NBER WORKING PAPER SERIES

VOLUNTARY DISCLOSURE AND PERSONALIZED PRICING

S. Nageeb Ali
Gregory Lewis
Shoshana Vasserman

Working Paper 26592
http://www.nber.org/papers/w26592

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

We thank Ben Golub, Nima Haghpanah, Kevin He, Navin Kartik, Vijay Krishna, Pepe Montiel Olea, Doron Ravid, Ron Siegel, Nikhil Vellodi, and Jidong Zhou for useful comments. We are grateful to Microsoft Research for hosting Ali and Vasserman. We thank Xiao Lin for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Voluntary Disclosure and Personalized Pricing
S. Nageeb Ali, Gregory Lewis, and Shoshana Vasserman
NBER Working Paper No. 26592
December 2019
JEL No. D4,D8

ABSTRACT

A concern central to the economics of privacy is that firms may use consumer data to price discriminate. A common response is that consumers should have control over their data and the ability to choose how firms access it. Since firms draw inferences based on both the data seen as well as the consumer's disclosure choices, the strategic implications of this proposal are unclear. We investigate whether such measures improve consumer welfare in monopolistic and competitive environments. We find that consumer control can guarantee gains for every consumer type relative to both perfect price discrimination and no personalized pricing. This result is driven by two ideas. First, consumers can use disclosure to amplify competition between firms. Second, consumers can share information that induces a seller—even a monopolist—to make price concessions. Furthermore, whether consumer control improves consumer surplus depends on both the technology of disclosure and the competitiveness of the marketplace. In a competitive market, simple disclosure technologies such as "track / do-not-track" suffice for guaranteeing gains in consumer welfare. However, in a monopolistic market, welfare gains require richer forms of disclosure technology whereby consumers can decide how much information they would like to convey.

S. Nageeb Ali
Penn State University
411 Kern Bldg.
University Park, PA 16802
nageeb@psu.edu

Shoshana Vasserman
Stanford Graduate School of Business
655 Knight Way
Stanford, CA 94305
svass@stanford.edu

Gregory Lewis
Microsoft Research New England
1 Memorial Drive
Cambridge, MA 02142
and NBER
glewis@microsoft.com

Contents

"Privacy is not the opposite of sharing—rather, it is control over sharing."

– Acquisti, Taylor, and Wagman (2016)

1 Introduction

Alice is shopping for a new widget on the internet. How much should online sellers know about her preferences? In some countries, regulations allow every seller to leave a cookie that tracks Alice's online behavior, but in other countries, sellers cannot track Alice at all without her consent. Tracking is controversial. On the one hand, it may allow sellers to personalize their offerings by showing more targeted advertisement. On the other hand, as these sellers learn Alice's preferences, they may set prices that exploit that knowledge. Should these sellers be allowed to track Alice by default or should she have a say in this matter? How much control should Alice have in deciding what firms learn about her?

These questions are at the forefront of an ongoing international debate that has precipitated action in both the public and private sectors. Regulators concerned that personalized pricing will harm consumers have focused on the importance of consumer consent, passing wide-reaching legislation on data storage and tracking. A prominent example, the General Data Protection Regulation (GDPR) passed by the European Union, requires firms to obtain consent from consumers before obtaining and processing their personal data.[1] In the United States, 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). Meanwhile, private sector firms have responded to consumer demand for privacy by designing commercial products and brands that are specifically developed to limit tracking.[2]

Against this backdrop, we study the market implications of consumer consent and control. We investigate what happens when consumers fully control their data—not only whether they are tracked, but what specific information is disclosed to firms. Each consumer's data is encoded in a verifiable form that she can partially or fully disclose to firms. Based on the information that is disclosed, each firm draws an inference about the consumer's type and charges her an equilibrium price based on her disclosure. Our motivating question is: *when consumers fully control their information, are they hurt or helped by personalized pricing?*

We pose this question in an environment in which products cannot be personalized, and so there is no match value from data. A classical intuition might suggest that con-

[1]Starting on January 1, 2020, California will enforce the California Consumer Privacy Act (CCPA), which has similar provisions to the GDPR.

[2]For example, Apple recently added a feature to its Safari browser that limits the ways in which its user's activities are tracked by third parties (Hern, 2018).

sumers would not benefit from being permitted to voluntarily disclose information. Because the market's *equilibrium inferences* are based both on 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 unraveling equilibria of Grossman (1981) and Milgrom (1981). Contrary to that intuition, we find that the combination of personalized pricing and consumer control is actually beneficial to consumers in both monopolistic and competitive markets. We construct simple equilibria of the consumers' disclosure game in which sharing data weakly increases consumer surplus for *every* consumer type, relative to the benchmark of no personalized pricing.

Two key ideas drive this result. 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, while consumers with a strong preference for the product of one particular firm can hide this preference.[3] Second, even in the absence of competition, consumers can benefit from sending coarse signals that pool their valuations. These pools are constructed in such a way that a monopolist finds it optimal to sell to every type within that pool, and therefore everyone within that pool pays the price of the consumer type that has the lowest valuation in that pool. Disclosures lead to price discounts that benefit every consumer type. The take-away is that offering consumers control—and possibly building tools that coordinate the sharing of data for consumer benefit—may make personalized pricing attractive *even in the absence of better matching.*

A Preview: We build on 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 "verifiable" disclosure game, as in Grossman (1981) and Milgrom (1981): before the monopolist sets her price, a consumer chooses what "evidence" or hard information to disclose about her valuation. We study both those disclosure environments in which evidence is **simple**, where a consumer can either speak "the whole truth" (reveal her type) or say nothing at all, as well as those in which evidence is **rich**, where a consumer can disclose facts about her type without having to reveal it completely.[4] We first study simple and rich evidence structures in a monopolistic environment, and then use those results to characterize behavior in a competitive market.

Here is the timing of our game: the consumer first observes her type and then chooses a message to disclose to the firm from the set of messages available to her, the firm then

[3]Of course, the firms will interpret this non-disclosure as a signal and update accordingly.

[4]We borrow this terminology from Hagenbach and Koessler (2017).

quotes a price, and the consumer then chooses whether to buy the product at that price. Neither the firm nor the consumer can commit to its strategic choices.

Our first conclusion in the monopolistic environment (Proposition 1) is that simple evidence never benefits the consumer 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 an unraveling equilibrium in which the monopolist extracts all surplus.

Our second conclusion is that once 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 constructs an equilibrium that improves the consumer surplus for almost all consumer types without reducing the surplus of any consumer type. In this equilibrium, all types are partitioned into segments on the basis of their willingness to pay, and trading is fully efficient. Because the consumer cannot commit ahead of time to her disclosure strategy, every consumer type must find that her equilibrium message induces a weakly lower price than that induced by any other message; our segmentation guarantees this property. Moreover, our segmentation ensures that for each segment, the monopolist's optimal price is the lowest willingness to pay in that segment. This "greedy algorithm" identifies a "Pareto-improving" equilibrium segmentation for every distribution of consumer types and identifies the equilibrium segmentation that maximizes ex ante consumer surplus for a class of distributions.

We use these insights to study competitive behavior in a model of Bertrand duopoly with horizontally differentiated products where the firms are uncertain of the consumer's location. The consumer can disclose information about her location to the firms, who then simultaneously make price offers to her. As before, we compare the outcomes when the consumer can disclose, either via simple or rich evidence, with a benchmark model in which there is no personalized pricing. Here, voluntary disclosure and personalized pricing is particularly beneficial to consumer surplus because of a new economic force: *information can be selectively disclosed to amplify competition.*[5]

More specifically, we show that if the distribution of consumer location has a symmetric and log-concave density, then an equilibrium in the game with simple evidence (where the consumer's disclosure strategies are all-or-nothing) improves consumer welfare for every type relative to the no-personalized-pricing benchmark. With rich evidence,

[5]We focus on horizontal differentiation because it is the minimal setting where this force appears. Without product differentiation, Bertrand competition reduces prices to marginal cost regardless of information disclosure. If differentiation is vertical, voluntary disclosure and personalized pricing lead to segmentations similar to the monopolistic setting (where the consumer does not benefit from simple evidence and the pooling equilibria with rich evidence are similar).

one can do even better by using the greedy segmentation strategy similar to that used in the monopolist's problem, with pools becoming progressively finer as one approaches the central type from either end.

Implications: From this stylized model, we draw two broad lessons for policy. First, voluntary disclosure facilitates price concessions in both monopolistic and competitive markets. Thus, there is something missing in the view that tracking involves a tradeoff between the benefits of personalized products and the costs of personalized prices. Even without the benefits of product personalization, a consumer can benefit from personalized pricing when she has control. Disclosure generates discounts and amplifies competition.

The second lesson is that whether a track / do-no-track regime (as evoked by the GDPR) suffices to give consumers *useful* control over their data depends on the competitiveness of the marketplace. In a monopolistic environment, richer forms of data sharing are necessary for the consumers to gain, but the same is not true in a competitive market. Useful control may thus involve a choice not only of whether to share information but also of how much information to share, and to whom.

While online communications between consumers and sellers are not yet as sophisticated as that envisioned in our rich-evidence setup, an important element of the digital economy is its increasing ability to verify information (Goldfarb and Tucker, 2019). These advances suggest that it may be technologically feasible for consumers to use intermediaries or platforms to verifiably disclose that their preferences or characteristics (e.g., income, age, address etc.) lie with a certain range without having to forfeit all of their information to online sellers.

