Product Variety and Market Segmentation*

Stefan Terstiege† Adrien Vigier‡

January 23, 2023

PRELIMINARY AND INCOMPLETE

Abstract

We study market segmentation in settings where a monopolistic seller possesses an inventory containing several variants of a given good and needs to decide which variant to offer. We fully characterize the producer-consumer surplus pairs induced by market segmentation as product variety becomes large, and prove that whether or not the seller can price discriminate is irrelevant. We show that, along the Pareto frontier, higher consumer surplus entails lower social welfare, but is compatible with greater privacy. We then study market segmentation arising from the sale of consumer data by intermediaries. Competition among data intermediaries results in lower match quality between consumers and products and lower social welfare.

Keywords: market segmentation, product variety, consumer privacy, data intermediaries.

JEL-Classification: D42, D83

*We thank...

†Maastricht University. Email: s.terstiege@maastrichtuniversity.nl.

‡University of Notthingam. Email: adrien.vigier@nottingham.ac.uk

1 Introduction

In this paper, we study how sellers' access to consumer data affects welfare when product variety is large. Large product variety is a salient feature of many online marketplaces. Already in the early 2000s, Amazon offered about 2.3 million book titles, whereas large brick-and-mortar stores held between 40,000 and 100,000 titles (Brynjolfsson, Hu, and Smith, 2003). Zalando, an online fashion retailer, has 63,965 items in the category "Men's T-Shirts & Polos" alone.[1] Similarly, online shoe retailers may offer over 50,000 distinct models, whereas traditional retailers usually stock at most a few thousand ones (Quan and Williams, 2018).

While the online shelf space is virtually unlimited, consumers' attention and time are limited. A key aspect of online retail is thus the ability to direct consumers towards the products which they are most likely to value. The fact that Netflix launched a million dollar competition (the "Netflix Prize") inviting computer scientists to outperform the company's recommender system gives a glimpse of the importance of the various tools which online sellers use to steer consumers' choices.

As an input to advertising and recommender systems, online retailers' information about consumers plays a crucial role. Which combinations of producer and consumer surpluses are attainable when sellers have access to consumer data? Can regulation enhance the privacy of consumers without sacrificing efficiency or consumer surplus? Why do online retailers often refrain from price discrimination although they have so much data? What are the welfare consequences of data intermediation, and which data do intermediaries supply? What are the effects of competition in the data market? These are the questions that we address in this paper.

In our model, a monopolistic seller possesses an inventory containing several variants of a given good, and each consumer wishes to buy a single good. The seller's information induces a market segmentation, as in Bergemann, Brooks, and Morris (2015). In particular, each market segment is represented by a probability distribution over the set of possible valuation vectors for the different products of the seller. The seller's problem

[1] `https://en.zalando.de/men-clothing-shirts/`; accessed January 6, 2023.

is to choose, for each market, which product to offer and at what price.

The central result of our paper (Theorem 1) characterizes the combinations of producer and consumer surpluses that result from market segmentation. This characterization rests on the observation that when product variety is large, the combinations of producer and consumer surpluses obtained through market segmentation approximately coincide with the producer-consumer surplus pairs obtained in a much simpler auxiliary setting. Specifically, in this auxiliary setting a designer chooses the distribution of consumers' valuations for a *single* product, subject to certain constraints.

We show that in the limit, when the number of products grows without bound, the surplus pairs that result from market segmentation in fact coincide with the surplus pairs induced by market segmentations under which the seller does not benefit from price discrimination (Proposition 1). Our analysis thus sheds light among other things on the use of price discrimination, and suggests that "search discrimination", that is, the practice by which different customers are steered towards different products, can make price discrimination redundant.[2] Indeed, whereas targeted ads and personalized product recommendations are common, overt price discrimination is relatively rare (see, e.g., Cavallo, 2017; DellaVigna and Gentzkow, 2019).

We explore the effect of market segmentation on social welfare. Efficiency requires each consumer to purchase the product that he values the most. But if the seller can identify for each consumer the product that he values the most then, in equilibrium, the prices at which products are sold must be high. By this logic, we show that along the Pareto frontier social welfare decreases when consumer surplus goes up (Proposition 2).

We examine the extent to which securing consumer privacy can increase the surplus of consumers, and identify a precise sense in which greater consumer privacy is associated with greater consumer surplus. Specifically, we show that any market segmentation possesses a payoff-equivalent segmentation such that consumer surplus may be increased by giving greater privacy to consumers (Proposition 3).

[2] A much publicized case of search discrimination was the steering of Mac users to more expensive hotels than PC users by the travel-booking site Orbitz.com; see `https://www.wsj.com/articles/SB10001424052702304458604577488822667325882`.

We then study market segmentation arising from the sale of consumer data by intermediaries. To this end, we augment the model by an initial stage in which data intermediaries propose data policies to a consumer. A data policy specifies which data will be made available to the seller, provided that the consumer gives his consent. The intermediaries that obtain the consumer's consent subsequently sell their data to the monopolistic seller of the products.

Comprehensive regulation in the European Union has made the data market more competitive. Most prominently, the General Data Protection Regulation prescribes that data be portable from one platform to another (Regulation (EU) 2016/679, Article 20). The Digital Markets Act strengthens this requirement, and furthermore prescribes that platforms do not combine, without additional consent, their personal data with those collected by subsidiaries (Regulation (EU) 2022/1925, Article 5). Moreover, antitrust authorities globally appear to take a more aggressive stance towards mergers and acquisition of data-driven businesses; an example is the recent lawsuit of the US Federal Trade Commission against Meta.[3]

We examine the welfare consequences of competition in the data market. Our analysis yields sharp predictions: as product variety becomes large, we precisely pin down producer and consumer surplus for a monopolistic and a competitive data market, respectively (Proposition 4). Competition between data intermediaries results in greater consumer surplus, but reduces social welfare because the seller ends up offering less suitable products. We also show that competition in the data market can result in greater privacy for consumers (Proposition 5).

The rest of the paper is organized as follows. The related literature is discussed below. The model is presented in Section 2. Section 3 states and proves the central theorem of the paper. We show in Section 4 that price discrimination is irrelevant for the set of producer-consumer surplus pairs feasible through market segmentation. In Section 5, we examine the effect of market segmentation on social welfare, as well as the

[3]The lawsuit alleges that Meta, the parent company of Facebook, systematically accumulated market power through its acquisitions of Instagram and WhatsApp. The court case, *Federal Trade Commission v. Meta Platforms, Inc.*, is ongoing.

relation between consumer privacy and welfare. Section 6 studies online markets with data intermediaries. Finally, Section 7 concludes.

1.1 Related Literature

We contribute to the extensive literature studying third-degree price discrimination starting with Pigou (1920), unified by Aguirre, Cowan, and Vickers (2010), and recently revived by Bergemann, Brooks, and Morris (2015).

Our paper builds on the seminal work of Ichihashi (2020) on market segmentation in multi-product monopolies. Ichihashi (2020) compares the properties of the consumer-optimal market segmentations in two pricing regimes: one in which the seller commits to prices before the market has been segmented, and one in which the seller sets prices after having observed the market segment to which the consumer belongs. Finding the consumer-optimal market segmentations in the latter regime is a hard problem, for which no general solution is known. However, an ingenuous two-step procedure enables the author to prove that letting the seller use information for pricing: (a) induces inefficient trade whereby the seller occasionally offers a product that is *not* the consumer's most-preferred product; (b) decreases producer surplus; (c) increases consumer surplus. In particular, insight (b) provides a possible explanation for the rare occurrence of price discrimination by online sellers.[4]

In contrast to Ichihashi (2020), our paper focuses entirely on the no-commitment regime (that is, on the case in which the seller can adjust prices depending on the market segment). We fully characterize the set of feasible combinations of consumer and producer surpluses. We prove the existence of a general trade-off between consumer surplus and social welfare, quantify the social welfare loss associated with consumer-optimal market segmentations, and identify a sense in which greater consumer privacy is associated with greater consumer surplus. Our findings also suggest an explanation complementing Ichihashi's own explanation for why price discrimination by online sellers is uncommon. Specifically, we show that, when product variety is large, the limits of

[4]See, e.g, Narayanan (2013) and DellaVigna and Gentzkow (2019).

price discrimination and uniform pricing approximately coincide. In particular, whatever profit sellers make by using information for pricing can also be achieved by setting one price for each product.

Several other important papers are closely linked in spirit to the core of our paper. Haghpanah and Siegel (2022) show that with remarkable generality a "simple" segmentation of the aggregate market improves welfare in the sense of Pareto. Haghpanah and Siegel (2021) obtains conditions under which the multi-product counterpart of the "surplus triangle" of Bergemann, Brooks, and Morris (2015) corresponds to the set of feasible consumer-producer surplus pairs. These conditions do not hold in the setting we study, and the set of feasible surplus pairs is different.

While closely related to our work, the models of Hidir and Vellodi (2021) and Pram (2021) exhibit important differences with the model we study. In our setting, a consumer is assigned to a specific market segment based on individual characteristics such as age, gender, nationality, or whatever information the seller possesses about this consumer. By contrast, both Hidir and Vellodi (2021) and Pram (2021) consider settings in which individual consumers exert a form of control over the market segment to which they belong: in the setting of the former study, each consumer chooses his preferred market segment through cheap-talk communication; in the latter study, each consumer communicates instead hard information, and thus chooses his preferred market segment from within a subset of segments.

Our paper also contributes to recent economic research on markets with data intermediaries. We study how competition between data intermediaries affects the type of information sold to downstream firms and, ultimately, consumer surplus and social welfare.[5] We show that competition benefits consumers but reduces social welfare.

In Hidir and Vellodi (2021), a single online platform provides consumer data to sellers, and each consumer faces an opportunity cost of participation. The authors show that a higher opportunity cost leads to lower product prices and lower match quality

[5] A different strand of literature shows how the combination of information externalities and coordination failure results in too much data sharing on the part of consumers (Choi, Jeon, and Kim, 2019; Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2022; Bergemann, Bonatti, and Gan, 2022).

between consumers and products. This key insight resonates with our results if one interprets the higher opportunity cost as being due to greater competition among online platforms. De Corniere and De Nijs (2016) consider a setting where online platforms auction advertising slots. To the extent that an increase in the number of slots plays the same role as a decrease in the number of bidders, competition between platforms benefits consumers by inducing lower prices in the product market. In Bounie, Dubus, and Waelbroeck (2022), data intermediaries first acquire costly information, and then choose the information sold to downstream firms. Competition between data intermediaries benefits consumers because it induces the former to acquire less information, which in turn reduces extraction by sellers in the product market.

Bergemann and Bonatti (2015) and Ichihashi (2021) offer different perspectives than the aforementioned papers, and put forth that competition need not benefit consumers. In Bergemann and Bonatti (2015), the key aspect is that raising the price at which information about one consumer is sold to downstream firms reduces the demand for information about all other consumers. By contrast, in Ichihashi (2021) data intermediaries can compensate customers through monetary transfers, and so the degree of competitiveness in the market for data leaves trade efficiency in the product market unaffected.

2 Model

Throughout the paper,

$$X := \{x_1, \ldots, x_m\}, \quad 0 < x_1 < \cdots < x_m,$$

and f is a distribution in ΔX that has full support.[6]

There is a seller (she) with an inventory containing n products, and a continuum of unit-demand consumers.[7] Any consumer's valuations for the n products can be repre-

[6]The notation ΔY indicates the set of all distributions with finite support over the set Y. Our assumptions that f has full support and that $x_1 > 0$ merely simplify the exposition.

[7]That is, the value attached by a consumer to any set of products is equal to the consumer's maximum valuation for a single item in this set.

sented by some vector $\mathbf{v} \in X^n$, with v_k (the k^{th} component of the vector $\mathbf{v}$) indicating this consumer's valuation for product k. We use the generic notation μ for a probability distribution over X^n, that is, $\mu \in \Delta X^n$; we refer to such a distribution as a *market*. The k-marginal of a market μ is denoted by μ_k.[8]

The proportion $\bar{\mu}(\mathbf{v})$ of consumers whose valuations are given by the vector $\mathbf{v}$ satisfies

$$\bar{\mu}(\mathbf{v}) = \prod_k f(v_k), \quad \forall \mathbf{v} \in X^n. \tag{1}$$

We refer to the market $\bar{\mu}$ defined through (1) as the *aggregate market*.

A typical element of $\Delta\Delta X^n$ is denoted by τ; if

$$\sum_\mu \tau(\mu)\mu(\mathbf{v}) = \bar{\mu}(\mathbf{v}), \quad \forall \mathbf{v} \in X^n, \tag{2}$$

then τ is called a *market segmentation*. For a fixed market segmentation τ, the problem of the seller is to choose for each market comprised in the support of τ, which product to offer and at what price.

