

Voluntary Disclosure and Personalized Pricing

S. NAGEEB ALI

Pennsylvania State University

GREG LEWIS

Microsoft Research

and

SHOSHANA VASSERMAN

Stanford University Graduate School of Business and NBER

First version received August 2020; Editorial decision April 2022; Accepted June 2022 (Eds.)

Firms have ever-increasing access to consumer data, which they use to personalize their advertising and to price discriminate. This raises privacy concerns. Policymakers have argued in response that consumers should be given control over their data, able to choose what to share and when. Since firms learn about a consumer's preferences both from what they do and do not disclose, the equilibrium implications of consumer control are unclear. We study whether such measures improve consumer welfare in monopolistic and in competitive markets. We find that consumer control can improve consumer welfare relative to both perfect price discrimination and uniform pricing. First, consumers can use disclosure to amplify competitive forces. Second, consumers can disclose information to induce even a monopolist to lower prices. Whether consumer control improves welfare depends on the disclosure technology and market competitiveness. Simple disclosure technologies suffice in competitive markets. When facing a monopolist, a consumer needs partial disclosure possibilities to obtain any welfare gains.

Key words: Privacy, Disclosure, Segmentation, Unravelling.

JEL Codes: D4, D8

“Privacy is not the opposite of sharing—rather, it is control over sharing.”
— Acquisti, Taylor and Wagman (2016)

1. INTRODUCTION

Consumers' online activity and behaviour are increasingly tracked through clicks and cookies, generating a bank of personalized data. These data are consolidated across sources by data owners and brokers and are used to infer what consumers are interested in and how much they are willing to pay. Data exist for literally hundreds of millions of consumers: as of 2014, one of the biggest data brokers, Axciom, had data on 700 million consumers worldwide, a number that has surely

The editor in charge of this paper was Andrea Galeotti.

grown over time (Federal Trade Commission, 2014). Firms use these data to target consumers for advertising or discounts. A recent survey from Deloitte found that of the 500 companies they surveyed that have adopted AI for personalization, 40% were using this technology for real-time promotions and pricing (Hogan, 2018).¹

How does all of this personalization affect consumer surplus? Targeted advertising and price offers can lead to efficiency gains if firms are able to accurately personalize their marketing efforts. Consumers may then see more relevant content and products, and incur lower search costs while doing so. But regulators are concerned that personalization may adversely affect the distribution of that surplus: if firms have detailed knowledge of consumers' willingness to pay, they may appropriate not only the efficiency gains of personalization but also much of the consumer surplus altogether. Consumers may then be strictly worse off than in a world without personalized pricing.

These concerns play a prominent role in the ongoing international debate about *consumer privacy*. Consumer advocates have pushed for greater *control* of how consumer data are shared, and their advocacy efforts have borne fruit across the globe. The General Data Protection Regulation (GDPR) passed by the European Union in 2016, requires firms to obtain consent from consumers before collecting and processing their personal data. In the US, the Federal Trade Commission recommends that “best practices include...giving consumers greater control over the collection and use of their personal data...” (Federal Trade Commission, 2012).² The private sector has also taken steps to give consumers more control. For instance, Apple's iOS 14.5 update, introduced in April 2021, defaults to asking users whether they want to allow an app to track them, a move that attracted attention in the mainstream media.³

Against this backdrop, we study the market implications of consumer consent and control. Although consumer control features prominently in regulatory debates and public rhetoric about online interactions, relatively little is understood about how to conceptualize its equilibrium effects and design effective regulations.⁴ We investigate what happens when consumers fully control their data—not only whether they are tracked, but what specific information is disclosed to firms. We treat data as verifiable hard information that consumers can partially or fully disclose to firms. In equilibrium, firms draw inferences about a consumer's preferences based on both what she shares and what she chooses not to share. Our motivating question is: *when consumers voluntarily disclose information about their preferences, are they hurt or helped by personalized pricing?*

A classical intuition suggests that consumers cannot benefit from being permitted to voluntarily disclose information. Because the market makes equilibrium inferences that are based both on the information that is disclosed and what is not being disclosed, giving consumers the ability to separate themselves may be self-defeating—as seen in the unravelling equilibria of Grossman (1981) and Milgrom (1981). This intuition suggests that consumer control may not, in equilibrium, protect consumer interests. By contrast, we find that the combination of personalized pricing and consumer control can be beneficial to consumers in both monopolistic and competitive markets.

1. See Organization for Economic Co-operation and Development (2018) for a discussion of the empirical evidence on personalized pricing. The consensus seems to be that empirical evidence of personalized pricing is “relatively limited”, but that this may be because firms are good at making such pricing behaviours opaque and hard to detect.

2. Consumer control is also increasingly featured in US state regulations, with support from both legislators and voters: the California Consumer Privacy Act (CCPA) was put into effect in January 2020 and expanded by a ballot proposition in November 2020 to have features similar to the GDPR.

3. For an excellent explanation of the current ad ecosystem, see Apple's own promotional material at https://youtu.be/Ihw_A14RNno. The Daily Show ran a segment entitled “Facebook and Apple battle over privacy”. See <https://youtu.be/rxXkPBWHDik>.

4. See also the report on “Algorithms: How they can reduce competition and harm consumers” by the UK Competition and Markets Authority, and Chapter 3 of the GDPR bill.

We construct equilibria of the consumers' disclosure game in which sharing data weakly increases consumer surplus for *every* type of consumer, relative to the benchmark of uniform pricing.

Two ideas drive these results. First, voluntary disclosure and personalized pricing together amplify competition between firms. Nearly indifferent consumers benefit from the ability to credibly communicate their flexibility, intensifying competition for their business. Meanwhile, consumers with a strong preference for the product of one particular firm can hide this preference by pooling together. Although firms do interpret this non-disclosure as a signal of strong preferences, the resulting prices are lower than those that obtain without the option to disclose information. Second, even in the absence of competition, consumers can benefit from sending coarse signals that pool their valuations. These pools are constructed so that the monopolist finds it optimal to sell to every type within each pool and offer almost all types a price discount.

1.1. *A preview*

To describe our work, we start with the problem of a monopolist choosing what price to charge a consumer whose valuation he does not know. We augment that classical problem with a disclosure stage: before the monopolist sets his price, the consumer chooses what *evidence* to disclose about her valuation. Evidence is *hard information* about the consumer's preferences, which can be verifiably disclosed to the firm subject to the consumer's disclosure decisions. We study two distinct kinds of evidence. The first is that of *simple evidence*, where a consumer can either reveal her type completely or say nothing at all. This corresponds to speaking "the whole truth and nothing but the truth", and is a stylized version of the track/do-not-track technology that features prominently in discussions of consumer control. The second is that of *rich evidence*, where a consumer can disclose some of her data or personal characteristics without having to reveal it completely. This is a partial disclosure setting where the consumer can speak "nothing but the truth", but is not constrained to having to speak the "whole truth".⁵ We take the evidence structure as a primitive of the game and consider the following timeline: (i) the consumer chooses a message to send to the firm, (ii) the firm quotes a price, and (iii) the consumer chooses whether to buy the product at that price. To characterize outcomes in this setting, we study the perfect Bayesian equilibria (PBE) of the resulting game.

Proposition 1 shows that simple evidence never benefits the consumer in a monopolistic market and potentially hurts her: there is no equilibrium in which *any type* of the consumer is better off relative to a setting without personalized pricing. Moreover, there are equilibria in which all consumer types are worse off, such as a fully revealing equilibrium in which the monopolist extracts all surplus. This result articulates that all-or-nothing forms of consumer control—such as track/do-not-track—do not benefit consumers when they are facing a monopolist.

However, if the evidence structure is rich—where consumers can partially disclose information without revealing all of it—all consumer types can benefit from disclosing information. Proposition 2 describes an efficient equilibrium that improves consumer surplus for almost all types without reducing the surplus of any type. We construct this equilibrium using an elementary approach: for a distribution of consumer types, find the segment whose willingness to pay exceeds the monopolist's optimal posted price for that distribution. This is then the highest market segment. Our "greedy" algorithm removes this segment and repeats the process for the truncated distribution. Continuing this process until all types whose valuation exceeds the cost of the product have been exhausted results in a Pareto improving efficient equilibrium segmentation. This equilibrium also maximizes *ex ante* consumer surplus for a class of distributions.

5. We borrow the "simple" and "rich" terminology from Hagenbach and Koessler (2017).

Generalizing this result, we show that although rich evidence suffices to generate Pareto gains, it is not necessary. As we describe in Proposition 5, it suffices to have partial disclosure possibilities that facilitate “group pricing”. That is, there may be Pareto gains from the combination of voluntary disclosure and personalized pricing so long as there is evidence that low-value consumer types can use to separate themselves from higher-value types. This result connects to group-pricing schemes seen in practice—to students, senior citizens, or low-income households—where a third-party verifies the consumer as being a low-value type and offers evidence that can be disclosed to obtain price discounts.⁶

Stronger conclusions emerge when there are multiple firms and there is some horizontal differentiation. We consider general competitive markets with differentiated products, in which the consumer demands a single unit. The firms are uncertain of the consumer’s preferences and the consumer can disclose information about her preferences to the firms, who then simultaneously make price offers to her. We compare the outcomes when consumers can disclose information—either via simple or rich evidence—against a benchmark model in which there is uniform pricing. Here, voluntary disclosure and personalized pricing may be beneficial to consumer surplus because of a new economic force: *information can be selectively disclosed to amplify competition*.

This force is sufficiently powerful that even with simple evidence, there is an equilibrium in which consumers benefit relative to the setting without personalized pricing. In this equilibrium, consumer types with strong brand loyalty to a firm conceal their types from that firm and pool with those who have less extreme brand loyalty. Meanwhile, consumers who have moderate preferences for the product fully reveal their preferences to all firms. More generally, we show that consumers are better off with voluntary disclosure and personalized pricing relative to uniform pricing in even the worst equilibrium (within a certain class).

We study several extensions of our results. Principally, we consider a model with limited control, in which firms have access to some consumer data (e.g. from their own past transactions with this consumer), though not enough to fully infer a consumer’s type. Even without disclosure, prices are no longer uniform in this setting since firms price discriminate on the basis of their available information. But as we show, if consumers are able to send messages that depend both on their type and on what the firm knows about them, they can still refine the firms’ beliefs and get a better deal.

1.2. Implications

Our aim is to inform the ongoing regulatory debate on consumer control by studying the market implications of consumer control. The last decade has seen a major push towards giving consumers control over where they are tracked, and this is reflected both in regulatory and market-based solutions. One way to view our study is that it describes *necessary* conditions for consumer control to benefit consumers and elucidates a subtle interaction between the technology by which consumers control their information and the degree of market competitiveness. Our results show that all-or-nothing forms of consumer control—such as track/do-not-track—do not benefit consumers when there is a monopolist. In this case, richer and more sophisticated technologies are needed for consumer control to benefit consumers. By contrast, if the market features both competition and some horizontal differentiation, then simple evidence may suffice for consumer gains. Moreover, as we highlight, even the worst equilibrium for consumers in this setting can

6. Group pricing has been discussed in the context of third-degree price discrimination where the firm prices on the basis of *publicly observable* characteristics; see Varian (1989) and Belleflamme and Peitz (2015). Here, we consider a setting in which the consumer endogenously discloses the group to which she belongs.

be better than uniform pricing. Simple forms of consumer control are therefore potentially more valuable in competitive settings than in a monopolistic market.

Our results also speak to the role of consumer control in the trade-offs between product personalization and price discrimination. One view is that sharing data with firms is beneficial because it better matches consumers with products, but information sharing comes at the cost of price discrimination. This view might suggest that without personalization, consumers do not gain from the combination of control and personalized pricing. Contrary to this view, our analysis shows that offering consumers control can make personalized pricing attractive *even in the absence of better matching*.

Finally, our analysis is germane to current technological solutions that enable greater consumer control. Pressure from regulators and consumers has led major platforms to search for richer ways to mediate the disclosure and use of user information. As of Q4 2021, Google has begun testing a new technology called “Federated Learning of Cohorts” in its Chrome browser that replaces third-party cookies—which allows a user to be tracked across multiple websites—with an assignment of each user to a cohort of similar users (i.e. a pool), thus preserving some individual privacy.⁷ Furthermore, Google and Facebook have released personalized privacy tools that allow users to see how their information is bucketed into demographic and interest bins and to opt out of sharing specific types of information or sharing information with specific advertisers or sellers. We elaborate further on the connection between our analysis and consumer control in practice in Section 5.

1.3. *Relationship to literature*

Our work belongs to the literature on privacy, information, and their implications for markets; see Acquisti *et al.* (2016) and Bergemann and Bonatti (2019) for recent surveys. We combine classical models of pricing and verifiable disclosure. Unlike the first analyses of verifiable disclosure (Grossman, 1981; Milgrom, 1981), unravelling is not the unique equilibrium outcome of the markets that we study. An observation central to our results is that a firm’s optimal price need not be monotone in its beliefs about the consumer’s willingness to pay. This observation permits us to pool high and low types without giving the low type an incentive to separate.⁸

For the case of a monopolistic seller, a couple of papers touch upon related issues. Sher and Vohra (2015) study monopolistic price discrimination where the seller commits to a schedule of evidence-contingent prices.⁹ We assume that sellers cannot commit and instead set prices that best respond to the evidence that has been presented. Closer to our work, Glode, Opp and Zhang (2018) and Pram (2021) show that a consumer can disclose rich evidence to obtain Pareto gains from a monopolistic seller, looking at both private and interdependent values. Glode *et al.* (2018) focus on an *ex ante* disclosure model, where a buyer commits to a disclosure plan before she learns her type. They show that these disclosure plans sustain efficient trade and that under certain conditions the optimal disclosure plan among those that are “monotone” proceeds along the greedy algorithm studied here. Pram (2021) studies a broader class of adverse-selection settings and characterizes necessary and sufficient conditions for the existence of a Pareto improving equilibrium.¹⁰ These results align with those in our article for

7. See https://en.wikipedia.org/wiki/Federated_Learning_of_Cohorts.

8. Prior analyses have highlighted other reasons for why markets may not unravel, in particular (i) uncertainty about whether the sender has evidence (Dye, 1985), (ii) disclosure costs (Jovanovic, 1982; Verrecchia, 1983), or (iii) the possibility for receivers to be naive (Hagenbach and Koessler, 2017).

9. McAdams (2011) studies a related model in which disclosure is costly.

10. Stahl and Strausz (2017) also study adverse selection and certification, focusing on complementary issues.

the case of rich evidence in a monopolistic setting. Distinct from this prior work, our motivation is to understand the interaction of consumer control and market competition, which is why we study both simple and rich evidence in both monopolistic and competitive markets. Our message of how simple evidence amplifies competitive forces but richer forms of evidence are needed with a monopolist does not feature in this prior work. Thus, we view our papers to be highly complementary shedding light on different aspects of verifiable disclosure in markets.

Our verifiable-disclosure approach to consumer control complements other ways for consumers to share information in markets. One is through information design, where an intermediary that knows the consumer's type commits to a segmentation strategy. With a monopolistic seller, that intermediary can achieve payoffs characterized by Bergemann, Brooks and Morris (2015). Another way is for the consumer to use cheap talk. Hidir and Vellodi (2021) show that if the product can be customized to the consumer's tastes, cheap talk can improve matching without resulting in the monopolist capturing the entire surplus. Verifiable disclosure offers a middle ground between information design and cheap talk and is relevant when a consumer can directly control the disclosure of her information to firms.

