Voluntary Disclosure and Personalized Pricing*

Preliminary and Incomplete

S. Nageeb Ali[†] Greg Lewis[‡] Shoshana Vasserman[§]

November 16, 2018

Abstract

A concern central to the economics of privacy is that firms may use consumer data to price discriminate. A common response to these concerns is that consumers should be given control over their own data, choosing which firms access it and what they access. Since the market draw inferences based on both the data it sees and the consumer's disclosure choice, the strategic implications of this proposal are unclear. We investigate whether such measures improve consumer welfare in both monopolistic and competitive environments, and find that a consumer can use verifiable information to create exclusive partial pools that guarantee gains relative to both perfect price discrimination and no personalized pricing.

*We thank Nima Haghpanah, Kevin He, Navin Kartik, Vijay Krishna, Andres Santos, Ron Siegel, Joel Sobel, and Glen Weyl for useful suggestions, and are grateful to Microsoft Research for hosting Ali (as a Visiting Researcher) and Vasserman (as an intern), during which time this research began.

[†]Department of Economics, Penn State. Email: nageeb@psu.edu.

[‡]Microsoft Research New England. Email: glewis@microsoft.com.

[§]Department of Economics, Harvard University. Email: svasserman@fas.harvard.edu.

1 Introduction

Motivation: Recent technological developments have enhanced firms' ability to learn about a consumer's characteristics, what she values, whether she is shopping around for better deals, etc., and to offer her a personalized price. Personalized pricing helps firms target the consumer with a price and product that matches her interest but raises the distributional concern that firms extract too much of her surplus. This prospect has spurred an ongoing debate on the role of privacy in markets, and the potential importance of protecting consumer privacy. Firms themselves have sometimes hesitated to personalize prices based on consumer data for fear that it might antagonize and alienate consumers (Wallheimer, 2018).

A leading response to this debate has been to suggest that consumers be given more control over their data, for them to choose what is shared with firms, and whether firms can personalize prices based on that data. This response is consistent with recent regulation, such as the European Union's General Data Protection Regulation (GDPR), that requires firms to store personal data in an anonymized way and process it only with consumer consent. As another example, Apple recently changed its Safari browser to limit the ways in which its user's activities were tracked by third parties (Hern, 2018). All of these efforts are an attempt to empower consumers. In this spirit, Acquisti, Taylor, and Wagman (2016), in their recent survey on the economics of privacy, argue that *"Privacy is not the opposite of sharing—rather, it is control over sharing."*

This paper investigates the economic implications of giving consumers control over information in monopolistic and competitive markets. We study a setting in which consumers have their characteristics encoded as hard information that they can disclose partially or fully to the marketplace. We investigate the implications of two forces. First, consumers cannot commit to a disclosure policy: in searching for the best possible deal, a consumer will disclose whatever information she can to secure the best possible price. Second, consumers face sophisticated firms who, as in Grossman (1981) and Milgrom (1981), draw inferences based not only what *is* disclosed to them, but what *isn't*. Namely, firms consider the strategic incentives of consumer to disclose what she disclosed and no more. For example, if a consumer opts out from being tracked, firms know that the consumer is of the kind that prefers to not be tracked. Our motivating question in this framework are: *when consumers control their information, are they hurt or helped by personalized pricing?*

We pose this question in simple and canonical models of monopolistic and competitive markets with horizontal differentiation. We show that whether consumers benefit from personalized pricing depends subtly on the interaction between market competitiveness and the technology through which consumers can disclose their information. But several robust and simple lessons emerge from our analysis.

First, personalized pricing benefits consumers when consumers can partially disclose information about their characteristics. In both monopolistic and competitive environments, the consumer-optimal equilibria feature intervals of consumer types who pool together in such a way that they obtain the same prices from the marketplace. These pools are constructed so that (i) firms have no incentive to exclude or screen consumer types with their prices, and (ii) no consumer type gains from deviating and sending different information about her characteristics. We first construct these pools in a monopolistic environment and then show that it emerges also in competitive markets with horizontal and vertical differentiation (and even if the consumer doesn't have evidence with positive probability, as in Dye 1985).

Second, personalized pricing benefits consumers in a strong way in competitive markets with horizontal differentiation. The key idea is that personalized pricing amplifies competition between firms: even an unraveling equilibrium (in which every consumer type reveals herself) is superior to the equilibrium of the environment in which personalized pricing is prohibited. In the consumer-optimal equilibrium, consumers who have strong preferences for one firm (or the other) pool from the perspective of the closer firm, and are able to retain significant consumer surplus as the more distant firm competes for their business. Our analysis formalizes the intuition, articulated below by Lars Stole (quoted in Wallheimer 2018), that personalized pricing and price-matching benefits 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."*

By contrast, banning personalized pricing dampens competition, and makes it incentive-compatible for firms to set high prices.

