

WORKING PAPERS

N° 1333

March 2024

“Personalized Pricing and Competition”

Andrew Rhodes and Jidong Zhou

Personalized Pricing and Competition*

Andrew Rhodes[†] Jidong Zhou[‡]

March 2024

Abstract

We study personalized pricing in a general oligopoly model. The impact of personalized pricing relative to uniform pricing hinges on the degree of market coverage. If market conditions are such that coverage is high (e.g., the production cost is low, or the number of firms is high), personalized pricing harms firms and benefits consumers, whereas the opposite is true if coverage is low. When only some firms have data to personalize prices, consumers can be worse off compared to when either all or no firms personalize prices.

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

JEL classification: D43, D82, L13

*We are grateful to the Editor Pietro Ortoleva and three anonymous referees, as well as Nageeb Ali, Mark Armstrong, Heski Bar-Isaac, Christoph Carnehl, Yongmin Chen, Ioana Chioveanu, Alexandre de Cornière, Joyee Deb, Florian Ederer, Makoto Hanazono, Christian Hellwig, Johannes Hörner, Bruno Jullien, Ilan Kremer, David Martimort, Leslie Marx, Barry Nalebuff, Martin Peitz, Ben Polak, Helen Ralston, Patrick Rey, Guillem Roig, David Ronayne, Larry Samuelson, Ali Shourideh, Alex Smolin, Steve Tadelis, Guofu Tan, Nikhil Vellodi, John Vickers, Chengsi Wang, Junjie Zhou, and various seminar and conference 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) and funding from the European Union (ERC, DMPDE, grant 101088307). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. 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

Firms increasingly have access to very rich data about individual consumers, including their (real-time) location, income, browsing and purchase history, and social media posts. Moreover, thanks to advances in information technology, firms are able to glean more and more information about consumers from their data, such as their preferences over different products. This in turn makes it easier for firms to do personalized pricing, i.e., charge different consumers different prices based on their preferences. Indeed, personalized pricing has already been documented in a wide range of industries including retail, travel, personal finance, mobile gaming and dating.¹ To avoid a consumer backlash, personalized prices are often sent by email or smartphone app in the form of personalized discounts.²

Personalized pricing has attracted a lot of interest from policymakers in recent years, and its use by firms raises a number of important questions.³ For example, it is natural to expect that personalized pricing will cause some consumers to pay more and others to pay less—but which consumers gain and which ones lose? And what is the impact on aggregate consumer surplus and firm profit? Moreover, in some industries, only certain firms have access to data that is needed to personalize prices. How does this data asymmetry affect market performance, and what would be the effect of policies that force data-rich firms to share data with their rivals?

In this paper we develop a general oligopoly model to address these and related questions. Personalized pricing is a form of very fine-tuned price discrimination, where small numbers of consumers with very similar preferences are grouped together and offered the same price. We focus on the limit case of personalized pricing, i.e., first-degree price discrimination, where firms can perfectly infer individual consumers’ preferences and tailor prices accordingly. Although firms may never fully reach this limit case, as they acquire increasingly rich data and access to powerful Artificial Intelligence (AI) algorithms, they are likely to get closer and closer to it.⁴

¹See <https://bit.ly/3A4Rk10> and <https://bit.ly/38Ygzq6> for a history of personalized pricing, and OECD (2018), Which? (2018) and <http://bit.ly/3E4nDBT> for more details of the examples.

²See, e.g., <https://bit.ly/370ftAc> for how Kroger uses its mobile app to send personalized coupons. Firms can conceal personalized prices in other ways too, such as “sticky targeting”, where prices are fixed for all consumers (including the targeted one) for a short period (Shiller, 2021), or personalized search results, which steer different consumers to products with different prices (see, e.g., Hannak et al., 2014).

³See, e.g., reports by OECD (2018), European Commission (2018), Ofcom (2020) and BEIS (2021).

⁴In this spirit, Council of Economic Advisers (2015) points out that the “increased availability of behavioral data has also encouraged a shift from third-degree price discrimination based on broad demographic categories towards personalized pricing.” Similarly, Varian (2018) writes that “Fully personalized pricing is unrealistic, but prices based on fine grained features of consumers may well be feasible, so the line between third degree and first degree is becoming somewhat blurred.”

Two benchmarks To set the scene, we first review the impact of personalized pricing in two well-known benchmarks from the literature, and then explain our contribution.

The first benchmark is the monopoly case studied by Pigou (1920). Suppose consumers wish to buy one unit of a product and have heterogeneous valuations. Personalized pricing enables the firm to charge each consumer with a valuation above marginal cost a personalized price equal to their valuation. As a result, total surplus is maximized but it is fully extracted by the monopolist. Compared to monopoly uniform pricing, personalized pricing therefore increases total welfare and firm profit but reduces consumer surplus.

The second benchmark is the linear Hotelling case studied by Thisse and Vives (1988). Suppose consumers are uniformly distributed along a unit-length Hotelling line, 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. Suppose the marginal production cost is normalized to zero. Under uniform pricing firms set the standard Hotelling price of 1. Under personalized pricing, except for consumers at $x = 1/2$, the firms engage in asymmetric Bertrand competition for each consumer individually. Specifically, for a consumer with $x < 1/2$ who prefers firm 1, firm 2 prices at cost, while firm 1 sells to her at a price $p_1(x) = v_1 - v_2 = 1 - 2x$. Similarly, for a consumer with $x > 1/2$, firm 1 prices at cost, while firm 2 sells at a price $p_2(x) = v_2 - v_1 = 2x - 1$.

Figure 1a depicts the uniform price (the solid curve) and personalized prices paid by different consumers (the dashed curve). It is clear that *every* consumer pays (weakly) less under personalized pricing because $p_1(x) \leq 1$ and $p_2(x) \leq 1$. (Intuitively, each firm tries to poach consumers on its rival’s “turf” with low prices, which then forces the rival to charge less even to consumers with a strong preference for its product.) Therefore going from monopoly to duopoly completely reverses the impact of personalized pricing—it now harms firms but benefits consumers. This insight that personalized pricing intensifies competition 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, including the burgeoning literature on data privacy and data brokers.

However, as observed in Armstrong (2007) and Ali, Lewis, and Vasserman (2023), the result that each personalized price is lower than the uniform price can easily be overturned.⁵ Suppose instead that consumers are distributed along the Hotelling line according to a symmetric and strictly log-concave (so single-peaked) density. Personalized prices are *independent* of the distribution and so remain unchanged, but the uniform price,

⁵Armstrong (2007, p. 114) observes this in an example with a specific non-uniform distribution, while Ali, Lewis, and Vasserman (2023, p. 560) show it for a symmetric and strictly log-concave distribution as discussed here.

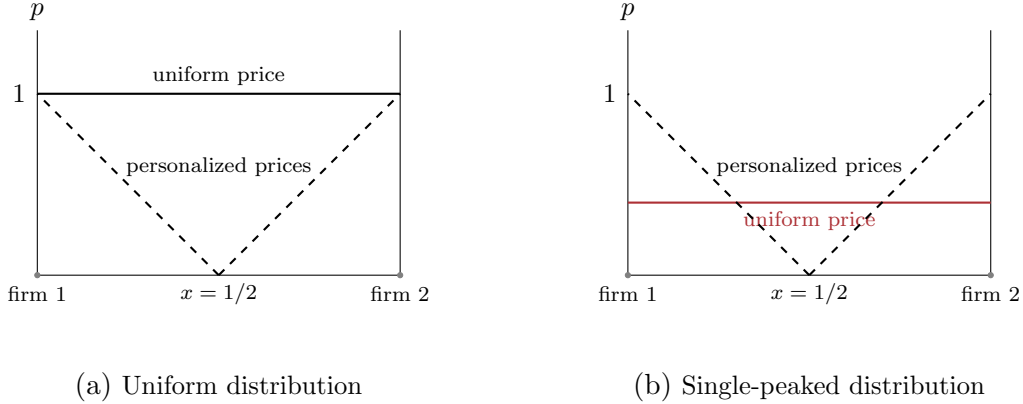


Figure 1: The impact of personalized pricing in the Hotelling model

which equals 1 over the density of consumers at $x = 1/2$, is now strictly below 1. Hence, as depicted in Figure 1b, consumers near the middle of the line (with relatively weak preferences) still pay less under personalized pricing, but those near the two ends of the line (with relatively strong preferences) now pay more. The impact of personalized pricing on industry profit and aggregate consumer surplus is then less clear.

Our contribution In this paper we develop a general oligopoly model to examine the impact of personalized pricing. We investigate when the Thisse and Vives (1988) insight that competitive personalized pricing reduces profit and raises *aggregate* consumer surplus holds and when it fails. In turn, this allows us to also reconcile the opposing impacts of personalized pricing under monopoly and (linear) Hotelling.

In more detail, in Section 2 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 the classic random-utility model developed in 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 do not have or are banned from using 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.

Section 3 compares market performance in these two regimes. Under a mild regularity condition, we show that personalized pricing benefits consumers who regard their two best products as close substitutes but harms others with a strong preference for one product. This generalizes the existing observation from the non-linear Hotelling model depicted in Figure 1b. Nevertheless, we then demonstrate that if the market is fully covered (i.e., if all

consumers buy) under uniform pricing, competitive personalized pricing lowers industry profit and increases aggregate consumer surplus under a log-concavity condition.⁶ We are therefore able to significantly generalize the aggregate welfare results from the classic Thisse and Vives (1988) paper.

