must be true if condition (a) holds. To see condition (b), notice that the log-concavity of the joint density \tilde{f} implies the log-concavity of the marginal density f and so $\frac{f'(z)}{f(z)} \geq \frac{f'(v)}{f(v)}$ for $v \geq z$. Therefore, a sufficient condition for (34) is

$$G(z|z)f(z) + \int_z^{\bar{v}} g(v|v)dF(v) + \int_z^{\bar{v}} G(v|v)f'(v)dv + \frac{G_2(z|z)}{G(z|z)} \int_z^{\bar{v}} G(v|v)dF(v) \geq 0 .$$

Applying integration by parts to the third term, we can rewrite the above condition as

$$f(\bar{v}) - \int_z^{\bar{v}} G_2(v|v)f(v)dv + \frac{G_2(z|z)}{G(z|z)} \int_z^{\bar{v}} G(v|v)dF(v) \geq 0 .$$

This holds if

$$\frac{G_2(z|z)}{G(z|z)} \geq \frac{G_2(v|v)}{G(v|v)}$$

for any $v \in [z, \bar{v}]$. This is true if $\frac{G_2(v|v)}{G(v|v)}$ is non-increasing in v . \square

Proof of Lemma 1. Here we establish existence and uniqueness of the equilibrium price. (The rest of the lemma follows from arguments in the text.)

Clearly $p - c < \phi(p)$ when $p = c$. Since $\phi(p)$ is non-increasing due to Assumption 1, it suffices to show that $p - c > \phi(p)$ when $p = \bar{v}$. The latter must be true if $\bar{v} = \infty$, because $\phi(p)$ is non-increasing and thus finite as $p \rightarrow \infty$. It also holds if $\bar{v} < \infty$ and $f(\bar{v}) > 0$, because in that case $\phi(\bar{v}) = 0$. Finally, then, consider $\bar{v} < \infty$ and $f(\bar{v}) = 0$, in which case $f(v)$ must be decreasing for v sufficiently close to \bar{v} . Notice that $\phi(p) \leq \frac{\int_p^{\bar{v}} f(v)dv}{G(p|p)f(p)}$, which for p close to \bar{v} is itself weakly less than $\frac{(\bar{v}-p)f(p)}{G(p|p)f(p)} = \frac{\bar{v}-p}{G(p|p)}$. This is clearly less than $\bar{v} - c$ when p is close to \bar{v} . \square

Remark. To ease the exposition we assumed that the joint distribution of valuations has full support on $[\underline{v}, \bar{v}]^n$. However Lemma 1 also holds under the weaker condition that if $G(\hat{v}|\hat{v}) = 0$ for some \hat{v} , then $G(v|v) = 0$ for any $v \leq \hat{v}$. (This is true, e.g., when the joint distribution has a convex support.) Define $\underline{v}^* \equiv \max\{v : G(v|v) = 0\}$. Then \underline{v} in the lemma can be replaced by \underline{v}^* .

Proof of Proposition 2. Suppose $c = \bar{v} - \varepsilon$ with $\varepsilon > 0$ close to zero. We first approximate the uniform price in equation (7), which in this case must be close to \bar{v} . One can check that $\phi(\bar{v}) = 0$ and $\phi'(\bar{v}) = -1$,⁴⁵ and so using a Taylor expansion we have $\phi(p) \approx \bar{v} - p$.

⁴⁵Notice that

$$\phi'(p) = \frac{-h_p(0)G(p|p)f(p) - \int_p^{\bar{v}} G(v|v)dF(v) \frac{\partial h_p(0)}{\partial p}}{h_p(0)^2} ,$$

and so $\phi'(\bar{v}) = -1$ given $h_{\bar{v}}(0) = f(\bar{v})$ and $G(\bar{v}|\bar{v}) = 1$.

The uniform price is therefore $p \approx \frac{1}{2}(\bar{v} + c) = \bar{v} - \frac{\varepsilon}{2}$, and the profit margin is $p - c \approx \frac{\varepsilon}{2}$. (Note that the approximated price is actually independent of the number of firms.) The equilibrium demand is $1 - F_{(n)}(p) \approx f_{(n)}(\bar{v})(\bar{v} - p) \approx f_{(n)}(\bar{v})\frac{\varepsilon}{2}$, where $f_{(n)}(\cdot)$ is the density of the highest order statistic. Hence industry profit under uniform pricing is

$$\Pi_U \approx f_{(n)}(\bar{v}) \frac{\varepsilon^2}{4}.$$

On the other hand, under perfect price discrimination recall from (18) that industry profit is $\Pi_D(c) = n \int_c^{\bar{v}} [G(v|\bar{v}) - G(v|v)F(v)]dv$. One can check that $\Pi'_D(\bar{v}) = 0$ and

$$\Pi''_D(\bar{v}) = n [f(\bar{v}) + G_2(\bar{v}|\bar{v})] = nf(\bar{v}),$$

where $G_2(\cdot|\cdot)$ is defined in (32) and we have used the fact $G_2(\bar{v}|\bar{v}) = 0$ given $G(\bar{v}|y) = 1$ for any y . We can therefore approximate industry profit as

$$\Pi_D \approx nf(\bar{v}) \frac{(\bar{v} - c)^2}{2} = nf(\bar{v}) \frac{\varepsilon^2}{2}.$$