Related to how voluntary disclosure amplifies competition in markets with differentiated goods is the innovative work of Thisse and Vives (1988). In a model of Bertrand duopoly, they study whether firms choose to personalize prices. In their model, the consumer's type is commonly known and they show that the unique equilibrium involves personalized pricing even though the firms' joint profits would be higher if they could commit to uniform pricing. An important distinction with our work is that all of the action in their model (and in the subsequent literature) is on the side of the firms: taking consumers as passive, these papers study whether firms personalize prices when they know (or can learn) consumer types.¹¹ In our model, the consumer actively chooses whether to disclose information and it is her voluntary disclosure that facilitates personalized pricing. Moreover, the ability to pool with other types is necessary for every consumer type to benefit from personalized pricing; otherwise, extreme types are worse off from personalized pricing. Thus, the welfare gains that we study would not emerge in a model in which consumer types are commonly known.

Several papers also study the role of information in competitive markets with differentiated products. Elliott, Galeotti, Koh and Li (2021) show that an information designer can segment the market so that consumers are allocated efficiently while nevertheless guaranteeing that consumers obtain no surplus. Armstrong and Zhou (2022) study optimal information structures for a consumer who does not know her tastes and show that the consumer-optimal signal may be quite uninformative so as to maintain price competition. Fainmesser and Galeotti (2016, 2020) study monopolistic and competitive price discrimination based on how consumers influence each other.

We abstract from a number of issues relevant to data markets. A prior literature studies how a consumer may distort her behaviour in dynamic settings if firms draw inferences about her tastes from her past choices.¹² Our analysis complements this agenda by studying how a consumer fares from directly controlling the flow of information rather than distorting her behaviour to influence the market's perception of her tastes. Recent work has focused on data externalities.¹³ We do not model data externalities directly, but it motivates our analysis of settings in which consumer

11. See Armstrong (2006) for a survey. We thank Jidong Zhou for drawing our attention to this work.

12. See Taylor (2004), Villas-Boas (2004), Acquisti and Varian (2005), Calzolari and Pavan (2006), Conitzer, Taylor and Wagman (2012), and Bonatti and Cisternas (2020).

13. See Choi, Jeon and Kim (2019), Acemoglu, Makhdoumi, Malekian and Ozdaglar (2021), Bergemann, Bonatti and Gan (2021), and Ichihashi (2021).

control is limited at best. Finally, we abstract from issues related to data collection, protection, and sale.¹⁴

1.4. Outline

We illustrate our main results in Section 2. Section 3 offers the general analysis for the monopolist setting and Section 4 studies competitive markets with product differentiation. Section 5 describes several practical implications of our work. Section 6 concludes.

2. TWO EXAMPLES

2.1. A monopolist

A consumer (“she”) is in the market for a single widget. Widgets are produced by a single firm (“he”), who thus is a monopolist. The consumer’s value for the widget is v drawn uniformly from $[0, 1]$. In interacting with the firm, she is able to provide evidence of this valuation v in the form of consumer data, such as her demographics, past browsing, and past purchase data (including past purchase prices). This evidence is communicated in the form of a message m , which takes the form of an interval $[a, b]$. The interpretation of a message $m = [a, b]$ is that it contains data that proves that v is in $[a, b]$.

We consider two possible evidence structures, *simple* and *rich*. Simple evidence is “all-or-nothing”—the consumer can either reveal her valuation (send a message $m = v$) or say nothing (send a message $m = [0, 1]$). This messaging system requires the consumer only to make a binary choice and has a real-world analogue in the decision by a consumer to either allow herself to be tracked by a firm (i.e. hand over her data) or to select “do-not-track” (i.e. retain all the data).

Rich evidence allows any truthful message—for example, a consumer with $v = 0.5$ could send the message $[0.3, 0.8]$. This evidence could be communicated by partial disclosure of consumer information—a subset of past purchases, say; or demographics only. By contrast to the simple evidence scheme which can be passively implemented by “track/do-not-track” choices, rich evidence requires consumers to actively choose their messages.

We want to know if disclosure hurts or harms consumers under both evidence structures. To determine this, we study the PBE of the following game: first the consumer observes her type v and chooses a message $m(v)$ to send; then the monopolist chooses a price $p \geq 0$ and the consumer decides whether to purchase. We compare the PBE of the disclosure games with a baseline model in which there are no disclosure possibilities. Without disclosure, this is a standard monopoly pricing problem: the monopolist optimally posts a uniform price of $p = \frac{1}{2}$, yielding an *ex interim* consumer surplus of $\max\{v - \frac{1}{2}, 0\}$, and a producer surplus of $\frac{1}{4}$.

We first ask whether the consumer can ever be better off relative to this baseline when she has simple evidence. Consistent with the concerns posed by privacy advocates, any consumer sending the perfectly revealing message $m = v$ will be charged exactly her valuation and obtain zero surplus. Nonetheless, there is an equilibrium where all types are compelled to disclose. In this equilibrium, the monopolist believes that non-disclosure ($m = [0, 1]$) indicates that $v = 1$ with probability 1 and charges $p = 1$. The consumer is then indifferent between revealing and paying her value to buy—which results in a payoff of zero—and not revealing or buying—which also results in a payoff of zero.

14. Jones and Tonetti (2020) show that, because data are non-rival, there are social gains from multiple firms using the same data simultaneously, and therefore, it is better to let consumers own and trade data. Fainmesser, Galeotti and Momot (2020) study tradeoffs between data collection and data protection.

But this is not the only equilibrium: there is also one in which the consumer never discloses ($m = [0, 1]$ for every type v), and the monopolist charges $1/2$. No one has an incentive to deviate because by disclosing they earn zero surplus, which is (weakly) worse for every type. Notice that this equilibrium results in the same payoffs, prices, and outcomes as the uniform pricing baseline. There is no equilibrium in which consumers can do better:

Observation 1. *With simple evidence, across all equilibria, the consumer’s interim payoff is no more than her payoff from uniform pricing, namely $\max\{v - 1/2, 0\}$.*

To see why Observation 1 holds, suppose that the monopolist charges price p_{ND} to all non-disclosing types. Any type with value $v > p_{ND}$ then prefers not to disclose, since otherwise they will be charged their valuation—and therefore the set of non-disclosing types must contain the set $(p_{ND}, 1]$. Now, the optimal price for a monopolist facing a set of types that includes the interval $(\bar{v}, 1]$ is $\frac{1}{2}$ if $\bar{v} \leq \frac{1}{2}$ and \bar{v} otherwise. So by optimal pricing, p_{ND} must be at least $\frac{1}{2}$, implying the bounds on interim payoffs noted above.

Observation 1 tells us that consumers do not benefit from simple evidence—at least when facing a monopolist. But as we show below, rich evidence can help.

Let us first start with a minimal elaboration: suppose that those consumer types whose valuation is below $1/2$ have evidence to show that their valuation is low (e.g. a senior citizen or student ID card). We can now consider an equilibrium with the following structure: the monopolist sells the widget at a sticker price of $\frac{1}{2}$, but offers a personalized group discount of 50% if the consumer proves that her valuation is below $1/2$. To see that this is an equilibrium, note that upon seeing evidence that the value is low, the monopolist would believe that the consumer’s type is uniformly distributed on $[0, 1/2]$, in which case his optimal price is $p = \frac{1}{4}$. If the monopolist does not see such evidence, he believes that the type is uniform on $(1/2, 1]$, in which case his optimal price is $p = \frac{1}{2}$. Meanwhile, the low-value consumer cannot benefit from deviating as not disclosing anything would lose her the discount and fully revealing v would lead to a payoff of zero. But now every consumer type in $[\frac{1}{4}, \frac{1}{2}]$ strictly benefits from the availability of evidence.

In fact, we can use rich evidence to construct equilibria with further gains for consumer surplus. To do this, we consider the countable grid $\{1, \frac{1}{2}, \frac{1}{4}, \dots\} \cup \{0\}$, which we refer to as *Zeno’s Partition*.¹⁵ We denote the $(k + 1)$ th element of this list, namely 2^{-k} , by a_k , and consider the following construction.

Observation 2. *With rich evidence, there is an equilibrium that generates Zeno’s Partition: the consumer’s disclosure strategy is*

$$m(v) = \begin{cases} [a_{k+1}, a_k] \text{ where } a_{k+1} < v \leq a_k & \text{if } v > 0, \\ \{0\} & \text{if } v = 0. \end{cases}$$

When the monopolist receives message $[a_{k+1}, a_k]$, he charges a_{k+1} thereby selling to that entire segment. Relative to uniform pricing, this equilibrium strictly improves consumer surplus for all v in $(0, 1/2]$, and leaves consumer surplus unchanged for all other types.

15. This is inspired by Zeno’s Paradox, summarized by Aristotle as “...that which is in locomotion must arrive at the half-way stage before it arrives at the goal...” See <https://plato.stanford.edu/entries/paradox-zeno/>.

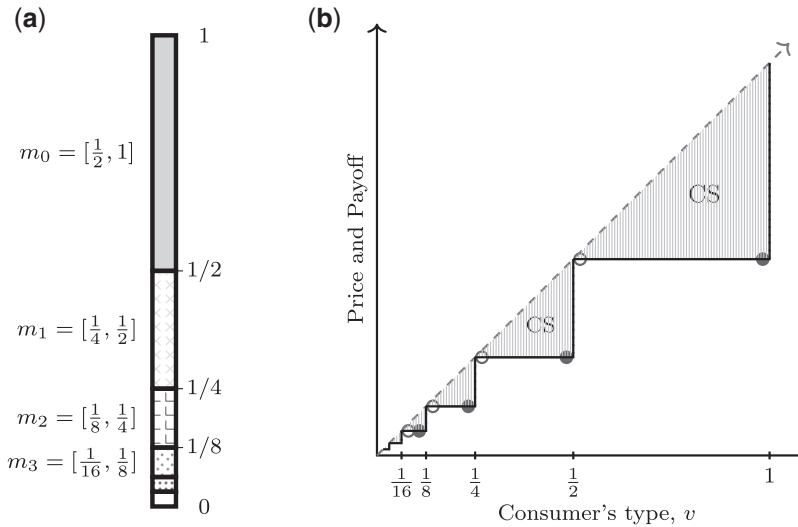


FIGURE 1

(a) Illustrates Zeno's Partition. (b) Illustrates prices and payoffs: for each consumer type v , the step function shows the equilibrium price that is charged, and the dashed 45° line shows the payoff from consumption. The shaded region illustrates the consumer surplus achieved by Zeno's Partition.

In this equilibrium, the highest market segment is composed of types in $(\frac{1}{2}, 1]$, all of which send the message $m_0 \equiv [\frac{1}{2}, 1]$; the next highest market segment comprises types in $(\frac{1}{4}, \frac{1}{2}]$, all of which send the message $m_1 \equiv [\frac{1}{4}, \frac{1}{2}]$, and so on and so forth. We depict this partition in Figure 1. Once the monopolist receives any message corresponding to each market segment, he believes that the consumer's value is uniformly distributed on it. His optimal strategy then is to price at the bottom of the segment. Therefore, trade occurs with probability 1, with each higher consumer type capturing some surplus.

This equilibrium generates an *ex ante* consumer surplus of $\frac{1}{6}$ and producer surplus of $\frac{1}{3}$. By contrast, uniform pricing—or the consumer-optimal equilibrium with simple evidence—results in an *ex ante* consumer surplus of $\frac{1}{8}$ and producer surplus of $\frac{1}{4}$. We interpret this rich-evidence equilibrium as involving a sticker price of $\frac{1}{2}$, and if a consumer can verify that her valuation belongs to a lower segment, she obtains a price discount. Using this scheme, personalized pricing improves the monopolist's profit and strictly improves the surplus for some consumer types without making any worse off.

How is Zeno's Partition supportable as an equilibrium? First, let us see what deters consumers from using messages that are not in Zeno's Partition. If the monopolist sees such a message, his off-path beliefs ascribe probability 1 to the highest type that could send such a message, which leads him to charge a price equal to that type. Such beliefs ensure that these off-path messages are not profitable deviations. How about deviations to other on-path messages? For every v in (a_{k+1}, a_k) , there is only one on-path message that it can send, and for every v on the boundaries of such messages, our strategy profile prescribes the message that results in the lower price. Thus, there are no profitable deviations for any consumer type. Finally, by construction, the monopolist is always setting a price that is optimal given its beliefs.

It is useful to understand why we do not see unravelling. In many disclosure models, the sender strictly prefers to induce the receiver to have higher (or lower) beliefs in the sense of first-order

stochastic dominance. Unravelling emerges as the unique equilibrium outcome as extreme types of the sender have a motive to separate from pools. By contrast, in our setting, there are beliefs $(\mu, \hat{\mu})$ that are ranked by FOSD such that the sender is indifferent between inducing μ and $\hat{\mu}$ because they result in the receiver taking the same action. For example, the monopolist charges the same price when he ascribes probability 1 to type $\{1/2\}$ as he does when his beliefs are $U[1/2, 1]$. Thus, higher types may be pooled with a lower type without giving that lower type an incentive to separate.

As we noted earlier, Zeno’s Partition is not the only equilibrium of this example. In fact, there are worse equilibria for consumer surplus: for example, a fully revealing equilibrium continues to exist with rich evidence. Accordingly, rich evidence, in this setting, is necessary for gains to consumer surplus but does not guarantee that consumers achieve such gains.

However, there are no equilibria in this example that achieve higher *ex ante* consumer surplus than Zeno’s Partition. To see why, let us consider only those equilibria that are fully efficient and take a partition form as above.¹⁶ Because consumers always purchase in a fully efficient equilibrium, maximizing consumer surplus is equivalent to minimizing the average price. For a monopolist to price at the bottom of an interval $[a, b]$ when v is uniformly distributed between a and b , it must be that $a \geq b/2$. Suppose that the consumer-optimal equilibrium involves types from $[\lambda, 1]$ forming the highest segment; by the logic of the previous sentence, λ is at least $1/2$. The monopolist charges a price of λ to that segment, and thus, its contribution to the *ex ante* expected price is $(1 - \lambda)\lambda$. The remaining population, $[0, \lambda]$, amounts to a λ -rescaling of the original problem, and so the consumer-optimal equilibrium after removing that highest segment involves replicating the same segmentation on a smaller scale. Thus, the consumer-optimal segmentation can be framed as a recursive problem where $P(\bar{v})$ is the lowest expected price generated by a partition when types are uniformly distributed on the interval $[0, \bar{v}]$:

$$P(1) = \min_{\lambda \geq \frac{1}{2}} (1 - \lambda)\lambda + \lambda P(\lambda) = \min_{\lambda \geq \frac{1}{2}} (1 - \lambda)\lambda + \lambda^2 P(1) = \min_{\lambda \geq \frac{1}{2}} \frac{(1 - \lambda)\lambda}{1 - \lambda^2} = \frac{1}{3},$$

where the first equality frames the problem recursively, the second is from $P(\lambda)$ being a re-scaled version of the original problem, and the remainder is algebra. Because Zeno’s Partition induces an expected price of $1/3$, no segmentation can generate higher consumer surplus.

Summarizing, we have shown that for the consumer to benefit from consumer control when facing a monopolist, simple evidence is insufficient. It never makes the consumer better off and in some equilibria, makes the consumer worse off relative to the baseline model without disclosure. With rich evidence, it is possible to construct equilibria that lead to efficient trade and strictly increase consumer surplus for almost every consumer type. The equilibrium that we construct pools consumers into groups such that all members of the group trade at the optimal price for that group and have no incentive to opt out. As we highlight in Section 3, it is not necessary that the rich evidence corresponds to partitioning into fine intervals; all that is necessary for a Pareto improvement is for some low-value consumers to have a verifiable way to separate themselves from high-value consumers. So, one may interpret the positive result as corresponding to group discounts often seen in practice.