We then turn to the (arguably more realistic) case where the market is not fully covered under uniform pricing. Here we show that the welfare impact of personalized pricing can be *reversed*: competitive personalized pricing can now increase industry profit and lower consumer surplus. For example, this always happens when valuations are independent across firms and exponentially distributed. For a more general valuation distribution, we prove that it happens when market coverage is sufficiently low due to a high production cost (or a good outside option). In addition, for the duopoly case with a family of generalized Pareto distributions, as well as in numerical simulations with many other common distributions, we find a cutoff result: personalized pricing benefits consumers in aggregate when cost is below a threshold (so coverage is relatively high) but otherwise harms consumers; the opposite is true (but with a lower cost threshold) for the impact on firm profits.^{7,8}

Intuitively, under personalized pricing competing firms fully pass cost increases through to consumers (who value at least two products above their production cost), whereas under uniform pricing they share some of the burden. Hence, as cost increases, personalized pricing becomes less favorable to consumers even if it helps expand the overall demand, but becomes more favorable to firms. Moreover, as cost increases further, eventually some consumers value only one product above cost; under personalized pricing firms have monopoly power over these consumers and fully extract their surplus, so this is an additional force by which personalized pricing can harm consumers and benefit firms.⁹

Finally, we discuss two extensions in Section 4. We show that the asymmetric case where only some firms can use data to price discriminate can be worse for consumers compared to the symmetric cases where all or no firms do personalized pricing. This suggests that it is sometimes desirable for consumers to force a seller with superior information to share its data with its competitors or prevent it from personalizing prices.

⁶The log-concavity condition ensures existence of a pure-strategy pricing equilibrium under uniform pricing, and also implies there are relatively more consumers with weak (rather than strong) preferences.

⁷For a cost between the two thresholds, personalized pricing benefits both firms and consumers. This is because when uniform pricing leads to only partial coverage, personalized pricing increases total surplus by expanding the market.

⁸Our numerical examples also show a similar cut-off result in terms of the number of firms. Personalized pricing harms firms and benefits consumers when the number of firms is large (so market coverage is high), but it can raise industry profit and harm consumers even for moderately high numbers of firms.

⁹However, as we discuss in more detail later, this additional “monopoly” force is not necessary for reversing the impact of personalized pricing relative to the full-coverage case.

We also discuss the case where the market structure is endogenously determined in a free-entry game. There we show that personalized pricing leads to the socially optimal firm entry, and so must benefit consumers relative to uniform pricing (if we ignore integer constraints). Section 5 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.)¹⁰ One exception is the study of spatial price discrimination, where firms can charge customers in different locations different prices. An influential paper within this literature is Thisse and Vives (1988) which, as explained earlier, can also be reinterpreted as a model of competitive personalized pricing.

The Thisse and Vives (1988) framework 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 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 data intermediary should sell data exclusively when firms use it to personalize prices, while 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.^{11,12}

Our paper is closely related to Anderson, Baik, and Larson (2023) (ABL henceforth), who also use a general discrete-choice framework to study competitive personalized pricing. One 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

¹⁰Since perfect price discrimination is the limit case of third-degree price discrimination, our paper is more related to competitive third-degree price discrimination. 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.

¹¹Thisse and Vives (1988) also consider the case where firms have different costs. The low-cost firm can then earn more than under uniform pricing, but within the parameter range considered by the authors, industry profit is still lower and consumer surplus is still higher under discriminatory pricing.

¹²Jullien, Reisinger, and Rey (2023) 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. Lu and Matsushima (2022) consider a Hotelling setup where consumers can buy from both firms. When the additional utility gain from buying a second product is sufficiently high, firms are close to being monopolists, and hence personalized pricing benefits firms and harms consumers.

personalized pricing. Another 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.¹³ 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. Moreover we also explore the impact of personalized pricing when firms are asymmetrically informed or when the market structure is endogenous.

Our paper is also related to Ali, Lewis, and Vasserman (2023), who investigate consumer privacy choice when firms can personalize prices. In Section 4 of their paper they study the competition case (with full market coverage). They show that if consumers can make different privacy choices across firms, there exists an equilibrium where consumers with weak preferences fully disclose their preference information to all firms, while consumers with strong preferences disclose to all firms except their favorite. This induces firms to compete so strongly that every consumer is better off relative to uniform pricing. Hence the Thisse and Vives (1988) result that every consumer benefits from personalized pricing can be restored with endogenous information disclosure. Our paper does not consider consumer privacy choice, but uses a more general demand setup which allows for partial market coverage to assess when the aggregate results of Thisse and Vives (1988) hold.

There is also growing empirical research on personalized pricing. One strand looks for evidence of personalized pricing. As discussed earlier, detecting personalized pricing is hard because sellers often disguise personalized offers. Nevertheless 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 (2024) 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; Shiller, 2020; Kehoe, Larsen, and Pastorino, 2020; and Dubé and Misra, 2023). For instance, Shiller (2020) shows that if Netflix could use rich consumer-level 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.

¹³Such randomized offers cause consumers to sometimes buy the wrong product, which harms match efficiency, and means that personalized pricing can make both firms and consumers worse off. In contrast, in our model personalized pricing (weakly) expands demand and hence can benefit both firms and consumers.

2 The Model and Pricing Equilibrium

There are n firms, 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 unit of one of the products. If a consumer buys nothing she obtains an outside option with zero surplus. Consumers perfectly know their own valuations for the n products, which are denoted by $\mathbf{v} = (v_1, \dots, v_n) \in \mathbf{R}^n$. 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.) We assume that the density function $\tilde{f}(\mathbf{v})$ is everywhere finite and differentiable. 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, where infinite valuation bounds are allowed.¹⁴ (Sometimes we focus on the IID case where the v_i 's are independent across products, in which case each v_i has a CDF F .) 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. To ensure an active market, we assume $c < \bar{v}$.

We consider two different pricing regimes. Under uniform pricing, firms set the same price for every consumer (e.g., because they do not have access to data on consumer preferences). 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. In either regime, after seeing the prices consumers choose the product with the highest surplus (if positive); if consumers are indifferent between several offers they choose the product with the highest valuation (so that total welfare is maximized). Firms and consumers are risk neutral. In both regimes we look for a Nash equilibrium where each firm maximizes its own profit.

Remarks. We will see that the degree of market coverage (i.e., how many consumers buy) affects the impact of personalized pricing. We have chosen to normalize the outside option, and to vary the degree of coverage via the marginal cost c . However the same qualitative insights obtain if instead we normalize the production cost and vary the outside option.

In the regime of personalized pricing, we assume that firms observe each consumer's vector of valuations $\mathbf{v} = (v_1, \dots, v_n)$. Later on we discuss an alternative information structure with "partial discrimination" where firm i observes only v_i .

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 $v_{n-1:n}$ denote the highest

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, there is at least one correlation structure over (v_1, v_2) that generates it.

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

$$F_{(n)}(v) = \tilde{F}(v, \dots, v) = \int_{\underline{v}}^v G(v|v_i) dF(v_i) , \quad (1)$$

and

$$F_{(n-1)}(v) = F_{(n)}(v) + n \int_{\underline{v}}^{\bar{v}} G(v|v_i) dF(v_i) . \quad (2)$$

(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 the associated density functions. In the IID case we have $G(v|v_i) \stackrel{\text{IID}}{=} F(v)^{n-1}$, $F_{(n)}(v) \stackrel{\text{IID}}{=} F(v)^n$ 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

$$x_z \equiv v_i - \max_{j \neq i} \{z, v_j\} , \quad (3)$$

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,

$$1 - H_z(x) = \Pr[v_i - x > \max_{j \neq i} \{z, v_j\}] = \int_{z+x}^{\bar{v}} G(v_i - x|v_i) dF(v_i) , \quad (4)$$

and

$$h_z(x) = G(z|z+x)f(z+x) + \int_{z+x}^{\bar{v}} g(v_i - x|v_i) dF(v_i) . \quad (5)$$

When z is irrelevant (i.e., when $z \leq \underline{v}$), 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

2.1 Uniform pricing

We first study the regime of uniform pricing. We focus on a symmetric pure-strategy pricing equilibrium where each firm charges price p .¹⁶ Using the definition of x_z and $H_z(x)$ in equations (3) and (4), firm i 's demand if it unilaterally deviates to price p_i is

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

¹⁵Anderson, Baik, and Larson (2023) 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 at least in the duopoly case (Caplin and Nalebuff, 1991) and in the IID case (Quint, 2014).

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 Online Appendix we report some primitive conditions under which this assumption holds. For example, it holds in the IID case with a log-concave f , and in the Hotelling case (see footnote 14) provided that $v_1 - v_2$ has a log-concave density. In the rest of the paper, we suppose that Assumption 1 holds whenever uniform pricing is involved.

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

$$p - c = \phi(p) , \tag{6}$$

where

$$\phi(p) \equiv \frac{1 - H_p(0)}{h_p(0)} = \frac{\int_p^{\bar{v}} G(v|v)dF(v)}{G(p|p)f(p) + \int_p^{\bar{v}} g(v|v)dF(v)} . \tag{7}$$

Note that $1 - H_p(0)$ is a firm's equilibrium demand (which can be also written as $\frac{1}{n}[1 - F_{(n)}(p)]$ due to firm symmetry), while $h_p(0)$ measures how many marginal consumers the firm loses if it slightly raises its price.

Given the second condition in Assumption 1, $\phi(p)$ is non-increasing (and constant for $p \leq \underline{v}$), so the equilibrium price is unique. We can then show the following result. (All omitted proofs in this section and Section 3 are available in the Appendix.)