Relationship to Literature: Our work belongs to a burgeoning literature on privacy, information, and their implications for markets; see Acquisti, Taylor, and Wagman (2016) and Bergemann and Bonatti (2019) for recent surveys. We view our paper as making two contributions. First, it formulates and investigates the economic implications of giving consumers control over their data. Our goal is to study the simplest possible model, abstracting from a number of details (e.g. the importance of product customization), in order to elucidate the strategic issues at the core of voluntary disclosure and personalized pricing. Second, our analysis shows that whether consumers benefit from controlling their information depends on a subtle interaction between the technology by which consumers disclose information and the degree of market competition.

Our work combines classical models of market pricing with the now classical study of verifiable disclosure. Unlike the first analyses of verifiable disclosure (Grossman, 1981;

Milgrom, 1981), unraveling is not the unique equilibrium outcome of the market interactions that we study. An observation at the core of our results is that the price charged by a firm need not be strictly increasing in his beliefs (in an FOSD sense) about the consumer's willingness to pay. This observation permits us to pool low and high types without giving the low type an incentive to separate itself from the pool.[6]

The literature on verifiable disclosure has had a recent resurgence,[7] and a closely related contribution therein is Sher and Vohra (2015). They study a general model of price discrimination with hard evidence in which the monopolist commits to a schedule of evidence-contingent prices. By contrast, we assume in both monopolistic and competitive settings that each seller cannot commit and instead sets a price that is a best-response to the evidence that has been presented.

Our approach to consumer control complements two important approaches studied in the literature. The first approach is that of "information design," which considers the actions of an intermediary that already knows the consumer's type and can commit to a segmentation strategy. In a monopolistic setting, that intermediary can achieve payoffs characterized by Bergemann, Brooks, and Morris (2015). Necessarily, any equilibrium outcome of our monopolist setting is attainable in their model but the converse is false. The reason is that an intermediary may pool a consumer type stochastically into different market segments that induce different prices. But a consumer who cannot commit would strictly prefer the market segment that offers the best price and would not randomize unless she were indifferent. The second approach is to envision that the consumer communicates her preferences using cheap talk. In the setting modeled here, cheap talk is ineffectual because every consumer that trades prefers to send the message that induces the lowest price, and thus, cheap talk alone cannot improve upon the benchmark of no-personalized-pricing. Hidir and Vellodi (2019) show that cheap talk is effective if the monopolist also matches a product to the consumer's tastes; then consumers can use cheap-talk to sort so that they are matched with a better product without being completely exploited by price discrimination.[8]

---

[6] Prior analyses have highlighted other reasons for why markets may not unravel, in particular (i) uncertainty about whether the sender has evidence (Dye, 1985; Shin, 1994), (ii) disclosure costs (Jovanovic, 1982; Verrecchia, 1983), or (iii) the possibility for receivers to be naive (Hagenbach and Koessler, 2017). Our setting does not have any of these features.

[7] A partial list of recent contributions is Kartik and Tercieux (2012), Ben-Porath and Lipman (2012), Hagenbach, Koessler, and Perez-Richet (2014), Hart, Kremer, and Perry (2017), Ben-Porath, Dekel, and Lipman (2017, 2019), and Koessler and Skreta (2019). Our work also relates to the study of certification in mitigating adverse selection in markets (Lizzeri, 1999; Stahl and Strausz, 2017; Glode, Opp, and Zhang, 2018), but we study a "private values" model rather than one with interdependent values.

[8] Their study extends beyond cheap-talk as they introduce and characterize the buyer-optimal incentive-compatible market segmentation in their setting. Ichihashi (2019) and Haghpanah and Siegel (2019) also study price discrimination with multiple products, using an information-design approach.

Verifiable disclosure complements and offers a middleground between information design and cheap talk. It is particularly germane when a consumer can use an intermediary to verify information about her type in her communication to firms without forfeiting control over the disclosure of that information. As the evolving digital economy balances an increasing ability to verify information cryptographically and public pressure for individual privacy, we believe it to be useful to complement the existing frameworks with a verifiable disclosure approach to the question of consumer control.

A large part of our motivation is to understand the role of voluntary disclosure and personalized pricing in competitive markets. Using a model of Bertrand duopoly with horizontal differentiation, we show that the consumer can use disclosure to amplify competitive forces. The notion that personalized pricing may amplify competitive forces is seen in the innovative work of Thisse and Vives (1988). They consider a setting where the consumer's type is commonly known, and show that the unique equilibrium involves firms adopting personalized pricing strategies even though their joint profits would be higher if they could commit to uniform pricing. An important difference is that the action in their model and in the subsequent literature is entirely 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.[9] By contrast, our analysis focuses on the consumer who chooses *ex interim* whether to disclose information and it is her voluntary disclosure that facilitates personalized pricing. Moreover, the ability to pool is necessary for every consumer type to benefit from personalized pricing; otherwise, extreme types would be worse off from personalized pricing (unless types are uniformly distributed). Thus, the welfare gains that we study would not emerge in a model where consumer types are commonly known.

Our results complement recent work that has shed light on the role of information in competitive markets with horizontal differentiation. Elliott and Galeotti (2019) show that information can be used to suppress competition: an information-designer can segment the market so that consumers are allocated efficiently while guaranteeing that consumers obtain no surplus. Armstrong and Zhou (2019) study firm-optimal and consumer-optimal information structures for an consumer that does not know her tastes. They show that the consumer-optimal signal may involve learning little so as to amplify price competition.

We focus on personalized pricing based on a consumer's tastes and whether a consumer would like to disclose information directly about those tastes. An important prior

[9]See Armstrong (2006) for a survey of this literature. Liu and Serfes (2004) analyze an information acquisition game, in which the consumer type space is partitioned, and firms can choose to learn about the location of the consumer. They show that if partitions are sufficiently fine, firms will acquire information and price discriminate. We thank Jidong Zhou for drawing our attention to this work.

literature investigates how a consumer may distort her behavior in dynamic settings if firms draw inferences about her tastes from her past choices. This literature has studied a broad range of issues, including whether firms prefer to commit to not personalize prices (Taylor, 2004; Villas-Boas, 2004; Acquisti and Varian, 2005; Calzolari and Pavan, 2006), whether consumers would like to remain anonymous (Conitzer, Taylor, and Wagman, 2012), and how anticipating future pricing may induce consumers to randomize (Bhaskar and Roketskiy, 2019). Bonatti and Cisternas (2019) study the welfare properties of aggregating consumers' past purchasing histories into scores and characterize optimal scoring rules. Our analysis complements this research by studying how a consumer fare from directly controlling the flow of information rather than distorting her behavior to influence the market's perception of her tastes.

Jones and Tonetti (2019) take a "macro approach" to whether consumers should control their own data. They show that, because data is non-rival, there are social gains from multiple firms using the same data simultaneously, and therefore, it is better to let consumers, rather than firms, own and trade data. Several recent papers (Choi, Jeon, and Kim, 2019; Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2019; Bergemann, Bonatti, and Gan, 2019) model a countervailing force where the data of some consumers is predictive about others, and each consumer does not internalize this externality, which induces excessive data-sharing. Fainmesser and Galeotti (2016, 2019) study monopolistic and competitive price discrimination based on consumers' influence on others and their susceptibility to influence. Our analysis abstracts from these important externalities.

Paper outline: Our paper proceeds as follows. Section 2 illustrates our main ideas in a simple example of a monopolist facing a consumer whose valuation is uniformly distributed on the unit interval. Section 3 considers the monopolist environment with a general model of consumer types. Section 4 studies the role of voluntary disclosure in a model of Bertrand competition with horizontal differentiation. Section 5 concludes.

2 Example

A monopolist ("he") sells a good to a single consumer ("she"), who demands a single unit. The consumer's value for that good is v, which is drawn uniformly from $[0, 1]$. If the consumer purchases the good from the monopolist at price p, her payoff is $v - p$ and the monopolist's payoff is p; otherwise, each party receives a payoff of 0. The consumer knows her valuation for the good and the monopolist does not.[10] In this setting, and without any disclosure, the monopolist optimally posts a uniform price of $\frac{1}{2}$, which induces an ex interim consumer surplus of $\max\{v - \frac{1}{2}, 0\}$, and a producer surplus of $\frac{1}{4}$.

We augment this standard pricing problem with voluntary disclosure on the part of the consumer. After observing her value, the consumer chooses a message m from the set of feasible messages for her. The set of *all* feasible messages is $\mathcal{M} \equiv \{[a, b] : 0 \le a \le b \le 1\}$, and we interpret a message $[a, b]$ as *"My type is in the set* $[a, b]$*."* When a consumer's type is v, the set of messages that she can send is $M(v) \subseteq \mathcal{M}$. The evidence structure is represented by the correspondence $M : [0, 1] \rightrightarrows \mathcal{M}$.

Here is the timeline for the game: first, the consumer observes her type v and chooses a message m from $M(v)$. The monopolist then observes the message and chooses a price $p \ge 0$. The consumer then chooses whether to purchase the good. Each party behaves sequentially rationally: we study Perfect Bayesian Equilibria (henceforth PBE) of this game. Our interest is in the implications of this model for simple and rich evidence structures, described below.

Simple evidence: An evidence structure is **simple** if for every v, $M(v) = \{\{v\}, [0, 1]\}$. In other words, each type v can either fully reveal her type using the message $m = \{v\}$ (which is unavailable to every other type), or not disclose anything at all, using the message $m = [0, 1]$ (which is available to every type). Such an evidence structure offers a stylized model for the dichotomy between "track" and "do-not-track": a consumer who opts into tracking will have all of her digital footprint observed by the buyer, whereas do-not-track obscures it entirely.