A Preview: Our approach combines two canonical economic problems. The first is the classic problem of a monopolist choosing what price to charge a consumer whose valuation he does not know. The second is the issue of "verifiable" disclosure, similar in spirit to Grossman (1981) and Milgrom (1981), where a consumer chooses what "evidence" or hard information to disclose about her valuation. We study both those disclosure environments in which evidence is **simple**—a consumer can either speak "the whole truth" (her type) or 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 her type perfectly. Our analysis begins with exploring simple and rich evidence structures in a monopolistic environment, and then using those results to characterize behavior in a competitive market.

The timing of the game we model is as follows: the consumer (she) first observes her valuation drawn from the support $[v_l, v_h]$, chooses a message to disclose to the firm (choosing from the set of messages available to her), then the firm chooses a price to charge her, and finally she chooses whether to buy the product at that price. Neither the firm nor the consumer can commit to disclosure, pricing, or purchasing strategies.

Our first conclusion in the monopolistic environment is that simple evidence never benefits the consumer, and potentially hurts her: there exists no equilibrium in which *any type* of the consumer is better off relative to the setting without personalized pricing. However, there exist equilibria in which all consumers are worse off, such as an "unraveling-equilibrium" in which the consumer fully reveals its type and the monopolist extracts all surplus. Simple evidence does not benefit consumers in monopolistic markets.

Our second conclusion for the monopolistic environment is that better equilibria for consumers exist once the evidence structure is rich. The rich evidence structure that we model are interval-messages: the set of messages that a consumer of type v can send is the set of all closed intervals that contain v. This partial disclosure technology nests the simple evidence case: the consumer can fully reveal her type, or send a fully uninformative message. But it also allows for the type space to be partitioned into a discrete set of consumer categories, where each consumer reveals their category but not their type.

This is in fact the structure of the consumer optimal equilibrium, which improves upon the benchmark of no-personalized-pricing. We prove that there always exists a consumer-optimal equilibrium in which the allocation is efficient, and the equilibrium involves breaking the market into segments, each of which is an interval of types. Each segment is constructed so that the monopolist's optimal price is at the bottom of each segment, and so trade occurs with probability 1. Characterizing the optimal segmentation in closed-form is challenging and so we develop a "greedy algorithm" that always identifies a segmentation equilibrium that improves upon no-personalized-pricing. For certain distributions (e.g., v is uniformly distributed), this algorithm identifies the consumer-optimal equilibrium.

We leverage these insights to extend to a model of Bertrand duopoly with horizontally differentiated products where the consumer faces a (linear) transportation cost and 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. In this setting, we find that voluntary disclosure and personalized pricing has a particularly strong effect on consumer surplus, as information can be selectively disclosed to induce competition from the market.

With a simple evidence structure, we show how higher consumer surplus is obtained using a private communication protocol. The consumer lets the firm farther from her observe her type. If the consumer is sufficiently close to the center, she also discloses her type to the firm closer to her, and benefits from the resulting price competition. Otherwise, she does not share her data with the firm closer to her. In equilibrium, each firm knows that a consumer who is not disclosing any information must be sufficiently close, but is uncertain about just how close. In this way, a pooling segment forms at the extremes, but all centrally located types disclose fully to both firms.

With rich evidence, we construct optimal equilibria that use the segmentation strategy similar to the monopolist's problem, with pools becoming progressively finer as one approaches the central type from either end. In this equilibrium, within each segment, the more distant firm charges a price of 0 while the closer firm charges a price that ensures that all types purchase from that closer firm.

One seemingly controversial feature of our model is that consumers can directly reveal a credible signal of their valuation to firms. In reality, the hard information that consumers are typically able to disclose consists of "features" such as their demographic information (age, gender, income), and search and purchase histories. So in the appendix we show that if there is a commonly-known deterministic and continuous mapping f from features into valuations ("a machine learning algorithm"), a model where the consumers send a closed convex subset of the feature space that contains their valuation (rather than an interval containing their type) is isomorphic to the model we solve.