2.2. Bertrand competition with horizontal differentiation

Now suppose that the consumer has the choice between two different firms that make widgets. These widgets are differentiated and the consumer prefers one widget to the other. Using her past

16. We show that this is without loss of generality: for every equilibrium, there is an interim payoff-equivalent equilibrium in which trade occurs with probability 1 and types segment into partitions.




FIGURE 2

The figure shows disclosure strategies for every type. Centrally located types fully reveal the location to both firms. Extreme types reveal location only to the distant firm and conceal it from the closer firm.

purchase data, she is able to provide evidence that verifiably communicates this preference. Can she use her data to get better prices than in the baseline without disclosure?

Suppose there are two firms, L and R , who compete to sell to a consumer who must purchase one unit of the good from either firm. Firm L is located at the point $\ell_L = -1$, firm R at the point $\ell_R = 1$. The consumer's location, ℓ , is drawn uniformly from $[-1, 1]$. The consumer knows her location but the firms do not. If the consumer purchases the good from firm i , then she pays the price p_i that is set by firm i , as well as a linear transportation cost $|\ell_i - \ell|$. In the baseline model, where the consumer has no option to disclose information, each firm would set a price of 2 in equilibrium. The consumer would then buy the good from the closer firm and incur a total expenditure of $2 + \min\{1 + \ell, 1 - \ell\}$.

Now, let us describe what can happen with simple evidence. The consumer at location ℓ can disclose one of two messages to each firm privately before prices are set: either she can send a message of $\{\ell\}$, which fully reveals her location, or a message of $[-1, 1]$, which fully conceals it. That the messages are sent privately reflects that the consumer can make two independent tracking decisions: she can allow firm L to track or not track her independently of whether she has allowed firm R to do so.

Figure 2 illustrates an equilibrium where the consumer selectively discloses evidence to amplify price competition and reduce her total expenditure. In this construction, the consumer conceals her location when she has a strong preference for the product of one of the firms. Thus, if her location is in $[1/2, 1]$, she reveals her location to the distant firm L but conceals it from R ; she plays a symmetric strategy if her location is in $[-1, -1/2]$. Only if her location is in $(-1/2, 1/2)$ does she reveal her location to both firms.

This disclosure strategy lowers market prices. Suppose that the consumer is located in $[1/2, 1]$, and hence she discloses her location to firm L but not to firm R . Firm R infers from this non-disclosure that the consumer must be located in $[1/2, 1]$ but does not know where in that interval. Firm L , on the other hand, knows the consumer's location, and to poach the consumer's business, it has to significantly reduce its price. In equilibrium, this powerful force pushes the firm L to lower its price all the way to its marginal cost of 0. Anticipating this outside option for the consumer, firm R solves for its optimal local monopoly price, which is 1. At these prices, the consumer purchases from firm R and therefore has a total expenditure of $2 - \ell$, which is strictly smaller than the total expenditure without personalized pricing.

What about all those interior types in $(-1/2, 1/2)$ that are fully revealing themselves to each firm? In this competitive market, these consumer types benefit from revealing their locations: each obtains a price of 0 from the distant firm and a price from the closer firm that makes it just indifferent. In equilibrium, each incurs a total expenditure of $\max\{1 + \ell, 1 - \ell\}$, which once again is below that without personalized pricing.

We see that consumer control, even through simple evidence, can be a powerful force for consumer gains. By allowing the distant firm to track her, the consumer motivates them to lower prices. This in turn forces the home firm to compete for her business. As we show in Section 4, this force manifests more broadly: it applies when there are more than two firms, and for both

general distributions and general forms of product differentiation. We also show there that if the consumer has access to rich evidence, she can improve upon this segmentation using a Zeno-like construction just as we did in the monopolist example. But an important lesson from this analysis is that in a competitive market with horizontal differentiation, rich evidence is not necessary for consumer gains.

We now move on to describing the monopoly model in full generality.

3. VOLUNTARY DISCLOSURE TO A MONOPOLIST

3.1. *Environment*

A monopolist (“he”) sells a good to a single consumer (“she”), who demands a single unit. The consumer’s valuation, denoted by v , is drawn according to a continuous and atomless cdf F on the support $V \equiv [\underline{v}, \bar{v}]$ where $\underline{v} \geq 0$. Payoffs are quasilinear: if the consumer purchases the good from the monopolist at price p when her type is v , her payoff is $v - p$ and the monopolist’s payoff is p ; otherwise, each player receives a payoff of 0.¹⁷

We append a disclosure game to this pricing problem. After observing her type, the consumer chooses a message m from the set of messages available to her. The set of *all* feasible messages is $\mathcal{M}^{\mathcal{F}}$, which denotes the collection of all intervals of the form $[a, b]$ that is a subset of V . We interpret the message $[a, b]$ as meaning “my type is in the set $[a, b]$ ”. When the consumer’s type is v , the set of messages that she can send is $\mathcal{M}(v) \subseteq \mathcal{M}^{\mathcal{F}}$. We study two forms of disclosure:

- the evidence technology is *simple* if for every v , $\mathcal{M}(v) = \{V, \{v\}\}$,
- the evidence technology is *rich* if for every v , $\mathcal{M}(v) = \{M \in \mathcal{M}^{\mathcal{F}} : v \in M\}$.

Each type of technology represents a medium through which the consumer can credibly convey information about her type. A simple evidence technology maps to the track/do-not-track dichotomy whereby a consumer can either share her personal characteristics and consumer data with the firm or not share anything at all. Our analysis here assumes that the firm can then translate the consumer’s data into the consumer’s valuation for the good. The rich evidence technology reflects a more sophisticated form of consumer control, through which the consumer can verifiably disclose true statements about her type without being compelled to reveal everything. This corresponds to the consumer revealing, say, a subset of her characteristics or being able to verifiably disclose that she is part of a certain “bucket” of consumers, without having to reveal her type entirely.

Our analysis here assumes that the firm can translate the consumer’s data into the consumer’s valuation for the good. We present our analysis in a unidimensional type space for simplicity, but as we elaborate in Section 3.5.3, our analysis applies to settings in which the consumer’s valuation depends on multiple characteristics, and a consumer has control over the information that is shared about these characteristics. In Section 5, we discuss several ways to interpret the rich evidence model through existing technologies in the public and private sectors.

We consider the following timeline. First, the consumer observes v and chooses a message M from $\mathcal{M}(v)$. The monopolist observes the message and chooses a price $p \geq 0$. The consumer then chooses whether to purchase the good. We study PBE of this game: the seller’s belief system is updated via Bayes’ Rule whenever possible, and both the consumer and the seller behave

17. We assume the monopolist has zero marginal costs. This assumption does not play an important role in our analysis.

sequentially rationally.¹⁸ For expository convenience, we assume that a consumer always breaks her indifference in favour of purchasing the good.

We compare the results of this model with a benchmark in which communication is infeasible—for instance, because personalized pricing is altogether prohibited by regulations. In this uniform pricing benchmark, the monopolist sets a price p that maximizes $p(1 - F(p))$. We denote the (lowest) revenue-maximizing price by p^* , so that the consumer's interim payoff is $\max\{v - p^*, 0\}$. Our analysis compares the benchmark with an alternative policy of allowing the consumer to control her data and letting the firm set prices based on what the consumer shares with the firm. If $p^* = v$, it would be impossible to improve upon the benchmark because the monopolist would be servicing the entire market. Accordingly, we assume that there is some scope for improving upon the benchmark.

Assumption 1. *The monopolist's optimal uniform price p^* is strictly higher than the lowest type v .*

Section 3.5.1 considers an alternative benchmark in which the consumer has limited control over their data, so that some information about their preferences “leaks” and is observed by the monopolist.

3.2. Simple evidence does not help consumers

In Section 2, we argued that simple evidence technologies may not be able to improve consumer surplus when the seller is a monopolist. In this section, we show that the result from our example extends more generally under our model. There are equilibria with simple evidence that make all consumer types worse off but there are no equilibrium in which *any* type is strictly better off.

To see that consumers may be worse off, consider an equilibrium in which the consumer fully reveals her type v with probability 1 and the monopolist charges a price of v . Off-path, the seller's beliefs are *maximally sceptical* in that he believes that the consumer's type has the highest possible valuation with probability 1. In this fully revealing equilibrium, the monopolist extracts all of the surplus, leaving consumers with nothing.

But there are also partially revealing equilibria in which only those types below a cutoff are revealed. For example, there is an equilibrium in which all types v such that $v \geq p^*$ send the message V and only types below p^* reveal themselves. As we show below, not only is this a consumer-optimal equilibrium, but the resulting payoffs for the consumer are also identical to those with uniform pricing.

Proposition 1. *With simple evidence, across all equilibria, the consumer's interim payoff is bounded above by $\max\{v - p^*, 0\}$.*

Thus, the consumer gains nothing, *ex ante* and *ex interim*, from the ability to disclose her type using simple evidence. We prove Proposition 1 by first observing that any equilibrium features a price, p_{ND} , that the monopolist charges when the consumer does not disclose her valuation. In equilibrium, any consumer whose valuation strictly exceeds p_{ND} must not disclose her type because disclosing her type results in a payoff of 0. Therefore, the set of valuations that do not disclose includes $(p_{ND}, \bar{v}]$. Given this observation, the non-disclosure price p_{ND} is at least p^* , establishing Proposition 1.

18. If the seller receives message m , his beliefs (both on- and off-path) have support contained in $\{v \in V : m \in \mathcal{M}(v)\}$.

This result suggests that consumer control in the form of track/do-not-track regulation does not benefit consumers in a monopolistic setting. There are equilibria in which consumers are much worse off and no equilibria in which the consumer is better off.

3.3. How rich evidence can help consumers

In this section, we show how giving the consumer the ability to disclose partial information about her type can make the consumer better off and result in Pareto gains. The idea is to use verifiable messages to segment the market into groups of consumers. The monopolist optimally charges each group a price equal to the minimum valuation among members of that group, so that all consumers purchase and the equilibrium allocation is efficient. The difficulty is in determining how the groups should be constructed so as to improve consumer welfare while allowing no profitable deviations. We use a “greedy” construction, in which we start with the consumers with the highest valuations and group as many of them together as possible subject to the constraint that the monopolist will not exclude anyone, and then proceed to the next group. This group segmentation based on verifiable disclosures often arises in practice; see Section 5 for examples.

To define the segmentation strategy, consider a sequence of prices $\{p_s\}_{s=0,1,2,\dots,S}$ where $S \leq \infty$, $p_0 = \bar{v}$, and for every s , where $p_{s-1} > \underline{v}$, p_s is the (lowest) maximizer of $p_s(F(p_{s-1}) - F(p_s))$. If $p_{s'} = \underline{v}$ for some s' , then we halt the algorithm and set $S = s'$; otherwise, $S = \infty$ and $p_\infty = \underline{v}$. Observe that $p_1 = p^*$ and by Assumption 1, $p_1 > \underline{v}$.

Proposition 2. *With rich evidence, there is a Pareto-improving equilibrium in which the consumer’s disclosure strategy is*

$$M^*(v) = \begin{cases} [p_s, p_{s-1}] & \text{if } v \in (p_s, p_{s-1}], \\ \{\underline{v}\} & \text{if } v = \underline{v}. \end{cases}$$

When receiving a message of this form, the firm charges a price of p_s . For any other feasible message M , the firm charges the price $\max_{v \in M} v$.

This segmentation generalizes Zeno’s Partition constructed in Section 2. The highest market segment comprises those consumer types whose valuations strictly exceed the monopolist’s optimal posted price, $p_1 = p^*$; these are the types who send message M_1 . The next highest market segment comprises those whose valuations exceed the optimal posted price, p_2 , for the *truncated distribution* that excludes the highest market segment; they send message M_2 . This iterative procedure continues either indefinitely (if $p_s > \underline{v}$ for every s) or halts once the monopolist has no incentive to exclude any type in the truncated distribution from trading. The idea that optimal prices need not be monotone in the firm’s beliefs plays an important role in this construction. When consumer types in $(p_s, p_{s-1}]$ are pooled, the monopolist prices at p_s . This ensures that even though low consumer types whose values are near p_s could separate from this larger pool, they have no incentive to do so.

Let us compare this equilibrium to the benchmark with uniform pricing. Consumer types in the highest market segment face a price of p^* but now consumers in other market segments can also purchase at prices that are (generically) below their willingness to pay. Firm profits are also increased and trading is efficient. Thus, the segmentation is a Pareto improvement.

How does this equilibrium compare with other equilibria of this game? As with simple evidence, there is a fully revealing equilibrium in which the consumer sends the message $\{v\}$ and purchases the good at a price that extracts all of her surplus. But, there also are many other equilibria. For example, one may interpolate between the fully revealing equilibrium and

that constructed in Proposition 2 to construct others. We simplify the comparison of equilibria by showing that from the standpoint of consumer surplus, we can restrict attention to efficient partitional equilibria. An equilibrium is *efficient* if trade occurs with probability 1. A collection of messages in \mathcal{M}^F denoted by \mathcal{P} is a partition of $[\underline{v}, \bar{v}]$ if $\bigcup_{m \in \mathcal{P}} m = [\underline{v}, \bar{v}]$ and for every distinct m, m' in \mathcal{P} , $m \cap m'$ is at most a singleton. For two messages m and m' , we say that $m \succeq_{\mathcal{M}} m'$ if for every $v \in m$ and $v' \in m'$, $v \geq v'$; argmin and argmax over a set of messages refers to this order. Given a partition \mathcal{P} , let $m^{\mathcal{P}}(v) \equiv \text{argmin}_{\{m \in \mathcal{P}: v \in m\}} m$. An equilibrium σ is *efficient partitional* if there is a partition \mathcal{P} such that the equilibrium message sent by type v , $m^{\sigma}(v)$ coincides with $m^{\mathcal{P}}(v)$, and for every m in \mathcal{P} , the equilibrium price charged by the firm $p^{\sigma}(m)$, is $\min_{v \in m} v$.

Proposition 3. *Given any equilibrium, there is an efficient partitional equilibrium that results in the same interim consumer surplus for almost every type.*

The proof proceeds in two steps. First, we show that any inefficient equilibrium can be converted into one that is efficient and results in the same interim consumer surplus for every type. The idea is that if a type is not purchasing the product, one could have that type fully reveals itself to the seller. Second, we show that for any efficient equilibrium, the set of types that are charged the same price is a connected set. This allows us to construct a payoff-equivalent efficient partitional equilibrium.

In comparing the equilibrium of Proposition 2 with other efficient partitional equilibria, we note that the equilibrium constructed in Proposition 2 is Pareto efficient: any partition that differs from it must raise the lowest type in at least one segment, which increases the price in that segment. However, it is not the unique interim Pareto efficient segmentation; for example, increasing the lowest type in the highest market segment, and then following a greedy construction on the truncated distribution would also result in a Pareto efficient equilibrium.

In terms of *ex ante* consumer surplus, the equilibrium constructed in Proposition 2 need not be optimal. A consideration that it ignores is that it may benefit average prices to exclude some high types from a pool, making those types pay a higher price, and pool intermediate types with low types. We illustrate this below.