In this game, there exists an equilibrium in which every type v fully reveals itself using the message $m = \{v\}$, and the monopolist extracts all surplus on the equilibrium path. Off-path, if the consumer sends the non-disclosure message, $m = [0, 1]$, the monopolist believes that $v = 1$ with probability 1, and charges a price of 1. In this equilibrium, all consumers are hurt by the possibility of voluntary disclosure and personalized pricing but the monopolist benefits from it.

But this is not the only equilibrium: there is also one in which every type sends the non-disclosure message $m = [0, 1]$, and the monopolist charges a price of $\frac{1}{2}$. No consumer type wishes to deviate because revealing her true type results in a payoff of 0. Here, both consumer and producer surplus are exactly as in the world without personalized pricing.

[10] While we think of v as the consumer's "valuation," our setting is compatible with v being the consumer's posterior expected value (after observing a signal), as in Roesler and Szentes (2017), and with the consumer learning no more than that. She then faces a choice of whether and how to disclose evidence about that posterior expected value.

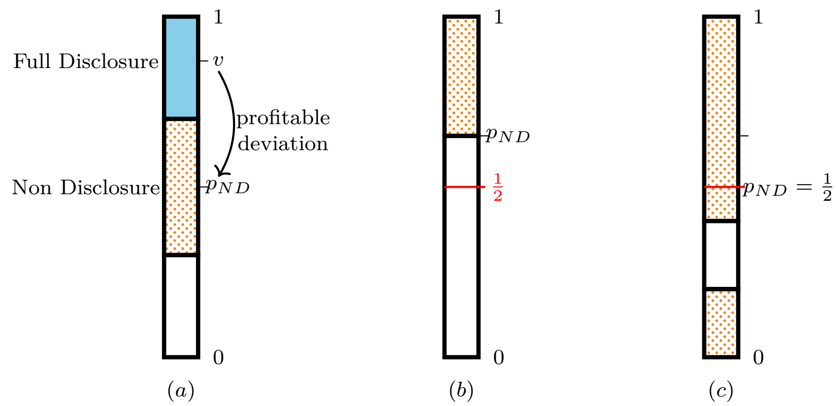

Figure 1: (a) shows that any disclosing type that is strictly higher than p_{ND} *has a profitable deviation* $\Rightarrow$ *the set of non-disclosing types includes* $(p_{ND}, 1]$*. (b) and (c) illustrate different equilibria where the shaded region is the set of non-disclosing types. Across equilibria,* $p_{ND} \geq 1/2$*.*

In fact, there are an uncountable number of equilibria. But *none* of them improve upon the benchmark of no-personalized-pricing from the perspective of *any* consumer type.

Observation 1. *With simple evidence, across all equilibria, the consumer's interim payoff is no more than her payoff without personalized pricing, namely* $\max\{v - 1/2, 0\}$.

The argument is illustrated in Figure 1. In an equilibrium where a positive mass send the non-disclosure message, suppose that the monopolist charges p_{ND} when he receives this message. Any type v that is strictly higher than p_{ND} must send this non-disclosure message because her other option—revealing herself—induces a price that extracts all of her surplus (this property is illustrated in Figure 1(a)). Hence, the set of types that send the non-disclosure message must include $(p_{ND}, 1]$. There are numerous configurations of dislosure and non-disclosure segments that are compatible with this requirement (illustrated in (b) and (c)), but across all of them, the monopolist's optimal non-disclosure price p_{ND} never goes below $\frac{1}{2}$, which is the price charged without personalized pricing.

Rich evidence: Observation 1 illustrates that simple evidence structures and personalized pricing do not benefit the consumer. Now we study how the consumer can do better if she can use a rich evidence structure. An evidence structure is **rich** if for every v, $M(v) = \{m \in \mathcal{M} : v \in m\}$; in other words, a type v can send any interval that contains v. With a rich evidence structure, all the equilibrium outcomes that can be supported using simple evidence are also supportable with this richer language. But now new possibilities emerge, some of which dominate the payoffs from no-personalized-pricing.

We describe an equilibrium that strictly improves consumer surplus for a positive measure of consumer types without making any type worse off. Inspired by Zeno's

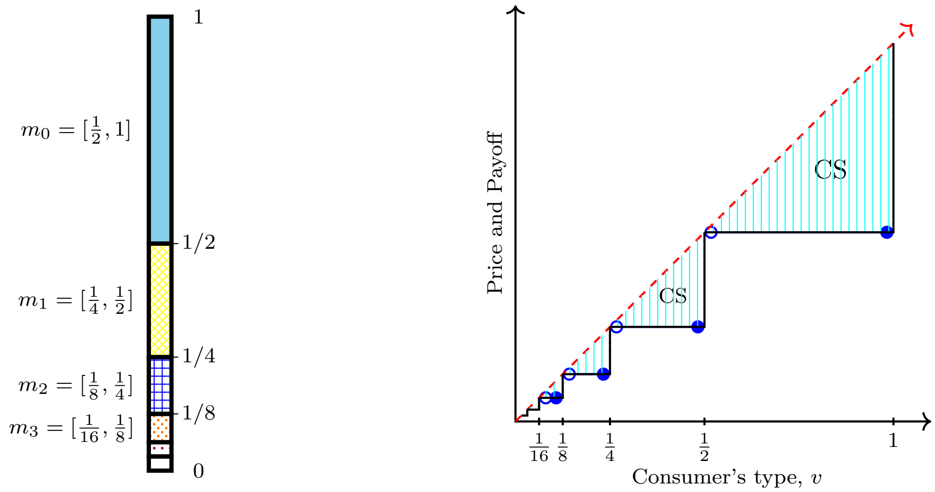

(a) (b)

Figure 2: (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.

Paradox,[11] consider the countable grid $\{1, \frac{1}{2}, \frac{1}{4}, \ldots\} \cup \{0\}$. We denote the $(k+1)^{th}$ element of this ordered list, namely 2^{-k}, by a_k, and the set $m_k \equiv [a_{k+1}, a_k]$. We use this partition to construct an equilibrium segmentation that improves consumer surplus.

Observation 2. *With rich evidence, there exists an equilibrium that generates Zeno's Partition: a consumer's reporting strategy is*

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

When the monopolist receives message m_k, he charges a_{k+1} thereby selling to that entire segment. Relative to no-personalized-pricing, this equilibrium strictly improves consumer surplus for all v in $(0, 1/2]$, and leaves consumer surplus unchanged for all other types.

In this equilibrium, the highest market segment is composed of types in $\left(\frac{1}{2}, 1\right]$, all of which send the message $m_0 \equiv \left[\frac{1}{2}, 1\right]$; the next highest market segment comprises types in $\left(\frac{1}{4}, \frac{1}{2}\right]$, all of which send the message $m_1 \equiv \left[\frac{1}{4}, \frac{1}{2}\right]$, and so on and so forth. We depict this partition in Figure 2. 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 price then is to price at the bottom of the segment. Therefore, trade occurs with probability 1, with each higher consumer type capturing some surplus.

[11] Zeno's Paradox is 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/.

This equilibrium generates an ex ante consumer surplus of $\frac{1}{6}$ and producer surplus of $\frac{1}{3}$, each of which is higher than what is achieved without personalized pricing. All types in $(1/2, 1]$ receive the same price that they would have if personalized pricing were infeasible, and almost every other type is strictly better off.[12] Thus, personalized pricing generates a Pareto improvement for the monopolist and each consumer type.

How is Zeno's Partition supportable as an equilibrium? We first describe how we deter consumers from using messages that are not in Zeno's Partition. We assume that if the monopolist sees such a message, he puts probability 1 on the highest type that could send such a message, and sets a price equal to that type in response to that off-path message. Such beliefs ensure that these off-path messages are not profitable deviations for any consumer type. How about deviations to other on-path messages? For every v in (a_{k+1}, a_k), there exists only one on-path message that it can send, and for every v on the boundaries of such messages, our strategy profile prescribes that they send the message that results in the lower price. Thus, there are no profitable deviations for any consumer type. Finally, we have already discussed how the monopolist's best-response after every equilibrium path message of the form $[a_{k+1}, a_k]$ is to charge a_{k+1}.[13]

It is useful to understand features of our setting that allow us to escape unraveling. 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. Unraveling then emerges as the unique equilibrium outcome as extreme types have a motive to separate from pools. By contrast, in our setting, there exist many pairs of 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 $U[1/2, 1]$. This observation permits us to build pools that do not give types the motive to separate themselves from the pool.

Zeno's Partition isn't the only possible equilibrium of this example. But it turns out to be the equilibrium that maximizes ex ante consumer surplus.[14] We prove in Section 3.4 that with a unidimensional type space, for every equilibrium, there exists an interim payoff-equivalent equilibrium in which trade occurs with probability 1 and types segment into partitions. Thus, it is without loss of generality to restrict attention to

[12] The consumer of type v obtains surplus $v - \left(\frac{1}{2}\right)^{\lfloor \frac{\log v}{\log(1/2)} \rfloor + 1}$ where $\lfloor \cdot \rfloor$ denotes the floor function. This expression is strictly positive for all $v > 0$.

[13] The logic of this market segmentation illustrates the role of hard information: even though types greater than $\frac{1}{2}$ would obtain a lower price by sending the message $\left[\frac{1}{4}, \frac{1}{2}\right]$, they are unable to do so.

[14] That Zeno's Partition is optimal implies that the best consumer-optimal equilibrium delivers payoffs below the consumer-optimal segmentation of Bergemann, Brooks, and Morris (2015).

equilibria that are fully efficient and partitional. Let us illustrate why Zeno's Partition is optimal when types are uniformly distributed using the following heuristic argument.