Notice that from

$$F_{(n)}(v) = \int_{[\underline{v}, v]^n} \tilde{f}(\mathbf{v}) d\mathbf{v},$$

we have

$$f_{(n)}(v) = n \int_{[\underline{v}, v]^{n-1}} \tilde{f}(v, \mathbf{v}_{-i}) d\mathbf{v}_{-i},$$

where we have used the exchangeability of \tilde{f} . Therefore, $f_{(n)}(\bar{v}) = nf(\bar{v})$.⁴⁶ With this result, we can claim that the approximated Π_D is greater than the approximated Π_U , i.e., personalized pricing improves profit when c is sufficiently large.

Now consider consumer surplus. Recall from (13) that consumer surplus under uniform pricing is $V_U(p) = \int_p^{\bar{v}} [1 - F_{(n)}(v)]dv$. Using $V'_U(\bar{v}) = 0$ and $V''_U(\bar{v}) = f_{(n)}(\bar{v})$, we can approximate it as

$$V_U \approx f_{(n)}(\bar{v}) \frac{(\bar{v} - p)^2}{2} \approx f_{(n)}(\bar{v}) \frac{\varepsilon^2}{8}.$$

On the other hand, recall from (16) that consumer surplus under price discrimination is $V_D(c) = \int_c^{\bar{v}} [1 - F_{(n-1)}(v)]dv$. Using $V'_D(\bar{v}) = 0$ and $V''_D(\bar{v}) = f_{(n-1)}(\bar{v})$, we can approximate it as

$$V_D \approx f_{(n-1)}(\bar{v}) \frac{(\bar{v} - c)^2}{2} = f_{(n-1)}(\bar{v}) \frac{\varepsilon^2}{2},$$

⁴⁶The same result holds if the support of \tilde{f} is not full on $[\underline{v}, \bar{v}]^n$ but is of full dimension. Due to the exchangeability, we can write $F_{(n)}(v) = n \int_{\underline{v}}^v \int_{S_{n-1}(v_i)} \tilde{f}(v_i, \mathbf{v}_{-i}) d\mathbf{v}_{-i} dv_i$, where $S_{n-1}(v_i) \subset [\underline{v}, v_i]^{n-1}$ is the support of \mathbf{v}_{-i} conditional on v_i and when all $v_{j \neq i} \leq v_i$. Then $f_{(n)}(v) = n \int_{S_{n-1}(v)} \tilde{f}(v, \mathbf{v}_{-i}) d\mathbf{v}_{-i}$. When $v = \bar{v}$, $S_{n-1}(\bar{v})$ becomes the whole support of \mathbf{v}_{-i} . Therefore, $f_{(n)}(\bar{v}) = nf(\bar{v})$. Intuitively, when \tilde{f} has a support of full dimension, when one product has the highest valuation, the conditional chance that another product also has the highest valuation is zero. Since it is equally likely for each product to have the highest valuation, we have $f_{(n)}(\bar{v}) = nf(\bar{v})$.

where $f_{(n-1)}(\cdot)$ is the density of the second highest order statistic. From (3) we have $f_{(n-1)}(v) = f_{(n)}(v) - nG(v|v)f(v) + n \int_{\underline{v}}^{\bar{v}} g(v|v_i)dF(v_i)$, and so $f_{(n-1)}(\bar{v}) = f_{(n)}(\bar{v}) - nf(\bar{v}) = 0$, where we have used $f_{(n)}(\bar{v}) = nf(\bar{v})$. In other words, V_D is at most of the magnitude of ε^3 when $c = \bar{v} - \varepsilon$. Therefore, when $f(\bar{v}) > 0$, personalized pricing harms consumers when c is sufficiently large. \square

Proof of Corollary 1. From the proof of Proposition 2, we derive that when $c = \bar{v} - \varepsilon$, each firm's profit under competitive personalized pricing is $\frac{1}{n}\Pi_D \approx f(\bar{v})\frac{\varepsilon^2}{2}$, and the monopoly profit under uniform pricing is $\Pi_U^m \approx f(\bar{v})\frac{\varepsilon^2}{4}$. Therefore, we have $\frac{1}{n}\Pi_D > \Pi_U^m$. For consumer surplus, when $c = \bar{v} - \varepsilon$, $V_U^m \approx f(\bar{v})\frac{\varepsilon^2}{8}$ while V_D is at most of the magnitude of ε^3 . \square