Lemma 1. (i) If $c \leq \underline{v} - \phi(\underline{v})$, the equilibrium uniform price satisfies

$$p - c = \phi(\underline{v}) = \frac{1/n}{\int_{\underline{v}}^{\bar{v}} g(v|v)dF(v)} \tag{8}$$

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

(ii) Otherwise, the equilibrium uniform price uniquely solves (6) and $p > \underline{v}$, such that the market is not fully covered in equilibrium.

Lemma 1 shows that the market is partially covered under uniform pricing whenever the marginal cost c is sufficiently high. Note that a sufficient (but by no means necessary) condition is $c \geq \underline{v}$, i.e., some consumers value a product less than marginal cost.

Remark. The literature on random utility oligopoly models usually assumes that the market is fully covered. For example, Anderson, Baik, and Larson (2023) derive (6) under that assumption. Most works further restrict attention to the IID case (e.g., Perloff and

Salop, 1985; Anderson, de Palma, and Nesterov, 1995; Gabaix et al., 2016). Section 4.2 of Zhou (2017) studies the IID case without full market coverage. In that case, using our notation, we have

$$\phi(p) \stackrel{\text{IID}}{=} \frac{[1 - F(p)^n]/n}{F(p)^{n-1}f(p) + \int_p^{\bar{v}} f(v)dF(v)^{n-1}}. \quad (9)$$

If full market coverage is assumed, $\phi(p)$ further simplifies to $1/[n \int_{\underline{v}}^{\bar{v}} f(v)dF(v)^{n-1}]$.

Industry profit under uniform pricing is

$$\Pi_U \equiv n(p - c)[1 - H_p(0)] = n \frac{[1 - H_p(0)]^2}{h_p(0)}, \quad (10)$$

where we have used the equilibrium price condition (6). Since all consumers buy their favorite product as long as it gives them a positive surplus, (aggregate) consumer surplus is

$$V_U \equiv \mathbb{E}[\max\{0, v_{n:n} - p\}] = \int_p^{\bar{v}} (v - p)dF_{(n)}(v) = \int_p^{\bar{v}} [1 - F_{(n)}(v)]dv, \quad (11)$$

where the last equality uses integration by parts. Note that the expressions for Π_U and V_U are valid regardless of whether or not the market is fully covered.

2.2 Personalized pricing

Now consider the regime where firms 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 standard asymmetric Bertrand competition for each consumer. It is well known that this pricing game can have many equilibria. We first rule out uninteresting equilibria where one or more firms price below cost. All the remaining equilibria are then outcome equivalent to the following one: (i) a firm wins a consumer if and only if its product is her favorite and is valued above cost, (ii) each "losing" firm charges c , and (iii) the "winning" firm (if any) charges the consumer a price equal to the difference between her valuation and that of the best alternative (which is either the outside option, or the second best product sold at marginal cost).¹⁷ We therefore focus on this equilibrium in the subsequent analysis:

Lemma 2. *Under personalized pricing, firm i 's equilibrium pricing schedule is:*

$$p(v_i, \mathbf{v}_{-i}) = \begin{cases} c + v_i - \underbrace{\max_{j \neq i} \{c, v_j\}}_{x_c} & \text{if } v_i \geq \max_{j \neq i} \{c, v_j\} \\ c & \text{otherwise,} \end{cases} \quad (12)$$

¹⁷For example, if a consumer values her top two products more than c , there are other (outcome equivalent) equilibria in which firms outside the top two charge more than c .

where \mathbf{v}_{-i} denotes a consumer's valuations for all products other than i . The market is fully covered if and only if $c \leq \underline{v}$.

Since all consumers buy their favorite product or take the outside option in equilibrium, the market is fully covered whenever $c \leq \underline{v}$, a weaker condition than under uniform pricing. Therefore, if the market is fully covered under uniform pricing, it is also fully covered under personalized pricing.

Notice from (12) that firm i sells to the consumer and earns margin $p(v_i, \mathbf{v}_{-i}) - c = x_c$ if and only if $x_c \geq 0$. Hence firm i 's equilibrium profit is $\int_0^\infty x dH_c(x)$, and industry profit is

$$\Pi_D = n \int_0^\infty x dH_c(x) = n \int_0^\infty [1 - H_c(x)] dx . \quad (13)$$

A consumer always buys her favorite product provided she values it above c . The equilibrium prices in (12) then ensure that she is indifferent between her favorite product and the next best option (i.e., $v_i - p(v_i, \mathbf{v}_{-i}) = \max_{j \neq i} \{0, v_j - c\}$ if firm i is the consumer's favorite). Therefore, consumer surplus under personalized pricing is

$$V_D \equiv \mathbb{E}[\max\{0, v_{n-1:n} - c\}] = \int_c^{\bar{v}} (v - c) dF_{(n-1)}(v) = \int_c^{\bar{v}} [1 - F_{(n-1)}(v)] dv . \quad (14)$$

Note that the expressions for Π_D and V_D are valid regardless of whether or not the market is fully covered. Comparing V_D and V_U , there is a trade-off: under personalized pricing it is *as if* consumers buy the second-best product at price c , while under uniform pricing consumers buy the best product at the uniform price $p > c$.

3 The Impact of Personalized Pricing

We now examine how a shift from uniform to personalized pricing affects market performance. Our first result shows that, under a mild regularity condition, the highest personalized price exceeds the uniform price, and hence personalized pricing harms some consumers. Recall that $h(x)$ is the density of $v_i - \max_{j \neq i} \{v_j\}$.

Lemma 3. *The highest personalized price exceeds the uniform price if $c \geq \underline{v}$, or if $c < \underline{v}$ and $h(x) < h(0)$ for $x > 0$.*

This result generalizes the observation in Armstrong (2007) and Ali, Lewis, and Vasserma (2023). When the cost is above \underline{v} , there are consumers who value one product at \bar{v} but all other products below c . They will be charged a personalized price \bar{v} which of course exceeds the uniform price p . When the cost is below \underline{v} , the required condition on $h(x)$ holds, for example, in the IID case with a log-concave f . However it fails in the

linear Hotelling model, because there $h(x)$ is constant everywhere; this explains why, in Thisse and Vives (1988), the highest personalized price exactly equals the uniform price.

Figure 2 illustrates the duopoly case with full market coverage, i.e., $p < \underline{v}$. Under personalized pricing, consumers with $v_1 > v_2$ buy from firm 1 and pay $v_1 - v_2 + c$, while consumers with $v_1 < v_2$ buy from firm 2 and pay $v_2 - v_1 + c$. Compared to the uniform-pricing regime with price p , consumers with strong relative preferences (high $|v_1 - v_2|$) in the northwest and southeast corners pay more, while those with weaker relative preferences (low $|v_1 - v_2|$) in the middle region pay less. The importance of *relative* rather than *absolute* valuations implies that the common wisdom that richer consumers pay more under personalized pricing may fail under competition. In particular, it seems reasonable to suppose that low valuations may be relatively concentrated among poorer consumers. Then, it is indeed the case that poor consumers with a low valuation for both products (the southwest corner) pay less under personalized pricing, and choosy consumers (the northwest and southeast corners) who are likely to be rich pay more. However, consumers with a high valuation for both products (the northeast corner), who also seem most likely to be rich, actually pay less.

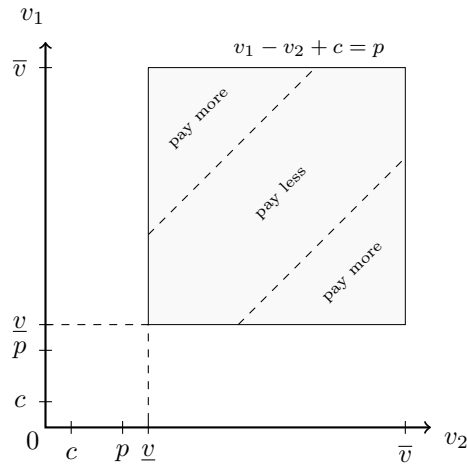


Figure 2: The impact of personalized pricing with full market coverage

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

3.1 The case of full market coverage

We first study the case where the market is fully covered under uniform pricing (and so is also fully covered under personalized pricing). Total welfare is therefore the same under either pricing regime, because in both cases every consumer buys her preferred product.

The following result reports the impact of personalized pricing on profit and aggregate consumer surplus.

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

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 from (10) and $1 - H(0) = \frac{1}{n}$, we can see that industry profit under uniform pricing is

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

while from (13) we can see that industry profit under personalized pricing 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. Therefore, firms suffer from personalized pricing. Since total welfare is unchanged, consumers benefit from personalized pricing. \square

Recalling our discussion of Figure 2, the log-concavity condition in Assumption 1 ensures that there are relatively more consumers with weak preferences who pay less under personalized pricing. Hence 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).¹⁸

3.2 The case of partial market coverage

We now turn to the case where the market is not fully covered under uniform pricing. From Lemma 1, we know this happens when $c > \underline{v} - \phi(\underline{v})$. Personalized pricing now expands total demand and strictly increases total welfare. The reason is that under uniform pricing a consumer buys if her best match is above the uniform price $p > c$, whereas under personalized pricing she buys if her best match is above c .

¹⁸This generalization of Thisse-Vives is also implied by Proposition 6 in Anderson, Baik, and Larson (2023) which does comparative statics in their advertising cost parameter. (Their CEPR working paper version has a more explicit discussion of this point.) Our proof is similar to that of their Proposition 7 which shows the opposite result when $1 - H$ is log-convex.