If consumers purchase with probability 1 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, $\lambda \geq 1/2$. 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 $\overline{P}(\bar{v})$ is the lowest expected price generated by a partition when types are uniformly distributed on the interval $[0, \bar{v}]$:

$$\begin{aligned}
\overline{P}(1) &= \min_{\lambda \geq \frac{1}{2}} (1-\lambda)\lambda + \lambda \overline{P}(\lambda) \\
&= \min_{\lambda \geq \frac{1}{2}} (1-\lambda)\lambda + \lambda^2 \overline{P}(1) \\
&= \min_{\lambda \geq \frac{1}{2}} \frac{(1-\lambda)\lambda}{1-\lambda^2} = \frac{1}{3},
\end{aligned}$$

where the first equality follows from framing the problem recursively, the second follows from $\overline{P}(\lambda)$ being a re-scaled version of the original problem, and the remaining corresponds to algebra. Because Zeno's Partition induces the same expected price, no alternative segmentation can generate higher consumer surplus.

3 Voluntary Disclosure to a Monopolist

3.1 Environment

The Pricing Problem. A monopolist ("he") sells a good to a single consumer ("she"), who demands a single unit. The consumer's type, denoted by t, is drawn according to a measure μ whose support is T. The type space T is a convex and compact subset of a finite-dimensional Euclidean space, $\Re^k$. Each of these k dimensions of a consumer's type reflect attributes that affect her valuation for the good according to $v : T \to \Re$. Payoffs are quasilinear: if the consumer purchases the good from the monopolist at price p when her type is t, her payoff is $v(t) - p$ and the monopolist's payoff is p; otherwise, each player receives a payoff of 0. We denote by F the induced CDF over valuations; in other words, $F(\tilde{v}) \equiv \mu(\{t \in T : v(t) \leq \tilde{v})$. We denote by $\underline{v}$ and $\overline{v}$ the highest and lowest valuations in the support. We simplify exposition by assuming that F is continuous and $F(\underline{v}) = 0$.[15]

Throughout our analysis, we assume that $v(t)$ is non-negative for every type $t \in \mathcal{T}$ and is quasiconvex.[16] A special leading case is where each dimension of t is a consumer characteristic (e.g. income) and $v(t)$ is linear; in this case, $v(t) = \sum_{i=1}^{k} \beta_i t_i$ where β_i is the coefficient on characteristic i. We order types based on their valuations: we say that $t \succeq t'$ if $v(t) \geq v(t')$, and we define $\succ$ and $\sim$ equivalently. When $t \succeq t'$, we refer to t as being a *higher* type.

Were communication infeasible, this pricing problem has a simple solution: the monopolist sets a price p that maximizes $p(1-F(p))$. Let p^* denote the (lowest) optimal price for the monopolist. The consumer's interim payoff is then no more than $\max\{v(t)-p^*, 0\}$.

The Disclosure Game. 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 $\mathcal{M}^{\mathcal{F}} \equiv \{M \subseteq T : M \text{ is closed and convex}\}$, and we interpret a message M in $\mathcal{M}^{\mathcal{F}}$ as meaning *"My type is in the set M."* When a consumer's type is t, the set of messages that she can send is $\mathcal{M}(t) \subseteq \mathcal{M}^{\mathcal{F}}$. We focus attention on the following two different forms of disclosure:

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

In both simple and rich evidence structures, the consumer has access to hard information about her type. In a simple evidence structure, the consumer can either disclose a "certificate" that fully reveals her type or say nothing at all. By contrast, in a rich evidence structure, the consumer can verifiable disclose true statements about her type without being compelled to reveal everything. The assumption that messages are convex sets implies that if types t and t' can disclose some common evidence, then so can any intermediate type $t'' = \alpha t + (1-\alpha)t'$ (for $\alpha \in (0, 1)$).

Timeline and Equilibrium Concept. First, the consumer observes her type t and chooses a message M from $\mathcal{M}(t)$. The monopolist then observes the message and chooses a price $p \geq 0$. The consumer then chooses whether to purchase the good. We study

[15] An equivalent assumption is that $\mu(\{t \in T : v(t) = v'\}) = 0$ for every v' in the range of $v(\cdot)$. Conditions that guarantee this property are that μ is absolutely continuous with respect to the Lebesgue measure, and $v(\cdot)$ is strictly monotone in each dimension.

[16] In other words, for every $\bar{v}$, the set $\{t \in \mathcal{T} : v(t) \leq \bar{v}\}$ is a convex set.

Perfect Bayesian Equilibria (henceforth PBE) of this game. For convenience, we assume that a consumer always breaks her indifference in favor of purchasing the good. We say that a PBE is **efficient** if trade occurs with probability 1, and is **consumer-optimal** if among equilibria, it maximizes ex ante consumer welfare.

3.2 Simple Evidence Does Not Help Consumers

Here, we show that when trading with a monopolist, consumers do not benefit from personalized pricing if the evidence structure is simple relative to a benchmark in which personalized pricing is impossible. As we described above, the interim payoff of each type t without personalized pricing is $\max\{v(t) - p^*, 0\}$ where p^* is the monopolist's optimal price. Relative to that benchmark, we show that there are equilibria with simple evidence that make all consumer types worse off, but no equilibrium in which *any* type is strictly better off. The argument that we use here generalizes that of Section 2.

To see how consumers may be worse off, consider the equilibrium in which the consumer fully reveals her type with probability 1 and the monopolist charges a price of $v(t)$ when the consumer reveals that her type is t; off-path, the seller's beliefs are *maximally skeptical* in that he believes that the consumer's type is one with the highest valuation with probability 1. In this equilibrium, the monopolist extracts all surplus, and so consumers are clearly worse off than without personalized pricing.

But there are also partially revealing equilibria in which only those types below a cutoff reveal themselves. For example, there exists an equilibrium in which all types t where $v(t) \geq p^*$ stay silent and only types below disclose; this results in payoffs for the consumer identical to that without personalized pricing. From an interim perspective, this is the best equilibrium for consumers.

Proposition 1. *With simple evidence, across all equilibria, the consumer's interim payoff is bounded above by* $\max\{v(t) - p^*, 0\}$.

In equilibrium, the seller extracts all surplus from any consumer who reveals her type. Each equilibrium can then be described by a price $\tilde{p}$ that is faced by those who do not. Of course, only those types whose valuation exceeds $\tilde{p}$ choose to purchase at that price. In equilibrium, $\tilde{p}$ is at least p^*; otherwise, the seller would wish to increase prices.

Thus, the consumer gains nothing, *ex ante* and *ex interim*, from the ability to disclose her type using simple evidence. If one takes the model of simple evidence as a stylized representation of track / do-not-track regulations, our analysis implies that this form of consumer protection does not benefit consumers in a monopolistic environment relative to a benchmark that prohibits personalized pricing. Instead, richer forms of verifiable disclosure are needed. We turn to constructing a segmentation with rich evidence that increases consumer surplus, both ex ante and ex interim.

3.3 A Pareto-Improving Segmentation with Rich Evidence

To improve consumer surplus using rich evidence, we develop a segmentation that generalizes that of Section 2. Each segment is constructed so that the monopolist's best-response to that segment is to sell to all consumer types in that segment. Consumers would profit if they could deviate "downwards" to a lower segment; our construction guarantees that this is impossible. Our construction is "greedy" insofar as we start with the highest segment and make each as large as possible without accounting for its effect on subsequent segments.

To define the segmentation strategy, consider a sequence of prices $\{p_s\}_{s=0,1,2,\dots,S}$ where $S \leq \infty$, $p_0 = \overline{v}$, 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 $p_\infty = \underline{v}$. We use these prices to construct sets of types, $(M_s)_{s=1,2,\dots,S} \cup M_\infty$:

$$M_s \equiv \{t \in T : v(t) \leq p_{s-1}\}.$$
$$M_\infty \equiv \{t \in T : v(t) = \underline{v}\}$$

Because v is quasiconvex and T is convex, M_s is a convex set for every $s = 0, 1, 2, \dots, S$, and therefore M_s is a feasible message. These messages segment the market.

Proposition 2. *With rich evidence, there exists a Pareto-improving equilibrium in which a consumer's reporting strategy is*

$$M^*(t) = \begin{cases} M_s & \text{if } p_s < v(t) \leq p_{s-1}, \\ M_\infty & \text{if } t \in M_\infty. \end{cases}$$

When receiving an equilibrium disclosure of the form M_s, the seller charges a price of p_s and sells to all types that send that message.

The segmentation described above generalizes the "Zeno Partition" constructed in Section 2. The highest market segment consists of 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.

Notice that in this segmentation, disclosures aren't taken at face value. Instead, the monopolist infers from receiving a message M_s that the consumer would have preferred to send message M_{s+1} but couldn't, and so her valuation must be in $(p_s, p_{s-1}]$. Notice also that the market segmentation is constructed so that given these beliefs about the consumer's valuation, the monopolist has no incentive to charge a price that excludes any type. In fact, this constraint for the seller binds in our *greedy segmentation* in that if types below p_s were included, the seller's optimal price would exclude those types.

This equilibrium segmentation is fully efficient—trade occurs with probability 1—and improves consumer surplus relative to the benchmark without personalized pricing. Consumer types in the highest market segment face the same price that they would without personalized pricing, but now consumers in other market segments can also purchase at prices that are (generically) below their willingness to pay. Thus, the segmentation is a Pareto improvement. One feature of the segmentation that is attractive is its simplicity: all that consumers have to disclose is information about their willingness to pay.