We use the generic notation ρ for a strategy of the seller, with $\rho_\mu(k, p)$ representing the probability that the seller offers product k at price p in market μ. We suppose that if a consumer's valuation for product k equals v_k then, when offered product k at price v_k, the consumer decides to buy. The producer surplus generated by the strategy ρ is[9]

$$\Pi_\tau(\rho) := \sum_\mu \tau(\mu) \sum_{k,p} \rho_\mu(k,p) p \sum_{x \geq p} \mu_k(x);$$

the corresponding consumer surplus is

$$U_\tau(\rho) := \sum_\mu \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \sum_{x \geq p} \mu_k(x)(x - p).$$

We say that a surplus pair (π, u) is *feasible* if there exist a market segmentation τ as well as a strategy $\rho^* \in \operatorname{argmax}_\rho \Pi_\tau(\rho)$ such that $\pi = \Pi_\tau(\rho^*)$ and $u = U_\tau(\rho^*)$. Finally, the set of feasible surplus pairs is denoted by S_n.

[8] Thus, $\mu_k \in \Delta X$, with $\mu_k(x) = \sum_{\mathbf{v}:\, v_k = x} \mu(\mathbf{v})$ for all $x \in X$.

[9] The mass of consumers is normalized to one, to save on notation.

Additional expository assumptions: we assume that $p \mapsto p\sum_{x\geq p} f(x)$ possesses a unique maximizer, which we denote by p_0. Then, letting

$$\pi_0 := p_0 \sum_{x\geq p_0} f(x),$$

we assume that $\pi_0 \in X$.

2.1 Discussion of the Model

A market segmentation could either depict geographically distinct markets, or summarize information available to an online seller, perhaps due to the seller's access to consumers' browsing histories or the use of cookies. An online seller might be able to determine, say, the age and nationality of each consumer. In this case, a market would represent the distribution of valuations within a given age group of a certain nationality. We return to this interpretation in Section 6, where we study online markets with data intermediaries.

The model supposes that the seller offers a single product in each market. We capture thereby situations in which a firm has a large inventory consisting of many different variants of a given good or service, and where the number of variants is far greater than the constraints imposed by consumers' limited attention or cognitive costs, thus forcing sellers to make strategic choices with regard to the products they offer in any given market. This feature is central to online retailing, among other things (Brynjolfsson, Hu, and Smith, 2003; Anderson, 2006). Note that, what we refer to as a product in the model might in practice represent a sub-category of products, such as "Italian movies from the 1960's", for example.

Finally, our model allows the seller to engage in third-degree price discrimination. This assumption seems realistic in online markets. At any rate, we show in Section 4 that, when product variety is large, whether or not the seller can price discriminate is inconsequential.

3 The Welfare Bounds of Market Segmentation

In this section, we characterize the set of feasible surplus pairs when the number of products is large.

We start with a couple of definitions. For $i \in \{1, \dots, m\}$, define $g_i \in \Delta X$ by

$$g_i(x_j) := \begin{cases} 0 & \text{if } j < i, \\ x_i/x_j - x_i/x_{j+1} & \text{if } i \le j < m, \\ x_i/x_m & \text{if } j = m. \end{cases}$$

It is readily checked that

$$p \sum_{x \ge p} g_i(x) = \begin{cases} p & \text{for all } p \in \{x_1, \dots, x_i\}, \\ x_i & \text{for all } p \in \{x_{i+1}, \dots, x_m\}. \end{cases} \tag{3}$$

To understand the significance of g_i, consider an auxiliary single-product setting without market segmentation in which consumers' valuations for the single product are distributed according to some distribution $\tilde{f} \in \Delta X$, and suppose that in said setting the seller obtains a surplus of x_i. Then $x_m \tilde{f}(x_m) \le x_i$, which we can rewrite as

$$\tilde{f}(x_m) \le g_i(x_m). \tag{4}$$

Similarly, $x_{m-1}\Big(\tilde{f}(x_{m-1}) + \tilde{f}(x_m)\Big) \le x_i$, which, if equality holds in (4) and $i \le m-1$, yields

$$\tilde{f}(x_{m-1}) \le g_i(x_{m-1}).$$

By pursuing the recursion we see that the valuation distribution g_i maximizes consumer surplus among all distributions in ΔX.[10]

Below, let

$$\overline{u}(x_i) := \sum_{x \ge x_i} g_i(x)(x - x_i) \tag{5}$$

[10] The class of distributions satisfying property (3) plays a key role in a wide variety of contexts. See, e.g., Neeman (2003), Bergemann and Schlag (2008), Bergemann, Brooks, and Morris (2015), and Condorelli and Szentes (2020).

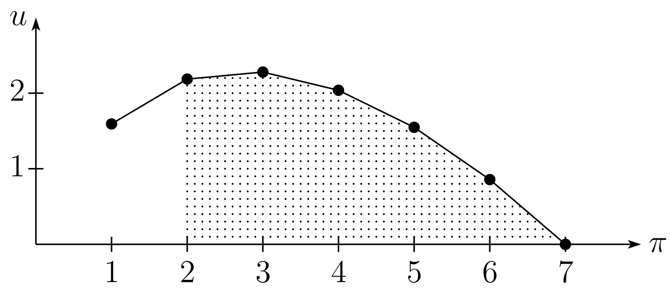

Figure 1: S (shaded area) and $\overline{u}$ (curve) for $X = \{1, 2, 3, 4, 5, 6, 7\}$, assuming $\pi_0 = 2$.

denote the aforementioned maximum consumer surplus in the auxiliary setting. Then, for $i = 1, \cdots, m-1$, extend (5) through

$$\overline{u}\Big((1-\lambda)x_i + \lambda x_{i+1}\Big) := (1-\lambda)\overline{u}(x_i) + \lambda\overline{u}(x_{i+1}), \quad \forall \lambda \in [0, 1]. \tag{6}$$

Finally, define

$$S := \Big\{(\pi, u) \in \mathbb{R}^2 \mid \pi \in [\pi_0, x_m], u \in [0, \overline{u}(\pi)]\Big\}.$$

We can now state our main theorem.

Theorem 1. *For every $n \in \mathbb{N}$, the set S_n of feasible surplus pairs is contained in the set S. Moreover, for every $(\pi, u) \in S$, there exists a sequence $\Big((\pi_n, u_n)\Big)_{n \in \mathbb{N}}$ such that $(\pi_n, u_n) \in S_n$ and $(\pi_n, u_n) \underset{n \to \infty}{\longrightarrow} (\pi, u)$.*

Figure 1 illustrates the set S for $X = \{1, 2, 3, 4, 5, 6, 7\}$ and $\pi_0 = 2$. Combining Theorem 1 with the discussion preceding it shows that, as product variety becomes large, the feasible surplus pairs approximately coincide with the surplus pairs attainable in a single-product setting without market segmentation, but where the valuation distribution is an object of design subject to the constraints that (a) the support of this distribution is contained in X, (b) the resulting surplus of the seller is not smaller than π_0.[11]

[11]This characterization of the feasible surplus pairs connects our analysis to Condorelli and Szentes (2020), which studies the problem of a buyer choosing the distribution of his valuation for a product supplied by a monopolistic seller. In particular, the authors show that in equilibrium the distribution chosen by the buyer generates a unit-elastic demand and that trade occurs with probablity 1.

Specifically, the first part of the theorem tells us that, regardless of the market segementation, the seller must obtain a surplus $\pi \in [\pi_0, x_m]$, while the surplus of the consumers is bounded from above by $\overline{u}(\pi)$.

The second part of the theorem tells us that, as product variety becomes large, any element of S may be approximately attained through market segmentation. The basic idea is as follows. Pick some $x_i \in X$ such that $x_i \geq \pi_0$. Next, for each $k \in \{1, \cdots, n\}$, divide the consumers in two groups, say k^+ and k^-, such that, in group k^+, the valuations for product k are distributed according to g_i. Let T^+ be the subset of individuals who belong to some group k^+, with $k \in \{1, \cdots, n\}$. Now segment the aggregate market into $n+1$ markets labelled $s_0, s_1, \cdots, s_n$, in the following way: if an individual belongs to T^+, place him in some market s_k with k chosen such that this individual belongs to the group k^+; then place all remaining individuals in market s_0. We show that in the market s_k, $k \neq 0$, the seller can do no better than to offer product k at a price of x_i. When n is large, the proportion of individuals in market s_0 is negligible. By segmentating the market in this way, we thus generate surplus approaching x_i for the seller and $\bar{u}(x_i)$ for consumers.

The next subsection contains the proof of Theorem 1. Readers uninterested in the technical details of the analysis can jump to Section 4 without loss.

3.1 Proof of Theorem 1

Our proof uses three lemmas.

Lemma 1. *If* $(\pi, u) \in S_n$ *then* $u \leq \overline{u}(\pi)$.

Proof. We treat below the case $\pi \in X$; the proof for the remaining case is similar, and therefore omitted. As $(\pi, u) \in S_n$, there exist a market segmentation τ and a strategy ρ of the seller that is optimal given τ, such that $\pi = \Pi_\tau(\rho)$ and $u = U_\tau(\rho)$. Then, define the distribution $h \in \Delta X$ by

$$h(x) := \sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \mu_k(x), \quad \forall x \in X.$$

Letting x_i be the element of X such that $\pi = x_i$, we have for any $q \in \{x_i, \cdots, x_m\}$:

$$\begin{aligned}\sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) q \sum_{x \geq q} \mu_k(x) &\leq \sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) p \sum_{x \geq p} \mu_k(x) \\ &= \Pi_\tau(\rho) = \pi = q \sum_{x \geq q} g_i(x).\end{aligned}$$

The inequality in the previous sequence follows from ρ being optimal given τ; the last equality follows from (3). Dividing through by q, we see that g_i first-order stochastically dominates S. Hence,

$$\begin{aligned}u &= \sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \sum_{x \geq p} \mu_k(x)(x-p) \\ &= \sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \sum_{x \geq p} \mu_k(x)x - \pi \\ &\leq \sum_{\mu} \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \sum_{x} \mu_k(x)x - \pi \\ &= \sum_{x} h(x)x - \pi \\ &\leq \sum_{x} g_i(x)x - \pi \\ &= \overline{u}(\pi).\end{aligned}$$

□

In what follows, we say that a distribution $\tau \in \Delta\Delta X^n$ is the product of distributions $\{\tau^k\}_{k=1}^n$ in ΔX, if $\tau(\mu) > 0$ implies

$$\mu(\mathbf{v}) = \prod_k \mu_k(v_k), \quad \forall \mathbf{v} \in X^n, \tag{7}$$

and

$$\tau(\mu) = \prod_k \tau^k(\mu_k). \tag{8}$$

Lemma 2. *Let τ be the product of $\{\tau^k\}_{k=1}^n$. If*

$$\sum_{\mu_k} \tau^k(\mu_k)\mu_k(x) = f(x) \quad \text{for all } k \in \{1, \cdots, n\} \text{ and all } x \in X, \tag{9}$$

then τ is a market segmentation.

The proof of this elementary result is relegated to the appendix. Our next lemma identifies points of S which may be approached by feasible surplus pairs as the number of products becomes large.

Lemma 3. *For every* $x_i \in \{\pi_0, \dots, x_m\}$ *and every* $p \in \{x_i, \dots, x_m\}$*, there exists a sequence* $\Big((\pi_n, u_n)\Big)_{n\in\mathbb{N}}$ *such that* $(\pi_n, u_n) \in S_n$ *for every* n*, and*

$$(\pi_n, u_n) \underset{n\to\infty}{\longrightarrow} \left(x_i, \sum_{x\geq p} g_i(x)(x-p)\right). \tag{10}$$

Proof. Let $x_i \in \{\pi_0, \dots, x_m\}$, and $p \in \{x_i, \dots, x_m\}$. Choose $\lambda \in (0,1)$ such that $\lambda g_i(x) \leq f(x)$ for all $x \in X$, and define

$$h(x) := \frac{f(x) - \lambda g_i(x)}{1-\lambda}. \tag{11}$$

Note that $h(x) \geq 0$ for all $x \in X$, and $\sum_x h(x) = 1$, whence $h \in \Delta X$. Moreover,

$$\lambda g_i(x) + (1-\lambda)h(x) = f(x), \quad \forall x \in X. \tag{12}$$

We claim that

$$\max_q q \sum_{x\geq q} g_i(x) = p \sum_{x\geq p} g_i(x) = x_i \geq \max_q q \sum_{x\geq q} h(x). \tag{13}$$

The equalities in (13) follow from (3). The fact that

$$x_i \geq q \sum_{x\geq q} h(x), \quad \forall q \leq x_i,$$

is immediate, as h is a distribution. Lastly, for all $q > x_i$:

$$x_i \;\geq \pi_0 \geq q \sum_{x\geq q} f(x) = \lambda q \sum_{x\geq q} g_i(x) + (1-\lambda) q \sum_{x\geq q} h(x) = \lambda x_i + (1-\lambda) q \sum_{x\geq q} h(x).$$

So

$$x_i \geq q \sum_{x\geq q} h(x), \quad \forall q > x_i,$$

which finishes the proof of (13).