Before investigating the impact of personalized pricing on firms and consumers, we offer an alternative formula to calculate industry profit under personalized pricing:

$$\Pi_D = n \int_c^{\bar{v}} \int_c^v G(x|v) dx dF(v) , \quad (15)$$

which is more convenient to use in some of the subsequent analysis.¹⁹ In the IID case, by integration by parts and using $G(x|v) = F(x)^{n-1}$, it simplifies to

$$\Pi_D \stackrel{\text{IID}}{=} \int_c^{\bar{v}} \frac{1 - F(v)}{f(v)} dF(v)^n . \quad (16)$$

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

$$\Pi_D = n \int_0^\infty [1 - H_c(x)] dx \leq n \frac{[1 - H_c(0)]^2}{h_c(0)} , \quad (17)$$

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 (17) binds and so we have $\Pi_D > \Pi_U$ whenever the market is not fully covered. That is indeed what we show next.

3.2.1 The case of the exponential distribution

It is illuminating to first consider the case of the exponential distribution. Suppose the v_i 's are independent and exponentially distributed with $F(v) = 1 - e^{-(v-\underline{v})}$ on $[\underline{v}, \infty)$. Then $\phi(p)$ defined in (9) equals 1,²⁰ and so the equilibrium price is $p = 1 + c$ regardless of

¹⁹To understand this alternative formula, notice that conditional on firm i winning a consumer and its product being valued at v_i , its expected profit margin is

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

where we have used (12) and integration by parts. Then industry profit under personalized pricing is $\Pi_D = n \int_c^{\bar{v}} m(v_i) G(v_i|v_i) dF(v_i)$, which is equal to (15).

²⁰Using integration by parts, the denominator in (9) 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 (9).

whether or not the market is fully covered (and irrespective of the number of firms). This means that under uniform pricing a fraction $F(1+c)^n$ of consumers are excluded from the market—and so industry profit is $\Pi_U = 1 - F(1+c)^n$. Meanwhile under personalized pricing, using (16) and the fact that $1 - F(v) = f(v)$ in this exponential example, we immediately have $\Pi_D = 1 - F(c)^n$. We then have the following observation:

Proposition 2. *Suppose valuations are IID exponential. Relative to uniform pricing, personalized pricing has no impact on firms or consumers if the market is fully covered under uniform pricing, but it benefits firms and harms consumers whenever the market is not fully covered.*

Proof. With full market coverage (which requires $1+c \leq \underline{v}$), it is immediate that $\Pi_D = \Pi_U = 1$, and since welfare is the same in both regimes, so must be consumer surplus.²¹

With partial coverage (meaning that $1+c > \underline{v}$), the profit result is also immediate, because $\Pi_D = 1 - F(c)^n > \Pi_U = 1 - F(1+c)^n$. To prove the consumer surplus result, note that

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

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

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

i.e., if the increase in profit under personalized pricing exceeds the welfare improvement. This condition must be true as $v-c < 1$ for $v \in (c, 1+c)$. \square

With partial market coverage, personalized pricing increases welfare by expanding demand, but it boosts profit so much that in aggregate consumers *suffer* from it. One way to see the intuition for this is as follows. Notice that under personalized pricing total demand is $1 - F(c)^n$, because a consumer makes a purchase unless all her valuations are below c . Since $\Pi_D = 1 - F(c)^n$ as well, it follows that under personalized pricing the *average* price that consumers pay is $1+c$, which is the same as the equilibrium 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. An alternative explanation will be provided after we examine the case with a more general distribution below.

²¹One can show that $1 - H(x)$ is log-linear in this exponential example, which explains why it is an edge case of Proposition 1.

3.2.2 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.

Proposition 3. *If $f(\bar{v}) > 0$ or in the IID case with a log-concave $f(v)$, there exists a \hat{c} such that when $c > \hat{c}$, personalized pricing benefits firms and harms consumers. More precisely,*

$$2 \leq \lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} \leq e \quad \text{and} \quad \lim_{c \rightarrow \bar{v}} \frac{V_D}{V_U} = 0 .$$

Notice that, in the limit case with large c , profit and consumer surplus in both regimes tend to zero. Meanwhile, Proposition 3 implies that for this limit case (i) profit in the discriminatory regime is at least twice as large as under uniform pricing, and (ii) consumer surplus tends to zero faster with discriminatory pricing.²² The intuition, as we will explain in detail later on, is that with a large c each firm acts almost like a local monopolist.

The contrast between Propositions 1 and 3 suggests a possible cutoff result with thresholds c_Π and c_V , such that personalized pricing benefits firms if and only if $c > c_\Pi$ and harms consumers if and only if $c > c_V$ (where $c_V > c_\Pi$ because personalized pricing raises total welfare). However, it seems difficult to prove such a cutoff result in general. In the following we first report an analytical result for a particular class of distributions when $n = 2$. We then report numerical simulations which suggest the cut-off result is true more generally. Finally, we provide an intuition for how the impact of personalized pricing varies with c .

Proposition 4. *Consider the IID duopoly case where valuations follow a generalized Pareto distribution $F(v) = 1 - [1 - a(v - \underline{v})]^{1/a}$ on $[\underline{v}, \underline{v} + \frac{1}{a}]$ with $a \in (0, 1]$. The impact of personalized pricing takes the above cut-off format when c varies.*

Notice that the generalized Pareto distribution covers the uniform and the exponential distribution as two polar cases. Such a cut-off result holds more generally as suggested by Figure 3, which plots the impact of personalized pricing on profit ($\Pi_D - \Pi_U$) and consumer surplus ($V_D - V_U$) for four common distributions in the IID duopoly case and for different values of c . (In the Online Appendix we report a similar cut-off result for $n > 2$ and other distributions.)

²²As detailed in the proof, the precise value of $\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U}$ depends on the tail behavior of $f(v)$.

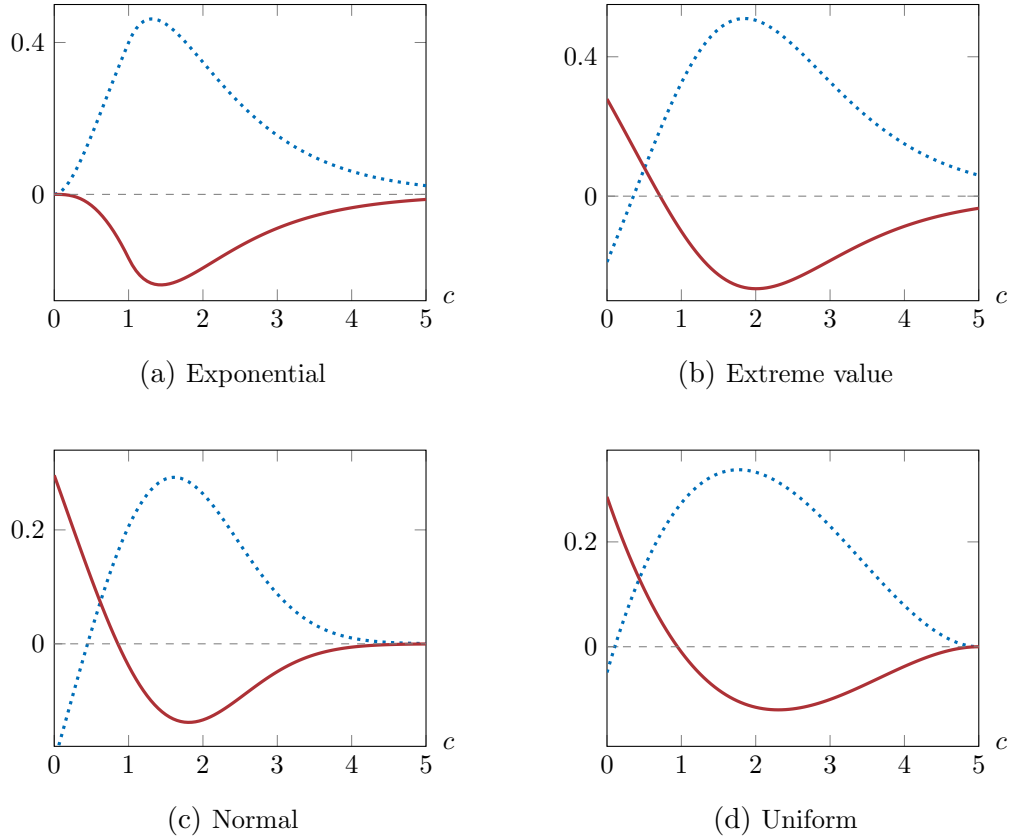


Figure 3: The impact of personalized pricing when $n = 2$, for different values of c (The dotted and solid lines represent, respectively, the change in industry profit and consumer surplus.)

Figure 3a considers the exponential case with $F(v) = 1 - e^{-(v-1)}$ on $[1, \infty)$. At $c = 0$, we have $p = 1$ and 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 pricing benefits firms and harms consumers. Figures 3b, 3c and 3d consider, respectively, the Extreme value distribution with $F(v) = e^{-e^{-(v-2)}}$, the Normal distribution with mean 2 and variance 1, and the uniform distribution with support $[0, 5]$. In each case, for low values of c (when coverage is high) personalized pricing harms firms and benefits consumers 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. These examples also show that we do not need a very low degree of market coverage to reverse the impact of personalized pricing. For instance, in the Extreme value example, personalized pricing benefits firms even if only around 11% of consumers are excluded under uniform pricing, and harms consumers if around 18% of consumers are excluded.²³

²³For the Normal example these thresholds are about 15% and 24% respectively, while for the uniform example they are about 18% and 29%.

