Consumer Control and Privacy Policies[*]

By S. Nageeb Ali, Greg Lewis, and Shoshana Vasserman[†]

Data collection and personalization are ubiquitous today, raising concerns of privacy and price discrimination. In response to these concerns, both regulatory authorities and the market emphasize the value of *consumer control*. In the public sphere, the European Union's GDPR and California's Consumer Privacy Act require websites to obtain consent before collecting browsing information and restrict the duration that consumer data can be retained. In the private sphere, Apple, Google, and other firms have rolled out product features that allow consumers to opt out of personalized tracking.

Although consumer control features prominently in discussions of privacy policies, relatively little is known about how to model consumer control and its effects on market outcomes. If consumers control the information possessed by sellers, is price discrimination beneficial? Or should it nevertheless be prohibited? Does it suffice for the consumer to be able to opt out from sharing information or is control at a more granular level needed? How does this relate to market competitiveness? Clearly, a strategic framework is necessary to assess the implications of privacy policies.

Our work (Ali, Lewis and Vasserman, 2022) offers such a framework. We view consumer control through the lens of *voluntary disclosure*. The consumer has certain verifiable characteristics—her age, income, or her data—that are correlated with her preferences. Rather than sellers having this information at the outset, the consumer chooses what to disclose to the market. From this perspective, opting in to a firm's tracking policy is tantamount to disclosing data pre-dictive of her preferences; opting out, by contrast, corresponds to choosing not to disclose information. These are but two choices; one may envision contexts that endow the consumer with control at a more granular level—e.g., the ability to disclose a student ID or a senior citizen card—without having to opt in entirely. In equilibrium, firms do not take disclosed information at face value; instead, they draw inferences both from what is said and what is left unsaid.

We use this framework to answer the questions above. In monopolistic markets, consumers do not benefit from personalized pricing if the only choice they have is to opt in or out from sharing information; more fine-tuned control is necessary for them to benefit from personalized pricing. By contrast, if the market is competitive, control even in the form of simple opt-in / opt-out policies are enough to assure consumer gains. Disclosure amplifies competition. Contrary to the view that firms should not price discriminate, our findings suggest that consumers may benefit from price discrimination if they control the flow of information.

We describe these findings in greater below and also pose new questions for which our approach may be useful. Section I studies the role of consumer control in a monopolistic market and Section II considers that for oligopoly with differentiated products. Section III shows how this framework offers a simple and direct resolution to the privacy paradox, namely the tendency for people to cede their privacy cheaply while stating that they value their privacy. Section IV revisits the limits of price discrimination through the lens of consumer control.

I. Disclosure to a Monopolist

We begin with the case of a monopolistic market. The monopolist faces a unit mass of consumers, each with unit demand and valuation v drawn from $V := [\underline{v}, \overline{v}]$ according to

a well-behaved CDF F, where $\underline{v} \geq 0$. The production cost is normalized to 0 and payoffs are quasilinear: if a consumer with value v purchases at price p, her payoff is $v - p$ and the monopolist's profit is p; otherwise, both parties obtain zero payoffs.

In the standard pricing problem, the monopolist chooses price p to maximize $p(1 - F(p))$. Denote the (lowest) maximizer by $\bar{p}$ and suppose $\bar{p} > \underline{v}$. The consumer's payoff in this *uniform-pricing* benchmark is $\max\{v - \bar{p}, 0\}$.

We append a disclosure game to this strategic interaction. Upon observing her value v, the consumer chooses a message M from the set of messages $\mathcal{M}(v)$ where $\mathcal{M} : V \rightrightarrows V$. The monopolist then chooses a price and the consumer chooses whether to purchase the good. We study Perfect Bayesian Equilibria of this game.

That the set of messages, $\mathcal{M}(v)$, varies with the consumer's type reflects the idea that the information that is shared is "hard" or verifiable, and not cheap talk. In other words, it is information—such as consumer characteristics or data—that can be shared by some types of consumers but not others.

A. *The Futility of All-or-Nothing Disclosure*

Suppose the consumer can choose to opt in or opt out from being tracked: if she opts in, the monopolist learns her value, and if she opts out, the monopolist learns nothing. Formally, the consumer with value v can send either the message $\{v\}$, which is fully revealing, or the completely uninformative message V. Disclosure is all or nothing.

This setting has multiple equilibria. First, there is a *fully revealing equilibrium* in which every consumer type opts in and the monopolist charges a personalized price of $p = v$; off-path, if the consumer opts out, the monopolist assumes with probability 1 that the consumer has the highest possible value $\bar{v}$ and charges that price. The setting also has a *fully concealing equilibrium* in which all consumer types opt out by sending message V and the firm charges $\bar{p}$. There are also a continuum of other equilibria, varying in the extent of revelation.