Example 1. *Suppose that the consumer's type is drawn from $\{1/3, 2/3, 1\}$, where $\Pr(t=1) = 1/6$, $\Pr(t=2/3) = 1/3 + \epsilon$, and $\Pr(t=1/3) = 1/2 - \epsilon$, where $\epsilon > 0$ is small. The greedy construction sets the highest segment as $\{2/3, 1\}$ and the next segment as $\{1/3\}$. This segmentation results in an average price of approximately $1/2$. A better segmentation for *ex ante* consumer surplus involves the segmentation $\{\{1/3, 2/3\}, \{1\}\}$, which has an average price of $4/9$.*

Generally, the optimal segmentation can be formulated as the solution to a constrained optimization problem over partitions that minimizes the average price subject to the constraint that the monopolist finds it optimal to price at the bottom of each segment. The greedy algorithm offers a simple program where that constraint binds in each segment. Identifying necessary and sufficient conditions on distributions under which these constraints necessarily bind is challenging because it requires understanding in detail how sharply the monopolist's optimal price responds to truncating the distribution at different points. This exercise is difficult for distributions where we cannot solve for the optimal price in closed form.¹⁹ A class of distributions where a closed-form

19. In certain cases, one may be able to show that the greedy algorithm is optimal without solving for the closed-form; for example, we show in the Supplementary Appendix that it is optimal if F is convex, and the optimal price on an interval $[\underline{v}, \bar{v}]$ —denoted by $p(\bar{v})$ —has a slope bounded above by 1 and is weakly concave.

solution is available is that of power distributions; for this class, the greedy algorithm identifies the consumer-optimal segmentation.

Proposition 4. *Suppose that $[\underline{v}, \bar{v}] = [0, 1]$ and the cdf on valuations, $F(v) = v^k$ for $k > 0$. Then the greedy segmentation is the consumer-optimal equilibrium segmentation.*

3.4. Partial disclosure and group pricing: how rich must evidence be?

In contrasting the positive message of Proposition 2 with the negative conclusion of Proposition 1, one may wonder how rich the evidence must be for consumers to gain. Here, we show that the evidence technology need not be so rich: Pareto improvements may obtain with a coarse technology where low-value consumers can send a message that high-value consumers cannot. Such equilibria resemble group-pricing schemes seen in practice where low-value consumers verifiably disclose that they belong to a group (e.g. students, retirees, or low-income households) that excludes high-value consumers.

To describe these schemes formally, recall that $\mathcal{M}(v)$ denotes the set of messages available to type v . Unlike our simple and rich evidence game, $\mathcal{M}(v)$ is a primitive of this framework. Accordingly, we assume that $\mathcal{M}(v)$ is well-behaved in that it is upper hemicontinuous in v and contains for each v the completely uninformative message, $[\underline{v}, \bar{v}]$.²⁰ We say that an evidence technology *facilitates group pricing* if there is a message \bar{M} that (1) cannot be sent by any consumer type whose valuation strictly exceeds p^* and (2) can be sent by a positive measure of types. We show that this form of evidence can lead to Pareto improvements relative to the benchmark of uniform pricing.

Proposition 5. *If an evidence technology facilitates group pricing, then there is a Pareto improving equilibrium.*

The idea of Proposition 5 is that a set of low-value consumers can use message \bar{M} to verifiably separate themselves from others. All other consumers send the uninformative message $[\underline{v}, \bar{v}]$. The monopolist charges a price of p^* to all consumers but if the monopolist receives the message \bar{M} , she offers a group discount to these low-value consumers. Such discounts induce those low-value consumers to purchase without hurting other consumers; it also increases the probability of trade. Therefore, this is a Pareto improvement relative to uniform pricing.

The logic of this argument highlights that we have imposed stronger than necessary conditions for this result. It is not necessary that every consumer type whose valuation exceeds p^* be unable to send \bar{M} ; the result obtains so long as the market segment that can send message \bar{M} induces a price below p^* , and that the segment that constitutes its complement does not induce a price above p^* . This is compatible with some high-value types being able to send message \bar{M} . Accordingly, in the Appendix, we prove Proposition 5 under these weaker conditions.

Schemes of this form arise in various contexts. It is often the case that buyers who do not disclose any evidence obtain the sticker price p^* , but disclosing evidence in the form of, for example, an Electronic Benefits Transfer (EBT) card, student ID, or senior citizen or Medicare card results in a group discount. In these contexts, a third party (a government, university, or association) verifies characteristics that would make one, on average, a low-value consumer. One may also interpret financial aid applications for private schools as being a form of disclosure that is used to separate types from high-value consumers and lead to a personalized discount.

20. These assumptions were satisfied for both the simple and rich evidence game.

3.5. Extensions

3.5.1. Limited control over information disclosure. In practice, firms may have access to their own forms of consumer information, such as records of the purchases that the consumer has made at that firm. In this sense, the consumer cannot truly prevent the firm from knowing anything about her type by choosing not to disclose.²¹ However, if the consumer has a good sense of what it is that the firm knows, then our results from Section 3 still apply.

To see this formally, consider an extension of our main model in which prior to any interaction, the firm has access to a signal x of the consumer's type. This signal may be multidimensional and complex, and in practice may correspond to the types of data that are captured by cookies from different platforms. We assume that the consumer valuation v and firm signal x are drawn from a joint distribution. To avoid unnecessary complications, we assume that v and x have a well-defined joint density, so that upon observing x , the firm is able to update its beliefs about the consumer's type to an atomless posterior cdf $F(v|x)$. This inference could be as simple as restricting the domain of v or as complex as learning a new conditional density for v for every "feature vector" x using a machine learning algorithm.²²

The game is now as follows. The consumer sends message $m(v, x) \in \mathcal{M}(v)$, where the choice of message may now depend on the commonly observed signal x . In the absence of personalized pricing, the firm chooses a price $p > 0$ that maximizes $p(1 - F(p|x))$. As before, we denote the (lowest) revenue-maximizing price by $p^*(x)$. When x is highly informative, the posterior distribution of v conditional on x may already induce a targeted price $p^*(x)$ that both increases the probability of purchase and reduces the consumer's *ex interim* surplus.

However, when there is still some exclusion, evidence technologies may be helpful. We formalize the exclusion assumption as follows:

Assumption 2. *For all x , the monopolist's optimal uniform price $p^*(x)$ is strictly higher than the lowest type \underline{v} in the support of the posterior distribution of v conditional on x , $F(v|x)$.*

This assumption is an extension of Assumption 1 above to hold pointwise in x . Given this assumption, we can make all the same arguments as in Propositions 1, 2, and 5 pointwise. Specifically, since the posterior cdf $F(v|x)$ satisfies the same conditions as the prior $F(v)$, the arguments in Section 3.2–3.4 hold. Thus, simple information cannot improve the consumer's interim payoff, but a rich information evidence structure can—whether through group messaging as in Proposition 5 or greedy messaging as in Proposition 2. Of course, this does not mean that the signal x has no impact. The posterior $F(v|x)$ may concentrate the mass of possible valuations in a way that elicits prices that are much higher or lower prices than what the monopolist would set with no information or disclosure. Indeed, it may be that for some x , Assumption 2 no longer holds, and so there cannot be any benefit to disclosure for individuals with that signal once it is observed. However, as Propositions 1, 2, and 5 hold for every x , our main results also hold *ex ante* (unconditional of x) so long as there is a positive mass of x that satisfies Assumption 2.

3.5.2. Robustness considerations. In this section, we discuss the degree to which the analysis here is robust or fragile to various issues abstracted from in our baseline analysis.

21. The firm may also learn about the consumer's information from other people's information, a point emphasized by the study of data externalities (Choi *et al.*, 2019; Acemoglu *et al.*, 2021; Bergemann *et al.*, 2021; Ichihashi, 2021).

22. See Bajari, Nekipelov, Ryan and Yang (2015) for an example of how machine learning can be used to derive posterior beliefs over consumers' willingness to pay.

Disclosure costs. One may envision disclosing evidence to be intrinsically costly, either as a pure cost or reflecting the cost of ceding one's privacy. To capture disclosure costs, suppose that the cost of sending the uninformative message $[\underline{v}, \bar{v}]$ is 0 but sending any other message involves a cost $c > 0$. In this setting, neither simple nor rich evidence could induce a consumer welfare improvement over uniform pricing: the unique equilibrium is for every type to send the uninformative message. The issue is that if in an equilibrium, consumer types were to disclose evidence and form a market segment, the monopolist would set a price so that at least some types in that segment obtain a price within c of their valuation. Such types would then have a strictly profitable deviation to concealing evidence. Therefore, even slight disclosure costs induce complete pooling and eliminate potential gains from voluntary disclosure. For this reason, it may be important to have a zero-disclosure-cost intermediary manage disclosure on behalf of consumers.

Concealment costs. Now suppose that it is costly for the buyer to "opt out" from disclosing evidence. With simple evidence, modelling this cost is straightforward: sending message $\{v\}$ is free but sending message $[\underline{v}, \bar{v}]$ costs $c > 0$. In this case, the unique equilibrium is the fully revealing equilibrium. The idea is that in any strategy profile in which a positive mass of types are opting out, the monopolist is setting a non-disclosure price p_{ND} that is within c of the valuation of some types that are opting out. Such types obtain a strictly negative equilibrium payoff and are better off revealing themselves. With rich evidence, it is less apparent as to how to model opt-out costs but, for simplicity, suppose that the only message that is costly is $[\underline{v}, \bar{v}]$. In this case, the equilibrium constructed in Proposition 2 survives since in equilibrium, no type opts out. Therefore, opt-out costs may amplify the difference between simple and rich evidence.

Dye-evidence structures. Suppose that with a positive chance, as in Dye (1985), the buyer lacks evidence to disclose, and is then forced to send the uninformative message $[\underline{v}, \bar{v}]$. In the simple evidence game, Proposition 1 would remain true: no equilibrium could improve upon the fully concealing equilibrium. In the rich evidence game, Proposition 2 remains completely unchanged as none of the consumer types who obtain evidence have any incentive to mimic those who do not.

3.5.3. Multi-dimensional types. Although we present our analysis in a unidimensional type space for simplicity, much of it applies to a multidimensional setting where the consumer's valuation depends on multiple characteristics. Suppose that the consumer's type t is drawn according to a measure μ with support T , which is a convex and compact subset of a finite-dimensional Euclidean space, \mathfrak{R}^k . Her valuation is specified by $v: T \rightarrow \mathfrak{R}$, where we assume that valuations are non-negative and v is quasiconvex. A special case is $v(t) = \sum_{i=1}^k \beta_i t_i$, where β_i reflects the degree to which attribute i influences her valuation. We denote by F the induced CDF over valuations, by \underline{v} and \bar{v} the highest and lowest valuations in the support, and we assume that F is atomless. In this setting, the set of all feasible messages is $\mathcal{M}^{\mathcal{F}} \equiv \{M \subseteq T : M \text{ is closed and convex}\}$. The set of messages available to type t is $\mathcal{M}(t) \subseteq \mathcal{M}^{\mathcal{F}}$. A simple evidence technology corresponds to $\mathcal{M}(t) = \{T, \{t\}\}$ for every t . A rich evidence technology corresponds to $\mathcal{M}(t) = \{M \in \mathcal{M}^{\mathcal{F}} : t \in M\}$ for every t . In the Supplementary Appendix, we show that the analogues of Propositions 1, 2, and 5 hold in this setting: namely, simple evidence is ineffectual, there is a greedy construction analogous to Zeno's partition, and a mild condition on the evidence technology guarantees the existence of Pareto improvements.

4. HOW DISCLOSURE AMPLIFIES COMPETITION

Our analysis above studies when consumers benefit from voluntary disclosure and personalized pricing when interacting with a monopolist. With simple evidence, consumers never benefit from

voluntary disclosure and are potentially harmed by it. With rich evidence, there is the potential for consumers to benefit. Both results rest on the idea that disclosure can generate consumer gains only if it is possible for consumers with similar valuations to pool together and obtain lower prices than they would have otherwise. As there are always fully revealing equilibria in which the consumer may be worse off relative to uniform pricing, however, these consumer gains are not assured.

This section investigates how voluntary disclosure and personalized pricing benefit consumers in competitive settings with differentiated products. Our analysis here focuses on a different force: *voluntary disclosure amplifies market competition*. As we articulate below, the rationale for voluntary disclosure is stronger than in the monopolistic setting. Consumers may benefit even with simple evidence, and even the worst equilibrium in the disclosure model leads to a higher (average) consumer surplus than the equilibrium with uniform pricing.

One of the new possibilities that emerges with competition is that consumers may send different messages to different firms. As it is easier to think of cases where consumers allow tracking by one firm and not another, as opposed to providing asymmetric pieces of rich evidence, we will emphasize simple evidence more in this section.

For expositional clarity, our main text focuses on the case of two firms with horizontal differentiation; the Supplementary Appendix shows that this analysis extends to more than two firms and general product differentiation.

4.1. *Bertrand duopoly with horizontal differentiation*

Firms L and R compete to sell to a single consumer with unit demand. The consumer's type is her *location*, denoted by t , which is drawn according to measure μ (and cdf F) with support T . We assume that $T \equiv [-1, 1]$ and that F is atomless with a strictly positive and continuous density f on its support. The firms L and R are located at the two end points, respectively -1 and 1 , and each firm i sets a price $p_i \geq 0$. The consumer has a value v for buying the good that is independent of her type t and faces a "transportation cost" when purchasing from firm i that is proportional to the distance between her location and that of the firm, ℓ_i . Thus, her payoff from buying the good from firm i at a price of p_i is $v - |t - \ell_i| - p_i$. We assume that v is sufficiently large that in the equilibria we study below, all types of the consumer purchase the good and no type is excluded from the market.²³

After observing her type, the consumer chooses a message M that is feasible and available for her to send to each of the firms. The set of feasible messages is $\mathcal{M}^{\mathcal{F}} \equiv \{[a, b] : -1 \leq a \leq b \leq 1\}$, where the message $[a, b]$ is interpreted as "my type is in the interval $[a, b]$ ". When the consumer's type is t , the set of messages that she can send is $\mathcal{M}(t) \subseteq \mathcal{M}^{\mathcal{F}}$.²⁴ We study two disclosure technologies:

- *simple* evidence messages for each type t , $\mathcal{M}(t) = \{-1, 1, \{t\}\}$.
- *rich* evidence messages for each type t , $\mathcal{M}(t) = \{[a, b] : a \leq t \leq b\}$.

Each evidence technology is identical to its counterpart in the monopolistic model. The main departure is that the consumer now sends two messages— M_L to firm L and M_R to firm R —and

23. For most of our analysis, it suffices for $v \geq 2$, so that a consumer is always willing to purchase the good from the more distant firm at a price of 0.

24. While we have treated t as the consumer's location, our analysis is also compatible with a setting where all that a consumer observes is a signal with her posterior expected location, like Armstrong and Zhou (2022), and chooses whether and how to disclose that expected location using simple or rich evidence.

each message is privately observed by its recipient. Both messages come from the same technology but are otherwise unrestricted. For example, a consumer of type t can reveal her type by sending the message $\{t\}$ to one firm while concealing it from the other firm using the message $[-1, 1]$. Our interpretation of being able to send two distinct messages is that it models consumers who can choose to have one firm track her without being forced to allow the other firm to also do so.

The timeline is: (1) the consumer observes her type t and chooses a pair of messages (M_L, M_R) , each from $\mathcal{M}(t)$; (2) firm i privately observes message M_i and sets price $p_i \geq 0$; (3) the consumer chooses which firm to purchase the good from, if any. We study pure strategy PBE in which the consumer breaks ties in favour of the closer firm, and henceforth refer to these as equilibria.

4.2. Constructing equilibria with simple and rich evidence

This section constructs equilibria of the disclosure game with simple and rich disclosure technologies for any distribution of consumer types. In both cases, we use the following strategic logic. Each consumer communicates to the firm that is more distant from her that she is “out of reach.” This distant firm then competes heavily for her business by setting a low price, which in equilibrium equals 0. Knowing that the consumer is obtaining a highly competitive price from the distant seller compels the closer seller to lower its price. With simple evidence, the consumer sends distinct messages to each firm; in equilibrium, each firm infers if it is closer or distant based on the message it receives. With rich evidence, it suffices for the consumer to send the same message to both firms.