We now explain how the impact of personalized pricing depends on the production cost c . We first consider the case where $c < \underline{v} < p$, such that the market is fully covered under personalized pricing, but only partially covered under uniform pricing. In this case, $\Pi_D - \Pi_U$ is strictly increasing while $V_D - V_U$ is strictly decreasing in c . That is, as c increases, personalized pricing becomes relatively more favorable for firms but less favorable for consumers. This is driven by the fact from equations (6) and (12) that under personalized pricing firms fully pass cost increases through to consumers, but that under uniform pricing firms share some of the burden and so $p'(c) \leq 1$.²⁴

This result can also be understood with the help of Figure 4a. (Compared to Figure 2, there is a new “expansion” region of consumers who are excluded under uniform pricing but who buy under personalized pricing.) An increase in c has two effects. First, the two dashed lines in Figure 4a with $|v_1 - v_2| + c = p$ move inwards, due to the different pass-through rates in the two pricing regimes. Hence more consumers are in the “pay more” regions and fewer are in the “pay less” region. Moreover, those who already lost from personalized pricing now lose more, while those who still gain from it now gain less. This harms consumers but benefits firms. (In the exponential example, the pass-through rate under uniform pricing is 1, so this first effect is absent.) Second, as c increases, the “expansion” region in Figure 4a grows, because additional consumers stop buying under uniform pricing and so switch from the “pay less” region to the “expansion” region. This clearly makes personalized pricing relatively more favorable to firms. For consumers, those who were already in the “expansion” region now benefit less due to the increased personalized prices; more surprisingly, those who switch into the “expansion” region also benefit less, because they have low product valuations and so the lost surplus from not buying under uniform pricing is outweighed by the lost surplus due to the price increase under personalized pricing.²⁵

It is worth mentioning that, in the case in Figure 4a, it is already possible for personalized pricing to benefit firms and harm consumers (even though it expands demand). Given that $\Pi_D - \Pi_U$ increases and $V_D - V_U$ decreases in c , this happens whenever $\Pi_D - \Pi_U$

²⁴In more detail, for the profit result, notice that Π_D is independent of c due to the full pass through of cost increases under personalized pricing, while Π_U decreases in c since both the uniform-pricing markup $p - c$ and the equilibrium demand decrease in c . For the consumer surplus result, notice that V_D decreases one-for-one in c when $c < \underline{v}$, while V_U decreases less than one-for-one in c since its derivative with respect to c is $-[1 - F_{(n)}(p)]p'(c) \in (-1, 0)$.

²⁵To see this more formally, suppose cost increases from c' to c'' , such that the uniform price increases from p' to p'' . Consider a consumer with $v_1 > v_2$ and also $p' < v_1 < p''$. The cost increase causes her to switch from the “pay less” region to the “expansion” region. Her loss from no longer buying under uniform pricing is $v_1 - p'$, while her loss under personalized pricing due to the price increase is $c'' - c'$. The former is smaller since $v_1 - p' < p'' - p' \leq c'' - c'$, where the first inequality is because the consumer has a relatively low valuation $v_1 < p''$ and the second is because $p'(c) \leq 1$.

is not too negative and $V_D - V_U$ is not too positive under full coverage. In particular, it happens in the exponential example where $\Pi_D - \Pi_U = V_D - V_U = 0$ with full coverage.

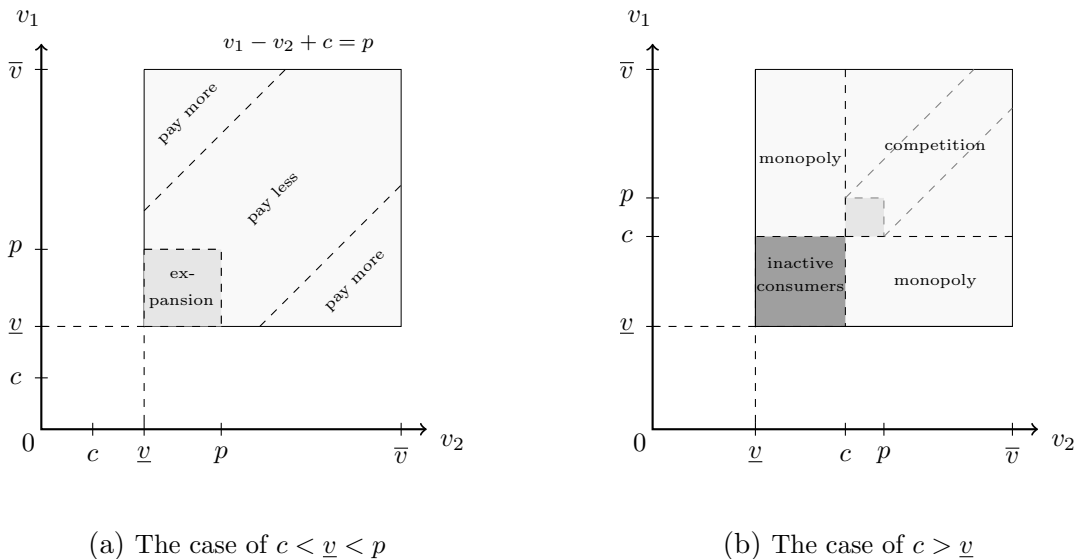


Figure 4: The impact of personalized pricing with partial market coverage

We now turn to the case when the production cost is higher, such that $c > \underline{v}$ as depicted in Figure 4b. In this case some consumers value only one firm’s product above cost, as indicated by the “monopoly” regions. Under personalized pricing, the firm is a monopolist over these consumers, and so it extracts all their surplus. This is an additional force for personalized pricing to benefit firms and harm consumers. Consumers in the “competition” region value both products above cost, and so for them the situation is the same as in Figure 4a. The size of the new monopoly force is, however, non-monotonic in c (e.g., it is zero at $c = \underline{v}$ or as c approaches \bar{v}). This makes it hard to compare it with the effect coming from the “competition” region, and hence prove the cut-off result for a general distribution.²⁶ Nevertheless, as c approaches \bar{v} , both the “monopoly” and “competition” regions shrink, but the latter is of second order relative to the former, so the impact of personalized pricing is qualitatively the same as under monopoly, as proved in Proposition 3. (Intuitively, in this case, conditional on a consumer valuing one product

²⁶More precisely, to understand the potential complication in combining the monopoly and competition effects, note that for a consumer with, say, $v_1 > v_2$, the impact of personalized pricing is $\max\{0, v_2 - c\} - \max\{0, v_1 - p\}$. For a consumer in the competition region, the first term is $v_2 - c$, in which case the impact weakly *decreases* in c given $p'(c) \leq 1$; for a consumer in the monopoly region, the first term is 0, in which case the impact is zero if $v_1 < p$ and otherwise *increases* in c given $p'(c) > 0$. Therefore, how the total impact varies with c now depends on which of the two regions is more important; moreover, as c changes, so do the boundaries of the two regions, and hence also their relative importance. A similar trade-off applies to the impact on firm profit.

above cost, it is very unlikely that she values another product above cost. Hence each firm is essentially a monopolist competing only against the outside option.) However, as noted already, personalized pricing can benefit firms and harm consumers even when there are no monopoly regions (as in Figure 4a).

3.2.3 The number of firms and market coverage

Another parameter which influences market coverage is the number of firms. When $n = 1$ we have the standard monopoly case; when n is large the best match should be relatively high, and so intuitively the impact of personalized pricing should be similar to the full-coverage case. (More rigorously, since the profit in both regimes often goes to zero as $n \rightarrow \infty$, it also matters how fast they converge to the full-coverage outcome as $n \rightarrow \infty$.) In the following, we first provide an analytical result regarding the case of large n , and then demonstrate by numerical examples that with partial coverage the impact of competitive personalized pricing can remain qualitatively similar as in the monopoly case for a relatively large range of n .

We deal with the case where n is large by approximating the equilibrium outcome. However, the approximation of the uniform equilibrium price when n is large is technically difficult. We rely on the approximation results for the IID and full-coverage case developed in Gabaix et al. (2016), and extend them to the case with partial coverage.

Proposition 5. *Consider the IID case with a log-concave $f(v)$, and let*

$$\gamma = \lim_{v \rightarrow \bar{v}} \frac{d}{dv} \left(\frac{1 - F(v)}{f(v)} \right) \quad (18)$$

denote the tail index of the valuation distribution of each product. If $\gamma \in (-1, 0)$, there exists \hat{n} such that when $n > \hat{n}$ personalized pricing harms firms and benefits consumers.

When f is log-concave, we must have $\gamma \in [-1, 0]$.²⁷ Unfortunately, our approximation in the proof is not precise enough for a meaningful comparison if $\gamma = -1$ or 0 . This rules out many common distributions such as the uniform, exponential, extreme value, and normal (see Table 1 in Gabaix et al., 2016). However, the numerical examples below demonstrate that our comparison results when n is large continue to hold in those examples. (See the Online Appendix for more simulations with other distributions.)

²⁷When f is log-concave, so is $1 - F$. Then $\frac{1-F}{f}$ is decreasing, 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 $\gamma = -1$.

Given that the impact of personalized pricing in Proposition 5 is the opposite of that under monopoly, this 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 5 below (the IID case with $c = 2$ and the same distributions used in Figure 3). Figure 5a is for the exponential distribution, and confirms our earlier analytic result that personalized pricing always benefits firms and harms consumers when the market is not fully covered. Figure 5b shows that for the Extreme value distribution, personalized pricing benefits firms if and only if $n < 10$, and harms consumers if and only if $n < 7$. A qualitatively similar pattern emerges in Figure 5c for the Normal distribution. Figure 5d shows that for the uniform distribution personalized pricing benefits firms for $n < 4$ and harms consumers for $n < 3$. Although the impact of personalized pricing can be non-monotonic in n , it goes to zero as n becomes large.