Proof of Proposition 3. In the IID case industry profit under uniform pricing can be written as

$$\Pi_U = \frac{[1 - F(p)^n]^2/n}{F(p)^{n-1}f(p) + \int_{\underline{v}}^{\bar{v}} f(v)dF(v)^{n-1}}.$$

Under the log-concavity condition, the uniform price p is decreasing in n , and so $F(p)^n$ must be of order $o(\frac{1}{n})$, i.e., $\lim_{n \rightarrow \infty} \frac{F(p)^n}{1/n} = 0$. Meanwhile, Theorem 1 in Gabaix et al. (2016), which approximates the Perloff-Salop price, has shown that as $n \rightarrow \infty$,

$$\int_{\underline{v}}^{\bar{v}} f(v)dF(v)^{n-1} \sim f(F^{-1}(1 - \frac{1}{n})) \cdot \Gamma(2 + \gamma),$$

where $\Gamma(\cdot)$ is the Gamma function. (Notice that $\Gamma(x)$ is non-monotonic in $x \in [1, 2]$ (decreases first and then increases) and is strictly positive but no greater than 1 in that range (with $\Gamma(1) = \Gamma(2) = 1$.) At the same time, notice that $\int_{\underline{v}}^p f(v)dF(v)^{n-1} < F(p)^{n-1} \times \max_{v \in [\underline{v}, p]} f(v)$, so it must be of order $o(\frac{1}{n})$ given f is finite. Therefore, as $n \rightarrow \infty$, we have

$$\Pi_U \sim \frac{[1 - o(\frac{1}{n})]^2/n}{o(\frac{1}{n}) + f(F^{-1}(1 - \frac{1}{n})) \cdot \Gamma(2 + \gamma)}.$$

Since the price is decreasing in n , π_U must be finite for any n . This implies that $\lim_{n \rightarrow \infty} nf(F^{-1}(1 - \frac{1}{n})) > 0$. Therefore, when n is large, those $o(\frac{1}{n})$ terms can be safely ignored. This yields

$$\Pi_U \sim \frac{1}{nf(F^{-1}(1 - \frac{1}{n})) \cdot \Gamma(2 + \gamma)}. \quad (35)$$

The case of personalized pricing is simpler. Industry profit in this case is

$$\Pi_D = \int_c^{\bar{v}} \frac{1 - F(v)}{f(v)} dF(v)^n = \int_{F(c)}^1 \frac{1 - t}{f(F^{-1}(t))} dt^n.$$

Proposition 2 in Gabaix et al. (2016) has shown that, as $n \rightarrow \infty$,⁴⁷

$$\mathbb{E}[v_{n:n} - v_{n-1:n}] = \int_0^1 \frac{1-t}{f(F^{-1}(t))} dt^n \sim \frac{\Gamma(1-\gamma)}{nf(F^{-1}(1-\frac{1}{n}))} .$$

Notice that

$$\Pi_D = \mathbb{E}[v_{n:n} - v_{n-1:n}] - \int_0^{F(c)} \frac{1-t}{f(F^{-1}(t))} dt^n ,$$

and the integrand in the second term is decreasing and so the second term is less than $\frac{F(c)^n}{f(v)}$ which is of order $o(\frac{1}{n})$. Therefore, the second term can be safely ignored when n is large, and so

$$\Pi_D \sim \frac{\Gamma(1-\gamma)}{nf(F^{-1}(1-\frac{1}{n}))} . \quad (36)$$

Comparing (35) and (36), we can claim that when n is sufficiently large, personalized pricing reduces profit (and so improves consumer surplus) if

$$\Gamma(1-\gamma)\Gamma(2+\gamma) < 1 ,$$

which is true when $\gamma \in (-1, 0)$.⁴⁸ □

Proof of Lemma 3. Notice that

$$\begin{aligned} \mathbb{E}[\max\{c, v_{n:n}\}] &= \frac{1}{n}\mathbb{E}[\max\{c, v_n\}|v_n > \max\{v_1, \dots, v_{n-1}\}] \\ &\quad + (1 - \frac{1}{n})\mathbb{E}[\max\{c, v_1, \dots, v_{n-1}\}|v_n < \max\{v_1, \dots, v_{n-1}\}] , \end{aligned}$$

and with Assumption 2 we also have

$$\begin{aligned} \mathbb{E}[\max\{c, \hat{v}_{n-1:n-1}\}] &= \frac{1}{n}\mathbb{E}[\max\{c, v_1, \dots, v_{n-1}\}|v_n > \max\{v_1, \dots, v_{n-1}\}] \\ &\quad + (1 - \frac{1}{n})\mathbb{E}[\max\{c, v_1, \dots, v_{n-1}\}|v_n < \max\{v_1, \dots, v_{n-1}\}] . \end{aligned}$$

Therefore, the match efficiency improvement in (23) is equal to

$$\frac{1}{n}\mathbb{E}[\max\{c, v_n\} - \max\{c, v_1, \dots, v_{n-1}\}|v_n > \max\{v_1, \dots, v_{n-1}\}] ,$$

which is just equal to (22).