Nevertheless, across all equilibria, the

combination of simple evidence and personalized pricing does not benefit consumers relative to uniform-pricing.

PROPOSITION 1: *Across all equilibria, the payoff of any consumer type v is bounded above by* $\max\{v - \bar{p}, 0\}$.

The logic of Proposition 1 is that any equilibrium involves a price, $\tilde{p}$, that is charged by the monopolist when the consumer chooses to opt out. In equilibrium, if a consumer's valuation v exceeds $\tilde{p}$, she must opt out; opting in results in the monopolist fully extracting her surplus while opting out results in her obtaining $v - \tilde{p}$. Therefore, at a price of $\tilde{p}$, the monopolist is selling at a uniform price to all types whose valuations exceed $\tilde{p}$. Hence, $\tilde{p}$ can be no lower than the (lowest) optimal uniform price, $\bar{p}$.

Proposition 1 implies that in the case of monopoly, prohibiting personalized pricing more effectively safeguards consumers than merely offering them the ability to opt out.

B. *Partial Disclosure Enables Group Pricing*

The conclusion above is overturned once consumers have access to richer forms of control. Suppose the consumer can disclose partial information without revealing her value. This form of consumer control benefits some consumers without hurting others.

To proceed formally, for a subset of consumer types $\tilde{V} \subseteq V$, let $p^*(\tilde{V})$ be the (lowest) optimal price when the monopolist knows the consumer's value is in $\tilde{V}$.[1] Also for a message M, let $\mathcal{M}^{-1}(M) = \{v : M \in \mathcal{M}(v)\}$ be the set of consumer types that can send message M.

DEFINITION 1: *An evidence technology enables group pricing if there exists a message M such that:*

(a) *the set $\mathcal{M}^{-1}(M)$ has positive measure and $p^*(\mathcal{M}^{-1}(M)) < \bar{p}$.*

(b) *$p^*(V \backslash \mathcal{M}^{-1}(M)) \leq \bar{p}$.*

Definition 1 stipulates that there is some message M such that the monopolist lowers prices if he learns that M is feasible for the

---

[1] In other words, $p^*(\tilde{V})$ maximizes $p(1 - F(p|v \in \tilde{V}))$.

consumer to send but does not raise prices when he infers otherwise. A special case is when consumers whose values are lower than the uniform price $\bar{p}$ can send evidence that distinguishes them from those whose values exceed $\bar{p}$. Examples of this kind of evidence are identification for senior citizens and students or Electronic Benefits Transfer cards; being able to produce such evidence is correlated with the consumer having a lower willingness to pay.

PROPOSITION 2: *If an evidence technology enables group pricing, there is a Pareto-improving equilibrium.*

The logic of Proposition 2 is that if an evidence technology enables group pricing, one market segment voluntarily discloses message M to obtain a price discount; the remainder send message V and obtain the price charged to the complement, which by Definition 1, is no higher than $\bar{p}$. The monopolist deters any other message by ascribing off-path beliefs that put probability 1 on the highest type that can send that message.

We view Proposition 2 as modeling group price schemes seen in practice. Evidence of having a lower willingness to pay, as in the examples described above, often result in lower prices. Other examples may be more subtle but just as pertinent; for instance, financial aid applications might be interpreted as disclosures that allow some consumers to verifiably distinguish themselves from those with a high WTP.

Proposition 2 describes a minimal form of evidence that allows for Pareto gains. But this equilibrium may be inefficient. In Ali, Lewis and Vasserman (2022), we describe a simple approach to construct an efficient equilibrium that Pareto improves on uniform pricing; our construction uses *rich evidence*, where the consumer can disclose any interval that contains her value.

II. Disclosure Amplifies Competition

We turn now to how consumer control may amplify competition. Our framework below emphasizes two intuitions: (i) personalized pricing in markets with differentiated products can lead to intense competition for consumer types that do not have strong brand loyalty, and (ii) extreme types can be pooled with those moderate types so that they too obtain these price discounts. These two intuitions together lead to a stark contrast with the monopolistic case: even if the consumer's only choices are to opt in or opt out of being tracked by each firm, she obtains significant gains relative to uniform pricing.

Ali, Lewis and Vasserman (2022) show that this intuition holds in a general discrete-choice framework with n firms; for intuition, we focus on the case of Bertrand duopoly with horizontal differentiation. Firms A and B produce differentiated products at a marginal cost of 0. The consumer's value for the product from firm i is $v_i > 0$. We treat the consumer's net value for product A as her *location*, denoted by $\ell := v_A - v_B$. To ease exposition, we assume that ℓ is drawn from $[-1, 1]$ distributed symmetrically around 0 with a density f that is strictly log-concave; the uniform distribution involves $f(0) = 1/2$ whereas a single-peaked distribution has $f(0) > 1/2$. We also focus on the "no exclusion case" where the consumer's value for each product is sufficiently high that the consumer always buys from either firm in the equilibria below.