Relationship to Literature: In studying voluntary disclosure and personalized pricing, we build on the burgeoning literature on the economics of privacy and its implications for markets.[1] We view our paper as making two main 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 potential multidimensionality of consumer characteristics, the importance of product customization, etc.), to elucidate the strategic issues at the core of voluntary disclosure and personalized pricing. These strategic issues are captured by combining classical models of markets with classical models of verifiable disclosure; perhaps surprisingly, we have not seen such a study undertaken in the prior literature.

Second, in deriving normative conclusions, our framework clarifies important interactions between the communication technology, market competition, and equilibrium

[1] See Acquisti, Taylor, and Wagman (2016) for a recent survey.

selection. Our results illustrate that the disclosure technology matters in monopolistic markets: if all that the consumer can do is either fully reveal or fully conceal her type (as in the binary choice between sharing one's entire browsing history and not doing so), she does not benefit from personalized pricing. But she does if she can disclose partial information about her type. By contrast, in competitive markets, personalized pricing benefits under both communication technologies.

An alternative approach to this problem uses an intermediary or platform to disclose information on behalf of the consumer. If that intermediary can commit to a disclosure strategy, it can segment the market, as in the elegant analysis of Bergemann, Brooks, and Morris (2015), to maximize consumer surplus. Necessarily, the outcome of the consumer-optimal equilibrium that we study is attainable by the intermediary, but the converse is not true. The issue is that an intermediary may pool a consumer type stochastically into different market segments that induce different prices, which would not be incentive-compatible for the consumer in our disclosure environment.[2] We complement their approach by starting with the premise that no party other than the consumer knows her valuation, and she chooses how to actively disclose it in a manner that is sequentially rational for her. Insofar as the electronic marketplace may involve both intermediaries and consumers with decision rights (who engage with the market actively), we view it to be useful to complement the *information-design* approach to this problem with a *verifiable disclosure* approach.[3]

Another strand to this literature has investigated the role of consumer privacy from the perspective of the "ratchet effect" where where *future* prices may be conditioned on a consumer's current consumption decisions. Taylor (2004), Villas-Boas (2004), and Acquisti and Varian (2005) show that if all consumers are sophisticated, firms are better off by committing to "confidential regimes" that do not price based on past behavior. Recently, Bonatti and Cisternas (2018) take an information-design approach to this problem, and characterize optimal scoring rules. Relatedly, Calzolari and Pavan (2006) show in a sequential contracting model that firms benefit from committing to share information about a consumer's preferences if she views the firms' goods to be complementary,

[2] Relatedly, the intermediary's segmentation also raises the tension, as alluded to in Footnote 1 of Bergemann, Brooks, and Morris (2015), that it can be done by "a benevolent intermediary who already knew consumers' valuations, but not by one who needed consumers to truthfully report their values...."

[3] Also taking an information-design approach, Ichihashi (2017) studies a model where an uninformed consumer chooses an information structure for a multi-product firm learns about her preferences, and commits to an information structure. He shows that the firm is better off by committing to not price discriminate based on what it learns, and such commitments encourages the consumers to choose more informative experiments. Less closely related, Roesler and Szentes (2017) study a setting where the buyer commits to an experiment on how much to learn about her preferences, and the seller makes a price offer without the buyer disclosing any information that she obtains.

and otherwise, the first firm is better off from committing to keeping the information secret. Conitzer, Taylor, and Wagman (2012) study a model where the consumer, not the firm, chooses whether to remain anonymous and faces a cost from doing so; they find that costs from being anonymous have a non-monotonic effect on consumer surplus.

Our work also connects to the study of the role of certification in mitigating adverse selection in markets; Lizzeri (1999), Stahl and Strausz (2017), and Glode, Opp, and Zhang (2018) identify conditions under which hard information avoids market breakdowns. While elements of our analysis may apply to lemons' markets, it isn't our focus; we focus on a "private values" model with voluntary disclosure to study whether personalized pricing benefits consumers in monopolistic and competitive markets.

2 Example

We illustrate the key ideas of our paper using the monopolistic model. 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. In this setting, and without any disclosure, the monopolist optimally posts a 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 $\mathcal{M}\equiv\{[a,b]:0\leq a\leq b\leq 1\}$, 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 $M:[0,1]\rightrightarrows\mathcal{M}$. The timeline for the game is: 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 $p\geq 0$. The consumer then chooses whether to purchase the good. Each party must behave sequentially rationally; we study Perfect Bayesian Equilibria (henceforth PBE) of this game. Our interest is in studying 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, every type can either fully disclose 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 corresponds to a stylized model for the dichotomy between "track" and "do-not-track," where when a consumer is tracked, potentially all of her digital footprint is observable whereas do-not-track obscures it entirely.