Finally, we note that this construction is robust to the possibility that the consumer does not have evidence with positive probability, an issue frequently considered in the verifiable disclosure literature (Dye, 1985; Shin, 1994). There exists an equilibrium in this expanded game where if the consumer does not have evidence, she is charged a price of $p_1 = p^*$, and all those with evidence behave as above. If the consumer has evidence, she does not gain from imitating those without evidence.

3.4 Optimal Equilibrium Segmentation

The previous section describes a disclosure strategy that is Pareto-improving and fully efficient. In this section, we explore conditions under which this is the optimal segmentation from the perspective of ex ante consumer surplus.

There are two reasons that this segmentation may not generally maximize ex ante consumer surplus. The first is that it ignores multidimensionality: even if two types have the same valuation, it may be optimal to separate them. The second reason that our segmentation may not be optimal is that even in a one-dimensional world, packing types greedily need not maximize consumer surplus. To elaborate on this second issue, we restrict attention to a one-dimensional model. We first prove that the consumer-optimal equilibrium necessarily features a partitional structure. We use an example to illustrate that the greedy partition may be suboptimal. We then prove that it is optimal for a specific class of distribution functions.

Let the set of types T be identical to the set of values $[\underline{v}, \overline{v}]$ and $v(t) = t$. Recall that $\underline{v} \geq 0$ and that F is an atomless CDF over valuations. The restriction that $M(t)$ is a closed convex set that includes t implies that here, $M(t) = \{[a, b] \subseteq [\underline{v}, \overline{v}] : a \leq t \leq b\}$; in other words, the set of all closed intervals that include t.

When applied to this setting, Proposition 2 identifies an equilibrium segmentation of the form $\{[0, p_s]\}_{s=0,1,2,\ldots,S}$ where p_s is the optimal price to when the distribution F is truncated to $[0, p_{s-1}]$. Because only types in $(p_{s+1}, p_s]$ send the message $[0, p_s]$, a payoff-equivalent segmentation is for a type t to send the message $[p_{s+1}, p_s]$ where $p_{s+1} < t \leq p_s$. This equilibrium is "partitional" in that types reveal the member of the partition to which they belong, and thus, these messages can be taken at "face value". We prove that for any equilibrium, there always exists a payoff-equivalent equilibrium that is partitional and involves the sale happening with probability 1.

Our characterization uses the following definitions. A PBE is efficient if trade occurs with probability 1. A collection of sets $\mathcal{P}$ is a **partition** of $[\underline{v}, \overline{v}]$ if $\mathcal{P}$ is a subset of $\mathcal{M}^F$ such that $\bigcup_{m \in \mathcal{P}} m = [\underline{v}, \overline{v}]$ and for every distinct m, m' in $\mathcal{P}$, $m \bigcap m'$ is at most a singleton. One message m dominates m' (i.e. $m \succeq_{\mathcal{M}} m'$) if for every $t \in m$ and $t' \in m'$, $t \geq t'$; arg min and arg max over a set of messages refers to this partial order. Given a partition $\mathcal{P}$, let $m^{\mathcal{P}}(t) \equiv \arg\min_{\{m \in \mathcal{P}: t \in m\}} m$. An equilibrium σ is **partitional** if there exists a partition $\mathcal{P}$ such that $m^\sigma(t) = m^{\mathcal{P}}(t)$, and for every m in $\mathcal{P}$, $p^\sigma(m) = \min_{t \in m} t$.

Proposition 3. *Given any equilibrium σ, there exists an efficient partitional equilibrium $\tilde{\sigma}$ that is payoff-equivalent for almost every type.*

The proof of Proposition 3 proceeds in two steps. First, we show that it is without loss of generality, from the perspective of consumer surplus, to look at efficient equilibria: for any equilibrium in which there exists a type that is not purchasing the product, there exists an interim payoff-equivalent equilibrium in which that type fully reveals itself to the seller. Second, we show that for any efficient equilibrium, there exists a partitional equilibrium that is payoff-equivalent for almost every type. In this step, we show that in any efficient equilibrium, prices must be (weakly) decreasing in valuation, and otherwise, some type has a profitable deviation.

Thus, it suffices to look at only partitional equilibria. How does the greedy segmentation compare to other partitional equilibria? The greedy partition is ex interim Pareto efficient: any partition that differs from the greedy partition must involve raising the lowest type in at least one segment, which increases prices for at least one type. But, from the perspective of ex ante welfare, 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 using a discrete type space.

Example 1. Suppose that the consumer's type is drawn from $\{1/3, 2/3, 1\}$ where $Pr(t = 1) = 1/6$, $Pr(t = 2/3) = 1/3 + \varepsilon$, and $Pr(t = 1/3) = 1/2 - \varepsilon$, where $\varepsilon > 0$ is small. The greedy construction sets the highest segment as $\{2/3, 1\}$—because the seller's optimal posted price here would be 2/3—and the next segment as $\{1/3\}$. This segmentation results in an average price of $\approx 1/2$. But a better segmentation for ex ante consumer surplus involves the high type perfectly separating as $\{1\}$, and the next highest segment being $\{1/3, 2/3\}$. This segmentation reduces the average price to 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 and Example 1 indicates that this may be sub-optimal. Identifying necessary and sufficient conditions on distributions when such 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 particularly difficult for distributions where we cannot solve for the optimal price in closed-form.[17] A class of distributions where a closed-form solution is available is that of power distributions. In this case, we show that the greedy algorithm identifies the consumer-optimal segmentation.

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

Summary: Our analysis concludes that in a monopolistic setting, (i) the combination of voluntary disclosure and personalized pricing does not benefit consumers if evidence is simple (Proposition 1), but (ii) generates an ex interim Pareto improvement if evidence is rich (Proposition 2). Thus, consumers' control over data benefits them when they can choose not only *whether* to communicate but also *what* to communicate.

[17] Without solving for the closed-form, we can verify that the greedy algorithm is optimal if (i) F is convex, and (ii) the optimal price on an interval $[0, \tilde{v}]$, denoted by $p(\tilde{v})$, has a slope bounded above by 1 and is weakly concave.

4 How Disclosure Amplifies Competition

Our analysis above identifies when and how consumers benefit from voluntary disclosure and personalized pricing when facing a monopolist. However, in many settings, consumers do not interact with only one seller but instead face a competitive market in which firms are differentiated. An important characteristic of each consumer then is her *location*, i.e., her tastes for the products made by each firm. In this section, we study the degree to which voluntary disclosure and personalized pricing benefits a consumer in a model of Bertrand duopoly with horizontally differentiated products.

Our analysis identifies a new strategic force absent in a monopolistic market: a consumer can use voluntary disclosure to amplify competition between firms. This force is sufficiently strong that all consumer types benefit from personalized pricing even with simple evidence. Richer evidence generates even stronger gains, where we can use an algorithm analogous to that of Proposition 2 to construct equilibria. Our analysis formalizes the intuition, articulated by Lars Stole (quoted in Wallheimer 2018), that targeting and personalized pricing benefit consumers in competitive markets:

> *"A competitor can quickly undercut a targeted price. Once you start doing this, you'll have companies in different markets matching those prices. You don't have much market power."*

We proceed as follows. Section 4.1 describes the market setting. Section 4.2 constructs equilibrium segmentations with simple and rich evidence. Section 4.3 compares the consumer's payoffs with those of a benchmark setting without personalized pricing.

4.1 Environment

Two firms, L and R, compete to sell to a single consumer who has unit demand. The type of the consumer is her *location*, denoted by t, which is drawn according to measure μ (and cdf F) with the 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 when purchasing from firm i, she faces a "transportation cost" that is proportional to the distance between her location and that of the firm's, ℓ_i.[18] Thus, her payoff from buying the good from firm i at a price of p_i is $V - |t - \ell_i| - p_i$. As is standard, we assume that

[18]Linear transportation costs simplify algebra, but our results do not hinge on it. Identical results emerge for quadratic distance costs.

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.[19]

Disclosure with Duopoly. After observing her type, the consumer chooses a message M that is feasible and available for her to send to each of the firms. As before, the set of feasible messages is $\mathcal{M}^{\mathcal{F}} \equiv \{[a,b] : -1 \leq a \leq b \leq 1\}$ where a message $[a,b]$ is interpreted as "*my type is in the interval* $[a,b]$." When a consumer's type is t, the set of messages that she can send is $\mathcal{M}(t) \subseteq \mathcal{M}^{\mathcal{F}}$. 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 when the type space is unidimensional. The novelty here is that the consumer now sends two messages—M_L to firm L and M_R to firm R—and 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 $[-1,1]$.

Timeline and Equilibrium Concept. The consumer first observes her type t and then chooses a pair of messages (M_L, M_R), each from $\mathcal{M}(t)$.[20] Each firm i privately observe its message M_i and sets price $p_i \geq 0$; price-setting is simultaneous. The consumer then chooses which firm to purchase the good from, if any.

We study Perfect Bayesian Equilibria of this game. As is well-known (Osborne and Pitchik, 1987; Caplin and Nalebuff, 1991), the price-setting game in Bertrand competition with horizontal differentiation may lack a pure-strategy equilibrium for general distributions. By contrast, we show constructively that pure-strategy equilibria always exist when this market setting is augmented with a disclosure game.

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

[19] See Osborne and Pitchik (1987), Caplin and Nalebuff (1991), Bester (1992), and Peitz (1997). For most of our analysis, it suffices for $V \geq 2$, so that a consumer is always willing to purchase the good from the most distant firm if that distant firm sets a price of 0.