We begin our analysis with a fully revealing equilibrium and then show how to improve upon it.

Proposition 6. *There is a fully revealing equilibrium in both simple and rich evidence games: each consumer type t sends message $\{t\}$ to each firm and purchases from the closer firm at a price of $2|t|$.*

Here is the logic. In an equilibrium where the consumer reveals her location to each firm, both firms do not charge her strictly positive prices. Standard Bertrand logic implies that the distant firm charges a price of 0 and the closer firm charges a price to make the consumer just indifferent; this tie is broken in favour of the closer firm.²⁵ If the consumer deviates by sending a message M that isn't a singleton to firm i , then the firm i believes that the consumer's type is the one in M closest to ℓ_i and that the consumer has revealed her location to firm j . This equilibrium involves each seller holding sceptical beliefs that the consumer is as close as possible (given the message that is sent).

This fully revealing equilibrium serves central types very well because they benefit from intense price competition. However, extreme types may suffer from the firm closer to them being able to charge a high price. Ideally, these extreme types would benefit from pooling with central types. The next result uses simple evidence to construct a partial pooling equilibrium that improves upon the fully revealing equilibrium for a strictly positive measure of types without making any type worse off.

Let p_1^i be the lowest maximizer of $p\ell_i(F(\ell_i) - F(p\ell_i/2))$, and let $t_1^i \equiv p_1^i \ell_i/2$. Observe that p_1^i is the (lowest) optimal price that firm i charges if he has no information about the consumer's type and firm j charges a price of 0. If firm j charges a price of 0 and firm i charges a price of $p > 0$, firm i expects to sell to the consumer with probability $\ell_i(F(\ell_i) - F(p\ell_i/2))$. When firm

25. Once types are revealed, these equilibrium prices coincide with those of Thisse and Vives (1988), where the consumer's type is commonly known.

i charges p_1^i , t_1^i is the most central type that purchases from firm i . We use these prices and types to describe an equilibrium in the setting with simple evidence. Moreover, this is the best pure-strategy equilibrium in this setting.

Proposition 7. *With simple evidence, there is an equilibrium in which the consumer's disclosure strategy is*

$$(M_L^*(t), M_R^*(t)) = \begin{cases} ([-1, 1], \{t\}) & \text{if } -1 \leq t \leq t_1^L, \\ (\{t\}, \{t\}) & \text{if } t_1^L < t < t_1^R, \\ (\{t\}, [-1, 1]) & \text{if } t_1^R \leq t \leq 1, \end{cases}$$

and the prices charged by firm i are

$$p_i^*(M) = \begin{cases} \max\{2t\ell_i, 0\} & \text{if } M = \{t\}, \\ p_1^i & \text{otherwise.} \end{cases}$$

In equilibrium, each consumer type purchases from the seller closer to her. This equilibrium results in a weakly higher payoff than any other pure-strategy equilibrium for each consumer type.

We view Proposition 7 as the main positive result for competitive markets: a simple track/do-not-track technology can make consumers better off when there is competition and horizontal differentiation. We depict this equilibrium for the special case of a uniform distribution in Figure 2 on p. 11. In this equilibrium, a centrally located consumer allows both firms to track her. Such consumers benefit from intense price competition. If the consumer is not centrally located, she allows the distant firm to track her but not the closer one. This guarantees that the distant firm offers an attractive price to poach her business. The firm that receives an uninformative (“do not track”) message infers that the consumer is located sufficiently close but does not learn where. That firm then chooses an optimal local monopoly price given the outside-option price of zero. By pooling extreme types with moderate types, this reduces prices relative to the fully revealing equilibrium. There may be multiple equilibria that feature such pooling; however that constructed in Proposition 7 has weakly lower prices for every type than any other pure-strategy equilibrium.

One facet of Proposition 7 is that the consumer exploits her ability to make different privacy choices with respect to each firm, allowing one to track her but not allowing the other to do so. This is an important feature since it permits a consumer with strong loyalties to one firm to attract low prices from the competing firm while maintaining some information rents with respect to the firm from which she eventually purchases the product. If it were not possible to make different privacy choices, and the consumer had to reveal her information either to both firms or no firms, then it would be difficult for extreme L -types to separate themselves from extreme R -types; in equilibrium, both firms might then charge high prices to these types. The ability to selectively disclose information is thus central to this analysis.

Consumer control with more sophisticated evidence technology can do better. Let us see how with rich evidence. Consider a segmentation analogous to Zeno's Partition but where type $t=0$ plays a role similar to the lowest type in the monopolistic setting. Define a sequence of types $\{t_s^i\}_{s=0,1,2,\dots}$, prices, and messages $\{p_s^i, M_s^i\}_{s=1,2,\dots}$ where for every firm i in $\{L, R\}$:

- $t_0^i = \ell_i$ and for every $s > 0$, $t_s^i = p_s^i \ell_i / 2$.
- p_s^i is the lowest maximizer of $p\ell_i(F(t_{s-1}^i) - F(p\ell_i/2))$.
- $M_s^i \equiv \{t \in [-1, 1] : t_s^i \ell_i \leq t \ell_i \leq t_{s-1}^i \ell_i\}$.




FIGURE 3

The figure shows a segmentation using rich evidence. The types in (t_3^L, t_3^R) are partitioned into countably infinitely many segments, and hence these segments are omitted.

Let $p_\infty^i = 0$ and let $M_\infty^i = \{0\}$. We use this notation to construct an equilibrium.

Proposition 8. *With rich evidence, there is an equilibrium in which the consumer’s disclosure strategy is to send message $M^*(t)$ to both firms where*

$$M^*(t) = \begin{cases} M_s^i & \text{if } t_s^i \ell_i < t \ell_i \leq t_{s-1}^i \ell_i \\ M_\infty^i & \text{if } t = 0. \end{cases}$$

When receiving an equilibrium disclosure of the form M_s^i , firm i charges a price of p_s^i and firm j charges a price of 0.

We depict this equilibrium in Figure 3. The left-most segment is M_1^L and comprises those consumer types who prefer purchasing the good from firm L at a price of p_1^L rather than from firm R at a price of 0. The price p_1^L is firm L ’s (lowest) optimal “local monopoly” price when it anticipates that firm R is charging 0. The construction proceeds by then removing the segment M_1^L , finding firm L ’s optimal monopoly price on the truncated distribution, and characterizing the set of types, M_2^L , willing to pay that price. We then proceed iteratively to 0. We follow the same procedure on types in $(0, 1]$ to find the right-most segment, M_1^R , proceeding iteratively to 0. This procedure generates an equilibrium, coupled with a “sceptical” off-path belief system. Notice that this equilibrium mirrors the construction of Zeno’s Partition in Proposition 2 but where the centrally located type $t = 0$ is the analogue of the consumer type whose valuation for the good is 0 in the monopolist setting.²⁶

The rich evidence game has other equilibria, including a fully revealing equilibrium. The equilibrium constructed in Proposition 8 uses a greedy algorithm, like our construction in Proposition 2. Our earlier results suggest that among pure-strategy equilibria, it is always interim Pareto efficient (for the same reasons as Proposition 3) and *ex ante* consumer optimal whenever the greedy algorithm is an optimal segmentation (e.g. for the uniform distribution). But our focus is not on whether this is the best equilibrium but whether and how it improves on uniform pricing, which we turn to next.

4.3. Benefits of personalized pricing in competitive markets

We compare our analysis above with the benchmark model where each firm i sets a uniform price p_i . Unfortunately, this benchmark model may lack a pure-strategy equilibrium (in prices), and there is no general characterization of the mixed strategy equilibria of this game. Accordingly,

26. We observe that although consumers share the same information with firms on the equilibrium path, this equilibrium construction nevertheless uses the property that each firm privately observes its message from the firm. Were the consumer to send a single message that is publicly observed by both firms, some consumer types may find it profitable to send off-path messages.

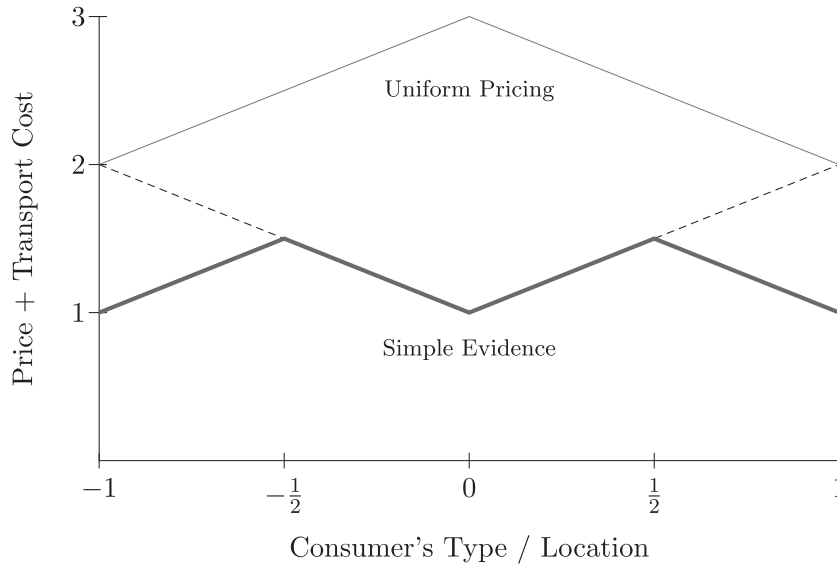


FIGURE 4

The figure compares the interim equilibrium cost (including price and transport cost) in the setting without personalized pricing with that of the equilibrium constructed in the simple evidence game (Proposition 7) for uniformly distributed types. Simple evidence reduces the expected cost by 50%.

we proceed with the approach standard in this literature, which is to assume that f is symmetric around 0 and is strictly log-concave. This assumption guarantees the existence and uniqueness of a symmetric pure strategy equilibrium (Caplin and Nalebuff, 1991) and is compatible with many standard distributions (Bagnoli and Bergstrom, 2005). In the symmetric pure-strategy equilibrium, each firm charges a price of $p^* = 1/f(0)$.²⁷ We compare this price to those of the equilibria constructed above.

Proposition 9. *If f is symmetric around 0 and log-concave, then every consumer type has a strictly higher payoff in the equilibria of the simple and rich evidence games constructed in Propositions 7 and 8 than in the benchmark setting with uniform pricing.*

The logic underlying Proposition 9 is that the price in the benchmark setting (p^*) is strictly higher than p_1^i , the price firm i charges to a consumer who conceals her type from firm i in the equilibrium of the simple evidence game (Proposition 7). The consumer must then be better off because this price (p_1^i) is strictly higher than all other equilibrium path prices both in this equilibrium and in the equilibrium that we construct in the rich evidence game. We illustrate the welfare gains from simple evidence in Figure 4 for the case of the uniform distribution.

Pooling facilitates interim Pareto gains that may not emerge in the fully revealing equilibrium: if $f(0) > 1/2$, extreme types pay a price of 2 in the fully revealing equilibrium but the uniform prices in the benchmark model are strictly less than 2. Our equilibrium constructions avoid this challenge by pooling these extreme types with moderate types so as to reduce the price they pay relative to the fully revealing equilibrium. Even simple evidence can induce beneficial pooling in this sense.

27. We assume that V is sufficiently high that all consumers purchase at these prices. It suffices that $V > (f(0))^{-1} + 1$.

The reader may wonder about other equilibria: might it be that consumers are worse off in other equilibria from voluntary disclosure and personalized pricing relative to uniform pricing? While it is infeasible to specify consumer surplus for all types across all equilibria, we argue that this is not the case in terms of *average* consumer surplus for a class of equilibria. This contrasts with the monopolist setting where the consumer may be worse off from the ability to disclose information.

Let us describe our argument. An equilibrium is efficient if, in equilibrium, a consumer purchases the good from the closer firm. An equilibrium involves *passive beliefs* if whenever firm i obtains an off-path message of the form $\{t\}$, firm i assumes that the consumer sent the *equilibrium path* message to the other firm. Passive beliefs are a standard restriction for games in which a single agent is interacting privately with a number of firms.²⁸ For brevity, we call such an equilibrium passive.

Proposition 10. *Suppose f is symmetric around 0 and log-concave. With both simple and rich evidence, every efficient passive equilibrium has higher average consumer surplus than the benchmark setting with uniform pricing.*

To prove this claim, we show that in every passive equilibrium, no consumer type's payoff is lower than the fully revealing equilibrium described in Proposition 6.²⁹ We then use a result of Rhodes and Zhou (2021), who show that under log-concavity, this equilibrium results in a higher average consumer surplus than uniform pricing.

A takeaway from these results is that under relatively weak conditions, consumers benefit from how voluntary disclosure amplifies competition. Even with simple evidence and relative to uniform pricing, every consumer type benefits in an equilibrium of the game, and average consumer surplus is higher even in the worst equilibrium (within a natural class).³⁰ Our analysis emphasizes the importance of allowing the consumer to selectively disclose information to firms. This idea builds on existing regulatory regimes that require each firm to individually obtain a consumer's consent to track her. Such consumer control would be impeded by centralized data brokers that may sell consumer data to multiple firms without the consumer's consent, in which case consumers may be unable to obtain highly competitive prices. Accordingly, assuring these gains requires the protection of consumer data so that it cannot be simultaneously shared with multiple firms without the consumer's permission.

5. VERIFIABLE VOLUNTARY DISCLOSURE IN PRACTICE

Our results throughout this article have relied on consumers having access to a verifiable disclosure technology that allows them to opt out of disclosing information if it benefits them to do so. In this section, we consider several examples of verifiable disclosure mechanisms that are currently in use. Our examples emphasize who administers the disclosure mechanism. This is because the administrator controls which messages can be (credibly) sent and may therefore exert some

28. For example, see Cremer and Riordan (1987), Horn and Wolinsky (1988), McAfee and Schwartz (1994), Rey and Vergé (2004), and Collard-Wexler, Gowrisankaran and Lee (2019).

29. Because this step involves issues relating to firm best-responses with continuous prices, we use a discrete price grid and make the grid arbitrarily fine.

30. An assumption that plays a role in our results is that the consumer's value for the object is sufficiently high that there is full coverage. Although we do not investigate how these results extend to settings without full coverage (and where there may be some captive consumers), we would expect that once transportation costs are sufficiently large, the setting converges to that of a monopoly, where every consumer is captive to a firm.