We need to further show that both the free-entry equilibrium and the socially optimal solution are unique. (Otherwise, a free-entry equilibrium could differ from a socially optimal solution due to a selection issue.) It suffices to show that (22) is decreasing

⁴⁷Note that $\int_0^1 t dt^n = 1 - \frac{1}{n}$, so the approximation is intuitive up to the adjustment $\Gamma(1-\gamma)$.

⁴⁸The equality holds when $\gamma = -1$ or 0 . Unfortunately, in these cases the approximations are not precise enough to help us compare profit meaningfully in the limit.

in n . To see that, it is more convenient to use the expression for Π_D in (15). Under Assumption 2, $x_c = v_i - \max_{j \neq i} \{c, v_j\}$ must become smaller in the sense of first-order stochastic dominance as one more firm is added, and so $1 - H_c(x)$ decreases in n for any x . This implies that $\frac{1}{n}\Pi_D$ decreases in n . \square

Omitted materials in Section 5. Here we report the details of the consumer surplus comparison across the three regimes when valuations are IID and follow the standard exponential distribution.

Lemma 7. *Suppose valuations are IID standard exponential and $k = 1$. Then there exists a \tilde{c} such that for $c < \tilde{c}$ we have $V_M < V_D \leq V_U$, and for $c > \tilde{c}$ we have $V_D < V_M < V_U$.*

Proof of Lemma 7. From (28), we have that with $k = 1$, $V_M = \mathbb{E}[\max\{0, v_{n-1:n-1} - p\}] = \int_p^{\bar{v}} [1 - F_{(n-1)}(v)] dv$. Notice also that with IID standard exponential valuations, a firm's optimal price is $1 + c$ regardless of the distribution of the outside option. Hence in both the uniform pricing and the mixed regime, the uniform price is $1 + c$. It then follows that

$$V_M = \int_{1+c}^{\infty} [1 - F(v)^{n-1}] dv < \int_{1+c}^{\infty} [1 - F(v)^n] dv = V_U .$$

We also proved on page 16 that $V_D \leq V_U$. Hence it suffices to show that $V_M < V_D$ for $c < \tilde{c}$, and $V_M > V_D$ for $c > \tilde{c}$, which we do in the remainder of the proof.

First, it is easy to show that $\lim_{c \rightarrow \infty} (V_D - V_M) = 0$, and also from l'hôpital's rule that

$$\lim_{c \rightarrow \infty} \frac{V_D - V_M}{V_M} = -1 < 0 .$$

That is, in the limit as $c \rightarrow \infty$, V_D approaches V_M from below.

Second, we claim that $V_D > V_M$ at $c = -1$.⁴⁹ To prove this, note that

$$\begin{aligned} V_D &= 1 + \int_0^{\infty} [1 - F(v)^{n-1}] dv - \int_0^{\infty} (n-1)[1 - F(v)]F(v)^{n-1} dv \\ &= 1 + V_M - \int_0^1 (n-1)t^{n-1} dt = V_M + \frac{1}{n} > V_M , \end{aligned}$$

where the second line uses the definition of V_M to rewrite the second term (from the first line), and the change of variable $t = F(v)$ to rewrite the third term (from the first line).

Third, note that for $c < -1$ we have that $V_D - V_M$ is constant in c , whereas for $c \in (-1, 0)$ we have that $V_D - V_M$ is decreasing in c . Moreover, $V_D - V_M$ is quasiconvex

⁴⁹We allow for negative costs because in this example we normalize the lowest valuation to equal 0.

in $c > 0$. To prove this, note that for $c > 0$ we can write

$$\begin{aligned}
\frac{\partial(V_D - V_M)}{\partial c} &= -F(1+c)^{n-1} + F(c)^{n-1} + (n-1)[1 - F(c)]F(c)^{n-1} \\
&= -[1 - e^{-1} + e^{-1}F(c)]^{n-1} + F(c)^{n-1} + (n-1)[1 - F(c)]F(c)^{n-1} \\
&\propto -\left[\frac{1 - e^{-1}}{s} + e^{-1}\right]^{n-1} + 1 + (n-1)(1-s) ,
\end{aligned} \tag{37}$$

where the second line uses $F(1+c) = 1 - e^{-1}[1 - F(c)]$, and the third line uses the change of variable $s = F(c)$. One can check that (37) is concave in s , negative as $s \rightarrow 0$ and zero as $s \rightarrow 1$, and decreasing in s as $s \rightarrow 1$. This establishes quasiconvexity of $V_D - V_M$.