Next, define $\tau^k \in \Delta X$ by

$$\tau^k(g_i) = \lambda = 1 - \tau^k(h),$$

and let τ be the product of $\{\tau^k\}_{k=1}^n$. By coupling (12) with Lemma 2, notice that τ is a market segmentation.

Now let ρ be a strategy of the seller with the following properties. For every market μ in the support of τ such that $\mu_k = g_i$ for some product k, offer any such product at price p. If $\mu_k = h$ for all products $k \in \{1, \cdots, n\}$, on the other hand, offer any product at some fixed price

$$q' \in \operatorname*{argmax}_q q \sum_{x \geq q} h(x).$$

By (13), the strategy ρ is optimal given τ. The resulting surplus of the seller is

$$\pi_n := \Pi_\tau(\rho) = (1 - (1 - \lambda)^n) x_i + (1 - \lambda)^n q' \sum_{x \geq q'} h(x);$$

the consumer surplus is

$$u_n := U_\tau(\rho) = (1 - (1 - \lambda)^n) q \sum_{x \geq p} g_i(x)(x - p) + (1 - \lambda)^n q' \sum_{x \geq q'} h(x)(x - q').$$

Then $(\pi_n, u_n) \in S_n$, and since $\lambda > 0$, the limit in (10) is established. □

We are now ready to prove the theorem.

Proof of Theorem 1. The strategy ρ given by $\rho_\mu(1, p_0) = 1$ for every μ yields surplus π_0 to the seller, so the seller can guarantee herself a surplus of π_0 regardless of the market segmentation. The first part of the theorem then follows from Lemma 1.

We now prove the second part of the theorem. Let (π', u') and (π'', u'') be arbitrary points in the set S. Suppose $(\pi'_n, u'_n) \in S_n$ for every n, with

$$(\pi'_n, u'_n) \underset{n \to \infty}{\longrightarrow} (\pi', u').$$

Similarly, suppose $(\pi''_n, u''_n) \in S_n$ for every n, with

$$(\pi''_n, u''_n) \underset{n \to \infty}{\longrightarrow} (\pi'', u'').$$

Let τ'_n and τ''_n be market segmentations inducing the surplus pairs (π'_n, u'_n) and (π''_n, u''_n), respectively. The set of market segmentations is evidently convex. Furthermore, note that for all $\zeta \in [0, 1]$, some optimal strategy of the seller given $(1 - \zeta)\tau'_n + \zeta\tau''_n$ yields a surplus of $(1 - \zeta)\pi'_n + \zeta\pi''_n$ for the seller and $(1 - \zeta)u'_n + \zeta u''_n$ for the consumers. We conclude that there exists a sequence $\big((\pi_n, u_n)\big)_{n \in \mathbb{N}}$ such that $(\pi_n, u_n) \in S_n$ for every n, and

$$(\pi_n, u_n) \underset{n \to \infty}{\longrightarrow} (1 - \zeta)(\pi', u') + \zeta(\pi'', u'').$$

Now, for all $x_i \in \{\pi_0, \ldots, x_m\}$, Lemma 3 gives us sequences $\big((\pi'_n, u'_n)\big)_{n\in\mathbb{N}}$ and $\big((\pi''_n, u''_n)\big)_{n\in\mathbb{N}}$ such that, firstly, (π'_n, u'_n) and (π''_n, u''_n) belong to S_n for every n, and secondly,

$$(\pi'_n, u'_n) \underset{n\to\infty}{\longrightarrow} \big(x_i, \overline{u}(\pi)\big) \quad \text{and} \quad (\pi''_n, u''_n) \underset{n\to\infty}{\longrightarrow} (x_i, 0).$$

We conclude using the previous observation that, for every $(\pi, u) \in S$, there exists a sequence $\big((\pi_n, u_n)\big)_{n\in\mathbb{N}}$ such that $(\pi_n, u_n) \in S_n$ for every n and $(\pi_n, u_n) \underset{n\to\infty}{\longrightarrow} (\pi, u)$. $\square$

4 Irrelevance of Price Discrimination

We say that a strategy ρ involves price discrimination if some product k is sold at different prices depending on the market in which this product is offered.[12]

It is easy to see that price discrimination may strictly benefit the seller. For example, suppose $X = \{x_1, x_2\}$, let τ be the market segmentation comprising 2^n markets separating consumers with different valuation vectors, and ρ^* some strategy of the seller that is optimal given τ. Now let μ^- denote the market in the support of τ in which every consumer's valuation vector equals $(x_1, \cdots, x_1)$, and μ^k the market in which every consumer values product k at x_2 and all other products at x_1. Then any product offered by the seller in market μ^- must be sold at a price of x_1, whence $\rho^*_{\mu^-}(k, x_1) > 0$ for some product $k \in \{1, \cdots, n\}$. On the other hand, the definition of the market μ^k implies $\rho^*_{\mu^k}(k, x_2) = 1$. So any strategy of the seller that is optimal given τ^* involves price discrimination.

We now show that price discrimination is irrelevant for the characterization of feasible surplus pairs in Theorem 1: in the limit, when the number of products grows without bound, any surplus pair that is feasible at all is also feasible without price discrimination.

Formally, we say that a surplus pair (π, u) is *feasible without price discrimination* if there exist a market segmentation τ, as well as a strategy $\rho^* \in \operatorname{argmax}_\rho \Pi_\tau(\rho)$, such that $\pi = \Pi_\tau(\rho^*)$, $u = U_\tau(\rho^*)$, and ρ^* does not price discriminate. The set of surplus pairs that are feasible without price discrimination is denoted by $\tilde{S}_n$.

[12]That is, formally, if there exist $\mu \neq \mu'$ and $p \neq p'$ such that $\rho_\mu(k, p) > 0$ and $\rho_{\mu'}(k, p') > 0$.

Proposition 1. *For every $n \in \mathbb{N}$, the set $\tilde{S}_n$ of surplus pairs that are feasible without price discrimination is contained in the set S. Moreover, for every $(\pi, u) \in S$, there exists a sequence $\big((\pi_n, u_n)\big)_{n\in\mathbb{N}}$ such that $(\pi_n, u_n) \in \tilde{S}_n$ and $(\pi_n, u_n) \underset{n\to\infty}{\longrightarrow} (\pi, u)$.*

Consider again the simple binary valuation example examined above; below, for concreteness, suppose $x_1 > x_2 f(x_2)$. We illustrate the underlying idea of the proposition by showing that, as product variety becomes large, the surplus pair $(x_2, 0)$ maximizing the surplus of the seller can almost be attained without price discrimination. To this end, let τ be the market segmentation with 2^{n-1} markets separating consumers whose vectors of valuations differ in some other component than the first one, and let μ^- denote the market in the support of τ in which every consumer values all products $k \neq 1$ at x_1. Now let ρ^* be some strategy of the seller such that:

- in market μ^-, the seller offers product 1 at a price of x_1;
- in any other market μ contained in the support of τ, the seller offers at a price of x_2 one of the products $k \neq 1$ to which all consumers in the market μ attach value x_2.

Notice that the strategy ρ^* is optimal given τ, and does not price discriminate. Furthermore, the proportion of consumers who belong to some market $\mu \neq \mu^-$ approaches 1 as n tends to infinity, whence $\big(\Pi_\tau(\rho^*), U_\tau(\rho^*)\big)$ approaches the surplus pair $(x_2, 0)$.

More generally, Proposition 1 shows that "search discrimination", the practice by which different customers are steered towards different products, makes price discrimination redundant when product variety is large.

5 Social Welfare and Consumer Privacy

In view of Theorem 1, we refer to the set of maximal elements of S as the *Pareto frontier*; that is, a surplus pair (π, u) belongs to the Pareto frontier if (i) $u = \overline{u}(\pi)$ and (ii) $\overline{u}(\pi) > \overline{u}(\pi')$ for all $\pi' \in (\pi, x_m]$. The *social welfare* at a surplus pair (π, u) is defined as $\pi + u$.

We start this section by showing that, along the Pareto frontier, increasing the surplus of consumers, or decreasing the producer surplus, implies lowering social welfare. To see why, notice that, for $1 \leq i \leq m-1$ and $p \in \{x_{i+1}, \cdots, x_m\}$, (3) gives

$$x_i = p \sum_{x \geq p} g_i(x) < x_{i+1} = p \sum_{x \geq p} g_{i+1}(x).$$

Thus

$$\sum_{x \geq p} g_i(x) < \sum_{x \geq p} g_{i+1}(x), \quad \forall p \in \{x_{i+1}, \cdots, x_m\},$$

whence g_{i+1} first-order stochastically dominates g_i. By Shaked and Shanthikumar (2007, Thm. 1.A.8), we conclude that

$$x_i + \overline{u}(x_i) = \sum_{x \geq x_i} x g_i(x) < \sum_{x \geq x_{i+1}} x g_{i+1}(x) = x_{i+1} + \overline{u}(x_i).$$

In other words, $x_i + \overline{u}(x_i)$ is strictly increasing in i. The previous remark establishes:

Proposition 2. *Along the Pareto frontier, increasing consumer surplus implies lowering social welfare.*

A central insight of Bergemann, Brooks, and Morris (2015) is that, in a single-product setting, market segmentation can be used as a tool to efficiently redistribute the gains from trade. Proposition 2 shows that, contrastingly, when the number of products is large, efficiently transferring surplus from the seller to the consumers through segmentation is infeasible. The broad idea is simple. Achieving efficiency is harder with product variety than without: whereas in a single-product setting efficiency obtains as long as trade occurs with probability 1, with product variety efficiency also requires each consumer to buy one of the products that he values the most. When product variety is large, the goal of achieving efficiency thus collides with that of inducing low prices. Along the Pareto frontier, the transfer of surplus from seller to consumers is achieved by segmenting the aggregate market in a way that leads the seller to occasionally offer products which do not accurately fit consumers' tastes.

The preceding discussion suggests that consumer privacy plays a key role in the determination of consumer surplus, and that securing the former might help to increase the latter. The remainder of this section examines the link between consumer privacy and welfare in greater details.

To formalize the notion of privacy, we build on Blackwell (1953). We say that a market segmentation τ' is *finer* than τ if there exists a function $\xi : \text{supp } \tau \to \Delta\Delta X^n$ such that, for every $\mu \in \text{supp } \tau$,

$$\mu(\mathbf{v}) = \sum_{\tilde{\mu}} \xi(\tilde{\mu} \mid \mu)\tilde{\mu}(\mathbf{v}), \quad \forall \mathbf{v} \in X^n,$$

and

$$\tau'(\tilde{\mu}) = \sum_{\mu} \tau(\mu)\xi(\tilde{\mu} \mid \mu).$$

Intuitively, we obtain τ' by splitting every market μ in the support of the market segmentation τ.[13] In this sense, τ gives greater privacy to consumers than τ'.

While greater consumer privacy evidently harms the seller, the effect of consumer privacy on the welfare of consumers is a lot more complex. On the one hand, privacy prevents the seller from extracting surplus through personalized prices. On the other hand, making detailed information available to the seller enables the latter to improve the match quality between consumers and products.[14] Consequently, whether increasing privacy benefits or harms consumers is generally ambiguous.

Our next result identifies a precise sense in which greater consumer privacy is associated with greater consumer surplus. Specifically, pick an arbitrary market segmentation τ: when product variety is large, some payoff-equivalent market segmentation $\tilde{\tau}$ is such that consumer surplus may be increased by giving greater privacy to consumers. The following proposition formalizes this insight.

Proposition 3. *Let (π, u) and (π', u') be two points in S, with $\pi' > \pi$. For every $n \in \mathbb{N}$, there exist market segmentations τ_n, τ'_n, where τ'_n is finer than τ_n, and strategies $\rho_n \in \text{argmax}_\rho \Pi_{\tau_n}(\rho)$ as well as $\rho'_n \in \text{argmax}_\rho \Pi_{\tau'_n}(\rho)$, such that*

$$\Big(\Pi_{\tau_n}(\rho_n), U_{\tau_n}(\rho_n)\Big) \underset{n\to\infty}{\longrightarrow} (\pi, u) \quad \textit{and} \quad \Big(\Pi_{\tau'_n}(\rho'_n), U_{\tau'_n}(\rho'_n)\Big) \underset{n\to\infty}{\longrightarrow} (\pi', u').$$

[13] Interpreting τ and τ' as distributions of posterior beliefs induced by Blackwell–experiments α and α', respectively, our notion corresponds to α' being "sufficient" for α, one of several equivalent definitions of "more informative" in Blackwell (1953).