We compare the case of consumer control to two benchmarks. The first is that of *uniform pricing* where firms have no information about the consumer's location. In the symmetric equilibrium therein, each firm charges a price of $1/f(0)$. The second benchmark is *personalized pricing without control*. Here, the consumer's location is commonly known. The unique equilibrium outcome is that the consumer who favors firm i purchases from it at a price of $2|\ell|$, rejecting a price offer of 0 from the other firm. Personalized pricing benefits consumers with moderate preferences ($\ell \approx 0$) as firms compete heavily; this is the intuition from Thisse and Vives (1988). But this intuition fails for consumers who have extreme preferences: if the distribution is single-peaked at 0, such consumers are better off with uniform prices.[2]

We show that combining consumer control and personalized pricing assures gains

---

[2] See Rhodes and Zhou (2022) for further elaboration of the Thisse and Vives logic.

Reveals to only B	Consumer reveals location to both firms	Reveals to only A

A | $([-1,1],\{\ell\})$ | $(\{\ell\},\{\ell\})$ | $(\{\ell\},[-1,1])$ | B

-1 ℓ^A 0 ℓ^B 1

FIGURE 1. MODERATE TYPES FULLY OPT IN AND EXTREME TYPES OPT OUT FROM SHARING WITH CLOSER FIRM.

for *all* consumer types relative to uniform pricing. Suppose the consumer can opt in or opt out from disclosing her location to each firm individually. In Figure 1, we depict an equilibrium of this setting: a moderate consumer type reveals her location to both firms whereas a consumer type with a strong preference for firm i's product shares her location only with the other firm. In equilibrium, moderate consumer types obtain the same prices as the personalized-pricing benchmark above since these consumers opt in to both firms. The interesting action comes from those who opt out. Firm i infers that a consumer that opts out must have a strong preference for its product; but as the location is shared with the other firm, it also expects that the other firm charges a price of 0. Firm i then charges its optimal (local) monopoly price given this competition. In equilibrium, type ℓ^i is just indifferent between buying from the two firms (and also indifferent between sharing her location with firm i and not); all types whose preferences are more extreme benefit from pooling with this type as they are charged the same price.

We depict these gains in Figure 2: the red line describes equilibrium prices from the benchmark of uniform pricing and the dashed line from that of personalized pricing without control. As seen in the figure, personalized pricing with consumer control (shown in blue) results in strictly lower prices for all consumer types relative to uniform pricing, and benefits extreme types relative to personalized pricing without control. We summarize below.

PROPOSITION 3: *Suppose f is symmetric around 0 and strictly log-concave. If the consumer can opt in and out of each firm's tracking individually, then there is an equilibrium in which every consumer type is strictly better off than with uniform pricing.*

This analysis shows that merely endowing consumers with the right to opt in or opt out from being tracked results in significant gains for all consumer types. As we show in Ali, Lewis and Vasserman (2022), richer forms of evidence can be used to construct more sophisticated equilibria that assure even higher gains. Moreover, subject to an equilibrium refinement, the consumer may be better off in terms of average surplus even in the worst equilibrium for consumers.

III. The Privacy Paradox

The "privacy paradox" considers the following puzzle: *if consumers value privacy, why is it that they cede it so easily?*[3] Here we show that strategic logic of optimal pricing and voluntary disclosure helps explain why: if maintaining privacy is even slightly costly, the consumer is unwilling to incur costs to protect it.

Suppose the set of feasible messages that can be sent by type v, $\mathcal{M}(v)$, is some subset of all (Borel) subsets of V that contain v, including the fully revealing message $\{v\}$. Messages are now potentially costly: a consumer of type v incurs cost $c(m,v) \geq 0$ to send message $m \in \mathcal{M}(v)$. We assume that it is costless for the consumer to fully reveal her type, i.e., $c(\{v\},v) = 0$. However, *every* other message entails at least some minimal cost. That is, there is some $\underline{c} > 0$ such that for every type v and message $m \neq \{v\}$, $c(m,v) > \underline{c}$. This assumption models the minimal cost that a consumer incurs in switching to a different browser or operating in "privacy mode."

PROPOSITION 4: *The consumer fully reveals her type in every equilibrium.*

[3]For discussions of the privacy paradox, see Acquisti, Taylor and Wagman (2016), Athey, Catalini and Tucker (2017), Goldfarb and Tucker (2019), and Acquisti, Brandimarte and Loewenstein (2020).

FIGURE 2. CONSUMER CONTROL LOWERS PRICES FOR ALL CONSUMERS RELATIVE TO UNIFORM PRICING.