Combining the above three steps then implies that there exists a critical \tilde{c} such that $V_D > V_M$ for $c < \tilde{c}$, and $V_D < V_M$ for $c > \tilde{c}$. \square

Proof of Lemma 4. The proof largely follows the literature on auctions with interdependent values. We look for a symmetric equilibrium where $b(v) = v - p(v)$ is the equilibrium surplus bidding function and $b(v)$ increases monotonically in v . When a firm observes consumer valuation v but deviates and bids according to valuation z , its expected profit is

$$[v - b(z) - c]G(z|v) . \tag{38}$$

Its derivative with respect to z is

$$-b'(z)G(z|v) + [v - b(z) - c]g(z|v) . \tag{39}$$

The deviation profit (38) should be maximized at $z = v$ in symmetric equilibrium, and so the first-order condition is

$$-b'(v)G(v|v) + [v - b(v) - c]g(v|v) = 0 ,$$

from which we derive a differential equation

$$b'(v) = [v - b(v) - c] \frac{g(v|v)}{G(v|v)} . \tag{40}$$

The natural boundary condition is $b(c) = 0$, which allows us to solve for

$$b(v) = \int_c^v (x - c) dL(x|v) = \int_c^v [1 - L(x|v)] dx ,$$

where $L(x|v)$ is defined in (29). Notice that $b'(v) = -\int_c^v \frac{\partial L(x|v)}{\partial v} dx > 0$ (where we have used the facts that $L(v|v) = 1$ and $L(x|v)$ decreases in v), so $b(v)$ is indeed increasing.

To check that the first-order condition is sufficient, substitute (40) into (39) to get

$$\begin{aligned} G(z|v) \left(-b'(z) + [v - b(z) - c] \frac{g(z|v)}{G(z|v)} \right) \\ = G(z|v) \left([v - b(z) - c] \frac{g(z|v)}{G(z|v)} - [z - b(z) - c] \frac{g(z|z)}{G(z|z)} \right) . \end{aligned}$$

Under Assumption 3 this is positive for $z < v$ and negative for $z > v$, and hence the first-order condition is indeed sufficient for defining the equilibrium. \square

Proof of Proposition 5. For part (i), note that under partial discrimination $b(v)$ is strictly positive and strictly increasing in v whenever $v > c$, and so a consumer buys the best-matched product whenever its valuation exceeds marginal cost. The same is true under full discrimination, so the two regimes yield the same total welfare.

For parts (ii) and (iii) let us compare profit. (The comparison of consumer surplus is just the opposite.) When a firm wins a consumer with valuation v , its profit

$$p(v) - c = v - b(v) - c = \int_c^v L(x|v) dx . \quad (41)$$

Recall that as derived in (17) the counterpart under full discrimination is

$$\frac{\int_c^v G(x|v) dx}{G(v|v)} . \quad (42)$$

Suppose first that $\frac{g(z|v)}{G(z|v)}$ increases in v . Then

$$\begin{aligned} L(x|v) &= \exp \left(- \int_x^v \frac{g(t|t)}{G(t|t)} dt \right) \geq \exp \left(- \int_x^v \frac{g(t|v)}{G(t|v)} dt \right) \\ &= \exp \left(- \int_x^v [\ln G(t|v)]' dt \right) = \frac{G(x|v)}{G(v|v)} . \end{aligned}$$

Therefore, (41) is greater than (42), i.e., firms earn more under partial discrimination. The opposite is true if $\frac{g(z|v)}{G(z|v)}$ decreases in v . In the IID case, $\frac{g(z|v)}{G(z|v)}$ is independent of v , so the equivalence result follows. \square

Proof of Lemma 5. The derivative of equation (30) with respect to p_i is proportional to

$$\frac{1 - H_{p^* - \Delta}(p_i - p^* + \Delta)}{h_{p^* - \Delta}(p_i - p^*)} - (p_i - \Delta - c) . \quad (43)$$

Firm i 's profit is quasiconcave provided this is decreasing in p_i . This is clearly true for the second term. The derivative of the first term with respect to p_i is proportional to

$$-\frac{h'_{p^*-\Delta}(p_i - p^*)}{h_{p^*-\Delta}(p_i - p^*)} - \frac{h_{p^*-\Delta}(p_i - p^* + \Delta)}{1 - H_{p^*-\Delta}(p_i - p^* + \Delta)} \leq \frac{h_{p^*-\Delta}(p_i - p^*)}{1 - H_{p^*-\Delta}(p_i - p^*)} - \frac{h_{p^*-\Delta}(p_i - p^* + \Delta)}{1 - H_{p^*-\Delta}(p_i - p^* + \Delta)} \leq 0,$$

where the two inequalities follow because $1 - H_{p^*-\Delta}(x)$ is logconcave in x . Quasiconcavity of firm i 's profit then follows.