[14] This trade-off plays a central role in the analysis of Ichihashi (2020) and Hidir and Vellodi (2021).

In plain words, if two points P' and P in the set S are such that producer surplus is greater at P' than at P, then P' and P can be approximately achieved by market segmentations τ' and τ with τ' finer than τ.

The proof of Proposition 3 rests on two basic ideas. Firstly, different market segmentations lead the seller to offer different products. Secondly, for a given market segmentation, the seller tends to offer the subset of products reagarding which the segmentation is "most informative" (that is, with regard to which the market segmentation best distinguishes consumers). Now suppose $\pi < \pi'$, and we want to find market segmentations τ and τ' respectively generating surplus π and π' for the seller. Proceed as follows. Firstly, partition the products in two subsets of equal size, say K and K'. Then construct τ by separating consumers *exclusively* with respect to their valuations for the products in K. Under the market segmentation τ, the seller offers products in the subset K and obtains surplus π. Finally, construct τ' by splitting every market in the support of τ according to consumers' valuations for the products in K'. Under this finer market segmentation τ', the seller offers products in the subset K' and obtains surplus $\pi' > \pi$.

6 Online Markets with Data Intermediaries

In this section, we study the sale of consumer data by data intermediaries.

The setting is as follows. There are a seller with an inventory comprising n products, and a unit-demand consumer. The consumer's valuation for product k is denoted by v_k, and the valuation vector by $\mathbf{v} = (v_1, \ldots, v_n)$. These valuations are initially unknown to all parties; the common prior probability assigned to $\mathbf{v} = \mathbf{x}$ is given by $\bar{\mu}(\mathbf{x})$, defined by (1).

The setting also comprises $l \geq 1$ data intermediaries, each of whom chooses a *data policy*, that is, a tuple (D, ϕ) where D is a set of signals and ϕ a mapping

$$\phi : X^n \to \Delta D.$$

Under data policy (D, ϕ), the signal $d \in D$ is drawn with probability $\phi(d \mid \mathbf{x})$ if the consumer has valuation vector $\mathbf{v} = \mathbf{x}$. The signals of different data intermediaries are drawn independently conditional on $\mathbf{v}$.

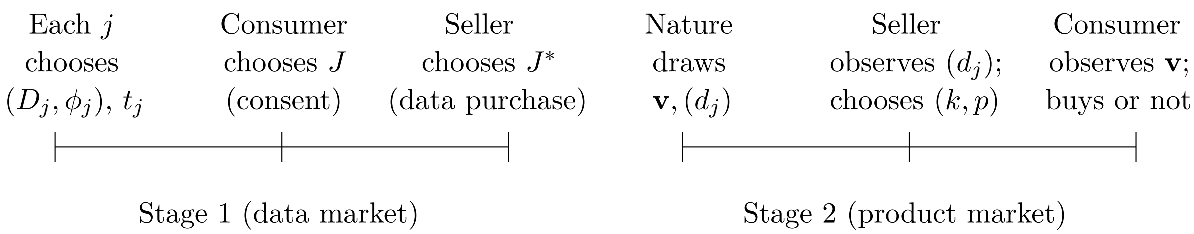

Figure 2: Timeline.

There are two stages: Stage 1 describes the data market, and Stage 2 describes the product market. The timeline is depicted in Figure 2.

Stage 1 (data market). First, every data intermediary $j = 1, \ldots, l$ chooses a data policy (D_j, ϕ_j), as well as a fee t_j at which it intends to sell the data d_j generated by this policy. At this point, the consumer selects a subset of data intermediaries, say $J \subseteq \{1, \ldots, l\}$, comprising all data intermediaries receiving his consent. The seller then purchases data from a subset of data intermediaries J^* selected from the set J.

Stage 2 (product market). First, the valuation vector $\mathbf{v}$ is drawn from the distribution $\bar{\mu}$. Then, the signals of the data intermediaries $j \in J^*$ are drawn according to their data policies. The seller observes these signals, and chooses which product to offer and at what price. Lastly, the consumer learns his valuation for the offered product, and decides whether to buy.

The payoff of a data intermediary is its revenue from selling data. If the consumer buys the product offered by the seller, his payoff equals his valuation minus the price; otherwise his payoff is zero. The payoff of the seller equals her revenue minus the cost $\sum_{j \in J^*} t_j$ of acquiring data.

The solution concept is perfect Bayesian equilibrium, with two refinements: firstly, the data intermediaries use pure strategies; secondly, the seller breaks ties in favor of the consumer, both when purchasing data and when choosing a product-price combination.

6.1 Discussion of the Model

A key assumption of the model is that the data intermediaries must obtain the consent of the consumer before selling information to the seller. This assumption is consistent with the EU General Data Protection Regulation (Regulation (EU) 2016/679, Articles 4 and 7), among other things.

Our assumption that the data intermediaries sell information to the seller *directly* accords with the business model of firms such as Acxiom, Nielsen, and Oracle, for example. On the other hand, online platforms acting as data intermediaries, such as Google and Facebook, do not sell information per se, but sell instead access to targeted consumer segments. To keep the analysis focused, we disregard in this paper the distinction between direct and indirect sale of information.[15]

The assumption that the data intermediaries know precisely the consumer's valuation vector evidently lacks realism, and merely ensures tractability. In particular, in practice one of the gains from having multiple data intermediaries may be that different intermediaries possess complementary information about consumers' preferences. Such considerations are beyond the scope of our analysis.

Finally, the model makes a number of technical assumptions. The assumption that a data intermediary simultaneously chooses its data policy and the fee at which it intends to sell its data simplifies the structure of the game, but is irrelevant for our results. Our assumption that the seller breaks ties in favor of the consumer ensures that each data market outcome at the end of Stage 1 induces both a unique expected revenue for the seller and a unique expected payoff for the consumer. Our focus on pure strategies circumvents possible miscoordination among data intermediaries.

[15]See Federal Trade Commission (2014) and Bergemann and Bonatti (2019) for details about the various business models of data intermediaries.

6.2 Data Intermediation and Welfare

We now characterize the equilibrium payoffs when the number of products is large.[16] Define

$$u_0 := \sum_{x \geq p_0} f(x)(x - p)$$

and

$$\pi_A := \max \Big\{ \pi \in [\pi_0, x_m] \mid \overline{u}(\pi) = u_0 \Big\}.$$

To simplify the statement of the next proposition, we assume that the function $\overline{u}$ possesses a unique maximizer in $[\pi_0, x_m]$,[17] which we denote by π_B.

Proposition 4. *For every $n \in \mathbb{N}$, fix some equilibrium. Let $(\pi_n, u_n)_{n \in \mathbb{N}}$ be the corresponding combinations of expected revenue of the seller and expected payoff of the consumer. If $l = 1$, then $(\pi_n, u_n) \underset{n \to \infty}{\longrightarrow} \Big(\pi_A, \overline{u}(\pi_A)\Big)$; if $l > 1$, then $(\pi_n, u_n) \underset{n \to \infty}{\longrightarrow} \Big(\pi_B, \overline{u}(\pi_B)\Big)$.*

The proposition can be understood as follows. A monopolistic data intermediary ($l = 1$) fully extracts the seller's gain from purchasing data. This results in a data policy that maximizes the seller's expected revenue, subject to the constraint that the consumer gives his consent. By contrast, when the data market is competitive ($l > 1$), the implemented payoff pair maximizes the consumer's expected payoff, subject to the constraint that the seller purchases the data.

We sketch here the main ideas of the proof. Every data market outcome at the end of Stage 1 induces a market segmentation τ. In the current setting, $\tau(\mu)$ represents the probability that the seller's posterior belief concerning $\mathbf{v}$ (after observing the signals) is equal to μ. Thus, every data market outcome induces an expected revenue π_n for the seller, and an expected payoff u_n for the consumer, such that (π_n, u_n) belongs to the set S_n of feasible surplus pairs defined in Section 2.[18] We then prove that S_n satisfies a

[16] We omit a proof of the existence of an equilibrium. Lemma B1 in Section B of the Appendix implies that both consumer-optimal and seller-optimal data policies exist. Based on this, it is straightforward to deduce existence of an equilibrium.

[17] This is the case if $\sum_{x \geq x_i} g_i(x)(x - x_i) \neq \sum_{x \geq x_{i+1}} g_{i+1}(x)(x - x_{i+1})$ for all $i = 1, \ldots, m - 1$.

[18] Specifically, (π_n, u_n) belongs to the subset of S_n that consists of the surplus pairs which are consistent with the seller breaking ties in favor of the consumer.

number of properties which enable us to pin down both the expected revenue of the seller and the expected payoff of the consumer in any equilibrium. Finally, an application of Theorem 1 yields the convergences stated in Proposition 4.

6.3 Data Intermediation and Consumer Privacy

We have focused so far on the implications of competition in the data market for welfare. We now extend the scope of our analysis, and investigate the effect of competition in the data market on consumer privacy. Precisely characterizing the data policies effectively used in equilibrium is a hard problem; to make progress, we thus focus in this subsection on the case of binary valuations: $X = \{x_1, x_2\}$.

We rank the privacy afforded by different data policies according to Blackwell–informativeness. A data policy (D, ϕ) is *more informative* than another data policy (D', ϕ') if there exists a function $\sigma : D \to \Delta D'$ such that

$$\phi'(d' \mid \mathbf{x}) = \sum_{d \in D} \phi(d \mid \mathbf{x})\sigma(d' \mid d), \quad \forall d' \in D', \forall \mathbf{x} \in X^n.$$

Thus, (D', ϕ') differs from (D, ϕ) by additional noise.

Below, say that the seller purchases data given by (D, ϕ) if in Stage 1 the seller purchases data from a single data intermediary, say j, and $(D_j, \phi_j) = (D, \phi)$. We can now state the main result of this subsection.

Proposition 5. *Let $X = \{x_1, x_2\}$. Fix $l' \in \mathbb{N}$ with $l' > 1$, and $n \in \mathbb{N}$ with $n \geq (\ln f(x_2) + \ln(x_2 - x_1) - \ln x_1)/\ln f(x_1)$. There exist two data policies, (D, ϕ) and (D', ϕ'), as well as an equilibrium for $l = 1$ and an equilibrium for $l = l'$, such that:*

- *(D, ϕ) is more informative than (D', ϕ');*
- *in the equilibrium for $l = 1$, the seller purchases data given by (D, ϕ);*
- *in the equilibrium for $l = l'$, the seller purchases data given by (D', ϕ');*

In the proof, we construct two equilibrium data policies (D, ϕ) and (D', ϕ'). They have a straightforward structure, and we illustrate them here by means of Figure 3 for

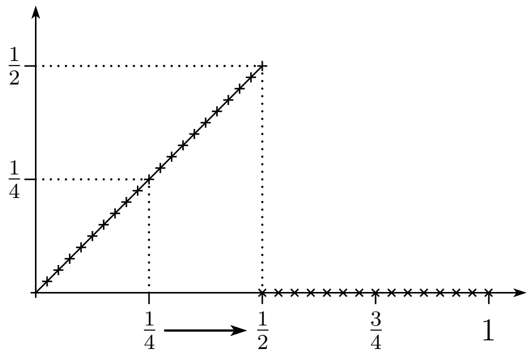

(a) Competition ($l = l'$)

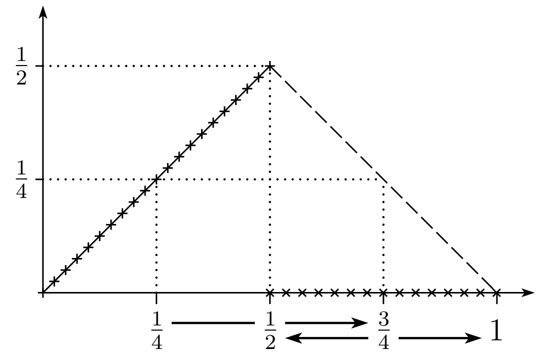

(b) Monopoly ($l = 1$)

Figure 3: Illustration of equilibrium data policies, assuming $x_1 = 1$, $x_2 = 2$, $f(x_2) = 1/4$. Horizontal axis: probability $\Pr[v_{k^*} = x_2]$; vertical axis: consumer's expected payoff.

parameters $x_1 = 1$, $x_2 = 2$, $f(x_2) = 1/4$. As a benchmark, note that in the absence of information, the seller offers an arbitrary product at a price of x_1.[19] In this case, $f(x_2)$ is the probability that the consumer's valuation for the offered product is equal to x_2, and the consumer obtains an expected payoff equal to

$$f(x_2)(x_2 - x_1) = \frac{1}{4}.$$

Panel (a) refers to the competitive data market. The seller purchases data enabling her to identify, with some noise, a product k^* for which the consumer's valuation is the highest across all products. More specifically, the seller obtains data making her indifferent between offering product k^* at a price of x_1, or at a price of x_2. She chooses the lower price x_1, resulting for the consumer in an expected payoff of

$$\frac{x_1}{x_2}(x_2 - x_1) = \frac{1}{2}.$$

Panel (b) refers to the monopolistic data market. Once again, the seller purchases data enabling her to identify, with some noise, a product k^* for which the consumer's valuation is the highest across all products. But, the data is more informative: the posterior probability that $v_{k^*} = x_2$ is now equal to $3/4 > x_1/x_2$. In principle, the seller would therefore offer product k^* at a price of x_2, leaving zero surplus to the consumer.