Let us argue why this is true. Towards a contradiction, suppose that in some equilibrium, some partially informative message m were sent on the equilibrium path. Let $\underline{v}(m)$ be the infimum of types that send that message. Observe that when receiving message m, the monopolist would charge a price no lower than $\underline{v}(m)$. But for $\varepsilon \in (0, \underline{c})$, all types in the ε-neighborhood of $\underline{v}(m)$ that send message m incur a cost of at least $\underline{c}$ to obtain a price discount no more than ϵ. Such types have a strictly profitable deviation to the fully revealing message, contradicting m being an equilibrium message.

Proposition 4 shows that if protecting one's privacy is even slightly costly, the consumer may be unwilling to incur those costs. In this light, default options matter: Apple's shift to a default-neutral setting for mobile app tracking has led to a considerably higher fraction of consumers opting out.

IV. Price Discrimination Revisited

Bergemann, Brooks and Morris (2015) identify the limits of price discrimination: the combinations of producer and consumer surplus that are achievable with some information structure in a monopolistic market. A feature of their setting is that the information structure is a primitive and not endogenously generated through choices of the consumer. Might consumer control limit the set of achievable outcomes?

To see how it may, consider a finite set of consumer types V for whom there are strict gains from trade. Define a *segment* as a distribution over consumer types and a *segmentation* as a distribution over segments that averages to the prior. A segmentation is efficient if the product is purchased with probability 1 and deterministic if there is only one segment containing any given type.

Our interest is in segmentations that result from the consumer choosing what to disclose to the seller. Suppose that $\mathcal{M}(v)$ comprises all subsets of V that contain v. Now, the segmentation is defined by the consumer's disclosure strategy: each on-path message M defines a segment corresponding to the monopolist's belief following that message. A segmentation is compatible with consumer control if it arises as an equilibrium of the disclosure game.

PROPOSITION 5: *An efficient segmentation is compatible with consumer control only if it is outcome equivalent to a deterministic segmentation.*

The logic of Proposition 5 is that in equilibrium, type v never mixes between two messages that result in different prices. Every segment that contains type v must then involve the same price. As this is true for all types, market segments that induce the same price can be pooled so that the outcome is equivalent to a deterministic segmentation.

Proposition 5 suggests that consumer control constrains the set of achievable payoffs; in particular, the consumer-optimal segmentation compatible with consumer control may be bounded away from that in Bergemann, Brooks and Morris (2015). This

preliminary observation raises several questions that we view to be of applied and theoretical interest. What is the set of achievable payoffs with consumer control? Which types are pooled in the consumer-optimal segmentation compatible with consumer control?

V. Conclusion

We have formulated a model of consumer control over data and derived its equilibrium implications for markets. Our analysis finds conditions under which giving consumers the ability to opt out allows them to benefit from personalized pricing.

The reader may wonder why consumer control is not self-defeating. Would the market not unravel, as in Grossman (1981) and Milgrom (1981)? An observation central to our results is that a firm's optimal price need not be monotone in its beliefs about the consumer's preferences. Therefore, extreme types may be pooled with moderate types without giving the moderate type any incentive to separate itself from the pool.

There are several other questions for which our approach may be useful. How does voluntary disclosure interact with the possibility to personalize the product to the consumer's tastes?[4] What is the consumer-optimal information and evidence design? If consumers intrinsically value privacy, how would this affect their disclosure choices?[5] Were consumers to use a private intermediary to control their disclosure, how might that intermediary collect their information voluntarily and what would it disclose to firms to obtain discounts?

REFERENCES

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature*, 54(2): 442–492.

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein.

[4]Ichihashi (2020) and Hidir and Vellodi (2021) study this issue in related frameworks.

[5]Lin (2022) shows, using a lab experiment, that the value for privacy can be decomposed into intrinsic and instrumental components.

2020. "Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age." *Journal of Consumer Psychology*, 30(4): 736–758.

Ali, S. Nageeb, Greg Lewis, and Shoshana Vasserman. 2022. "Voluntary disclosure and personalized pricing." *Review of Economic Studies*.

Athey, Susan, Christian Catalini, and Catherine Tucker. 2017. "The digital privacy paradox: Small money, small costs, small talk." Working Paper.

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

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature*, 57(1): 3–43.

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics*, 24(3): 461–483.

Hidir, Sinem, and Nikhil Vellodi. 2021. "Personalization, Discrimination and Information Revelation." *Journal of European Economic Association*, 119: 1342–1363.

Ichihashi, Shota. 2020. "Online privacy and information disclosure by consumers." *American Economic Review*, 110(2): 569–95.

Lin, Tesary. 2022. "Valuing intrinsic and instrumental preferences for privacy." *Marketing Science*, 41(4): 663–681.

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics*, 12(2): 380–391.

Rhodes, Andrew, and Jidong Zhou. 2022. "Personalized Pricing and Competition." Working Paper.

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