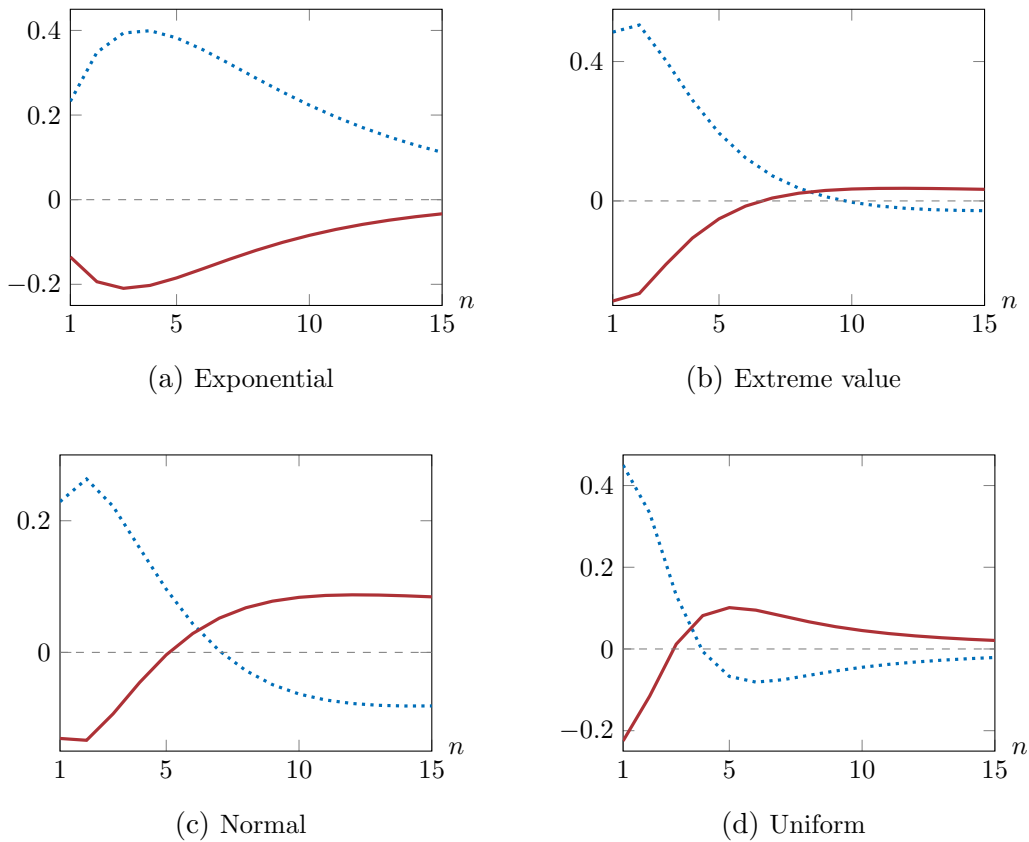


Figure 5: The impact of personalized pricing when $c = 2$, for different values of n (The dotted and solid lines represent, respectively, the change in industry profit and consumer surplus.)

3.3 Discussions

Valuation correlation and dispersion. We have examined how the impact of personalized pricing varies with c and n , but another important primitive of our model is the distribution of product valuations. It would be interesting to examine how a change in the valuation distribution which reflects, say, a change in product differentiation, shapes the impact of personalized pricing. While a general treatment of this issue is challenging, the Online Appendix investigates it using a few examples. Intuitively, at a high level, as product valuations become more correlated or less dispersed, product differentiation falls, and this tends to lessen the impact of personalized pricing because price competition under both uniform pricing and personalized pricing becomes fiercer.²⁸

Alternative information structures. We have assumed that under personalized pricing each firm observes a consumer’s valuations for all the n products. A natural alternative information structure is that each firm only observes a consumer’s valuation for its own product. This case of “partial discrimination” resembles a first-price auction,²⁹ while the case studied so far resembles a second-price auction. Therefore, if valuations are IID across products, the well-known revenue equivalence theorem from auction theory implies that these two information structures lead to the same market outcome, and consequently the impact of personalized pricing under the alternative information structure remains unchanged. (We note, however, that uniform pricing has no counterpart in the auctions literature.) In the Online Appendix we formally investigate this alternative information structure, including beyond the IID case.³⁰

4 Extensions

We now discuss how our model can be extended to a situation where only some firms have access to data and can do personalized pricing, or where the number of firms is endogenous and determined as part of a free-entry equilibrium. We focus on a few main results and their intuition; further details and proofs are in the Online Appendix.

²⁸This is broadly speaking what we find in the examples in the Online Appendix. Nevertheless, we note that the impact of product differentiation can also be more subtle than this, because it affects not only equilibrium prices, but also market coverage given fixed prices.

²⁹This is easy to see once we interpret a firm as “bidding” surplus of $v - p(v)$ to a consumer who values its product at v . This is also analogous to price competition among firms that sell homogeneous goods but have privately-known production costs as studied in, for example, Spulber (1995).

³⁰Of course it would also be interesting to consider other information structures and investigate how the impact of price discrimination changes with the amount of information that firms have access to. See the Online Appendix for a further discussion of this issue.

4.1 Asymmetrically informed firms

So far we have analyzed the cases where all firms do uniform pricing or all firms do personalized pricing. However, in certain markets, some firms have access to more data and better technology—and hence are more able to do personalized pricing—compared to other firms. For example, Amazon possesses lots of information about customer shopping behavior, and in principle can use this 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. We now discuss such a “mixed” case where only some firms can price discriminate.

Suppose that k firms have consumer data and can price discriminate, while the other $n - k$ firms do not have such data and thus have to offer a uniform price. When $k = 0$ all firms do uniform pricing as in Section 2.1, and when $k = n$ all firms do personalized pricing as in Section 2.2. When $0 < k < n$ (i.e., the “mixed” case) a subtle technical issue arises: if all firms set prices simultaneously, as explained in the Online Appendix, there is no pure-strategy pricing equilibrium,³¹ and the mixed-strategy equilibrium is rather complicated to characterize. To avoid this problem, we assume that the $n - k$ firms with no data simultaneously set their uniform prices first, and after seeing those prices the other k firms use their data to simultaneously offer personalized prices. This timing captures the idea that firms with lots of data often also have better pricing technology and so can adjust prices more frequently.

We can compare the three regimes analytically when the production cost c is large. Specifically, if $f(\bar{v}) > 0$, there exists a \hat{c} such that for $c > \hat{c}$ the mixed regime is ranked in the middle for industry profit, consumer surplus, and total welfare. Intuitively, recall that when c is large, each firm approximately acts like a monopolist. Hence, as more firms can do personalized pricing, profit and welfare increase but consumers are made worse off.

Analytical results are also available for a general c when valuations are IID exponential. If c is relatively high, the mixed case is ranked in the middle for each welfare measure; however, if c is relatively low, the mixed case is (weakly) the best for industry profit (tied with discriminatory pricing) but the worst for both total welfare and consumer surplus. Numerical simulations suggest that qualitatively similar patterns emerge when product valuations are drawn from other distributions (except that for low values of c the mixed regime is now strictly best for industry profit). Figure 6 illustrates this for consumer surplus in the duopoly case when valuations are Extreme value or Normal. (Corresponding plots for total welfare and industry profit are in the Online Appendix.)

³¹This issue is well known in the literature, and the usual approach to avoid it is to consider a sequential pricing game as we do here. See, e.g., the duopoly analysis of Thisse and Vives (1988) when only one firm can price discriminate.

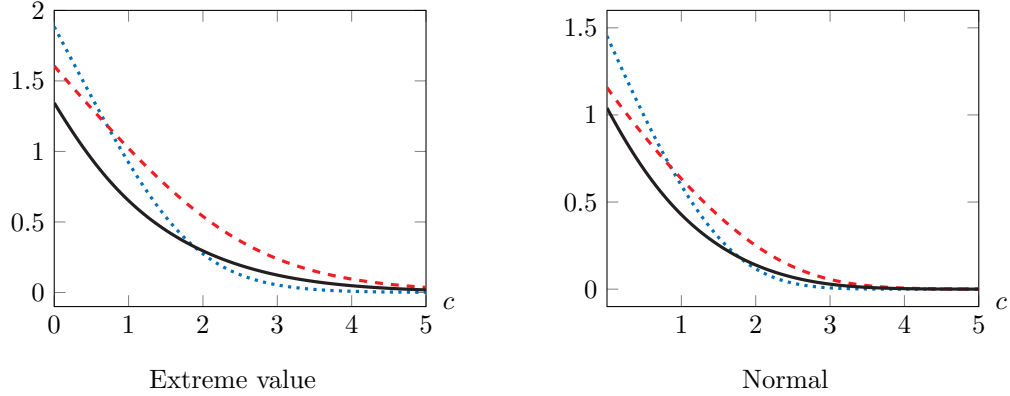


Figure 6: Consumer surplus in the asymmetric case *vs* symmetric cases when $n = 2$, for different values of c

(The solid, dashed, and dotted curves are respectively the mixed, uniform, and discriminatory cases.)

Notice that at low c the mixed regime is worst for consumers. Intuitively, when c is low, market coverage is high irrespective of how many firms can personalize. However, in the mixed case, firms that are able to personalize can “poach” some consumers for whom they are not the consumer’s favorite product via a low personalized price; this harms match efficiency and reduces total welfare compared to the other two regimes. At the same time, the asymmetry in information and timing allows the firms to better segment the market, increasing industry profit. Consequently, consumer surplus ends up being lowest in the mixed regime.