[19]For $X = \{x_1, x_2\}$, the assumption $\pi_0 \in X$ implies $p_0 = \pi_0 = x_1$.

To secure the consumer's consent, the intermediary provides the seller with yet more information: in essence, the seller receives one of two additional signals, say d^+ and d^-. Conditional on d^+, the seller knows that $v_{k^*} = x_2$; however, conditional on d^-, the probablity that $v_{k^*} = x_2$ is again x_1/x_2, and so the seller offers product k^* at a price of x_1. The resulting expected payoff of the consumer is

$$\frac{1}{4} = f(x_2)(x_2 - x_1).$$

7 Conclusion

TO BE ADDED.

Appendix

A Omitted Proofs for Sections 3–5

Proof of Lemma 2. Since τ is the product of $\{\tau^k\}_{k=1}^n$, notice that $\mu \in \text{supp}\ \tau$ if and only if (7) holds and $\mu_k \in \text{supp}\ \tau^k$ for every k. Then, using (7), (8), and (9) gives

$$\begin{aligned}\sum_{\mu} \tau(\mu)\mu(\mathbf{v}) &= \sum_{\substack{\mu_1 \in \text{supp}\ \tau^1, \\ \cdots, \\ \mu_n \in \text{supp}\ \tau^n}} \prod_k \tau^k(\mu_k)\mu_k(v_k) \\ &= \prod_k \left(\sum_{\mu_k \in \text{supp}\ \tau^k} \tau^k(\mu_k)\mu_k(v_k) \right) \\ &= \prod_k f(v_k) = \bar{\mu}(\mathbf{v}),\end{aligned}$$

for all $\mathbf{v} \in X^n$. □

Proof of Proposition 1. The first part of the proposition follows from Theorem 1 because $\tilde{S}_n \subseteq S_n$.

We next prove the second part. We treat below the case $\pi = x_i \in X$ and $u = \zeta\overline{u}(x_i)$, where $\zeta \in [0, 1]$; the proof for the remaining case is similar, and therefore omitted.

Let $\lambda \in (0, 1)$, and $h \in \Delta X$ be given by (11). We claim that

$$x_i \geq \pi_0 \geq \max_p \ p \sum_{x \geq p} h(x). \tag{A.1}$$

The first inequality holds because $(\pi, u) \in S$. We next show the second inequality. For all $p \leq x_i$,

$$\pi_0 \geq p \sum_{x \geq p} f(x) = \lambda p \sum_{x \geq p} g_i(x) + (1 - \lambda)p \sum_{x \geq p} h(x) = \lambda p + (1 - \lambda)p \sum_{x \geq p} h(x).$$

Hence, $(1 - \lambda)p \sum_{x \geq p} f(x) \geq p \sum_{x \geq p} f(x) - \lambda p = (1 - \lambda)p \sum_{x \geq p} h(x)$. For all $p > x_i$,

$$\pi_0 \geq p \sum_{x \geq p} f(x) = \lambda p \sum_{x \geq p} g_i(x) + (1 - \lambda)p \sum_{x \geq p} h(x) = \lambda x_i + (1 - \lambda)p \sum_{x \geq p} h(x).$$

In either case, we obtain $\pi_0 \geq p \sum_{x \geq p} h(x)$, which finishes to prove (A.1).

Next, let K^{x_i} and K^{x_m} be two disjoint subsets of $\{1, \cdots, n\}$, each containing $(n-1) \div 2$ elements (where $\div$ denotes division with remainder). Let $k_0 \in \{1, \cdots, n\} \setminus (K^{x_i} \cup K^{x_m})$. For each $p \in \{x_i, x_m\}$, define $\tau_k^p \in \Delta X$ by

$$\tau_k^p(g_i) = \lambda = 1 - \tau_k^p(h), \quad \forall k \in K^p,$$

and

$$\tau_k^p(f) = 1, \quad \forall k \in \{1, \ldots, n\} \setminus K^p.$$

Lastly, let τ^p be the product of $\{\tau_k^p\}_{k=1}^n$. By coupling (12) with Lemma 2, notice that τ^p is a market segmentation. Consequently, the mixture

$$\tau := \zeta \tau^{x_i} + (1 - \zeta) \tau^{x_m}$$

is a market segmentation too.

Each market μ in the support of τ satisfies exactly one of the following conditions:

(a) $\mu_k = g_i$ for some $k \in K^{x_i}$;

(b) $\mu_k = g_i$ for some $k \in K^{x_m}$;

(c) $\mu_k \in \{h, f\}$ for all $k \in K^{x_i} \cup K^{x_m}$.

Now let ρ be a strategy of the seller with the following properties:

- for every market $\mu \in \text{supp}\, \tau$ satisfying (a), offer one of the products $k \in K^{x_i}$ for which $\mu_k = g_i$ at a price of x_i;
- for every market $\mu \in \text{supp}\, \tau$ satisfying (b), offer one of the products $k \in K^{x_m}$ for which $\mu_k = g_i$ at a price of x_m;
- for every market $\mu \in \text{supp}\, \tau$ satisfying (c), offer product k_0 at price p_0.

By (A.1), the strategy ρ is optimal given τ. Furthermore, note that ρ does not price discriminate: the products in K^{x_i} (respectively, K^{x_m}) are always offered at a price of x_i (respectively, x^m), and product k_0 is always offered at price p_0.

The resulting surplus of the seller is

$$\pi_n := \Pi_\tau(\rho) = \Big(1 - (1 - \lambda)^{(n-1) \div 2}\Big) x_i + (1 - \lambda)^{(n-1) \div 2} \pi_0;$$

the consumer surplus is

$$u_n := U_\tau(\rho) = \zeta\Big(1-(1-\lambda)^{(n-1)\div 2}\Big)\sum_{x\geq x_i} g_i(x)(x-x_i) + (1-\lambda)^{(n-1)\div 2}\sum_{x\geq p_0} f(x)(x-p_0).$$

So $(\pi_n, u_n) \in \tilde{S}_n$, and as $\lambda > 0$,

$$(\pi_n, u_n) \underset{n\to\infty}{\longrightarrow} \Big(x_i, \zeta\overline{u}(x_i)\Big). \qquad \square$$

Proof of Proposition 3. We will use the following lemma. Its proof is analogous to the proof of Lemma 2, and therefore omitted.

Lemma A1. *For every $k \in \{1, \ldots, n\}$, let $\tau^k \in \Delta X$ and $\xi^k : \Delta X \to \Delta\Delta X$ satisfy*

$$\sum_{\mu_k} \tau^k(\mu_k)\mu_k(x) = f(x), \quad \forall x \in X,$$

$$\sum_{\tilde{\mu}_k} \xi^k(\tilde{\mu}_k \mid \mu_k)\tilde{\mu}_k(x) = \mu_k(x), \quad \forall x \in X, \forall \mu_k \in \Delta X.$$

Define $\langle \tau^k, \xi^k\rangle \in \Delta X$ by

$$\langle \tau^k, \xi^k\rangle(\tilde{\mu}_k) = \sum_{\mu_k} \tau^k(\mu_k)\xi^k(\tilde{\mu}_k \mid \mu_k), \quad \forall \tilde{\mu}_k \in \Delta X. \tag{A.2}$$

Let τ be the product of $\{\tau^k\}_{k=1}^n$, and $\hat{\tau}$ be the product of $\{\langle \tau^k, \xi^k\rangle\}_{k=1}^n$. Then both τ and $\hat{\tau}$ are market segmentations, and $\hat{\tau}$ is finer than τ.

We can now prove the proposition. We treat below the case $(\pi, u) = (x_i, \overline{u}(x_i))$ and $(\pi', u') = (x_j, \overline{u}(x_j))$ with $x_i, x_j \in X$ and $x_i < x_j$; the proof for the remaining cases is similar, and therefore omitted.

Let $\lambda_i \in (0, 1)$, and define $h_i \in \Delta X$ by

$$h_i(x) := \frac{f(x) - \lambda_i g_i(x)}{1-\lambda_i},$$

as in the proof of Lemma 3. Moreover, let $\lambda_j \in (0, 1)$, and define $h_j \in \Delta X$ analogously.

Next, for $n > 1$, define $\tau_n^k \in \Delta X$ by

$$\tau_n^k(g_i) = \lambda_i = 1 - \tau_n^k(h_i), \quad \forall k \in \{1, \ldots, n \div 2\}$$

(where $\div$ denotes division with remainder), and

$$\tau_n^k(f) = 1, \quad \forall k \in \{(n \div 2) + 1, \ldots, n\}.$$

Moreover, define $\xi_n^k : \Delta X \to \Delta\Delta X$ by

$$\xi_n^k(g_j \mid f) = \lambda_j = 1 - \xi_n^k(h_j \mid f), \quad \forall k \in \{(n \div 2) + 1, \dots, n\},$$

and $\xi_n^k(\mu_k \mid \mu_k) = 1$ if $k \in \{1, \dots, n \div 2\}$ or $\mu_k \neq f$. Lastly, let τ_n be the product of $\{\tau_n^k\}_{k=1}^n$, and $\hat{\tau}_n$ the product of $\{\langle \tau_n^k, \xi_n^k \rangle\}_{k=1}^n$, where $\langle \tau_n^k, \xi_n^k \rangle$ was defined in (A.2). By Lemma A1, both τ_n and $\hat{\tau}_n$ are market segmentations, and $\hat{\tau}_n$ is finer than τ_n.

Every market μ in the support of τ_n satisfies $\mu_k \in \{g_i, h_i\}$ for all $k \in \{1, \dots, n \div 2\}$, and $\mu_k = f$ for all $k \in \{(n \div 2) + 1, \dots, n\}$. By (13), there exists a strategy $\rho_n \in \operatorname{argmax}_\rho \Pi_{\tau_n}(\rho)$ with the following property:

> For every market $\mu \in \operatorname{supp} \tau_n$ satisfying $\mu_k = g_i$ for some $k \in \{1, \dots, n \div 2\}$, offer a product k for which $\mu_k = g_i$ at a price of x_i.

Because the probability that $\mu_k = g_i$ for some $k \in \{1, \dots, n \div 2\}$ is $1 - (1 - \lambda_i)^{n \div 2}$, which tends to 1 as n grows without bound, it holds that

$$\lim_{n\to\infty} \Pi_{\tau_n}(\rho_n) = \lim_{n\to\infty} (1 - (1 - \lambda_i)^{n \div 2}) x_i,$$
$$\lim_{n\to\infty} U_{\tau_n}(\rho_n) = \lim_{n\to\infty} (1 - (1 - \lambda_i)^{n \div 2}) \sum_{x \geq x_i} g_i(x)(x - x_i).$$

Consequently,

$$\lim_{n\to\infty} \Big(\Pi_{\tau_n}(\rho_n), U_{\tau_n}(\rho_n)\Big) = \Big(x_i, \overline{u}(x_i)\Big).$$

Every market μ in the support of $\hat{\tau}_n$ satisfies $\mu_k \in \{g_i, h_i\}$ for all $k \in \{1, \dots, n \div 2\}$, and $\mu_k \in \{g_j, h_j\}$ for all $k \in \{(n \div 2) + 1, \dots, n\}$. By (13), and since $x_j > x_i$, there exists a strategy $\hat{\rho}_n \in \operatorname{argmax}_\rho \Pi_{\hat{\tau}_n}(\rho)$ with the following property:

> For every market $\mu \in \operatorname{supp} \hat{\tau}_n$ satisfying $\mu_k = g_j$ for some $k \in \{(n \div 2) + 1, \dots, n\}$, offer a product k for which $\mu_k = g_j$ at a price of x_j.

Then, as above,

$$\lim_{n\to\infty} \Big(\Pi_{\hat{\tau}_n}(\hat{\rho}_n), U_{\hat{\tau}_n}(\hat{\rho}_n)\Big) = \Big(x_j, \overline{u}(x_j)\Big). \qquad \square$$

B Proofs for Section 6

The proofs of Propositions 4 and 5 build on a characterization of the expected revenues of the seller and the expected payoffs of the consumer that can result from arbitrary data policies. We state this characterization in the following subsection.