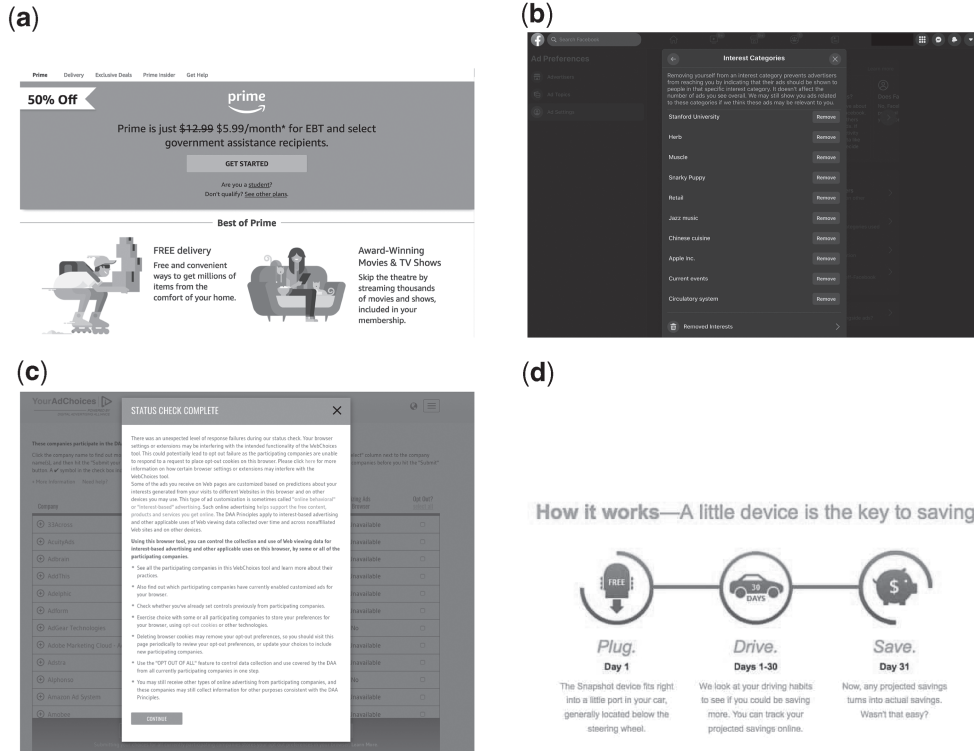


FIGURE 5

Screenshots from examples of verifiable disclosure interactions as of 2021. (a) EBT discount for Amazon Prime. (b) Facebook opt-out menu for interest-based ads. (c) Opt-out menu for Google advertisers. (d) Overview of an auto-insurer's monitoring program.

control over the equilibrium that arises. Since there are often multiple equilibria of the disclosure game, achieving the consumer-optimal equilibrium may require an intermediary to facilitate coordination.

As different types of administrators may prioritize different outcomes, we group our examples accordingly. Despite the differences, we argue that in each case, the administrator has palpable incentives to promote consumer welfare—and is likely to continue to as technology advances and richer verification becomes more universally accessible.

5.1. Government verification

A natural administrator for the verification of hard information is a local or federal authority. Governments routinely collect highly sensitive and detailed personal information and provide credible attestations like ID cards and immunization certificates that can be used to access goods and services. In many cases, government-vetted signals can also be used to obtain personalized pricing—typically either through group verification or full disclosure (as with a simple evidence technology). For instance, low-income individuals in the US can present an EBT card (used for food stamps) or Medicaid enrolment card to obtain steep discounts on everything from phone and internet service to Amazon prime membership. Figure 5(a) presents an example of a 50%

discount on Amazon Prime that is offered upon proof of an EBT card.³¹ Students hoping to attend elite universities can submit their parents' complete tax returns in order to obtain personalized financial aid offers. These schemes already mirror the group-pricing schemes that we discuss in Section 3.4. As record-keeping and commercial interactions become increasingly digital, it seems likely that opportunities to allow for richer forms of government-verified signals will grow. Government consumer protection bureaus may thus be well positioned to create tools for verified disclosure that increase consumer welfare.

5.2. *Private intermediary verification*

Third party intermediaries already provide verification for private information. For example, credit rating agencies provide discrete grades or score ranges that consumers use to secure loans. Banks provide verification of asset holdings, used for a range of durable financial arrangements. More recently, firms that observe granular user behaviour through their products such as Google³², Facebook,³³ and General Motors³⁴ use their data to group consumers into pools, which allows advertisers to target their ads and promotions to these groups.

As consumer concerns about privacy as well as regulatory oversight has grown—particularly in the tech sector, where “Big Tech” has faced increasing congressional scrutiny—intermediaries have slowly increased consumer control over what information is shared, and with which advertisers/sellers. For instance, as of October 2021, Facebook allows users to opt out of the groups they have been assigned to (called “interests”) that are used to generate personalized ads/discounts. Google allows consumers to opt out of engaging with specific advertisers. Figure 5(b) and (c) demonstrates how these choices look from the consumer's perspective. Meanwhile, in Q1 of 2021, Apple set the default option on its products to opt out of third party personalized advertising so that users need to actively opt in in order to be tracked.³⁵ This is evidence of a trend towards giving consumers increasing control over what they share and with whom.

Most intermediaries have two potentially competing design goals. On the one hand, they need user engagement with their products in order to thrive, and so they have an incentive to design these pools in such a way that consumers see relevant ads and get good promotions. On the other hand, they want to segment the market in a way that is informative to advertisers and firms, so they can sell them services. As the demand for consumer control from users and regulators grows, some intermediaries may choose to align their incentives with consumers as a competitive move to attract a wider user base. The resulting design problem closely resembles the ones studied in the article. Sections 3.3, 3.4, and 4.3 show that allowing users to segment themselves by valuation blocks—as with a rich evidence structure—can be an effective tool to generate effective personalized pricing when interacting in monopolistic settings. Indeed, Proposition 1 argues that rich segmentation along these lines is necessary to provide a benefit to consumers. In competitive settings (Section 4), intermediaries may see further gains from allowing consumers to control not only what data firms see but also which firms see it.

Moreover, our extensions discussed in Section 3.5.2 highlight an interesting subtlety as private intermediaries move from opt-out to opt-in policies for tracking (as has been the case for Apple).

31. See <https://lowincomerelief.com/ebt/> for other detailed examples.

32. <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox>.

33. <https://www.facebook.com/business/ads/ad-targeting>.

34. GM observes both GPS and radio browsing data. See <https://www.theverge.com/2018/10/17/17990052/gm-radio-listen-tracking-habits-advertising-future>.

35. See <https://www.cnet.com/tech/services-and-software/this-one-iphone-setting-change-keeps-ads-from-following-you-across-the-internet> for details.

Our analysis of concealment costs suggests that in markets where there is a monopolist and the evidence is simple, opt-out policies would involve no one opting out. Shifting to opt-in policies, by contrast, results in no one opting in, and leads to complete pooling (as suggested by our analysis of disclosure costs). Thus, the shift from opt-out to opt-in would improve consumer surplus, but nevertheless does not lead to any gains beyond uniform pricing, leaving “money on the table”.

5.3. Firm verification

In some cases, the seller may provide a mechanism for verifying consumers’ behaviour himself. For example, as of 2017, auto-insurers serving over 60% of drivers in the US offered a program through which consumers could have their driving behaviour tracked in order to prove that they warrant a discounted price.³⁶ Figure 5(d) summarizes an example of a tracking program of this sort as it is presented to consumers.

This is an example of the simple evidence message technology of Section 4: competitive consumers may voluntarily opt-in to fully disclosing their driving patterns—and consequently, the value of their insurance contracts. The remaining consumers do not disclose anything and receive a price based on public information as in Section 3.5.1. Although the economics of this setting is more complex than our present model allows—monitoring has been shown to not only deliver information but also to reduce moral hazard—we view this as a nice example of how firms may offer consumers the opportunity to self-disclose in a competitive setting.

In practice, most consumers opt-out of tracking and purchase insurance at a price that is competitive with other firms based on public information. The consumers who stand to benefit from tracking discounts then opt-in, share their full driving data with the insurer, and receive a substantial share of the efficiency rents generated by the tracking program. Although these programs vary in implementation, recent studies have suggested that—given competitive pressures and regulatory constraints—these programs have been broadly beneficial for consumers. Jin and Vasserman (2021) show that the introduction of a monitoring program in their setting induced an *ex ante* Pareto improvement for consumers,³⁷ while Reimers and Shiller (2019) show that profit advantages from tracking programs have eroded over time as entrants adopt tracking programs of their own.

6. CONCLUSION

As the digital economy matures, policymakers and industry leaders are working to establish norms and regulations to govern data ownership and transmission. In light of the privacy concerns and the distributional concerns that this issue raises, we set out to study the question: *do consumers benefit from personalized pricing when they have control over their data?* Although this question looms large in regulatory and public debates, relatively little has been done to conceptualize how it manifests in equilibrium. We argue that the answer depends on both competition and the disclosure technology. Under monopoly, disclosure benefits consumers only when the disclosure technology allows low-valuation consumers to pool together and receive a group price. By contrast, under oligopoly, disclosure intensifies competition and is generally beneficial.

Our model is designed to analyse a wide class of disclosure technologies. Many specific instantiations already exist. Track/do-not-track is a part of GDPR; a simple slider system in which consumers report the group or range that they belong to on different dimensions is not far

36. See the 2017 annual report of the National Association of Insurance Commissioners for details.

37. Jin and Vasserman (2021) show that consumers who entered a tracking program earned substantial discounts on average, while consumers who chose not to enter the program faced an unchanged menu of prices.

from the privacy control menus offered by Google and Facebook. Giving consumers control over their data when such technologies are available can generate substantial gains to welfare.

One challenge of implementing consumer control is assuring consumers that their data is not leaked across markets and firms. If a consumer is buying a range of goods and services, she may worry that information that she discloses in one market is leaked to another market. In a competitive market, if a consumer's data are being acquired by a single firm, that firm may not have incentives to share that information with competing firms; by contrast, if the data are being acquired by a data broker, the broker may sell that information to multiple firms within the market. In the latter case, the consumer may lose her ability to selectively disclose information—which is central to the consumer gains we outlined in Section 4. In this light, our analysis suggests the importance of coupling consumer control within a market with the protection of that data across markets.

As we show, consumer gains are also contingent on which equilibrium is selected. This gives rise to a coordination problem. Government agencies, third-party intermediaries, and sellers themselves are well positioned to resolve this coordination problem through their choice of message space. These parties are not necessarily incentivized to act on behalf of consumers or society at large, and regulation may be necessary to correctly align incentives. In this light, we hope that our article may provide some guidelines for when and how technologies that enable voluntary disclosure can be beneficial to the public interest.

Acknowledgements. We thank the co-editor and six referees for their constructive feedback. We have benefited from suggestions made by Wioletta Dziuda, Tan Gan, Nima Haghpanah, Zi Yang Kang, Navin Kartik, Vijay Krishna, Wenhao Li, Alessandro Pavan, Martin Peitz, Doron Ravid, Andy Skrzypacz, Yossi Spiegel, Nikhil Vellodi, Jidong Zhou, and various conference and seminar audiences. We thank Microsoft Research for hosting Ali and Vasserman. Xiao Lin provided excellent research assistance.

A. APPENDIX

Proof of Proposition 1. Consider an equilibrium. Let \tilde{V} be the set of types that in equilibrium send the non-disclosure message V . Thus, every type in $V \setminus \tilde{V}$ sends a message that fully reveals itself. Sequential rationality demands that the monopolist charges a price that leads to an interim payoff of 0 for each such type. We prove below that the non-disclosure message induces a price that is no less than p^* .

Suppose towards a contradiction that it leads to a price $\tilde{p} < p^*$. If the consumer's value v exceeds \tilde{p} , she must send the non-disclosure message V (because the message $\{v\}$ leads to a payoff of 0). Therefore, $\tilde{V} \supseteq \{v \in V : v > \tilde{p}\} \supseteq \{v \in V : v \geq p^*\}$. By charging a price of \tilde{p} , the firm's payoff is

$$\begin{aligned} \tilde{p} \Pr(\{v \in \tilde{V} : v \geq \tilde{p}\}) &\leq \tilde{p}(1 - F(\tilde{p})) \\ &< p^*(1 - F(p^*)) \\ &= p^* \Pr(\{v \in \tilde{V} : v \geq p^*\}), \end{aligned}$$

where the weak inequality follows from $\tilde{V} \subseteq V$, the strict inequality follows from p^* being the (lowest) optimal price, and the equality follows from $\{v \in V : v \geq p^*\} \subseteq \tilde{V}$. Therefore, the monopolist gains from profitably deviating from charging \tilde{p} to a price of p^* when receiving the non-disclosure message, thereby generating a contradiction. \square

Proof of Proposition 2. We describe the seller's off-path belief system: if the seller receives a message $M \notin (\cup_{s=1, \dots, S} M_s) \cup M_\infty$, she puts probability 1 on $\max_{v \in M} v$.

Let us show that neither the seller nor the consumer has a profitable deviation. For the seller, for each (on- or off-path) message, the equilibrium price is optimal given beliefs induced by that message. For the consumer, let us consider on-path messages first. Consider a type v that is prescribed to send message $[p_s, p_{s-1}]$ where $p_s < v \leq p_{s-1}$. The only potentially feasible on-path message is $[p_{s-1}, p_{s-2}]$, but this results in a strictly higher and therefore is not a profitable deviation. If $v = v$, her equilibrium payoff is 0, and sending any other message results in a weakly higher price. Thus, the consumer has no profitable deviation to any other on-path message. There is also no profitable deviation to any off-path message: because for any set M that contains v , $v \leq \max_{v' \in M} v'$, any off-path message is guaranteed to result in a payoff of 0. \square

Proof of Proposition 3. Consider an equilibrium σ . Let $m^\sigma(v)$ denote the message reported by type v , let $F_m^\sigma \in \Delta[0,1]$ denote the firm's belief when receiving message m and $\underline{v}^\sigma(m)$ be the lowest type in the support of that belief, and let $p^\sigma(m)$ be the sequentially rational price that he charges. In equilibrium, $p^\sigma(m) \geq \underline{v}^\sigma(m)$, because otherwise the firm has a profitable deviation. We say that a message is an equilibrium-path message if there exists at least one type that sends it, and a price is an equilibrium-path price if there exists at least one equilibrium-path message that induces the firm to charge that price.

Lemma 1 (Efficiency lemma) *For any equilibrium σ , there exists an equilibrium that is efficient that results in the same payoff for every consumer type.*

Proof. Consider an equilibrium σ . Define a strategy profile $\tilde{\sigma}$ in which

$$m^{\tilde{\sigma}}(v) = \begin{cases} m^\sigma(v) & \text{if } v \geq p^\sigma(m^\sigma(v)), \\ \{v\} & \text{otherwise,} \end{cases}$$

$$p^{\tilde{\sigma}}(m) = p^\sigma(m).$$

In this strategy profile, a type v that does not buy in equilibrium σ is fully revealing in $\tilde{\sigma}$. Because σ is an equilibrium, and the pricing strategy remains unchanged, such a type purchases in $\tilde{\sigma}$ at price v .

We argue that $\tilde{\sigma}$ is an equilibrium. Because σ is an equilibrium, and we have not changed the payoff for any message, no type has a motive to deviate. We now argue that the monopolist has no incentive to change prices. Because $p^\sigma(m)$ is an optimal price for the firm to charge in the equilibrium σ when receiving message m ,

$$p^\sigma(m)(1 - F_m^\sigma(p^\sigma(m))) \geq p(1 - F_m^\sigma(p)) \text{ for every } p. \quad (\text{A.1})$$

After receiving message m in $\tilde{\sigma}$, the monopolist's payoff from setting a price of $p^{\tilde{\sigma}}(m)$ is $p^{\tilde{\sigma}}(m) = p^\sigma(m)$ (because that price is accepted for sure), and the payoff from setting a higher price is $p(1 - F_m^{\tilde{\sigma}}(p))$. But observe that by Bayes' Rule, for every $p \geq p^{\tilde{\sigma}}(m)$,

$$1 - F_m^{\tilde{\sigma}}(p) = \frac{1 - F_m^\sigma(p)}{1 - F_m^\sigma(p^\sigma(m))}.$$

Thus, (A.1) implies that $p^{\tilde{\sigma}}(m) \geq p(1 - F_m^{\tilde{\sigma}}(p))$ for every $p > p^{\tilde{\sigma}}(m)$, and clearly the monopolist has no incentive to reduce prices below $p^{\tilde{\sigma}}(m)$. Therefore, the monopolist has no motive to deviate. \square

Lemma 2 (Partitional lemma) *For every efficient equilibrium σ , there is a partitional equilibrium $\tilde{\sigma}$ that results in the same payoff for almost every type.*

Proof. In any equilibrium, $p^\sigma(m) \geq \underline{v}^\sigma(m)$. But in an efficient equilibrium, trade occurs with probability 1, which implies that $p^\sigma(m) = \underline{v}^\sigma(m)$.