Now consider part (i) of the lemma. Substitute $p_i = p^*$ into equation (43) and set it to 0 to get equation (31). To prove that this equation has a unique solution, rewrite it as

$$p^* - \Delta - c - \frac{1 - H_{p^*-\Delta}(\Delta)}{1 - H_{p^*-\Delta}(0)} \frac{1 - H_{p^*-\Delta}(0)}{h_{p^*-\Delta}(0)} = 0. \quad (44)$$

The lefthand side of (44) is strictly negative when evaluated at any $p^* \leq c + \Delta$. It is strictly positive when evaluated at $p^* = \bar{v}$. To see this, recall from equation (4) the definition of the random variable $x_z \equiv v_i - \max_{j \neq i} \{z, v_j\}$, which has cdf $H_z(x)$ and pdf $h_z(x)$. Then notice that by assumption $\bar{v} - \Delta - c > \max\{c, \underline{v}\} - c \geq 0$, and also by assumption $\bar{v} - \Delta > \max\{c, \underline{v}\} \geq \underline{v}$ and thus $\sup x_{\bar{v}-\Delta} = \bar{v} - \max\{\underline{v}, \bar{v} - \Delta\} = \Delta$ which in turn implies that $1 - H_{\bar{v}-\Delta}(\Delta) = 0$ and $h_{\bar{v}-\Delta}(0) > 0$. The lefthand side of (44) is also strictly increasing in p^* . Specifically, the second part of Assumption 4 implies that $\frac{1 - H_{p^*-\Delta}(0)}{h_{p^*-\Delta}(0)}$ decreases in p^* , and it also implies that $\frac{1 - H_{p^*-\Delta}(\Delta)}{1 - H_{p^*-\Delta}(0)}$ decreases in p^* .⁵⁰

Now consider part (ii) of the lemma. The market is fully covered if and only if $p^* - \Delta \leq \underline{v}$. Equivalently, since we have just shown that the lefthand side of (44) is increasing in p^* , the market is fully covered if and only if the lefthand side of (44) is positive when evaluated at $p^* = \underline{v} + \Delta$ —which gives the condition stated in the lemma.

Now consider part (iii) of the lemma. Differentiating (31) gives:

$$\frac{\partial p^*}{\partial \Delta} = \frac{1 - \frac{h_{p^*-\Delta}(\Delta)}{h_{p^*-\Delta}(0)} - \frac{\partial}{\partial z} \left(\frac{1 - H_{p^*-\Delta}(\Delta)}{h_{p^*-\Delta}(0)} \right)}{1 - \frac{\partial}{\partial z} \left(\frac{1 - H_{p^*-\Delta}(\Delta)}{h_{p^*-\Delta}(0)} \right)} \in (0, 1), \quad (45)$$

since $\frac{\partial}{\partial z} \left(\frac{1 - H_{p^*-\Delta}(\Delta)}{h_{p^*-\Delta}(0)} \right) \leq 0$. Hence p^* increases and $p^* - \Delta$ decreases in Δ . \square

Proof of Proposition 6. Start with part (i). Using equation (30) a firm's equilibrium profit under a covered market is equal to

$$(p^* - c)[1 - H(\Delta)] + \int_0^\Delta (p^* - \Delta x - c) dH(x).$$

⁵⁰To see the second point, notice that logsubmodularity implies that for each $x' > x''$ and $z' > z''$,

$$[1 - H_{z'}(x')][1 - H_{z''}(x'')] \leq [1 - H_{z'}(x'')][1 - H_{z''}(x')].$$

Letting $x' = 0$ and $x'' = \Delta$ yields the desired result.

Its derivative with respect to Δ equals

$$\frac{\partial p^*}{\partial \Delta}[1 - H(\Delta)] + [H(\Delta) - H(0)] \left[\frac{\partial p^*}{\partial \Delta} - 1 \right] = 1 - H(\Delta) - [1 - H(0)] \frac{h(\Delta)}{h(0)} \leq 0 ,$$

where the equality uses $\frac{\partial p^*}{\partial \Delta} = 1 - \frac{h(\Delta)}{h(0)}$ (from equation (45) evaluated at $p^* - \Delta < \underline{v}$), and the inequality uses the logconcavity of $1 - H(x)$ from Assumption 4. Hence profit decreases in Δ . Since the market is fully covered total welfare is invariant to Δ , consumer surplus increases in Δ .