B.1 Preliminaries

Every data market outcome at the end of Stage 1 induces a distribution $\tau \in \Delta\Delta X^n$ of posterior beliefs $\mu \in \Delta X^n$; furthermore, this τ is a market segmentation. Conversely: every market segmentation τ is the distribution of posterior beliefs induced by some data policy (see Kamenica and Gentzkow, 2011).

Next, the profile of data policies $\Big((D_j, \phi_j)\Big)_{j \in J^*}$ has the same informational content as the "aggregate" data policy (D', ϕ') with $D' = \prod_{j \in J} D_j$ and ϕ' given by

$$\phi'\Big((d_j)_{j \in J^*} \mid \mathbf{x}\Big) = \prod_{j \in J} \phi_j(d_j \mid \mathbf{x}).$$

Hence, we can represent any profile of data policies by a single data policy.

Combining the previous observations shows that every data market outcome at the end of Stage 1 induces a subgame in Stage 2 in which the seller obtains an expected revenue π_n and the consumer obtains an expected payoff u_n such that (π_n, u_n) belongs to the set S_n of feasible surplus pairs defined in Section 2. In fact, since here the seller breaks ties in favor of the consumer, the previous payoffs are uniquely pinned down by the aggregate data policy (D', ϕ'). We thus say that (D', ϕ') *implements* (π_n, u_n). The set of pairs (π, u) which can be implemented by some data policy (D', ϕ') will be denoted by $\hat{S}_n$.

The following two lemmas provide a characterization of the set $\hat{S}_n$. Their proofs are relegated to Section OA.1 of the Online Appendix.

Lemma B1. *Define* $\pi_n^{\max} := \sum_{\mathbf{w}} \bar{\mu}(\mathbf{w}) \max_k w_k$. *For every* $n \in \mathbb{N}$, *it holds that:*

a) $\{\pi \in \mathbb{R} \mid$ *there exists* $u \in \mathbb{R}$ *s.t.* $(\pi, u) \in \hat{S}_n\} = [\pi_0, \pi_n^{\max}]$;

b) $\{u \in \mathbb{R} \mid (\pi, u) \in \hat{S}_n\}$ *has a greatest element* $\overline{u}_n(\pi)$ *for every* $\pi \in [\pi_0, \pi_n^{\max}]$;

c) $\overline{u}_n : [\pi_0, \pi_n^{\max}] \to \mathbb{R}$ *is concave and continuous, and* $\overline{u}_n(\pi_n^{\max}) = 0$.

Lemma B2. *For every* $n \in \mathbb{N}$, $\hat{S}_n \subseteq S$. *Moreover, for every* $(\pi, \overline{u}(\pi)) \in S$, *there exists a sequence* $\big((\pi_n, u_n)\big)_{n\in\mathbb{N}}$ *such that* $(\pi_n, u_n) \in \hat{S}_n$ *and* $(\pi_n, u_n) \underset{n\to\infty}{\longrightarrow} (\pi, \overline{u}(\pi))$.

B.2 Proofs of Propositions 4 and 5

Proof of Proposition 4. We abbreviate "data intermediary" to "DI".

(i) Suppose there is a single DI. Fix some $n \in \mathbb{N}$. In every perfect Bayesian equilibrium, the consumer's expected payoff is at least u_0, which the consumer obtains if he does not give consent to the proposed data policy. Hence, the seller's expected gross payoff is at most

$$\max\{\pi \in \mathbb{R} \mid \text{there exists } u \geq u_0 \text{ s.t. } (\pi, u) \in \hat{S}_n\}.$$

By Lemma B1, the maximum exists and is equal to

$$\pi_n^* := \max\{\pi \in [\pi_0, \pi_n^{\max}] \mid \overline{u}_n(\pi) = u_0\}.$$

In every perfect Bayesian equilibrium, the DI chooses fee $\pi_n^* - \pi_0$ and a data policy that implements (π_n^*, u_0), the consumer consents, and the seller purchases the data.

It remains to show that $\lim_{n\to\infty} \pi_n^* = \pi_A$. There are two cases. Case (i): π_A coincides with π_B, the unique maximizer of $\overline{u}$. Then $\overline{u}(\pi_B) = \overline{u}(\pi_A) = u_0$, which implies $\pi_n^* = \pi_B = \pi_A = \pi_0$ for all n and hence $\lim_{n\to\infty} \pi_n^* = \pi_A$. Case (ii): $\pi_A \neq \pi_B$. Then $\overline{u}$ is strictly decreasing at π_A. By the second part of Lemma B2, we can find $n \in \mathbb{N}$ such that $(\pi, u) \in \hat{S}_n$ with π in any neighborhood of π_A and $u > \overline{u}(\pi_A) = u_0$. Consequently, $\liminf_{n\to\infty} \pi_n^* \geq \pi_A$. On the other hand, $\limsup_{n\to\infty} \pi_n^* \leq \pi_A$ because $\hat{S}_n \subseteq H$ by the first part of Lemma B2. Thus, $\lim_{n\to\infty} \pi_n^* = \pi_A$

(ii) Suppose there is more than one DI. Fix some $n \in \mathbb{N}$. First, we show that the consumer's expected payoff in every perfect Bayesian equilibrium is equal to

$$\max\{u \in \mathbb{R} \mid \text{there exists } \pi \in \mathbb{R} \text{ s.t. } (\pi, u) \in \hat{S}_n\}.$$

By Lemma B1, the maximum exists and is

$$u_n^* := \max\{\overline{u}_n(\pi) \mid \pi \in [\pi_0, \pi_n^{\max}]\}.$$

By contradiction, suppose the consumer's expected payoff in some perfect Bayesian equilibrium is $u < u_n^*$. Let π be the seller's expected revenue in this equilibrium. Then, the sum of the expected payoffs of the DIs is at most $\pi - \pi_0$. Consequently, there exists a DI j whose expected payoff is strictly smaller than $(\pi - \pi_0)/2$. By Lemma B1, part c), we can find $\left(\pi', \overline{u}_n(\pi')\right) \in \hat{S}_n$ with $\pi' > (\pi + \pi_0)/2$ and $\overline{u}_n(\pi') > u$. Suppose DI j chooses fee $\pi' - \pi_0$ and a data policy that implements $\left(\pi', \overline{u}_n(\pi')\right)$. We show that the consumer would give consent to j, and the seller would purchase j's data.

Indeed, if the consumer gives consent to DI j alone, then the seller must purchase j's data, given our equilibrium restriction that the seller breaks ties in favor of the consumer when purchasing data. The consumer's expected payoff is then $\overline{u}_n(\pi') > u$. Suppose the consumer gives consent to a subset of DIs not including j. This choice was also possible when j did not deviate, and it would have resulted in the same expected payoff for the consumer. Hence, the consumer's expected payoff is at most $u < \overline{u}_n(\pi')$ in this case. Suppose finally the consumer gives consent to a subset J of DIs that includes j. If the seller then does not purchase j's data, the consumer's expected payoff is the same as if he gives consent just to the DIs $J \setminus \{j\}$.

Thus, if DI j chooses chooses fee $\pi' - \pi_0$ and a data policy that implements $\left(\pi', \overline{u}_n(\pi')\right)$, the consumer consents and the seller purchases the data. With this deviation, j's payoff is $\pi' - \pi_0 > (\pi - \pi_0)/2$, contradicting the hypothesis that j earns $(\pi - \pi_0)/2$ in equilibrium. Hence, the consumer obtains $u = u_n^*$ in every perfect Bayesian equilibrium.

Choose $\hat{\pi}_n \in [\pi_0, \pi_n^{\max}]$ such that $\overline{u}_n(\hat{\pi}_n) = u_n^*$. We show $\lim_{n\to\infty}(\hat{\pi}_n, \overline{u}(\hat{\pi}_n)) = (\pi_B, \overline{u}(\pi_B))$. Because π_B is the unique maximizer of $\overline{u}$, $\overline{u}_n(\hat{\pi}_n) \leq \overline{u}(\pi_B)$ by the first part of Lemma B2. Hence, $\limsup_{n\to\infty} \overline{u}_n(\hat{\pi}_n) \leq \overline{u}(\pi_B)$. By the second part of Lemma B2, $\liminf_{n\to\infty} \overline{u}_n(\hat{\pi}_n) \geq \overline{u}(\pi_B)$. Thus, $\lim_{n\to\infty} \overline{u}_n(\hat{\pi}_n) = \overline{u}(\pi_B)$. By contradiction, suppose $\liminf_{n\to\infty} \hat{\pi}_n = \pi' < \pi_B$ or $\limsup_{n\to\infty} \hat{\pi}_n = \pi' > \pi_A$. Then, $\overline{u}_n(\pi') \leq \overline{u}(\pi') < \overline{u}(\pi_B)$, contradicting $\lim_{n\to\infty} \overline{u}_n(\hat{\pi}_n) = \overline{u}(\pi_B)$. Thus, $\lim_{n\to\infty} \hat{\pi}_n = \pi_B$. □

Proof of Proposition 5. Note that for $X = \{x_1, x_2\}$, the assumption $p_0 \in X$ implies $\pi_0 = p_0 = x_1$. Thus,

$$x_1 > f(x_2)x_2 \quad \iff \quad f(x_2) < \frac{x_1}{x_2}.$$

Moreover, $u_0 = f(x_2)(x_2 - x_1)$, and

$$\overline{u}(\pi) = \frac{x_2 - \pi}{x_2 - x_1}\frac{x_1}{x_2}(x_2 - x_1), \quad \forall \pi \in [x_1, x_2].$$

Lastly,

$$n \geq \frac{\ln f(x_2) + \ln(x_2 - x_1) - \ln x_1}{\ln f(x_1)} \implies (f(x_1))^n \leq f(x_2)\frac{x_2 - x_1}{x_1}. \tag{B.1}$$

The proof consists of three steps.

Step 1: an equilibrium for $l = 1$. Suppose $l = 1$. If the data intermediary proposes a data policy that implements (π, u), it is optimal for the consumer to give his consent if $u \geq u_0$, and it is optimal for the seller to purchase the data if the fee is at most $\pi - \pi_0$.

By Lemma B2, the set of (π, u) that can be implemented by data policies is $\hat{S}_n \subseteq S$. In S, the seller's expected revenue is maximal at $\big(\pi_A, \overline{u}(\pi_A)\big)$, where $\pi_A = \max\big\{\pi \in [x_1, x_2] \mid \overline{u}(\pi) = u_0\big\}$. For future reference, note that

$$\overline{u}(\pi_A) = u_0 \quad \iff \quad \pi_A = x_2 - f(x_2)\frac{x_2 - x_1}{x_1}x_2.$$

We present a data policy (D, ϕ) that implements $\big(\pi_A, \overline{u}(\pi_A)\big)$. It then follows that there exists an equilibrium in which the data intermediary proposes (D, ϕ) at fee $\pi_A - \pi_0$, the consumer gives his consent, and the seller purchases the data.

Consider the data policy (D, ϕ), where $D = \{1, \ldots, n\} \times X$ and ϕ is defined as follows. With probability

$$\frac{\pi_A/x_2 - f(x_2)}{1 - (f(x_1))^n - f(x_2)},$$

the first component of the signal $k \in \{1, 2\}$ is drawn uniformly at random from the set $\operatorname{argmax}_{k'} v_{k'}$. This is indeed a number between zero and one by (B.1). With the remaining probability, k is drawn uniformly at random from $\{1, \ldots, n\}$. Thus, the first component displays with some noise a product for which the consumer has the highest valuation across all products. Based on this information, the posterior probability that $v_k = x_2$ when k is displayed is

$$\frac{\pi_A/x_2 - f(x_2)}{1 - (f(x_1))^n - f(x_2)}(1 - (f(x_1))^n) + \left(1 - \frac{\pi_A/x_2 - f(x_2)}{1 - (f(x_1))^n - f(x_2)}\right) f(x_2) = \frac{\pi_A}{x_2}.$$

If the first component of the signal is k and $v_k = x_2$, then the second component is $x = x_2$ with probability

$$\frac{\pi_A - x_1}{x_2 - x_1}\frac{x_2}{\pi_A},$$

and $x = x_1$ with the remaining probability. If $v_k = x_1$, on the other hand, then $x = x_1$ with probability one. Thus, the second component potentially reveals the consumer's valuation for the displayed product k if the valuation is equal to x_2. The posterior probability that $v_k = x_2$ based on the two components of the signal is one if $x = x_2$, and

$$\frac{\pi_A}{x_2}\left(1 - \frac{\pi_A - x_1}{x_2 - x_1}\frac{x_2}{\pi_A}\right) \Big/ \left(1 - \frac{\pi_A}{x_2}\frac{\pi_A - x_1}{x_2 - x_1}\frac{x_2}{\pi_A}\right) = \frac{x_1}{x_2}$$

if $x = x_1$.