Overall, our results suggest that when coverage is relatively high (i.e., c is relatively low), policies which force large firms to share their data, or which prevent those firms from personalizing their prices, can benefit consumers and overall welfare.

4.2 Free entry and endogenous market structure

Since the ability to do personalized pricing affects firm profits, it may also influence firms’ entry decisions and thus the market structure. We now examine the impact of personalized pricing with an endogenous market structure. To do this, we use a standard free-entry game, where firms first decide whether or not to pay a fixed cost to enter, and then after entering compete in prices. (As often assumed in the literature, we consider a large number of potential entrants, and use a sequential entry game to avoid coordination problems.)

Throughout this section, we assume that entry of a new firm does not affect consumers’ valuations for existing products. (Formally, as usual let v_i denote the valuation for product i when there are n firms in the market, and let \hat{v}_i denote the same when there are $n - 1$

firms in the market. We assume that the distribution of $(\hat{v}_1, \dots, \hat{v}_{n-1})$ is the marginal distribution of (v_1, \dots, v_{n-1}) . This assumption is trivially satisfied in the IID case, and can also hold with correlated valuations (e.g., if we fix a grand set of firms with a certain joint valuation distribution).³²

With the above assumption, we can show that the free-entry equilibrium under personalized pricing must be socially optimal. The intuition is simple: Suppose $n - 1$ firms are already in the market, and consider the entry of an n^{th} firm. Amongst consumers with $v_n \leq \max_{j \leq n-1} \{c, v_j\}$, 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 under personalized pricing it fully extracts this incremental surplus via Bertrand competition. As a result, the incentives of the social planner and this new firm are perfectly aligned.³³ (Using the terminology introduced by Mankiw and Whinston, 1986, here the business-stealing effect and the product-diversity effect exactly cancel each other out.)

On the other hand, the free-entry equilibrium under uniform pricing is generically not socially efficient.³⁴ We then deduce that if there is no integer constraint, in a free-entry market consumer surplus is higher under personalized pricing than under uniform pricing. This is simply because absent integer constraints, free entry drives industry profit to zero under both our pricing regimes (after accounting for fixed entry costs). Therefore, since total welfare is maximized under personalized pricing, so is consumer surplus.

However, we emphasize that the result that personalized pricing outperforms uniform pricing for consumers may not hold anymore when the number of firms in the industry is restricted to be an integer. The reason is that although personalized pricing maximizes total welfare, with integer constraints it may also lead to higher industry profit, in which case consumer surplus can be lower compared to under uniform pricing. For instance,

³²This assumption fails if entry of a new firm induces existing firms to reposition their products. This happens, for example, in the Salop circle model. (Contrary to our observation below, Stole, 2007 shows that entry is socially excessive in the Salop model under perfect price discrimination.) However this assumption can hold in other spatial models, such as Chen and Riordan (2007) where entry of a new firm does not cause existing ones to reposition.

³³This result is in the same spirit as Spence (1976). He shows that in a competitive market with perfect discrimination, each firm's choice of quantity or product characteristic is socially optimal because its profit is equal to its marginal contribution to total surplus. See also Bhaskar and To (2004) for a similar observation in a game where firms enter, then choose product characteristics, then set prices. (We thank John Vickers for drawing our attention to these papers.)

³⁴Let n^* denote the socially optimal number of firms in the free-entry market. Entry under uniform pricing is then excessive if $\Pi_U > \Pi_D$ at $n = n^*$, but insufficient if $\Pi_U < \Pi_D$ at $n = n^*$. In particular, if the market is fully covered, then entry must be excessive under uniform pricing. (This was previously shown in the IID case by Anderson, de Palma, and Nesterov, 1995 using a different approach to ours.)

it is possible to construct examples where there is a monopoly (i.e., exactly one firm enters) in both regimes, such that consumers are better off with uniform pricing. (See a more detailed discussion of this issue in the Online Appendix. There we also provide examples where considering the integer constraint can instead strengthen the advantage of personalized pricing for consumers. In addition, we discuss the possibility that a natural monopoly can arise due to economies of scale in data collecting.)

5 Conclusion

We have investigated the impact of personalized pricing—a form of price discrimination which is becoming increasingly relevant in the digital economy—in a general oligopoly model. We find that competitive personalized pricing harms firms and benefits consumers under a standard log-concavity condition if the market is fully covered. However, the impact can be reversed in the (arguably more realistic) case without full market coverage; in particular, when the production cost is sufficiently high or the outside option is sufficiently good, personalized pricing raises industry profit and decreases consumer surplus. When some firms can use consumer data to price discriminate while others cannot, this asymmetric 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 benefit consumers.

Two important issues remain unaddressed in this paper. First, do firms have incentives to adopt personalized pricing? This of course depends on how costly it is to acquire consumer data and develop the technology needed for personalized pricing, as well as the extent to which firms can commit to a pricing strategy before making offers. If it is costless to implement personalized pricing, and if firms simultaneously choose whether to do personalized pricing and what prices to offer, then the only equilibrium is that all firms do personalized pricing because a discriminatory pricing schedule includes uniform pricing as a special case. Second, do consumers have incentives to allow their data to be collected and then used for personalized pricing? For example, 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. Following Ali, Lewis, and Vasserman (2023) and Anderson, Baik, and Larson (2023), we are investigating this question in some ongoing work.

Finally, throughout the paper we have focused on a standard IO setting. However our model also applies to other markets—such as competing employers that offer personalized wages according to workers’ preferences for different job positions, or competing schools that offer personalized scholarships according to students’ preferences and family incomes.

Appendix: Omitted Proofs for Sections 2 and 3

Proof of Lemma 1. We prove the existence and uniqueness of the equilibrium uniform price. (The rest of the lemma follows from arguments in the text.) Clearly $p - c < \phi(p)$ at $p = c$. Since $\phi(p)$ is non-increasing due to Assumption 1, it suffices to show that $p - c > \phi(p)$ at $p = \bar{v}$. This 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 must be less than $\bar{v} - c$ when p is close to \bar{v} . \square

Proof of Lemma 3. Using equation (12) the highest personalized price is $p_{\max} = c + \bar{v} - \max\{c, \underline{v}\}$. If $\underline{v} \leq c$, then $p_{\max} = \bar{v} > p$. If $\underline{v} > c$, then $p_{\max} = c + \bar{v} - \underline{v}$, and so $p < p_{\max}$ if and only if $p - c < \bar{v} - \underline{v}$. Under Assumption 1, $\phi(p)$ is constant for $p \leq \underline{v}$ and non-increasing for $p > \underline{v}$, and so the uniform price satisfies $p - c = \phi(p) \leq \phi(\underline{v}) = \frac{1}{nh(0)}$. Meanwhile,

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

where the equality is from the fact that $\Pr[v_i \geq \max_{j \neq i} \{v_j\}] = \frac{1}{n}$ due to firm symmetry, 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

Proof of Proposition 3. We first state a few lemmas (whose proofs can be found in the Online Appendix):

Lemma 4. *Under uniform pricing, the equilibrium pass-through rate $p'(c)$ at $c \rightarrow \bar{v}$ is*

- (i) $p'(\bar{v}) = \frac{1}{2}$ if $f(\bar{v}) > 0$;
- (ii) $p'(\bar{v}) = \frac{2}{3}$ if $f(\bar{v}) = 0$ and $f'(\bar{v}) < 0$;
- (iii) $p'(\bar{v}) \in [\frac{1}{2}, 1]$ if $f(\bar{v}) = 0$, $f'(\bar{v}) = 0$, and $f(v)$ is log-concave.

Lemma 5. *If $f(\bar{v}) > 0$ or in the IID case,*

$$\lim_{p \rightarrow \bar{v}} \frac{\phi(p)f(p)}{1 - F(p)} = 1 .$$

Lemma 6. *In the IID case, if $f(v)$ is log-concave and $f(\bar{v}) = 0$, then*

$$e^{2p'(\bar{v})-1} \leq \lim_{c \rightarrow \bar{v}} \frac{f(c)}{f(p)} \leq e^{2-\frac{1}{p'(\bar{v})}} .$$

Profit result. Recall that $\Pi_U = (p - c)[1 - F_{(n)}(p)] = n(p - c) \int_p^{\bar{v}} G(v|v)dF(v)$ and $\Pi_D = n \int_c^{\bar{v}} \int_c^v G(x|v)dx dF(v)$. Both go to zero as $c \rightarrow \bar{v}$. Then L'hôpital's rule implies that

$$\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} = \lim_{c \rightarrow \bar{v}} \frac{\int_c^{\bar{v}} G(c|v)dF(v)}{[1 - p'(c)] \int_p^{\bar{v}} G(v|v)dF(v) + \phi(p)G(p|p)f(p)p'(c)}, \quad (19)$$

where we have used the equilibrium price condition $p - c = \phi(p)$. Divide both the numerator and denominator by $1 - F(p)$. Notice that

$$\lim_{c \rightarrow \bar{v}} \frac{\int_p^{\bar{v}} G(v|v)dF(v)}{1 - F(p)} = \lim_{c \rightarrow \bar{v}} \frac{G(p|p)f(p)}{f(p)} = 1.$$

This, together with Lemma 5, implies that the denominator of (19) divided by $1 - F(p)$ converges to 1. Therefore,

$$\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} = \lim_{c \rightarrow \bar{v}} \frac{\int_c^{\bar{v}} G(c|v)dF(v)}{1 - F(p)}. \quad (20)$$

Consider first the general case with $f(\bar{v}) > 0$. L'hôpital's rule implies that (20) equals

$$\lim_{c \rightarrow \bar{v}} \frac{-G(c|c)f(c) + \int_c^{\bar{v}} g(c|v)dF(v)}{-f(p)p'(c)} = \frac{1}{p'(\bar{v})} = 2,$$

where we have used $p'(\bar{v}) = \frac{1}{2}$ from Lemma 4(i). Therefore, in this case, $\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} = 2$.