[20] 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 (2019), and chooses whether and how to disclose that expected location using simple or rich evidence.

strategic logic. Each consumer reveals her type to the firm that is more distant from her, indicating 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. The seller who does not obtain a fully revealing message infers that the consumer is closer to his location. Based on that inference, this seller sets a profit-maximizing price subject to the consumer having the option to buy from the other seller at a price of 0. We use the assumption that $V \geq 2$ to guarantee that the consumer weakly prefers purchasing the good from the distant firm at a price of 0 to not purchasing it at all. We begin our analysis with a fully revealing equilibrium in both simple and rich evidence environments, and then show how to improve upon it.

Proposition 5. *There exists a fully revealing equilibrium in both simple and rich evidence games: every type of consumer t sends the message $\{t\}$ to each firm, and purchases from the firm nearer to her at a price of $2|t|$.*

The logic of Proposition 5 is straightforward. In an equilibrium where the consumer reveals her location to each firm, both firms do not charge her strictly positive prices in equilibrium. Standard Bertrand logic implies that the distant firm must charge her a price of 0 and the closer firm charges her the highest price that it can subject to the constraint that the consumer finds it incentive-compatible to purchase from the closer firm at that price.[21] If the consumer deviates by sending a message M that isn't a singleton to firm i, then 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, thus, involves each seller holding skeptical 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 suffer from the firm closer to them being able to charge a high price. Ideally, types that are located close to firm i would benefit from pooling with types more distant from firm i. 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.

Our construction uses the following notation. Let p_1^i be the lowest maximizer of $p\ell_i(F(\ell_i) - F(p\ell_i/2))$, and let $t_1^i \equiv p_1^i \ell_i/2$. To provide some intuition, 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; in other words, this is firm i's optimal *local monopoly price*

[21] Once types are revealed, these equilibrium prices necessarily coincide with those of Thisse and Vives (1988), where the consumer's type is common knowledge.

against an outside option where firm j charges a price of 0. At those prices, firm i expects a probability of trade $\ell_i(F(\ell_i) - F(p\ell_i/2))$ and t_1^i is the most distant type from firm i that still purchases from firm i. It is necessarily the case that $-1 < t_1^L < 0 < t_1^R < 1$. We use these types to describe our equilibrium.

Proposition 6. *With simple evidence, there exists a partially pooling equilibrium in which the consumer's reporting strategy is*

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

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, every consumer type purchases from the seller nearer to her.

An intuition for Proposition 6 is as follows. If the consumer is centrally located—i.e., in (t_L^1, t_R^1)—she disclose her type ("track") to both firms. These consumers then benefit from intense price competition, exactly as in the fully revealing equilibrium of Proposition 5. If the consumer is not centrally located, she reveals her location to the firm farther from her but not to the nearer one. This private messaging strategy guarantees that the distant firm prices at zero and offers an attractive outside option. The firm that receives an uninformative ("don't 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. This pool of extreme types improves consumer welfare by guaranteeing that extreme consumer types can pool with moderate types thereby decreasing type-contingent prices relative to the fully revealing equilibrium.[22] We depict this disclosure strategy in Figure 3.

One can do even better with rich evidence by using a segmentation that is analogous to the "Zeno Partition" constructed in the monopolistic market. In this case, the central type, $t = 0$, obtains equilibrium prices of 0 from each firm, and plays a role similar to the lowest type in the monopolistic setting. Accordingly, one sees a segmentation that

[22] When types are uniformly distributed, this equilibrium has a particularly intuitive form. The cutoff types in this case are symmetrically $t_L^1 = -1/2$ and $t_R^1 = 1/2$. Consequently, types $t \in [-1, -1/2]$ and $t \in [1/2, 1]$ purchase the good at a price of 1, rather than $2|t|$.

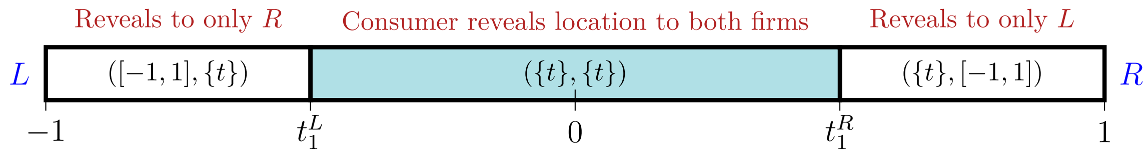

Figure 3: The figure shows disclosure strategies for every type. Centrally located types fully reveal location to both firms. Extreme types reveal location only to the distant firm and conceal it from the closer firm

goes from the extremes to the center, and becomes arbitrarily fine as one approaches the center. To develop notation for this argument, let us define a sequence of types $\{t_s^i\}_{s=0,1,2,\dots}$ and 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 $p\ell_i(F(t_{s-1}^i) - F(p\ell_i/2))$.
- $M_s^i \equiv \{t \in [-1,1] : t_s^i \ell_i \leq t\ell_i \leq t_{s-1}^i \ell_i\}$.

Let $p_\infty^i = 0$ and let $M_\infty^i = \{0\}$. We have thus defined a sequences of cutoffs, prices, and messages where at every stage, we are constructing segments greedily: given a segment M_s^i, firm i is charging the price that is the optimal local monopoly price (assuming that the other firm charges a price of 0), and at this price, firm i is servicing all consumer types in M_s^i. Because rich evidence allows consumers to disclose intervals directly, our disclosure strategy need not be asymmetric (unlike our analysis of the segmentation with simple evidence): a consumer of type t can send the message M_s^i that contains t to both firms. We use this notation to prove our result below.

Proposition 7. *With rich evidence, there exists a segmentation equilibrium in which a consumer's reporting 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.

This equilibrium construction highlights the versatility of rich evidence disclosure. While the competitive environment differs from the monopolistic setting in many ways, the logic of the "Zeno Partition" strategy follows in much the same way. Consumers with the highest willingness to pay for the good from firm i are segmented together and send messages M_1^i. That message induces a price of 0 by firm j and given that outside option, firm i charges a price that makes indifferent the consumer type in M_1^i with the lowest willingness to pay for firm i's product. Prices diminish as the consumer types become

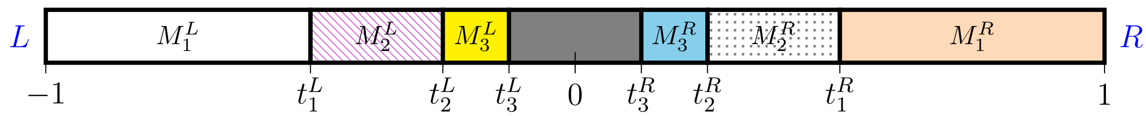

Figure 4: The figure shows a segmentation using rich evidence. The types in (t_3^L, t_3^R) are partitioned into countably infinitely many segments, and hence these segments are omitted.

closer to the center. As such, the segmentation follows iteratively from both sides of 0 exactly as in "Zeno."[23] We depict this segmentation strategy in Figure 4.

We have constructed equilibria with simple and rich evidence but we do not argue that these equilibria are consumer-optimal for an important reason: an equilibrium segmentation (with either simple or rich evidence) may generate segments that induce each firm to randomize in its pricing strategy. Characterizing or bounding prices across mixed strategy equilibria across segments appears intractable.[24] Instead, we compare these equilibria to that of a benchmark model without personalized pricing and show that these equilibria generate strict interim Pareto gains.

4.3 Benefits of Personalized Pricing in Competitive Markets

The benchmark is the standard model of Bertrand pricing with horizontal differentiation: each firm i sets a uniform price p_i and the consumer buys from one of the firms. Unfortunately, this game may lack a pure-strategy equilibrium (in prices), and characterizing the mixed strategy equilibria is challenging. Accordingly, we impose a distributional assumption that is standard in this setting, namely 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 a range of distributions, including uniform and Beta distributions (Bagnoli and Bergstrom, 2005).

With this assumption, a symmetric pure-strategy equilibrium in this benchmark setting exists and involves each firm charging a price of p^* where

$$p^* \equiv \arg\max_p pF\left(\frac{p^* - p}{2}\right) = \frac{2F(0)}{f(0)},$$

where the first equality is firm L's profit maximization problem, and the second comes from solving its first-order condition and substituting $p = p^*$.[25] Our welfare result com-

[23] For the uniform distribution, the construction mirrors that in Section 2 where $t_s^i = \ell_i(1/2)^s$.

[24] Restricting attention to segmentations that generate pure-strategy equilibria, we conjecture that the equilibrium that we construct in Proposition 6 is consumer-optimal in the game with simple evidence; similarly, we conjecture the same regarding the equilibrium constructed in the game with rich evidence whenever the greedy algorithm yields an optimal segmentation.

[25] As before, we assume that V is sufficiently high that all consumers purchase at these prices. It

Figure 5: The figure compares the interim equilibrium cost (incl. price and transport cost) in the setting without personalized pricing with that of the equilibrium constructed in the simple evidence game (Proposition 6) for uniformly distributed types. Simple evidence reduces the expected cost by 50%.

pares this price to those of the equilibria constructed in Section 4.2.

Proposition 8. *If f is symmetric around* 0 *and log-concave, then every type has a strictly higher payoff in the equilibria of the simple and rich evidence games constructed in Propositions 6 and 7 than in the benchmark setting without personalized pricing.*

The logic of Proposition 8 is that the price in the benchmark setting (p^*) is strictly higher than p_1^i, the price charged by firm i to a consumer who conceals her type from firm i in the equilibrium of the simple evidence game (Proposition 6). 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 5.