After signal $d = (k, x_2)$, the seller knows for sure that she can sell product k at price x_2; doing this is optimal. After signal $d = (k, x_1)$, the posterior probability that $v_k = x_2$ is x_1/x_2. For any product $k' \neq k$, by contrast, the posterior probability that $v_{k'} = x_2$ is bounded by the prior probability $f(x_2) < x_1/x_2$ because k' may not belong to the products for which the consumer's valuation is the highest across all products. Hence, it is again optimal to offer product k. Furthermore, x_1 is an optimal price because $x_1 = x_1/x_2 \cdot x_2$.

The signals $d \in \{(1, x_1), \ldots, (n, x_1)\}$ have total probability

$$1 - \frac{\pi_A}{x_2}\frac{\pi_A - x_1}{x_2 - x_1}\frac{x_2}{\pi_A} = \frac{x_2 - \pi_A}{x_2 - x_1}.$$

Consequently, the expected revenue of the seller is

$$\frac{x_2 - \pi_A}{x_2 - x_1}x_1 + \left(1 - \frac{x_2 - \pi_A}{x_2 - x_1}\right)x_2 = \pi_A,$$

and the expected payoff of the consumer is

$$\frac{x_2 - \pi_A}{x_2 - x_1}\frac{x_1}{x_2}(x_2 - x_1) = \overline{u}(\pi_A).$$

Thus, (D, ϕ) implements $\big(\pi_A, \overline{u}(\pi_A)\big)$.

Step 2: an equilibrium for $l = l'$. Let $l = l'$. In S, the consumer's expected payoff is at most

$$\max_{\pi} \overline{u}(\pi) = \overline{u}(\pi_B) = \overline{u}(x_1).$$

Suppose there exists a data policy (D', ϕ') that implements $\big(x_1, \overline{u}(x_1)\big)$. Then, there exists an equilibrium in which every data intermediary proposes (D', ϕ'), along with a fee of zero, the consumer gives his consent to exactly one intermediary, selected uniformly at random, and the seller purchases its data. To see this, note that the seller is indifferent whether to purchase the data, and the consumer cannot benefit by giving his consent to more than one intermediary. If an intermediary deviates to a data policy that does not implement $\big(x_1, \overline{u}(x_1)\big)$, it would be optimal for the consumer to give his consent to another intermediary, and at any fee strictly greater than zero the seller would not purchase the data.

We present a data policy (D', ϕ') that implements $\big(x_1, \overline{u}(x_1)\big)$. Let $D' = \{1, \ldots, n\}$. For every $(k, x) \in D$, define $\sigma(\cdot \mid (k, x)) \in \Delta D'$ as follows. With probability

$$\frac{x_1/x_2 - f(x_2)}{\pi_A/x_2 - f(x_2)},$$

$\sigma(\cdot \mid (k, x))$ draws (k, x). With the remaining probability, $\sigma(\cdot \mid (k, x))$ draws k' uniformly at random from D'. Define ϕ' by

$$\phi'(d' \mid \mathbf{v}) = \sum_{d \in D} \phi(d \mid \mathbf{v}) \sigma(d' \mid d), \quad \forall d' \in D', \forall \mathbf{v} \in X. \tag{B.2}$$

When (D', ϕ') displays signal $d' = k$, the posterior probability that $v_k = x_2$ is

$$\frac{x_1/x_2 - f(x_2)}{1 - (f(x_1))^n - f(x_2)}(1 - (f(x_1))^n) + \left(1 - \frac{x_1/x_2 - f(x_2)}{1 - (f(x_1))^n - f(x_2)}\right) f(x_2) = \frac{x_1}{x_2}.$$

For any product $k' \neq k$, by contrast, the posterior probability that $v_{k'} = x_2$ is again bounded by the prior probability $f(x_2) < x_1/x_2$ because k' may not belong to the products for which the consumer's valuation is the highest across all products. Hence, it is optimal to offer product k, and x_1 is an optimal price.

Consequently, the expected revenue of the seller is x_1. The expected payoff of the consumer is

$$\frac{x_1}{x_2}(x_2 - x_1) = \overline{u}(x_1).$$

Thus, (D', ϕ') implements $(x_1, \overline{u}(x_1))$.

Step 3: (D, ϕ) is more informative than (D', ϕ'). This holds by (B.2). $\square$

Online Appendix

OA.1 Proofs of Lemmas B1 and B2

To prove the lemmas, we first provide a formal statement of the set $\hat{S}_n$. We will use the generic notation ρ for the restriction of the seller's strategy to the problem of choosing, for each posterior belief μ, which product to offer and at what price, analogous to Section 2. We call ρ a strategy for short. Given a market segmentation τ and a strategy ρ, the expected gross payoff of the seller is $\Pi_\tau(\rho)$ and the expected payoff of the consumer is $U_\tau(\rho)$. We defined these expected payoffs in Section 2.

For convenience, we define ρ on the entire set of posterior beliefs ΔX^n, rather than just for the beliefs that have positive probability under the relevant market segmentation. A strategy ρ is optimal for the seller if

$$\forall \mu \in \Delta X^n: \quad (k^*, p^*) \in \text{supp } \rho_\mu \implies (k^*, p^*) \in \operatorname*{argmax}_{(k,p)} p \sum_{x \geq p} \mu_k(x), \tag{OA.1}$$

and it breaks ties in favor of the consumer if furthermore

$$(k', p') \in \operatorname*{argmax}_{(k,p)} p \sum_{x \geq p} \mu_k(x) \implies \sum_{x \geq p^*} \mu_{k^*}(x)(x - p^*) \geq \sum_{x \geq p'} \mu_{k'}(x)(x - p'). \tag{OA.2}$$

Thus,

$$\hat{S}_n = \Big\{ \Big(\Pi_\tau(\rho), U_\tau(\rho) \Big) \mid \ \rho \text{ satisfies (OA.1) and (OA.2)} \Big\}.$$

Proof of Lemma B1. a) We first show that $\hat{S}_n$ is convex. Let $(\pi', u') \in \hat{S}_n$ and $(\pi'', u'') \in \hat{S}_n$. Thus, there exist market segmentations $\tau', \tau'' \in \Delta\Delta X^n$, and strategies ρ', ρ'' that satisfy (OA.1) and (OA.2), such that $\Big(\Pi_{\tau'}(\rho'), U_{\tau'}(\rho')\Big) = (\pi', u')$ as well as $\Big(\Pi_{\tau''}(\rho''), U_{\tau''}(\rho'')\Big) = (\pi'', u'')$. Then for $\lambda \in (0, 1)$, the mixture $\tau = \lambda\tau' + (1 - \lambda)\tau''$ is another market segmentation, and $\Big(\Pi_\tau(\rho'), U_\tau(\rho')\Big) = \Big(\Pi_\tau(\rho''), U_\tau(\rho'')\Big) = \lambda(\pi', u') + (1 - \lambda)(\pi'', u'')$. Thus, $\hat{S}_n$ is convex.

Next, we show that $\Big\{\pi \in \mathbb{R} \mid \text{there exists } u \in \mathbb{R} \text{ s.t. } (\pi, u) \in \hat{S}_n\Big\} \subseteq [\pi_0, \pi_n^{\max}]$. Let τ be any market segmentation and ρ any optimal strategy. Then

$$\Pi_\tau(\rho) = \sum_\mu \tau(\mu) \max_{(k,p)} p \sum_{x \geq p} \mu_k(x)$$

and, letting k' be any product, we obtain

$$\Pi_\tau(\rho) \geq p_0 \sum_{x \geq p_0} \sum_\mu \tau(\mu)\mu_{k'}(x) = p_0 \sum_{x \geq p_0} f(x) = \pi_0,$$

and

$$\Pi_\tau(\rho) \leq \sum_\mu \tau(\mu) \sum_{\mathbf{v}} \mu(\mathbf{v}) \max_k v_k = \sum_{\mathbf{v}} \bar{\mu}(\mathbf{v}) \max_k v_k = \pi_n^{\max}.$$

Now, let τ be such that $\tau(\bar{\mu}) = 1$. Then, $\Pi_\tau(\rho) = \pi_0$ if ρ is optimal. Let τ' be the market segmentation that is supported on the Dirac measures of ΔX^n. That is, $\operatorname{supp} \tau' = \{\delta^{\mathbf{w}} \in \Delta X^n \mid \mathbf{w} \in X^n\}$, where $\delta^{\mathbf{w}}$ assigns probability 1 to $\mathbf{w} \in X^n$, and $\tau'(\delta^{\mathbf{w}}) = \bar{\mu}(\mathbf{w})$. Then, $\max_{(k,p)} p \sum_{x \geq p} \delta^{\mathbf{w}}_k(x) = \max_k w_k$, implying $\Pi_{\tau'}(\rho) = \pi_n^{\max}$ at an optimal ρ. Part a) now follows, as $\hat{S}_n$ is convex.

b) We start with preliminaries. Define on $\Delta X^n \times \{1, \cdots, n\} \times X$ the functions $(\mu, k, p) \mapsto p \sum_{x \geq p} \mu_k(x)$ and $(\mu, k, p) \mapsto \sum_{x \geq p} \mu_k(x)(x - p)$. For fixed (k, p), $\mu \mapsto p \sum_{x \geq p} \mu_k(x)$ and $\mu \mapsto \sum_{x \geq p} \mu_k(x)(x - p)$ are continuous. Because $\{1, \cdots, n\} \times X$ is finite, it follows that $(\mu, k, p) \mapsto p \sum_{x \geq p} \mu_k(x)$ and $(\mu, k, p) \mapsto \sum_{x \geq p} \mu_k(x)(x - p)$ are continuous. Define the value function $a : \Delta X^n \to \mathbb{R}$ by

$$a(\mu) := \max_{(k,p) \in \{1, \cdots, n\} \times X} p \sum_{x \geq p} \mu_k(x),$$

and the correspondence $\phi : \Delta X^n \rightrightarrows \{1, \cdots, n\} \times X$ of maximizers by

$$\phi(\mu) := \Big\{(k, p) \in \{1, \cdots, n\} \times X \mid p \sum_{x \geq p} \mu_k(x) = a(\mu)\Big\}.$$

By the continuity of $(\mu, k, p) \mapsto p \sum_{x \geq p} \mu_k(x)$, the Maximum Theorem (Aliprantis and Border, 2006, Thm. 17.31) implies that a is continuous and ϕ upper hemicontinuous with nonempty compact values. Moreover, define the value function $b : \Delta X^n \to \mathbb{R}$ by

$$b(\mu) := \max_{(k,p) \in \phi(\mu)} \sum_{x \geq p} \mu_k(x)(x - p).$$

As $(\mu, k, p) \mapsto \sum_{x \geq p} \mu_k(x)(x - p)$ is continuous and ϕ upper hemicontinuous with nonempty compact values, b is upper semicontinuous (see Aliprantis and Border, 2006, Lem. 17.30).

If τ is a market segmentation and ρ a strategy that satisfies (OA.1) and (OA.2), then

$$\Pi_\tau(\rho) = \sum_\mu \tau(\mu)a(\mu) \quad \text{and} \quad U_\tau(\rho) = \sum_\mu \tau(\mu)b(\mu).$$

Fix some $\pi \in [\pi_0, \pi_n^{\max}]$ for the rest of the proof. The problem of finding a greatest element in $\{u \in \mathbb{R} \mid (\pi, u) \in \hat{S}_n\}$ can be stated as maximzing $\sum_\mu \tau(\mu)b(\mu)$ over all market segmentations τ such that $\sum_\mu \tau(\mu)a(\mu) = \pi$.