Step 1: We prove that the set of types being charged an equilibrium-path price p is a connected set. Suppose that types v and $v' > v$ are sending (possibly distinct) equilibrium-path messages m and m' such that $p^\sigma(m) = p^\sigma(m')$. Because $p^\sigma(m) = \underline{v}^\sigma(m)$ and $p^\sigma(m') = \underline{v}^\sigma(m')$, it follows that $\underline{v}^\sigma(m) = \underline{v}^\sigma(m') < v < v'$. Because types arbitrarily close to $\underline{v}^\sigma(m')$ and v' are both sending the message m' , the message m' contains the interval $[\underline{v}^\sigma(m'), v']$.

Consider any type v' in $[v, v']$: because $[v, v'] \subseteq [\underline{v}^\sigma(m'), v'] \subseteq m'$, it follows that m' is a *feasible message* for type v' . Therefore, denoting m' as the equilibrium-path message of type v' , type v' does not have a profitable deviation to sending message m' only if $p^\sigma(m') \leq p^\sigma(m)$.

We argue that this weak inequality holds as an equality. Suppose towards a contradiction that $p^\sigma(m') < p^\sigma(m)$. Then it follows from $p^\sigma(m') = \underline{v}^\sigma(m')$ that $\underline{v}^\sigma(m') < \underline{v}^\sigma(m) \leq v \leq v'$. Therefore, the interval $[\underline{v}^\sigma(m'), v']$ is both a subset of m' and contains v , and hence, m' is a feasible message for type v . But then, type v has an incentive to deviate from her equilibrium-path message m to m' , which is a contradiction.

Step 2: For every equilibrium-path price p , let

$$M^\sigma(p) \equiv \{m \in \mathcal{M} : p^\sigma(m) = p \text{ and } m \text{ is an equilibrium-path message}\},$$

$$V^\sigma(p) \equiv \{v \in V : p^\sigma(m^\sigma(v)) = p\}.$$

Observe that for every message m in $M^\sigma(p)$, the monopolist's optimal price is p . Because the monopolist's payoff from charging any price is linear in his beliefs, and the belief induced by knowing that the type is in $V^\sigma(p)$ is a convex combination of beliefs in the set $\bigcup_{m \in M^\sigma(p)} \{F^\sigma(m)\}$, it follows that the monopolist's optimal price remains p when all he knows is that the type is in $V^\sigma(p)$.

Now consider the collection of sets

$$\mathcal{P}^\sigma \equiv \{m \in \mathcal{M} : m = cl(V^\sigma(p)) \text{ for some equilibrium-path price } p\},$$

where $cl(\cdot)$ is the closure of a set. We argue that \mathcal{P}^σ is a partition of $[\underline{v}, \bar{v}]$: clearly, $[\underline{v}, \bar{v}] \subseteq \bigcup_{m \in \mathcal{P}^\sigma} m$, and because each of $V^\sigma(p)$ and $V^\sigma(p')$ are connected for equilibrium-path prices p and p' , $cl(V^\sigma(p)) \cap cl(V^\sigma(p'))$ is at most a singleton. For a type v let $m^{\mathcal{P}^\sigma}(v)$ be the lowest member of \mathcal{P}^σ that contains v .

Consider a strategy-profile $\tilde{\sigma}$ where each type v sends the message $m^{\mathcal{P}^\sigma}(v)$. Fix such a message m generated by $\tilde{\sigma}$; there exists a price p that is on the equilibrium path (in the equilibrium σ) such that $m = cl(V^\sigma(p))$. Because the prior is atomless, the monopolist's optimal price when receiving message m in $\tilde{\sigma}$ is equivalent to setting the optimal price when knowing that the type is in $V^\sigma(p)$, which as established above, is p . If any other message $m = [a, b]$ is reported, the monopolist believes that the consumer's type is b with probability 1.

We argue that this is an equilibrium. We first consider deviations to other messages that are equilibrium-path for $\tilde{\sigma}$. For any type v such that there exists a unique element in \mathcal{P}^σ that contains v , there exists no other feasible message that is an equilibrium-path message for $\tilde{\sigma}$. For any other type v , the strategy of sending the message $m^{\mathcal{P}^\sigma}(v)$ ensures that type v is sending the equilibrium-path message that induces the lower price. Finally, no type gains from sending an off-path message. Observe that all but a measure-0 set of types are charged the same price in $\tilde{\sigma}$ as they are in σ . \square

Proof of Proposition 4. Since all partitional equilibria involve trade with probability 1, a partitional equilibrium σ has higher *ex ante* consumer welfare than the partitional equilibrium $\tilde{\sigma}$ if the average price in σ is lower than that in $\tilde{\sigma}$:

$$\int_0^1 p^\sigma(m^\sigma(v))dv \leq \int_0^1 p^{\tilde{\sigma}}(m^{\tilde{\sigma}}(t))dv.$$

Thus, it suffices to prove that the greedy segmentation attains the lowest average price attainable by any partitional equilibrium.

We first describe the greedy segmentation. For a truncation of valuations $[0, v]$ where $v \leq 1$, let $p(v)$ solve $pf(p) = F(v) - F(p)$, which implies that $p(v) = \frac{v}{\sqrt[k]{k+1}}$; let us denote the denominator of $p(v)$ by γ , and note that $\gamma > 1$. The greedy segmentation divides the $[0, 1]$ interval into sets of the form $\{0\} \cup \bigcup_{\ell=0}^\infty S_\ell$ where $S_\ell \equiv \left[\frac{1}{\gamma^{\ell+1}}, \frac{1}{\gamma^\ell}\right]$.

We prove that no partitional equilibrium generates a lower average price on the segment S_ℓ than $\frac{1}{\gamma^{\ell+1}}$. Consider an arbitrary $\ell \geq 0$. Consider some type $\tilde{v} \in \left(\frac{1}{\gamma^{\ell+1}}, \frac{1}{\gamma^\ell}\right)$. Consider an upper segment $\left(\tilde{v}, \frac{1}{\gamma^\ell}\right]$, which is charged \tilde{v} by the monopolist. Let us then pool all types in $\left[\frac{1}{\gamma^{\ell+1}}, \tilde{v}\right]$ with types below $\frac{1}{\gamma^{\ell+1}}$ so as to create a segment that is charged the lowest possible price. The lowest possible price is $\frac{\tilde{v}}{\gamma}$, which is achieved if all types in $\left[\frac{1}{\gamma^{\ell+1}}, \tilde{v}\right]$ are pooled with types in $\left[\frac{\tilde{v}}{\gamma}, \frac{1}{\gamma^{\ell+1}}\right]$. The resulting average price in the segment S_ℓ is

$$\begin{aligned} \bar{P}(\tilde{v}) &\equiv (F(\tilde{v}) - F(1/\gamma^{\ell+1})) \frac{\tilde{v}}{\gamma} + (F(1/\gamma^\ell) - F(\tilde{v}))\tilde{v} \\ &= (\tilde{v}^k - \gamma^{-k(\ell+1)}) \frac{\tilde{v}}{\gamma} + (\gamma^{-k\ell} - \tilde{v}^k)\tilde{v}, \end{aligned}$$

where the first equality substitutes $F(v) = v^k$. Taking derivatives,

$$\frac{d^2 \bar{P}}{d\tilde{v}^2} = (k+1)k\tilde{v}^{k-1} \left(\frac{1}{\gamma} - 1\right) < 0,$$

where the inequality follows from $\gamma > 1$. Therefore, \bar{P} is concave in \tilde{v} . The boundary condition that $\bar{P}(\gamma^{-\ell}) = \bar{P}(\gamma^{-(\ell+1)}) = \gamma^{-(\ell+1)}$ coupled with concavity of \bar{P} implies that $\bar{P}(\tilde{v}) \geq \gamma^{-(\ell+1)}$ for every $\tilde{v} \in \left(\frac{1}{\gamma^{\ell+1}}, \frac{1}{\gamma^\ell}\right)$. Therefore, no partitional equilibrium generates a lower average price than $\gamma^{-(\ell+1)}$ for the set of types in S_ℓ . Because the greedy segmentation attains this lowerbound pointwise on every interval S_ℓ for every ℓ , it is the consumer-optimal partitional equilibrium. \square

Proof of Proposition 5. We first generalize the notion of facilitating group pricing.

Definition 1. An evidence technology facilitates group pricing if

- (a) The set of types that can send message M is of positive measure and induces a price strictly below p^* , viz., the lowest maximizer of $\int_{v: \bar{M} \in \mathcal{M}(v)} p \mathbf{1}\{v \geq p\} dF$ is strictly less than p^* .
- (b) The set of types that cannot send message M induces a price no more than p^* , viz., the lowest maximizer of $\int_{v: \bar{M} \notin \mathcal{M}(v)} p \mathbf{1}\{v \geq p\} dF$ is no more than p^* .

Observe that this definition is compatible with the special case described in the main text: if $\{v: \bar{m} \in \mathcal{M}(v)\} \subseteq [\underline{v}, p^*]$, then the message \bar{M} induces a price strictly below p^* but not being able to send message \bar{M} does not induce a price exceeding p^* . But, we note that Definition 1 also encompasses cases where some high-valuation types can also send message \bar{M} .

Let p_l be the lowest price that maximizes $\int_{v: \bar{M} \in \mathcal{M}(v)} p \mathbf{1}\{v \geq p\} dF$, and p_h be the lowest price that maximizes $\int_{v: \bar{M} \notin \mathcal{M}(v)} p \mathbf{1}\{v \geq p\} dF$. By Definition 1, $p_l < p^*$ and $p_h \leq p^*$. Let $\bar{v}(M) \equiv \max_{v: M \in \mathcal{M}(v)} v$. Consider a strategy profile σ in which the messaging strategy is

$$M^{\bar{\sigma}}(v) = \begin{cases} \bar{M} & \text{if } \bar{M} \in \mathcal{M}(v), \\ [\underline{v}, \bar{v}] & \text{otherwise.} \end{cases}$$

and the monopolist's pricing strategy is

$$p^{\bar{\sigma}}(M) = \begin{cases} p_h & \text{if } M = [\underline{v}, \bar{v}], \\ p_l & \text{if } M = \bar{M}, \\ \bar{v}(M) & \text{otherwise.} \end{cases}$$

We consider a belief-system in which whenever the monopolist obtains a message $M \notin \{[\underline{v}, \bar{v}], \bar{M}\}$, she believes that it was sent by the highest type that could send that message. Given this belief system, the monopolist's pricing strategy is sequentially rational. Moreover, given this pricing strategy, the message strategy is the best response. Observe that this equilibrium results in a Pareto improvement. \square

Proof of Proposition 6. For a message M , let $\tau(i, M) \equiv \arg \min_{t \in M} |t - \ell_i|$ denote the closest type in M to seller i ; this is well-defined because M is closed. Let δ_t be the probability distribution that puts probability 1 on type t . We use this notation to describe a fully revealing equilibrium:

- The consumer of type t sends message $\{t\}$.
- If seller i receives message M , his beliefs are $\delta_{\tau(i, M)}$ and that the other seller has received a fully revealing message.
- If seller i holds belief $\delta_{\tau(i, M)}$, he charges a price $p_i(M) = \max\{2\tau(i, M)\ell_i, 0\}$.
- If $v - p_i - |t - \ell_i| > v - p_j - |t - \ell_j|$ and $v - p_i - |t - \ell_i| \geq 0$, then the consumer purchases from firm i .
- If $v - p_L - |t - \ell_L| = v - p_R - |t - \ell_R| \geq 0$, the consumer purchases from firm L if and only if $t \leq 0$, and otherwise, the consumer purchases from firm R .

We argue that this is an equilibrium. Observe that each seller's on-path beliefs are consistent with Bayes' rule, since $t = \tau(i, \{t\})$. In the case of an off-path message M , Bayes' rule does not restrict the set of possible beliefs, and therefore, the above off-path belief assessment is feasible.

Let us see that no firm has a profitable deviation. Suppose that firm i receives message M . He believes with probability 1 (on or off-path) that the consumer's type is $\tau(i, M)$ with probability 1 and that the other firm j has received a message $\{\tau(i, M)\}$. Denote this type by t .

- If $2t\ell_i > 0$ then the consumer is closer to firm i and so $2t\ell_j < 0$. In this case, firm i believes that firm j is charging a price of 0. Charging a price strictly higher than $2t\ell_i$ leads to a payoff of 0 (because the consumer will reject such an offer), and charging a price p weakly below $2t\ell_i$ leads to a payoff of p (because the consumer always breaks ties in favour of the closer firm).
- If $2t\ell_i \leq 0$, then the consumer is located closer to the other firm j and is being charged a price equal to $2t\ell_j$. In this case, charging any strictly positive price leads to a payoff of 0.

Therefore, in each case, firm i has no incentive to deviate.

We argue that the consumer has no incentive to deviate. If $t \leq 0$, the consumer obtains an equilibrium price of 0 from firm R , which is the lowest possible price. Sending any other message $M \in M(t)$ to firm L induces a weakly higher price because for any feasible message $M \in M(t)$, $\tau(L, M) \leq t$, and therefore, $2\tau(L, M)\ell_L \geq 2t\ell_L$. Thus, the consumer has no profitable deviation. An analogous argument applies if the consumer's type is greater than 0. \square

Proof of Proposition 7. We first describe the incentives of the consumer.

- Suppose that $t \in (t_1^L, t_1^R)$, or in other words, $t\ell_i < t_1^i\ell_i$ for each i . In equilibrium, the consumer sends the message $\{t\}$ to each firm and is offered a price of $\max\{2t\ell_i, 0\}$ by firm i . If the consumer deviates and sends message $[-1, 1]$ to firm i , she induces a price of $p_i^j = 2t_1^i\ell_i$, which is higher.

- Suppose that $t\ell_i \geq t_1^i \ell_i$ for some i . In equilibrium, the consumer sends message $[-1, 1]$ to firm i and $\{t\}$ to firm j . The equilibrium price offered to the consumer by firm j is 0, the lowest possible price. The price offered by firm i is $2t_1^i \ell_i$, and deviating leads to a higher price of $2t\ell_i$.

Therefore, in either case, the consumer does not have a profitable deviation.

We now consider whether firm i has an incentive to deviate. It follows from the proof of Proposition 6 that the prices are optimal whenever firm i receives an (equilibrium-path) message of $\{t\}$ for $t \in (t_1^L, t_1^R)$. An identical argument applies when firm i receives an (off-path) message of $\{t\}$ for $t\ell_i \geq t_1^i \ell_i$: in this case, firm i believes that firm j is charging a price of 0, and thus, the optimal price is $2t\ell_i$ (because the consumer always breaks ties in favour of firm i). When firm i receives an (equilibrium-path) message of $\{t\}$ for $t\ell_j \geq t_1^j \ell_j$, firm i believes that firm j is charging a price of $2t_1^j$. The equilibrium prescribes that firm i charges a price of 0, which leads to a payoff of 0 (because the consumer breaks ties in favour of firm j), and any strictly positive price also leads to a payoff of 0. Finally, consider the case when firm i receives an (equilibrium-path) message of $[-1, 1]$. Firm i infers that $t\ell_i \geq t_1^i \ell_i$ and believes that firm j is charging a price of 0. Because p_1^i is, by definition, a profit-maximizing price in response to a price of 0, firm i has no strictly profitable deviation.