We note that the ability to pool is needed for these interim Pareto gains and these gains would not generally emerge if the consumer fully revealed her type: apart from the uniform distribution, any symmetric log-concave density must involve $f(0) > 1/2$ and therefore, prices in the benchmark are strictly less than 2. By contrast, in the fully revealing equilibrium, the extreme types pay a price of 2. Thus, these interim Pareto gains emerge when consumers can disclose or conceal evidence about their type, and not necessarily when types are commonly known.

suffices that $V > (f(0))^{-1} + 1$.

5 Conclusion

As the digital economy matures, policymakers and industry leaders alike are working to establish norms and regulations to govern data ownership and transmission. In light of the privacy and 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?* We frame and answer this question using the language of voluntary disclosure, building on a rich theoretical literature on evidence and hard information.

Our initial instinct was that voluntary disclosure would not help. As the market draws inferences based on information that is *not* disclosed, giving consumers the ability to separate themselves would seem to be self-defeating. To put it differently, if the market necessarily unravels as in Grossman (1981) and Milgrom (1981), consumers retain no surplus and may be worse off with personalized pricing. We show that this conclusion is incorrect because it omits two important strategic forces present in market interactions.

First, one can construct pools in both monopolistic and competitive settings in which the consumer lacks an incentive to separate herself from the pool. These pools are simple, do not require commitment, and depend only on the willingness-to-pay rather than on intricate details of the type space. Second, when facing multiple firms, voluntary disclosure and personalized pricing amplify competitive forces. By revealing features of one's preferences to the market, the consumer obtains a significant price concession from a less competitive firm that forces the more competitive firm to also lower her price.

We have examined these basic strategic considerations through the lens of a stylized model. A separate question is about the technological feasibility of various data-sharing arrangements. As argued by Goldfarb and Tucker (2019) and others, an important element in the ongoing evolution of the digital economy is its increasing ability to verify information. These advances suggest that it may be technologically feasible for consumers to use intermediaries or platforms to verifiably disclose aspects of their preference, or at the very least, decide how much they would like to be tracked, and by whom. Offering consumers control over their data—and giving them tools to coordinate their sharing of data—may make personalized pricing attractive and improve consumer welfare.

References

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2019): "Too Much Data: Prices and Inefficiencies in Data Markets," Working Paper.

Acquisti, A., C. Taylor, and L. Wagman (2016): "The Economics of Privacy," *Journal of Economic Literature*, 54, 442–492.

Acquisti, A. and H. R. Varian (2005): "Conditioning prices on purchase history," *Marketing Science*, 24, 367–381.

Armstrong, M. (2006): *Recent Developments in the Economics of Price Discrimination*, Cambridge University Press, vol. 2 of *Econometric Society Monographs*, 97–141.

Armstrong, M. and J. Zhou (2019): "Information Structure and Price Competition," Working Paper.

Bagnoli, M. and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic theory*, 26, 445–469.

Ben-Porath, E., E. Dekel, and B. L. Lipman (2017): "Disclosure and choice," *The Review of Economic Studies*, 85, 1471–1501.

——— (2019): "Mechanisms with evidence: Commitment and robustness," *Econometrica*, 87, 529–566.

Ben-Porath, E. and B. L. Lipman (2012): "Implementation with partial provability," *Journal of Economic Theory*, 147, 1689–1724.

Bergemann, D. and A. Bonatti (2019): "Markets for Information: An Introduction," *Annual Review of Economics*, 11, 85–107.

Bergemann, D., A. Bonatti, and T. Gan (2019): "The Economics of Social Data," Working Paper.

Bergemann, D., B. Brooks, and S. Morris (2015): "The limits of price discrimination," *The American Economic Review*, 105, 921–957.

Bester, H. (1992): "Bertrand Equilibrium in a Differentiated Duopoly," *International Economic Review*, 33, 433–448.

Bhaskar, V. and N. Roketskiy (2019): "Consumer Privacy and Serial Monopoly," Working Paper.

Bonatti, A. and G. Cisternas (2019): "Consumer Scores and Price Discrimination," *Review of Economic Studies*.

Calzolari, G. and A. Pavan (2006): "On the optimality of privacy in sequential contracting," *Journal of Economic theory*, 130, 168–204.

Caplin, A. and B. Nalebuff (1991): "Aggregation and imperfect competition: On the existence of equilibrium," *Econometrica*, 59, 25–59.

Choi, J. P., D.-S. Jeon, and B.-C. Kim (2019): "Privacy and personal data collection with information externalities," *Journal of Public Economics*, 173, 113–124.

Conitzer, V., C. R. Taylor, and L. Wagman (2012): "Hide and seek: Costly consumer privacy in a market with repeat purchases," *Marketing Science*, 31, 277–292.

Dye, R. A. (1985): "Disclosure of nonproprietary information," *Journal of accounting research*, 123–145.

Elliott, M. and A. Galeotti (2019): "Market segmentation through information," Working Paper.

Fainmesser, I. P. and A. Galeotti (2016): "Pricing network effects," *The Review of Economic Studies*, 83, 165–198.

——— (2019): "Pricing network effects: Competition," *American Economic Journal: Microeconomics*.

Federal Trade Commission (2012): "Protecting Consumer Privacy in an Era of Rapid Change," https://tinyurl.com/y6euao98.

Glode, V., C. C. Opp, and X. Zhang (2018): "Voluntary disclosure in bilateral transactions," *Journal of Economic Theory*, 175, 652–688.

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

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 F. Koessler (2017): "Simple versus rich language in disclosure games," *Review of Economic Design*, 21, 163–175.

Hagenbach, J., F. Koessler, and E. Perez-Richet (2014): "Certifiable Pre-Play Communication: Full Disclosure," *Econometrica*, 82, 1093–1131.

Haghpanah, N. and R. Siegel (2019): "Pareto Improving Segmentation of Multi-product Markets," Working Paper.

Hart, S., I. Kremer, and M. Perry (2017): "Evidence games: Truth and commitment," *American Economic Review*, 107, 690–713.

Hern, A. (2018): "No tracking, no revenue: Apple's privacy feature costs ad companies millions," https://tinyurl.com/y9a9u4jm.

Hidir, S. and N. Vellodi (2019): "Personalization, Discrimination and Information Revelation," Working Paper.

Ichihashi, S. (2019): "Online Privacy and Information Disclosure by Consumers," *American Economic Review*.

Jones, C. I. and C. Tonetti (2019): "Nonrivalry and the Economics of Data," Working Paper.

Jovanovic, B. (1982): "Truthful disclosure of information," *The Bell Journal of Economics*, 36–44.

Kartik, N. and O. Tercieux (2012): "Implementation with evidence," *Theoretical Economics*, 7, 323–355.

Koessler, F. and V. Skreta (2019): "Selling with evidence," *Theoretical Economics*, 14, 345–371.

Liu, Q. and K. Serfes (2004): "Quality of Information and Oligopolistic Price Discrimination," *Journal of Economics & Management Strategy*, 13, 671–702.

Lizzeri, A. (1999): "Information revelation and certification intermediaries," *The RAND Journal of Economics*, 214–231.

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

Osborne, M. J. and C. Pitchik (1987): "Equilibrium in Hotelling's Model of Spatial Competition," *Econometrica*, 55, 911–922.

Peitz, M. (1997): "Differentiated Bertrand duopoly with variable demand," *Research in Economics*, 51, 85–100.

Roesler, A.-K. and B. Szentes (2017): "Buyer-optimal learning and monopoly pricing," *American Economic Review*, 107, 2072–80.

Sher, I. and R. Vohra (2015): "Price discrimination through communication," *Theoretical Economics*, 10, 597–648.

Shin, H. S. (1994): "The burden of proof in a game of persuasion," *Journal of Economic Theory*, 64, 253–264.

Stahl, K. and R. Strausz (2017): "Certification and market transparency," *The Review of Economic Studies*, 84, 1842–1868.

Taylor, C. R. (2004): "Consumer privacy and the market for customer information," *RAND Journal of Economics*, 631–650.

Thisse, J.-F. and X. Vives (1988): "On The Strategic Choice of Spatial Price Policy," *The American Economic Review*, 78, 122–137.

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," *The RAND Journal of Economics*, 35, 486–501.

Wallheimer, B. (2018): "Are you ready for personalized pricing?" *Chicago Booth Review*.

A Appendix

Proof of Proposition 1 on p. 14. Consider an equilibrium. Let $\tilde{T}$ be the set of types that in equilibrium send the non-disclosure message, T. Thus, every type in $T \backslash \tilde{T}$ sends a message that fully reveals itself. Sequential rationality demands that the monopolist charges a price of $v(t)$ to every such type, leading to an interim payoff of 0. We prove below that the non-disclosure message must induce a price that is no less than p^*.

Suppose towards a contradiction that it leads to a price $\tilde{p}$ that is strictly less than p^*. In equilibrium, if $v(t) > \tilde{p}$, the consumer must be sending the non-disclosure message T (because sending the message $\{t\}$ leads to a payoff of 0). Therefore, in equilibrium,

$$\tilde{T} \supseteq \{t \in T : v(t) > \tilde{p}\} \supseteq \{t \in T : v(t) \geq p^*\}.$$

By charging a price of $\tilde{p}$, the firm's payoff is

$$\begin{aligned}\tilde{p}\mu(\{t \in \tilde{T} : v(t) \geq \tilde{p}\}) &\leq \tilde{p}\mu(\{t \in T : v(t) \geq \tilde{p}\}) \\ &< p^*\mu(\{t \in T : v(t) \geq p^*\}) \\ &= p^*\mu(\{t \in \tilde{T} : v(t) \geq p^*\}),\end{aligned}$$

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

Proof of Proposition 2 on p. 15. We augment the description of the strategy-profile with the off-path belief system where when the seller receives an off-path message $M \notin (\cup_{s=1,\dots,S} M_s) \cup M_\infty$, she puts probability 1 on a type in M with the highest valuation (i.e. a type in $\arg\max_{t \in M} v(t)$), and charges a price equal to that valuation.