We momentarily enlarge the choice set of this problem so as to obtain a compact set. Let $\tilde{\Delta}\Delta X^n$ be the set of all Borel probability measures ζ on ΔX^n.[20] Let $Z \subset \tilde{\Delta}\Delta X^n$ be the subset of probability measures ζ that average to the prior belief $\bar{\mu}$,

$$\int \mu(\mathbf{w}) \mathrm{d}\zeta(\mu) = \bar{\mu}(\mathbf{w}), \quad \forall \mathbf{w} \in X^n. \tag{OA.3}$$

We endow $\tilde{\Delta}\Delta X^n$ with the weak* topology. Because ΔX^n is compact and metrizable, the space $\tilde{\Delta}\Delta X^n$ is compact (see Aliprantis and Border, 2006, Thm. 15.11). Being a closed subset, it follows that Z is compact. By the continuity of a, $\zeta \mapsto \int a(\mu)\mathrm{d}\zeta(\mu)$ is continuous. Hence, $\{\zeta \in Z \mid \int a(\mu)\mathrm{d}\zeta(\mu) = \pi\}$ is compact. Furthermore, by the upper semicontinuity of b, $\zeta \mapsto \int b(\mu)\mathrm{d}\zeta(\mu)$ is upper semicontinuous (see Aliprantis and Border, 2006, Thm. 15.5). It follows that there exists a maximizer ζ^* for the problem

$$\max_{\zeta \in Z} \int b(\mu)\mathrm{d}\zeta(\mu) \quad \text{s.t.} \quad \int a(\mu)\mathrm{d}\zeta(\mu) = \pi.$$

It remains to show that there exists a market segmentation $\tau \in \Delta\Delta X^n$ such that

$$\sum_\mu \tau(\mu)b(\mu) = \int b(\mu)\mathrm{d}\zeta^*(\mu) \quad \text{and} \quad \sum_\mu \tau(\mu)a(\mu) = \pi. \tag{OA.4}$$

The tuple $\left(\bar{\mu}, \pi, \int b(\mu)\mathrm{d}\zeta^*(\mu)\right)$ lies in the convex hull of

$$\left\{(\mu, r_1, r_2) \in \Delta X^n \times \mathbb{R}^2 \mid (r_1, r_2) = \left(a(\mu), b(\mu)\right)\right\}.$$

Because the dimension of this set is finite, Caratheodory's Theorem allows us to express $\left(\bar{\mu}, \pi, \int b(\mu)\mathrm{d}\zeta^*(\mu)\right)$ as a convex combination of finitely many elements. Denote a generic such element by (μ^y, r_1^y, r_2^y), and let $z^y > 0$ be the corresponding weight. Then, $\tau^* \in \Delta\Delta X^n$ with $\tau^*(\mu^y) = z^y$ is a market segmentation at which (OA.4) holds.

[20] Thus, in contrast to $\tau \in \Delta\Delta X^n$, the support of $\zeta \in \tilde{\Delta}\Delta X^n$ need not be finite.

c) The concavity of $\overline{u}_n$ follows from the convexity of $\hat{S}_n$, which we showed in the proof of part a). Being concave, $\overline{u}_n$ is continuous at every $\pi \in (\pi_0, \pi_n^{\max})$, and $\lim_{\pi \to \pi_0} \overline{u}_n(\pi) \geq \overline{u}_n(\pi_0)$ and $\lim_{\pi \to \pi_n^{\max}} \overline{u}_n(\pi) \geq 0$. It only remains to show that these weak inequalities hold with equality.

By contradiction, suppose $\lim_{\pi \to \pi_0} \overline{u}_n(\pi) > \overline{u}_n(\pi_0)$. We use again the notation from the proof of part b). Let $(\pi^s)_{s\in\mathbb{N}}$ be a sequence with $\pi^s > \pi_0$ for all s and $\lim_{s\to\infty} \pi^s = \pi_0$. Let ρ be a strategy that satisfies (OA.1) and (OA.2), and let $(\tau^s)_{s\in\mathbb{N}}$ be a sequence of market segmentations such that $\Pi_{\tau^s}(\rho) = \pi^s$ and $U_{\tau^s}(\rho) = \overline{u}_n(\pi^s)$ for all s. Then, $\tau^s \in Z$ for all s. As Z is compact and metrizable (see Aliprantis and Border, 2006, Thm. 15.11), there exists a subsequence $\left(\tau^{s(t)}\right)_{t\in\mathbb{N}}$ that converges to some $\zeta' \in Z$. By the continuity of $\zeta \mapsto \int a(\mu)\mathrm{d}\zeta(\mu)$ and the upper semicontinuity of $\zeta \mapsto \int b(\mu)\mathrm{d}\zeta(\mu)$,

$$\pi_0 = \lim_{t\to\infty} \Pi_{\tau^{s(t)}}(\rho) = \lim_{t\to\infty} \sum_{\mu} \tau^{s(t)}(\mu)a(\mu) = \int a(\mu)\mathrm{d}\zeta'(\mu),$$

$$\limsup_{t\to\infty} U_{\tau^{s(t)}}(\rho) = \limsup_{t\to\infty} \sum_{\mu} \tau^{s(t)}(\mu)b(\mu) \leq \int b(\mu)\mathrm{d}\zeta'(\mu) \leq \overline{u}_n(\pi_0).$$

As in the proof of part b), there exists a market segmentation τ such that $\sum_\mu \tau(\mu)a(\mu) = \pi_0$ and $\sum_\mu \tau(\mu)b(\mu) = \int b(\mu)\mathrm{d}\zeta'(\mu)$. This yields a contradiction to $\lim_{\pi\to\pi_0} \overline{u}_n(\pi) > \overline{u}_n(\pi_0)$. Hence, $\lim_{\pi\to\pi_0} \overline{u}_n(\pi) = \overline{u}_n(\pi_0)$.

By contradiction, suppose $\lim_{\pi\to\pi_n^{\max}} \overline{u}_n(\pi) = \eta > 0$. Then, there exist $\epsilon, \delta > 0$ and π such that $\pi_n^{\max} - \pi < \delta$, $|\eta - \overline{u}_n(\pi)| < \epsilon$, and $\epsilon + \delta < \eta$. Let ρ be a strategy that satisfies (OA.1) and (OA.2), and let τ be a market segmentation such that $\Pi_\tau(\rho) = \pi$ and $U_\tau(\rho) = \overline{u}_n(\pi)$. Then, $\Pi_\tau(\rho) + U_\tau(\rho) > \pi_n^{\max} - \delta + \eta - \epsilon > \pi_n^{\max}$. But

$$\begin{aligned}
\Pi_\tau(\rho) + U_\tau(\rho) &= \sum_\mu \tau(\mu) \sum_{k,p} \rho_\mu(k,p) \sum_{x\geq p} \mu_k(x)x \\
&\leq \sum_\mu \tau(\mu) \max_k \sum_x \mu_k(x)x \\
&\leq \sum_\mu \tau(\mu) \sum_{\mathbf{w}} \mu(\mathbf{w}) \max_k w_k = \sum_{\mathbf{w}} \bar{\mu}(\mathbf{w}) \max_k w_k = \pi_n^{\max};
\end{aligned}$$

a contradiction to $\Pi_\tau(\rho) + U_\tau(\rho) > \pi_n^{\max}$. Hence, $\lim_{\pi\to\pi_n^{\max}} \overline{u}_n(\pi) = 0$.

It remains to show $\overline{u}_n(\pi_n^{\max}) = 0$. By contradiction, suppose $\overline{u}_n(\pi_n^{\max}) > 0$. Let ρ be a strategy that satisfies (OA.1) and (OA.2), and let τ be a market segmentation such that $\Pi_\tau(\rho) = \pi_n^{\max}$ and $U_\tau(\rho) = \overline{u}_n(\pi)$. Then, $\Pi_\tau(\rho) + U_\tau(\rho) > \pi_n^{\max}$, which is impossible as shown above. $\square$

Proof of Lemma B2. The first part of the lemma holds because

$$\hat{S}_n \subseteq \Big\{ \Big(\Pi_\tau(\rho), U_\tau(\rho) \Big) \mid \ \rho \text{ satisfies (OA.1)} \Big\} = S_n \subseteq S.$$

Next, let $\Big(\pi, \overline{u}(\pi)\Big) \in S$, where we may assume $\pi < x_m$. By Theorem 1, there exists a sequence $\Big((\pi_n, u_n)\Big)_{n \in \mathbb{N}}$ such that $\Big(\pi_n, u_n\Big) \in S_n$ and $\lim_{n\to\infty}(\pi_n, u_n) = \Big(\pi, \overline{u}(\pi)\Big)$. Because $\lim_{n\to\infty} \pi_n^{\max} = x_m$, we have $\pi_n \in [\pi_0, \pi_n^{\max}]$ for n sufficiently large, say $n > n'$. Consider the sequence $\Big((\tilde{\pi}_n, \overline{u}_n(\tilde{\pi}_n))\Big)_{n \in \mathbb{N}}$, where $\tilde{\pi}_n = \pi_0$ for $n \leq n'$ and $\tilde{\pi}_n = \pi_n$ for $n > n'$. By construction, $\Big(\tilde{\pi}_n, \overline{u}_n(\tilde{\pi}_n)\Big) \in \hat{S}_n$. Furthermore, $\lim_{n\to\infty} \tilde{\pi}_n = \pi$. It remains to show that $\lim_{n\to\infty} \overline{u}_n(\tilde{\pi}_n) = \overline{u}(\pi)$. Because $\hat{S}_n$ differs from S_n only by the additional condition (OA.2), according to which the seller breaks ties in favor of the consumer, it holds that $\overline{u}_n(\tilde{\pi}_n) = \max\{u \in \mathbb{R} \mid (\tilde{\pi}_n, u) \in S_n\} \geq u_n$ for $n > n'$. Hence, $\liminf_{n\to\infty} \overline{u}_n(\tilde{\pi}_n) \geq \lim_{n\to\infty} u_n$. On the other hand, $\hat{S}_n \subseteq S$ and the continuity of the function u imply $\limsup_{n\to\infty} \overline{u}_n(\tilde{\pi}_n) \leq \lim_{n\to\infty} \overline{u}(\tilde{\pi}_n) = \overline{u}(\pi) = \lim_{n\to\infty} u_n$. Thus, $\lim_{n\to\infty} \overline{u}_n(\tilde{\pi}_n) = \overline{u}(\pi)$. $\square$

References

ACEMOGLU, D., A. MAKHDOUMI, A. MALEKIAN, AND A. OZDAGLAR (2022): "Too much data: Prices and inefficiencies in data markets," *American Economic Journal: Microeconomics*, 14(4), 218–56.

AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): "Monopoly price discrimination and demand curvature," *American Economic Review*, 100(4), 1601–15.

ALIPRANTIS, C. D., AND K. C. BORDER (2006): *Infinite Dimensional Analysis*. Springer, Berlin, third edn.

ANDERSON, C. (2006): *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, New York.

BERGEMANN, D., AND A. BONATTI (2015): "Selling cookies," *American Economic Journal: Microeconomics*, 7(3), 259–94.

——— (2019): "Markets for Information: An Introduction," *Annual Review of Economics*, 11, 85–107.

BERGEMANN, D., A. BONATTI, AND T. GAN (2022): "The economics of social data," *The RAND Journal of Economics*.

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The Limits of Price Discrimination," *American Economic Review*, 105(3), 921–957.

BERGEMANN, D., AND K. H. SCHLAG (2008): "Pricing without priors," *Journal of the European Economic Association*, 6(2-3), 560–569.

BLACKWELL, D. (1953): "Equivalent Comparison of Experiments," *Annals of Mathematical Statistics*, 24(2), 265–272.

BOUNIE, D., A. DUBUS, AND P. WAELBROECK (2022): "Competition and Mergers with Strategic Data Intermediaries," *working paper*.

Brynjolfsson, E., Y. J. Hu, and M. D. Smith (2003): "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49(11), 1580–96.

Cavallo, A. (2017): "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers," *American Economic Review*, 107(1), 283–303.

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.

Condorelli, D., and B. Szentes (2020): "Information Design in the Hold-Up Problem," *Journal of Political Economy*, 128(2), 681–709.

De Corniere, A., and R. De Nijs (2016): "Online advertising and privacy," *The RAND Journal of Economics*, 47(1), 48–72.

DellaVigna, S., and M. Gentzkow (2019): "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 134(4), 2011–2084.

Federal Trade Commission (2014): "*Data brokers: a call for transparency and accountability*," Washington, DC.

Haghpanah, N., and R. Siegel (2021): "The limits of multi-product price discrimination," *American Economic Review: Insights*.

——— (2022): "Pareto improving segmentation of multi-product markets," Discussion paper, Working paper.

Hidir, S., and N. Vellodi (2021): "Privacy, personalization, and price discrimination," *Journal of the European Economic Association*, 19(2), 1342–1363.

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

——— (2021): "Competing data intermediaries," *The RAND Journal of Economics*, 52(3), 515–537.

Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101(6), 2590–2615.

Narayanan, A. (2013): "Online price discrimination: conspicuous by its absence," *URL https://33bits. wordpress. com/2013/01/08/online-price-discrimination-conspicuous-by-its-absence*.

Neeman, Z. (2003): "The effectiveness of English auctions," *Games and economic Behavior*, 43(2), 214–238.

Pram, K. (2021): "Disclosure, welfare and adverse selection," *Journal of Economic Theory*, 197, 105327.

Quan, T. W., and K. R. Williams (2018): "Product variety, across-market demand heterogeneity, and the value of online retail," *The RAND Journal of Economics*, 49(4), 877–913.

Shaked, M., and J. G. Shanthikumar (2007): *Stochastic Orders*, Springer Series in Statistics. Springer New York.