With a simple evidence structure, there exists a full unraveling equilibrium in which every type v sends the fully revealing message $m = \{v\}$, and the monopolist extracts all surplus on the equilibrium path; off-path, if the consumer ever sends the non-disclosure message, the monopolist believes (off-path) that $v = 1$ with probability 1, and charges a price of 1. In this equilibrium, all consumers are hurt by the possibility of personalized pricing and the monopolist benefits from it.

But this is not the unique equilibrium: it is also an equilibrium for every type v to send the *non-disclosure* message $m = [0, 1]$, and the monopolist charges a price of $\frac{1}{2}$. No consumer type has a motive to deviate and reveal herself because doing so results in a payoff of 0. Here, both consumer and producer surplus are exactly as in the world without personalized pricing.

In fact, there are an uncountable number of equilibria. But the important observation is that *none* of them improve upon the benchmark of no-personalized-pricing from the perspective of consumers, ex ante and ex interim.

Observation 1. *With simple evidence, across all equilibria, the consumer's interim payoff is no more than her payoff without personalized pricing.*

The logic of Observation 1 is straightforward and geometric, as 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 then send this non-disclosure message because her other option—revealing herself—results in the monopolist extracting all of her surplus. Therefore, the set of types in $(p_{ND}, 1]$ must all be sending the non-disclosure message. Given this consideration, one can show that the monopolist's optimal price is never below $\frac{1}{2}$.

Observation 1 illustrates that simple evidence structures and personalized pricing do not benefit the consumer. Below, we show how the consumer can do better if she can use a rich evidence structure.

Rich evidence: 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 described above are supportable with this richer language, but now new possibilities emerge, some of which dominate the payoffs from no-personalized-pricing.

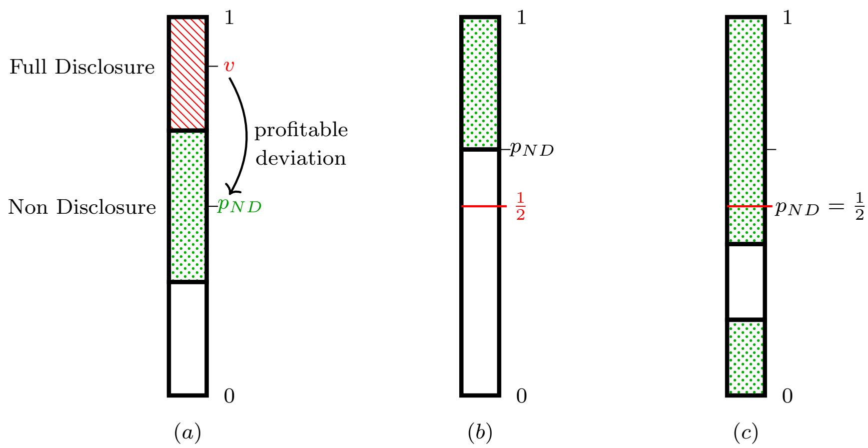

Figure 1: (a) shows that any disclosing type that is strictly higher than p_{ND} has a profitable deviation $\Rightarrow$ the set of non-disclosing types includes $(p_{ND}, 1]$. (b) and (c) illustrate different equilibria where the shaded region is the set of non-disclosing types. Across equilibria,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

We describe the equilibrium that maximizes consumer surplus. Inspired by Zeno's Paradox,[4] consider the countable grid $\left\{1, \frac{1}{2}, \frac{1}{4}, \ldots\right\} \cup \{0\}$. 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 prove that the consumer-optimal equilibrium generates this market segmentation, and that this segmentation induces the monopolist to price at the bottom of each segment.

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

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

For every equilibrium path message of the form $[a_{k+1}, a_k]$, the monopolist charges a_{k+1} after receiving that message, therefore selling to the entire segment sending that message.

In this equilibrium, the highest market segment is composed of types in $\left(\frac{1}{2}, 1\right]$, all of whom 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 whom send the message $m_1 \equiv \left[\frac{1}{4}, \frac{1}{2}\right]$, and so on and so forth. We illustrate this partition in Figure 2. For each market segment, the monopolist believes that the consumer's value is uniformly distributed on it, and his optimal price is to price at the bottom of the segment. Therefore, trade occurs with probability 1, with the consumer capturing some fraction of the surplus.

This equilibrium generates an ex ante consumer surplus of $\frac{1}{6}$ and producer surplus of

[4] 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/.