When $f(\bar{v}) = 0$, we focus on the IID case. Then (20) simplifies to

$$\lim_{c \rightarrow \bar{v}} \frac{G(c)[1 - F(c)]}{1 - F(p)} = \frac{1}{p'(\bar{v})} \lim_{c \rightarrow \bar{v}} \frac{f(c)}{f(p)}.$$

If $f(\bar{v}) = 0$ and $f'(\bar{v}) < 0$, L'hôpital's rule implies that the above limit equals $\frac{1}{p'(\bar{v})^2}$. Using $p'(\bar{v}) = \frac{2}{3}$ from Lemma 4(ii), we have $\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} = \frac{9}{4} \in (2, e)$.

If $f(\bar{v}) = 0$, $f'(\bar{v}) = 0$, and f is log-concave, then Lemma 6 implies that

$$\frac{e^{2p'(\bar{v})-1}}{p'(\bar{v})} \leq \lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} \leq \frac{e^{2-\frac{1}{p'(\bar{v})}}}{p'(\bar{v})}.$$

From Lemma 4(iii), we know $p'(\bar{v}) \in [\frac{1}{2}, 1]$. One can check that in this range both the upper bound and lower bound are increasing functions of $p'(\bar{v})$. Therefore, in this case $\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} \in [2, e]$. (The above bounds result also implies that $\lim_{c \rightarrow \bar{v}} \frac{\Pi_D}{\Pi_U} = e$ if $p'(\bar{v}) = 1$. This is the case for many distributions such as exponential, extreme value, and normal.)

Consumer surplus result. To prove the consumer surplus result, we need one more lemma (whose proof is also in the Online Appendix):

Lemma 7. *Recall that $f_{(n)}(v)$ is the density of $F_{(n)}(v)$. Then $f_{(n)}(\bar{v}) = nf(\bar{v})$.*

Recall that $V_U = \int_p^{\bar{v}} [1 - F_{(n)}(v)] dv$ and $V_D = \int_c^{\bar{v}} [1 - F_{(n-1)}(v)] dv$. Both go to zero as $c \rightarrow \bar{v}$. Therefore, by L'hôpital's rule we have

$$\lim_{c \rightarrow \bar{v}} \frac{V_D}{V_U} = \frac{1}{p'(\bar{v})} \lim_{c \rightarrow \bar{v}} \frac{1 - F_{(n-1)}(c)}{1 - F_{(n)}(p)} = \frac{1}{p'(\bar{v})^2} \lim_{c \rightarrow \bar{v}} \frac{f_{(n-1)}(c)}{f_{(n)}(p)}. \quad (21)$$

Consider first the general case with $f(\bar{v}) > 0$. From the definition of $F_{(n-1)}(v)$ in (2), we obtain

$$f_{(n-1)}(v) = f_{(n)}(v) - n[G(v|v)f(v) - \int_v^{\bar{v}} g(v|v_i)dF(v_i)].$$

Therefore, $f_{(n-1)}(\bar{v}) = f_{(n)}(\bar{v}) - nf(\bar{v}) = 0$, where the second equality uses Lemma 7. On the other hand, from the definition of $F_{(n)}(v)$ in (1), $f_{(n)}(v) = G(v|v)f(v) + \int_v^{\bar{v}} g(v|v_i)dF(v_i)$, and so $f_{(n)}(\bar{v}) > 0$ given $f(\bar{v}) > 0$. Therefore, in this general case (21) equals 0.

When $f(\bar{v}) = 0$, we consider the IID case with a log-concave f . Then $f_{(n-1)}(c) = n(n-1)(1-F(c))F(c)^{n-2}f(c)$ and $f_{(n)}(p) = nF(p)^{n-1}f(p)$. Therefore,

$$\lim_{c \rightarrow \bar{v}} \frac{f_{(n-1)}(c)}{f_{(n)}(p)} = (n-1) \lim_{c \rightarrow \bar{v}} [1 - F(c)] \frac{f(c)}{f(p)},$$

and so $\lim_{c \rightarrow \bar{v}} \frac{V_D}{V_U} = 0$ if $\lim_{c \rightarrow \bar{v}} \frac{f(c)}{f(p)}$ is finite.

If $f(\bar{v}) = 0$ and $f'(\bar{v}) < 0$, then

$$\lim_{c \rightarrow \bar{v}} \frac{f(c)}{f(p)} = \lim_{c \rightarrow \bar{v}} \frac{f'(c)}{f'(p)p'(c)} = \frac{1}{p'(\bar{v})} = \frac{3}{2},$$

where the first equality is from L'hôpital's rule, the second uses $f'(\bar{v}) < 0$, and the third uses Lemma 4(ii). This is of course finite.

If $f(\bar{v}) = 0$ and $f'(\bar{v}) = 0$, Lemma 4(iii) and Lemma 6 jointly imply that $\lim_{c \rightarrow \bar{v}} \frac{f(c)}{f(p)}$ is finite. This completes the whole proof. □

Proof of Proposition 4. Please see the Online Appendix. □

Proof of Proposition 5. Please see the Online Appendix. □

References

- ALI, N., G. LEWIS, AND S. VASSERMAN (2023): “Voluntary Disclosure and Personalized Pricing,” *Review of Economic Studies*, 90(2).
- ANDERSON, S., A. BAIK, AND N. LARSON (2019): “Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts,” *CEPR Discussion Paper No. 13793*. CEPR Press, Paris & London.
- (2023): “Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts,” *Review of Economic Studies*, 90(5).
- 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 (2024): “The Pricing Strategies of Online Grocery Retailers,” *Quantitative Marketing and Economics*, 22, 1–21.
- ARMSTRONG, M. (2007): “Recent Developments in the Economics of Price Discrimination,” in *Advances in Economics and Econometrics (Theory and Applications, Ninth World Congress, Vol. II)*, 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.
- BEIS (2021): “Personalised Pricing and Disclosure,” Discussion paper, Department for Business, Energy & Industrial Strategy (United Kingdom).
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105(3), 921–957.
- BHASKAR, V., AND T. TO (2004): “Is Perfect Price Discrimination Really Efficient? An Analysis of Free Entry,” *RAND Journal of Economics*, 35(4), 762–776.
- BOUNIE, D., A. DUBUS, AND P. WAELBROECK (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., 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.
- DUBÉ, J.-P., AND S. MISRA (2023): “Personalized Pricing and Consumer Welfare,” *Journal of Political Economy*, 131(1), 131–189.
- ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2021): “Market Segmentation Through Information,” mimeo.
- EUROPEAN COMMISSION (2018): “Consumer Market Study on Online Market Segmentation Through Personalised Pricing/Offers in the European Union,” Final Report.
- 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.
- ICHIHASHI, S., AND A. SMOLIN (2022): “Data Collection by an Informed Seller,” mimeo.
- JULLIEN, B., M. REISINGER, AND P. REY (2023): “Personalized Pricing and Distribution Strategies,” *Management Science*, 69(3), 1687–1702.
- KEHOE, P., B. LARSEN, AND E. PASTORINO (2020): “Dynamic Competition in the Era of Big Data,” mimeo.
- LIU, Q., AND K. SERFES (2004): “Quality of Information and Oligopolistic Price Discrimination,” *Journal of Economics & Management Strategy*, 13(4), 671–702.

- LU, Q., AND N. MATSUSHIMA (2022): “Personalized Pricing When Consumers Can Purchase Multiple Items,” mimeo.
- MANKIW, G., AND M. WHINSTON (1986): “Free Entry and Social Inefficiency,” *RAND Journal of Economics*, 17(1), 48–58.
- 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.
- OFCOM (2020): “Personalised Pricing for Communications: Making Data Work for Consumers,” Discussion paper, Office of Communications (United Kingdom).
- PERLOFF, J., AND S. SALOP (1985): “Equilibrium with Product Differentiation,” *Review of Economic Studies*, 52(1), 107–120.
- PIGOU, A. (1920): *The Economics of Welfare*. Macmillan.
- QUINT, D. (2014): “Imperfect Competition with Complements and Substitutes,” *Journal of Economic Theory*, 152, 266–290.
- 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.
- SPENCE, M. (1976): “Product Selection, Fixed Costs, and Monopolistic Competition,” *Review of Economic Studies*, 43(2), 217–235.
- SPULBER, D. (1995): “Bertrand Competition When Rivals’ Costs Are Unknown,” *Journal of Industrial Economics*, 43(1), 1–11.
- STOLE, L. (2007): “Price Discrimination and Competition,” in *Handbook of Industrial Organization (Vol. 3)*, ed. by M. Armstrong, and R. Porter. North Holland.

- THISSE, J.-F., 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.
- VARIAN, H. (2018): “Artificial Intelligence, Economics, and Industrial Organization,” NBER working paper 24389.
- WALDFOGEL, J. (2015): “First Degree Price Discrimination Goes to School,” *Journal of Industrial Economics*, 63(4), 569–597.
- WHICH? (2018): “Control, Alt or Delete? The Future of Consumer Data,” Policy Report.
- ZHOU, J. (2017): “Competitive Bundling,” *Econometrica*, 85(1), 145–172.