Observe that in this equilibrium, each consumer type t incurs a total cost no more than that of Proposition 6, which is $1 + |t|$. We argue that this equilibrium results in a higher payoff for every type than any other pure strategy equilibrium. Consider another pure strategy equilibrium σ and suppose that it improves the payoff of some type $t < 0$. Observe that it cannot be that type t purchases from firm R because the price charged by firm R is at least 0. Therefore, type t must purchase from firm L . It also cannot be that type t discloses itself to both firms because then it incurs a total cost of $1 + |t|$. Therefore t must send the message $[-1, 1]$ to firm L . In equilibrium σ , let $T_L^\sigma \equiv \{t < 0 : M_L(t) = [-1, 1]\}$ be the set of types sending message $[-1, 1]$ to firm L . If type t reveals its location to firm R , the equilibrium price from firm R is 0; indeed, this must be true for every type in $T^{\sigma L}$. For this to not be a profitable deviation, it implies that in equilibrium, each of these types is being charged a price of 0 by firm R . Therefore, it is without loss of generality to suppose that $M_R(t) = \{t\}$. Let p_L^{ND} be the equilibrium price charged by firm L when it obtains the non-disclosure message $[-1, 1]$. Observe that if $2|t| > p_L^{ND}$, then the consumer of type t has a strictly profitable deviation to the message $[-1, 1]$. Therefore, $T_L^\sigma = [-1, \tilde{t}]$ for some \tilde{t} .³⁸ Observe that there is no pure strategy equilibrium in which $\tilde{t} > t_1^L$ since otherwise, the firm would charge a price such that all types in (t_1^L, \tilde{t}) purchase from firm R at a price of 0. But then firm R would have a profitable deviation to increasing its price above 0. Therefore, $\tilde{t} \leq t_1^L$. For all such values of \tilde{t} , all consumer types with $t < 0$ are better off in the equilibrium constructed in Proposition 7. An analogous argument holds for types where $t \geq 0$. \square

Proof of Proposition 8. We use an off-path belief system where if firm i receives an off-path message M , she holds degenerate beliefs $\delta_{\tau(i, M)}$ that put probability 1 on type $\tau(i, M)$ where recall that $\tau(i, M)$ is defined as the type in M that is located closest to firm i (this was defined in the proof of Proposition 6). Given such beliefs, the firm charges a price $p_i(M) = \max\{2\tau(i, M)\ell_i, 0\}$ for an off-path message M .

First, we prove that given the pricing strategies, no consumer has an incentive to deviate. Consider a consumer type t such that $t_1^s \ell_i < t\ell_i \leq t_{s-1}^i \ell_i$ for some $s = 1, 2, \dots$. Such a consumer should be sending message M_s^i to both firms. Such a message induces a price of 0 from firm j and $p_s^i = 2t_1^s \ell_i$ from firm i . No message can induce a lower price from firm j . Therefore, any strictly profitable deviation must induce a strictly lower price from firm i . We show that this is not possible.

We first argue that the consumer does not have a profitable deviation to any other equilibrium-path message. Suppose that $t\ell_i < t_{s-1}^i \ell_i$. In this case, M_s^i is the only equilibrium-path message that type t can send to firm i . If $t\ell_i = t_{s-1}^i \ell_i$, then type t can send either message M_s^i or M_{s-1}^i but because $p_s^i \leq p_{s-1}^i$, this is not a strictly profitable deviation.

We now argue that the consumer does not have a profitable deviation to any off-path message. Any feasible message $M \in \mathcal{M}(t)$ satisfies the property that the closest type in M to firm i is at least as close as t to firm i ; or formally: $t\ell_i \leq \tau(i, M)\ell_i$. In that case, the price that the consumer is charged is $2\tau(i, M)\ell_i \geq 2t\ell_i > 2t_1^s \ell_i = p_s^i$. Therefore, this deviation is not strictly profitable.

Finally, we argue that the firms have no incentive to deviate in their pricing strategies. For any equilibrium-path message, the prices charged by firms are (by construction) equilibrium prices. For any off-path message M , each firm assumes that the consumer sent the equilibrium-path message to the other firm. If $\tau(i, M)\ell_i > 0$ then firm i assumes that firm j is charging a price of 0, and then charging a price of $2\tau(i, M)\ell_i$ is a best-response (assuming that the consumer breaks ties in favour of the closer firm). If $\tau(i, M)\ell_i \leq 0$, then firm i believes that the consumer is being charged a price p_s^j by firm j for some s where $t_1^s \ell_i < \tau(i, M)\ell_i \leq t_{s-1}^i \ell_i$. Because the consumer breaks ties in favour of the closer firm, firm i anticipates that the consumer will reject any strictly positive price. \square

Proof of Proposition 9. We show that $p^* > p_1^i$ for every i . Observe that

$$p_1^i = \frac{2F(-p_1^i/2)}{f(-p_1^i/2)} < \frac{2F(0)}{f(0)} = p^*,$$

38. Type \tilde{t} is indifferent, and it makes no difference as to how one breaks its tie.

where the first equality follows from the first-order condition that p_1^t solves, the inequality follows from F being strictly log-concave, and the second equality follows from the definition of p^* . A symmetric argument shows that $p_1^R < p^*$.

Now, we prove that all consumers are better off in the equilibrium we construct in the game with simple evidence (Proposition 7). All types where $t\ell_i \geq t_1^i \ell_i$ are buying the good at a lower price because $p_1^i < p^*$. Consider any other type, i.e., where $t\ell_i < t_1^i \ell_i$ for every $i \in \{L, R\}$. Suppose that $t\ell_i > 0$. That type in equilibrium buys the good from firm i at the price $2t\ell_i < 2t_1^i \ell_i$, which equals p_1^i . Therefore, it obtains the good at a lower price than p^* . Finally, if $t=0$, that type obtains the good at a price equal to 0.

An analogous argument ranks prices relative to the equilibrium constructed in the game with rich evidence (Proposition 8). All types in that equilibrium pay a price that is less than p_1^i , and therefore, buy the good at a price lower than p^* . \square

Proof of Proposition 10. Suppose that prices are chosen from a price grid $\{0, \Delta, 2\Delta, \dots\}$ where our interest is in small $\Delta > 0$, and then taking $\Delta \rightarrow 0$. For a price p , let $\lfloor p \rfloor_\Delta$ be the highest multiple of Δ that is lower than p .

Let us describe equilibrium path prices in a fully revealing equilibrium. Suppose $t < 0$. Then $p_R(t) \in \{0, \Delta\}$ and $p_L(t) = \lfloor 2|t| \rfloor_\Delta + p_R(t)$. Analogously, if $t \geq 0$, then $p_L(t) \in \{0, \Delta\}$ and $p_R(t) = \lfloor 2|t| \rfloor_\Delta + p_L(t)$. Observe that as $\Delta \rightarrow 0$, these prices converge to those described in Proposition 6.

Consider an efficient passive equilibrium. Consider a consumer type where $t < 0$. Because the equilibrium is efficient, the type purchases from firm L . Suppose the equilibrium prices it faces from firms L and R are p_L and p_R . We argue that

$$p_L \leq \lfloor 2|t| \rfloor_\Delta + \Delta.$$

Suppose otherwise. Observe that this implies that $p_R > \Delta$. But then in equilibrium, if the consumer sends a message of $\{t\}$ to firm R , firm R sets a price \tilde{p} such that the consumer is strictly better off purchasing from firm R at \tilde{p} than from firm L at price p_L .³⁹ Therefore, the consumer has a strictly profitable deviation, which is a contradiction. A similar argument applies for types $t \geq 0$. Taking $\Delta \rightarrow 0$ implies that no consumer type obtains worse payoffs than in Proposition 6.

Our result then follows from Proposition 1 of Rhodes and Zhou (2021), who show that if f is symmetric around 0 and log-concave, average consumer surplus is higher in the equilibrium studied in Proposition 6 than with uniform pricing; formally they show that

$$\int_0^1 2tf(t)dt \leq \frac{2F(0)}{f(0)},$$

where the LHS is the average price paid by types in $[0, 1]$ in the equilibrium studied in Proposition 6 and the RHS is that with uniform pricing. \square

REFERENCES

- ACEMOGLU, D., MAKHDOUNI, A., MALEKIAN, A. and OZDAGLAR, A. (2021), "Too Much Data: Prices and Inefficiencies in Data Markets", *American Economic Journal: Microeconomics*, forthcoming.
- ACQUISTI, A., TAYLOR, C. and WAGMAN, L. (2016), "The Economics of Privacy", *Journal of Economic Literature*, **54**, 442–492.
- ACQUISTI, A. and VARIAN, H. R. (2005), "Conditioning Prices on Purchase History", *Marketing Science*, **24**, 367–381.
- ARMSTRONG, M. (2006), "Recent Developments in the Economics of Price Discrimination", in Blundell, R., Newey, W. K. and Persson, T. (eds) *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, vol. 2 of Econometric Society Monographs* (Cambridge University Press) 97–141.
- ARMSTRONG, M. and ZHOU, J. (2022), "Consumer Information and the Limits to Competition", *American Economic Review*, **112**, 534–577.
- BAGNOLI, M. and BERGSTROM, T. (2005), "Log-Concave Probability and its Applications", *Economic Theory*, **26**, 445–469.
- BAJARI, P., NEKIPELOV, D., RYAN, S. P. and YANG, M. (2015), "Machine Learning Methods for Demand Estimation", *American Economic Review*, **105**, 481–85.
- BELLEFLAMME, P. and PEITZ, M. (2015), *Industrial Organization* (Cambridge, UK: Cambridge University Press).
- BERGEMANN, D. and BONATTI, A. (2019), "Markets for Information: An Introduction", *Annual Review of Economics*, **11**, 85–107.
- BERGEMANN, D., BONATTI, A. and GAN, T. (2021), "The Economics of Social Data", *Rand Journal of Economics*, forthcoming.
- BERGEMANN, D., BROOKS, B. and MORRIS, S. (2015), "The Limits of Price Discrimination", *American Economic Review*, **105**, 921–957.
- BONATTI, A. and CISTERNAS, G. (2020), "Consumer Scores and Price Discrimination", *Review of Economic Studies*, **87**, 750–791.

39. Recall that consumers break ties in favour of the closer firm.

- CALZOLARI, G. and PAVAN, A. (2006), "On the Optimality of Privacy in Sequential Contracting", *Journal of Economic Theory*, **130**, 168–204.
- CAPLIN, A. and NALEBUFF, B. (1991), "Aggregation and Imperfect Competition: On the Existence of Equilibrium", *Econometrica*, **59**, 25–59.
- CHOI, J. P., JEON, D.-S. and KIM, B.-C. (2019), "Privacy and Personal Data Collection with Information Externalities", *Journal of Public Economics*, **173**, 113–124.
- COLLARD-WEXLER, A., GOWRISANKARAN, G. and LEE, R. S. (2019), "Nash-in-Nash Bargaining: A Microfoundation for Applied Work", *Journal of Political Economy*, **127**, 163–195.
- CONITZER, V., TAYLOR, C. R. and WAGMAN, L. (2012), "Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases", *Marketing Science*, **31**, 277–292.
- CREMER, J. and RIORDAN, M. H. (1987), "On Governing Multilateral Transactions with Bilateral Contracts", *Rand Journal of Economics*, **18**, 436–451.
- DYE, R. A. (1985), "Disclosure of Nonproprietary Information", *Journal of Accounting Research*, **23**, 123–145.
- ELLIOTT, M., GALEOTTI, A., KOH, A. and LI, W. (2021), "Market Segmentation Through Information" (Working Paper).
- FAINMESSER, I. P. and GALEOTTI, A. (2016), "Pricing Network Effects", *Review of Economic Studies*, **83**, 165–198.
- (2020), "Pricing Network Effects: Competition", *American Economic Journal: Microeconomics*, **12**, 1–32.
- FAINMESSER, I. P., GALEOTTI, A. and MOMOT, R. (2020), "Digital Privacy" (Working Paper).
- FEDERAL TRADE COMMISSION (2012), "Protecting Consumer Privacy in an Era of Rapid Change" <https://tinyurl.com/y6eua098>.
- (2014), "Data Brokers: A Call for Transparency and Accountability" <https://tinyurl.com/wa8z6n>.
- GLODE, V., OPP, C. C. and ZHANG, X. (2018), "Voluntary Disclosure in Bilateral Transactions", *Journal of Economic Theory*, **175**, 652–688.
- GROSSMAN, S. J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality", *Journal of Law and Economics*, **24**, 461–483.
- HAGENBACH, J. and KOESSLER, F. (2017), "Simple versus Rich language in Disclosure Games", *Review of Economic Design*, **21**, 163–175.
- HIDIR, S. and VELLODI, N. (2021), "Personalization, Discrimination and Information Revelation", *Journal of European Economic Association*, **119**, 1342–1363.
- HOGAN, K. (2018), "Consumer Experience in the Retail Renaissance: How Leading Brands Build a Bedrock with Data" <https://tinyurl.com/6n5vd2ax>.
- HORN, H. and WOLINSKY, A. (1988), "Bilateral Monopolies and Incentives for Merger", *Rand Journal of Economics*, **19**, 408–419.
- ICHIHASHI, S. (2021), "The Economics of Data Externalities", *Journal of Economic Theory*, **196**, 105316.
- JIN, Y. and VASSERMAN, S. (2021), "Buying Data from Consumers: The Impact of Monitoring Programs in US Auto Insurance" (Working Paper).
- JONES, C. I. and TONETTI, C. (2020), "Nonrivalry and the Economics of Data", *American Economic Review*, **110**, 2819–2858.
- JOVANOVIC, B. (1982), "Truthful Disclosure of Information", *Bell Journal of Economics*, **13**, 36–44.
- MCADAMS, D. (2011), "Discounts for Qualified Buyers Only" (Working Paper).
- MCAFEE, R. P. and SCHWARTZ, M. (1994), "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity", *American Economic Review*, **84**, 210–230.
- MILGROM, P. R. (1981), "Good News and Bad News: Representation Theorems and Applications", *Bell Journal of Economics*, **12**, 380–391.
- ORGANIZATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT (2018), "Personalised Pricing in the Digital Era" <https://www.justice.gov/atr/page/file/1312741/download>.
- PRAM, K. (2021), "Disclosure, Welfare and Adverse Selection", *Journal of Economic Theory*, **197**, 105327.
- REIMERS, I. and SHILLER, B. R. (2019), "The Impacts of Telematics on Competition and Consumer Behavior in insurance", *Journal of Law and Economics*, **62**, 613–632.
- REY, P. and VERGÉ, T. (2004), "Bilateral Control with Vertical Contracts", *RAND Journal of Economics*, **35**, 728–746.
- RHODES, A. and ZHOU, J. (2021), "Personalized Pricing and Privacy Choice" (Working Paper).
- SHER, I. and VOHRA, R. (2015), "Price Discrimination Through Communication", *Theoretical Economics*, **10**, 597–648.
- STAHL, K. and STRAUSZ, R. (2017), "Certification and Market Transparency", *Review of Economic Studies*, **84**, 1842–1868.
- TAYLOR, C. R. (2004), "Consumer Privacy and the Market for Customer Information", *RAND Journal of Economics*, **35**, 631–650.
- THISSE, J.-F. and VIVES, X. (1988), "On the Strategic Choice of Spatial Price Policy", *American Economic Review*, **78**, 122–137.
- VARIAN, H. R. (1989), "Price Discrimination", *Handbook of Industrial Organization*, **1**, 597–654.
- VERRECCHIA, R. E. (1983), "Discretionary Disclosure", *Journal of Accounting and Economics*, **5**, 179–194.
- VILLAS-BOAS, J. M. (2004), "Price Cycles in Markets with Customer Recognition", *Rand Journal of Economics*, **35**, 486–501.