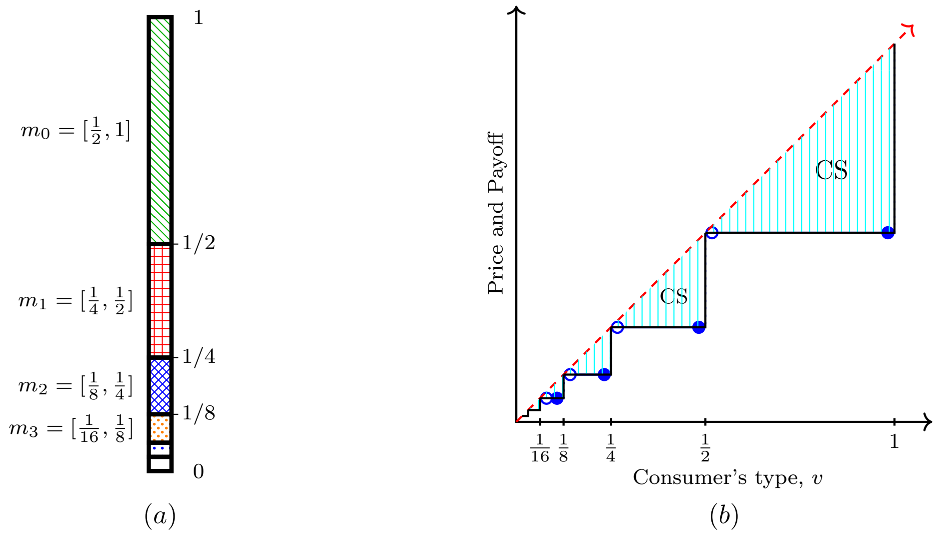

Figure 2: (a) illustrates Zeno's Partition. (b) illustrates prices and payoffs: for each consumer-type v, the step-function shows the equilibrium price that is charged and the dashed $45°$ line shows the payoff from consumption. The shaded region illustrates the consumer surplus achieved by Zeno's Partition.

$\frac{1}{3}$, each of which is higher than what is obtained without personalized pricing.[5] Thus, the possibility for personalized pricing with rich evidence is a Pareto improvement.

How is Zeno's Partition supportable as an equilibrium segmentation? 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.

How about "on-schedule" deviations? For every v in (a_{k+1}, a_k), there exists only one "on-schedule" 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}.

The logic of this market segmentation illustrates the role of partial hard information: the message $\left[\frac{1}{4}, \frac{1}{2}\right]$ ensures that even though types greater than $\frac{1}{2}$ would love to join this pool (to obtain a price of $1/4$), they are unable to do so. Because these higher types are excluded, the monopolist finds it optimal to not raise prices.

Proving that Zeno's Partition is the consumer-optimal equilibrium is less straightforward. Lemmas 2 and 3 in the full paper show that generally, a consumer-optimal equi-

[5]Moreover, every type of the consumer obtains surplus $v - \left(\frac{1}{2}\right)^{\lfloor \frac{\log v}{\log(1/2)} \rfloor + 1}$, which is strictly positive for all but a countable set of values ($\lfloor \cdot \rfloor$ denotes the floor function). All types in $(1/2, 1]$ are no worse off from personalized pricing and almost every other type is strictly better off.

librium involves both trade occurring with probability 1 and each equilibrium message being sent by an interval of types. But it is generally difficult to obtain a closed-form solution for the consumer-optimal segmentation. Instead, we offer a "greedy algorithm" that generates an equilibrium that *always* improves upon the benchmark of no-personalized pricing, and under certain conditions (such as uniformly distributed types) results in the consumer-optimal equilibrium. Below, we exposit a heuristic argument to illustrate why Zeno's Partition is the optimal segmentation in this example.

Because consumers purchase with probability 1, 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, where $\lambda \geq 1/2$. 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 must involve replicating the same segmentation but on a smaller scale. Thus, one can frame the consumer-optimal segmentation 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 $[0, \bar{v}]$. Therefore,

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

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.

References

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.

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

Bonatti, A. and G. Cisternas (2018): "Consumer Scores and Price Discrimination," Working Paper.

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

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.

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

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

Hern, A. (2018): "No tracking, no revenue: Apple's privacy feature costs ad companies millions," https://www.theguardian.com/technology/2018/jan/09/apple-tracking-block-costs-advertising-companies-millions-dollars-criteo-web-browser-

Ichihashi, S. (2017): "Online Privacy and Information Disclosure by Consumers," Working paper.

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.

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

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.

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*.