Now consider part (ii). Using equation (30) a firm's equilibrium profit is

$$(p^* - c)[1 - H_{p^*-\Delta}(\Delta)] + \int_0^\Delta (p^* - \Delta + x - c) dH_{p^*-\Delta}(x) .$$

Its derivative with respect to Δ around the point $\Delta = 0$ is

$$\frac{\partial p^*}{\partial \Delta}[1 - H_{p^*}(0)] + (p^* - c) \frac{\partial[1 - H_{p^*}(0)]}{\partial p^*} \left[\frac{\partial p^*}{\partial \Delta} - 1 \right] > 0, \quad (46)$$

where the inequality uses $\frac{\partial p^*}{\partial \Delta} \geq 0$, $\frac{\partial p^*}{\partial \Delta} - 1 < 0$, and $\frac{\partial[1 - H_{p^*}(0)]}{\partial p^*} < 0$.

Finally, consider consumer surplus. Total surplus is $\int_{p^*-\Delta}^{\bar{v}} (v - c) dF_{(n)}(v)$ since a consumer buys if the value of the best product exceeds the lowest possible price. Its derivative with respect to Δ around $\Delta = 0$ is then

$$(p^* - c) f_{(n)}(p^*) \left[1 - \frac{\partial p^*}{\partial \Delta} \right]. \quad (47)$$

Using $\frac{\partial[1 - H_{p^*}(0)]}{\partial p^*} = -G(p^*|p^*)f(p^*)$, and multiplying equation (46) by the number of firms n yields the derivative of industry profit with respect to Δ around $\Delta = 0$:

$$n \left\{ \frac{\partial p^*}{\partial \Delta}[1 - H_{p^*}(0)] + (p^* - c)G(p^*|p^*)f(p^*) \left[1 - \frac{\partial p^*}{\partial \Delta} \right] \right\}. \quad (48)$$

Starting from $\Delta = 0$, a small increase in Δ reduces consumer surplus if and only if (48) exceeds (47). Given $\frac{\partial p^*}{\partial \Delta} \in (0, 1)$, a sufficient condition for that is $nG(p^*|p^*)f(p^*) \geq f_{(n)}(p^*)$. This is not always true, but in the IID case, we have $nG(p^*|p^*)f(p^*) = nF(p^*)^{n-1}f(p^*) = f_{(n)}(p^*)$. The claimed result then follows. □

References

- ADACHI, T. (2022): “A Sufficient Statistics Approach for Welfare Analysis of Oligopolistic Third-Degree Price Discrimination,” mimeo.
- AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): “Monopoly Price Discrimination and Demand Curvature,” *American Economic Review*, 100(4), 1601–1615.
- ALI, N., G. LEWIS, AND S. VASSERMAN (2020): “Voluntary Disclosure and Personalized Pricing,” mimeo.
- ANDERSON, S., A. BAIK, AND N. LARSON (2021): “Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts,” mimeo.
- ANDERSON, S., A. DE PALMA, AND Y. NESTEROV (1995): “Oligopolistic Competition and the Optimal Provision of Products,” *Econometrica*, 63(6), 1281–1301.
- APARICIO, D., Z. METZMAN, AND R. RIGOBON (2021): “The Pricing Strategies of Online Grocery Retailers,” mimeo.
- ARMSTRONG, M. (2007): “Recent Developments in the Economics of Price Discrimination,” in *Advances in Economics and Econometrics (Theory and Applications, Ninth World Congress)*, ed. by R. Blundell, W. Newey, and T. Persson. Cambridge University Press.
- ARMSTRONG, M., AND J. VICKERS (2001): “Competitive Price Discrimination,” *RAND Journal of Economics*, 32(4), 579–605.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105(3), 921–957.
- BOUNIE, D., A. DUBUS, AND P. WAELEBROECK (2021): “Selling Strategic Information in Digital Competitive Markets,” *RAND Journal of Economics*, 52(2), 283–313.
- CAPLIN, A., AND B. NALEBUFF (1991): “Aggregation and Imperfect Competition: On the Existence of Equilibrium,” *Econometrica*, 59(1), 25–59.
- CHEN, Y., AND G. IYER (2002): “Consumer Addressability and Customized Pricing,” *Marketing Science*, 21(2), 197–208.
- CHEN, Y., J. LI, AND M. SCHWARTZ (2021): “Competitive Differential Pricing,” *RAND Journal of Economics*, 52(1), 100–124.