Observe that the seller has no incentive to deviate from this strategy-profile because for each (on- or off-path) message, the price that he is prescribed to charge in equilibrium is her optimal price given the beliefs that are induced by that message.

We consider whether the consumer has a strictly profitable deviation. Let us consider on-path messages first. Consider a consumer type t that is prescribed to send message M_s where $p_s < v(t) \leq p_{s-1}$. Sending any message of the form $M_{s'}$ where $s' < s$ results in a higher price and therefore is not a profitable deviation. All messages of the form $M_{s'}$ where $s' > s$ are infeasible because $t \notin M_{s'}$ for any $s' > s$. Finally, if the type t is such that she is prescribed to send message M_∞, 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 t, $v(t) \leq \max_{t' \in M} v(t')$, any off-path message is guaranteed to result in a payoff of 0. □

Proof of Proposition 3 on p. 17. Consider an equilibrium σ. Let $m^\sigma(t)$ denote the message reported by type t, let $F^\sigma_m \in \Delta[0,1]$ denote the firm's belief when receiving message

m and $\underline{t}^\sigma(m)$ be the lowest type in the support of that belief, and let $p^\sigma(m)$ be the sequentially rational price that he charges. In any equilibrium, $p^\sigma(m) \geq \underline{t}^\sigma(m)$, 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

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

In this disclosure strategy profile, a consumer-type that doesn't buy in equilibrium σ is fully revealing herself in $\tilde{\sigma}$. Because σ is an equilibrium, and the pricing strategy remains unchanged, such a type purchases in $\tilde{\sigma}$ at price v, and thus, efficiency is guaranteed without a change in payoffs.

We argue that $\tilde{\sigma}$ is an equilibrium. Note that because σ is an equilibrium, and we have not changed the price for any message, no consumer-type has a motive to deviate. We also 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,

$$p^\sigma(m)(1 - F^\sigma_m(p^\sigma(m))) \geq p(1 - F^\sigma_m(p)) \text{ for every } p. \tag{1}$$

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 $p(1 - F^{\tilde{\sigma}}_m(p))$. But observe that by Baye's Rule, for every $p \geq p^{\tilde{\sigma}}(m)$,

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

Thus (1) implies that $p^{\tilde{\sigma}}(m) \geq p(1 - F^{\tilde{\sigma}}_m(p))$ 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.

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

Proof. In an efficient equilibrium σ, trade occurs with probability 1. Therefore, for every equilibrium-path message, m, the price charged by the monopolist after that message, $p^\sigma(m)$, must be no more than the lowest type in the support of his beliefs after receiving message m, $\underline{t}^\sigma(m)$ (recall that $v(t) = t$). Sequential rationality of the monopolist demands that $p^\sigma(m)$ is at least $\underline{t}^\sigma(m)$ (because charging strictly below can always be improved), and therefore, in an efficient equilibrium, $p^\sigma(m) = \underline{t}^\sigma(m)$.

Step 1: We first prove that the set of types being charged an equilibrium-path price p is a connected set. Suppose that types t and $t'' > t$ are sending (possibly distinct) equilibrium-path messages m and m'' such that $p^\sigma(m) = p^\sigma(m'')$. Because $p^\sigma(m) = \underline{t}^\sigma(m)$ and $p^\sigma(m'') = \underline{t}^\sigma(m'')$, it follows that $\underline{t}^\sigma(m) = \underline{v}^\sigma(m'') < v < v''$. Because types arbitrarily close to $\underline{t}^\sigma(m'')$ and v'' are both sending the message m'', the message m'' contains the interval $[\underline{t}^\sigma(m''), t'']$.

Consider any type t' in $[t, t'']$: because $[t, t''] \subseteq [\underline{t}^\sigma(m''), v''] \subseteq m''$, it follows that m'' is a *feasible message* for type t'. Therefore, denoting m' as the equilibrium-path message of type t', type t' 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 $p^\sigma(m') = \underline{t}^\sigma(m')$ that $\underline{t}^\sigma(m') < \underline{t}^\sigma(m) \leq t \leq t'$. Therefore, the interval $[\underline{t}^\sigma(m'), t']$ is both a subset of m' and contains t, and hence, m' is a feasible message for type t. But then, type t 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

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

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 $T^\sigma(p)$ 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 $T^\sigma(p)$.

Now consider the collection of sets

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

where $cl(\cdot)$ is the closure of a set. We argue that $\mathcal{P}^\sigma$ is a partition of $[\underline{v}, \overline{v}]$: clearly, $[\underline{v}, \overline{v}] \subseteq \bigcup_{m \in \mathcal{P}^\sigma} m$, and because each of $T^\sigma(p)$ and $T^\sigma(p')$ are connected for equilibrium-path prices p and p', $cl(T^\sigma(p)) \bigcap cl(T^\sigma(p'))$ is at most a singleton.

Consider a strategy-profile $\tilde{\sigma}$ where each type t sends the message $m^{\mathcal{P}^\sigma}(t)$. 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 $m = cl(T^\sigma(p))$. 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 $T^\sigma(p)$, 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 t such that there exists a unique element in $\mathcal{P}^\sigma$ that contains t, there exists no other feasible message that is an equilibrium-path message for $\tilde{\sigma}$. For any other type t, the strategy of sending the message $m^{\mathcal{P}^\sigma}(t)$ ensures that type t 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 on p. 18. 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 $\tilde{\sigma}$:

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

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 $\{0\} \bigcup_{\ell=0}^{\infty} S_\ell$ 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 dividing S_ℓ into two segments $\left[\frac{1}{\gamma^{\ell+1}}, \tilde{v}\right]$ and $\left(\tilde{v}, \frac{1}{\gamma^\ell}\right]$ for some $\tilde{v} \in \left(\frac{1}{\gamma^{\ell+1}}, \frac{1}{\gamma^\ell}\right)$. The higher segment is charged $\tilde{v}$. The lowest possible price that the lower segment is charged is $\frac{\tilde{v}}{\gamma}$, which is achieved if all types in $\left[\frac{\tilde{v}}{\gamma}, \tilde{v}\right]$ send the same message. 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. □

Proof of Proposition 5 on p. 21. Given 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 type is well-defined because M is closed. Let δ_t denote the degenerate probability distribution that places probability 1 on type t. We use this notation to construct a fully revealing equilibrium:

- The consumer of type t always 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.

To see that each firm does not wish to deviate from charging the above prices, 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 in location to firm i and therefore $2t\ell_j < 0$. 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 purchase instead from the other firm), and charging a price p weakly below $2t\ell_i$ leads to a payoff of p (because the consumer always breaks ties in favor of the closer firm). Therefore, firm i has no incentive to deviate. 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 (because the consumer will purchase the good from the other firm). Therefore, in either case, firm i has no incentive to deviate.

Finally, we argue that the consumer has no incentive to deviate. By sending a fully revealing message, $\{t\}$, the consumer obtains an equilibrium payoff of $V - (|t| + 1)$. If $t \leq 0$, the consumer obtains a price of 0 from firm R, which is the lowest possible price. Therefore, there is no incentive to send any other message to firm R. 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 strictly profitable deviation from sending any other message $M \in M(t)$ to firm L. An analogous argument implies that if the consumer's type is $t > 0$, she also does not gain from deviating to any other feasible disclosure strategy.

□

Proof of Proposition 6 on p. 22. We first show given the pricing strategies that the consumer has no incentive to deviate.

Consider a consumer type t such that $t \in (t_1^L, t_1^R)$, or in other words, $t\ell_i < t_1^i\ell_i$. The equilibrium strategies are that the consumer sends the message $\{t\}$ to each firm. Given these equilibrium strategies, the consumer is quoted 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 $p_1^i = 2t_1^i\ell_i$, which is strictly higher. Therefore, this deviation is not strictly profitable.

Now suppose that $t\ell_i \geq t_1^i\ell_i$. The equilibrium strategies are that the consumer sends message $[-1, 1]$ to firm i and $\{t\}$ to firm j. Because the consumer, in equilibrium, is quoted a price of 0 by firm j, sending the other message cannot lower that price. Given the equilibrium message, the consumer is quoted a price of $2t_1^i\ell_i$ by firm i, and deviating leads to a weakly higher price of $2t\ell_i$. Therefore, this deviation is not strictly profitable.

We now consider whether firm i has an incentive to deviate. It follows from the proof of Proposition 5 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 favor 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 favor 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. □

Proof of Proposition 7 on p. 23. 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 5). 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 $t_s^i \ell_i < t\ell_i \leq t_{s-1}^i \ell_i$ for some $s = 1, 2, \ldots$. 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 $p_s^i = 2t_s^i \ell_i$ 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 $2\tau(i, M)\ell_i \geq 2t\ell_i > 2t_s^i \ell_i = p_s^i$. 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 favor 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 $t_s^i \ell_i < \tau(i, M)\ell_i \leq t_{s-1}^i \ell_i$. Because the consumer breaks ties in favor of the closer firm, firm i anticipates that the consumer will reject any strictly positive price. □

Proof of Proposition 8 on p. 25. We show that $p^* > p_1^i$ for every i. Observe that

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

where the first equality follows from the first-order condition that p_1^L 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 6). 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 7). 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^*.

□