- CHEN, Y., AND M. RIORDAN (2007): “Price and Variety in the Spokes Model,” *The Economic Journal*, 117(522), 897–921.
- CHEN, Z., C. CHOE, AND N. MATSUSHIMA (2020): “Competitive Personalized Pricing,” *Management Science*, 66(9), 4003–4023.
- CORTS, K. (1998): “Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment,” *RAND Journal of Economics*, 29(2), 306–323.
- COUNCIL OF ECONOMIC ADVISERS (2015): “Big Data and Differential Pricing,” Discussion paper, The White House.
- DELOITTE (2018): “Consumer Experience in the Retail Renaissance: How Leading Brands Build a Bedrock with Data,” Deloitte Digital.
- DUBE, J.-P., AND S. MISRA (2019): “Personalized Pricing and Customer Welfare,” mimeo.
- ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2021): “Market Segmentation Through Information,” mimeo.
- FINANCIAL CONDUCT AUTHORITY (FCA) (2019): “General Insurance Pricing Practices,” Market Study MS18/1.2.
- FUDENBERG, D., AND J. M. VILLAS-BOAS (2007): “Behavior-Based Price Discrimination and Customer Recognition,” in *Handbook on Economics and Information Systems (Vol. 1)*, ed. by T. Hendershott. Elsevier.
- GABAIX, X., D. LAIBSON, D. LI, H. LI, S. RESNICK, AND C. G. DE VRIES (2016): “The Impact of Competition on Prices with Numerous Firms,” *Journal of Economic Theory*, 165, 1–24.
- HANNAK, A., G. SOELLER, D. LAZER, A. MISLOVE, AND C. WILSON (2014): “Measuring Price Discrimination and Steering on E-Commerce Web Sites,” in *Proceedings of the 2014 Conference on Internet Measurement Conference*, p. 305–318.
- HOLMES, T. (1989): “The Effects of Third-Degree Price Discrimination in Oligopoly,” *American Economic Review*, 79(1), 244–250.
- ICHIHASHI, S., AND A. SMOLIN (2022): “Data Collection by an Informed Seller,” mimeo.
- JULLIEN, B., M. REISINGER, AND P. REY (2022): “Personalized Pricing and Distribution Strategies,” *Management Science*, forthcoming.

- KEHOE, P., B. LARSEN, AND E. PASTORINO (2020): “Dynamic Competition in the Era of Big Data,” mimeo.
- LEVIN, D., AND J. SMITH (1994): “Equilibrium in Auctions with Entry,” *American Economic Review*, 84(3), 585–599.
- LIU, Q., AND K. SERFES (2004): “Quality of Information and Oligopolistic Price Discrimination,” *Journal of Economics & Management Strategy*, 13(4), 671–702.
- MILGROM, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- MILGROM, P. R., AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089–1122.
- MONTES, R., W. SAND-ZANTMAN, AND T. VALLETTI (2019): “The Value of Personal Information in Online Markets with Endogenous Privacy,” *Management Science*, 65(3), 1342–1362.
- OECD (2018): “Personalised Pricing in the Digital Era,” Discussion paper, Organisation for Economic Co-operation and Development.
- PERLOFF, J., AND S. SALOP (1985): “Equilibrium with Product Differentiation,” *Review of Economic Studies*, 52(1), 107–120.
- QUINT, D. (2014): “Imperfect Competition with Complements and Substitutes,” *Journal of Economic Theory*, 152, 266–290.
- RHODES, A., AND J. ZHOU (2022): “Personalization and Privacy Choice,” under preparation.
- SCHMALENSEE, R. (1981): “Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination,” *American Economic Review*, 71(1), 242–247.
- SHAFFER, G., AND Z. J. ZHANG (2002): “Competitive One-to-One Promotions,” *Management Science*, 48(9), 1143–1160.
- SHILLER, B. R. (2020): “Approximating Purchase Propensities and Reservation Prices from Broad Consumer Tracking,” *International Economic Review*, 61(2).
- (2021): “Optimized Sticky Targeted Pricing,” mimeo.
- STOLE, L. (2007): “Price Discrimination and Competition,” in *Handbook of Industrial Organization (Vol. 3)*, ed. by M. Armstrong, and R. Porter. North Holland.

- TAN, G., AND J. ZHOU (2021): “Consumer Heterogeneity and Inefficiency in Oligopoly Markets,” mimeo.
- THISSE, J.-F. T., AND X. VIVES (1988): “On The Strategic Choice of Spatial Price Policy,” *American Economic Review*, 78(1), 122–137.
- VARIAN, H. (1989): “Price Discrimination,” in *Handbook of Industrial Organization (Vol. 1)*, ed. by R. Schmalensee, and R. Willig. North-Holland.
- WALDFOGEL, J. (2015): “First Degree Price Discrimination Goes to School,” *Journal of Industrial Economics*, 63(4), 569–597.
- ZHOU, J. (2017): “Competitive Bundling,” *Econometrica*, 85(1), 145–172.