Manipulative consumers

MICHAEL RICHTER, Baruch College, CUNY, USA and Royal Holloway, University of London, UK NIKITA ROKETSKIY, University College London, UK

We study optimal monopoly pricing with evasive consumers. The monopolist uses consumer data to estimate demand and menu pricing to optimally screen the residual uncertainty about consumers' preferences. Third degree price discrimination encourages data-conscious consumers to manipulate their observable attributes (at a cost). This reduces the precision of demand estimation, sometimes rendering the consumer data useless. We derive the monopolist's gains from using data and characterize the optimal investigation strategy. Large number of observable consumer attributes results in small overall value of data. Randomly restricting monopolist's access to consumer data increases profit.

1 INTRODUCTION

Sellers' use of consumer data for price discrimination is as old as the hills, but, recently, the amount and variety of available data has rapidly multiplied. This would spell trouble for consumers, if not for the fact that many of them are aware of these practices and have some control over their own data. Moreover, with a few exceptions, consumers are not liable for falsifying or manipulating the data that is harvested by sellers.

Thus, one of the key privacy-related questions is this: How are the spoils of trade divided between data-hungry sellers and data-conscious consumers? We answer this question using a canonical monopoly framework in which the seller has access to a full price-discrimination toolkit.

In our model, the seller's optimal pricing is a combination of the two instruments that are commonly used in practice: consumer profiling and menus. The seller employs consumer data to organize consumers into specific categories that share common characteristics. We refer to these categories as market segments. Because consumer data does not always perfectly explain variation in preferences, the seller screens the remaining intrasegment heterogeneity using a variety of vertically differentiated goods.

When the monopolist employs consumer data to segment the market, he solves a problem akin to a regression of market demand on consumer attributes. The explained part of variation in this regression represents the value of data for the monopolist. As such, the monopolist is interested in maximizing the part of variation that is explained by the consumer data. However, the consumers' interests are opposed. When some parts of their data become overly informative of the demand, the consumers muddle them. This limits the sellers' gains from segment-specific pricing. It is costly for the consumers to manipulate their data, hence some explanatory power always remains.

Because the data is endogenous to pricing, the derived value of the data is not a function of prior beliefs, but rather a function of the structure of the data and the cost of manipulation. We show that the seller gains from using the data, but this gain is vanishing when the data becomes very rich—i.e., when it contains a large number of (conditionally) independent variables.¹ The data-driven gain in profit comes at a cost of disproportionately large reduction

¹In our context, we abstract away from the problems associated with finite samples. We assume that seller can sample data from the continuum of consumers and the term "rich" refers to the number of variables, not observations.

in consumer surplus. In contrast with the seller's value of data, the consumer surplus loss may remain substantial when data becomes rich.

The dimensionality of consumer data takes a center stage in this result. Consumers manipulate their data to get a better deal from the seller. The more high-dimensional the data is, the more combinations of attribute values will correspond to desirable deals. Thus, the consumers can rely less on the number of manipulated attributes and more on selecting which attributes to manipulate—i.e., less on the costly and more on the costless aspect of data manipulation.

The hide and seek nature of monopolist-consumer relations means that the seller would benefit from limiting consumers' understanding of how data affects prices. One way to achieve this is to use a random, and therefore unpredictable, subset of variables that are contained in the original data. We show that such a tactic would resolve the issue of vanishing value when the data is rich, but implementing it would require commitment power that sellers may not possess.

2 RELATED LITERATURE

The general problem of inference from muddled data is studied by Frankel and Kartik (2019, 2022) (also, see earlier related work by Kartik, 2009). Ball (2022) studies scoring of strategic agents in which the data for scoring is provided by an intermediary. The latter may serve as a valuable commitment provider to the agency that develops the scoring rule. Ball (2022) shows that randomization by the data intermediary inhibits data manipulation because the target of manipulation becomes more difficult to identify. A distantly related result in the the context of information elicitation with verification is obtained by Carroll and Egorov (2019).

Milli, Miller, Dragan and Hardt (2019) study welfare costs associated with threshold binary classifiers when subjects can manipulate classifiers inputs strategically at a cost. They show that the pursuit to counter strategic manipulations via redesigning the classifier brings a disproportionately high cost on the subjects. A similar problem is studied by Cunningham and Moreno De Barreda (2022), Perez-Richet and Skreta (2022) and Hardt, Megiddo, Papadimitriou and Wootters (2016).

Incentives to misreport data are not unique to classification problems. Eliaz and Spiegler (2022) study incentives under regression estimators and Caner and Eliaz (2021) investigate the same question with an additional issue of variable selection and regularization.

Deneckere and Severinov (2022) and Severinov and Tam (2018) study costly misreporting in classic asymmetric information frameworks of signaling and screening. Liang and Madsen (forthcoming) study the use of observables in provision of effort when subjects productivity is private but correlated with these observables. Dana, Larsen and Moshary (2023), Tan (2023) and Perez-Richet and Skreta (2023) adopt a mechanism design approach in which the agents must report their type (e.g., the consumers must report their individual demand) and they incur a cost when their reports are not truthful. Our interpretation of manipulable consumer data is more literal than theirs. We assume that the seller has to use statistical tools to infer the true demand from the data.

Hu, Immorlica and Vaughan (2019) study strategic manipulation of data by multiple agents. They highlight an externality one subject imposes on others by manipulating their own record. In our framework, similar externality considerations are present.

The use of data by sellers is tightly related to consumer privacy (see Acquisti, Taylor and Wagman (2016) for review of the literature on consumer privacy). Bonatti and Cisternas (2020) investigate how the seller can condition current prices on the consumers' past choices via

an aggregator that assigns a score to each consumer. Bhaskar and Roketskiy (2021) consider unrestricted use of past purchases for equilibrium optimal pricing. Conitzer, Taylor and Wagman (2012) study how consumers' control over their past records affects their welfare.

Bonatti, Huang and Villas-Boas (2023) draw a connection between the value of privacy and the concavity of the indirect utility as a function of market beliefs. Their approach is complementary to ours. In their papers, privacy results in pooling of consumption for various types of consumers. In our case, the data is valuable to the seller if it allows him to differentiate his offering. A natural way to measure this value is via the dispersion: an expectation of the quadratic (convex) function of the price.

Acknowledging that some degree of privacy is desirable, Eilat, Eliaz and Mu (2021) suggest a Bayesian measure of privacy protection and use this measure to derive optimal privacypreserving pricing. They show that an optimal privacy-preserving menu always contains a finite number of alternatives.

We study how the seller uses data to segment the market and tailor prices to the demand in each segment. Our definition of market segment is closely related to the one in Yang (2022). In our setting, the consumer data is available to the seller at a negligible cost. When sellers use consumer data for both pricing and product design, Ichihashi (2020) and Hidir and Vellodi (2021) show that it is possible to find a segmentation that relies on consumer, volunteering the private information on their valuations. Ali, Lewis and Vasserman (2022) study how consumers' control of personal data affects consumer surplus, industry profits and overall welfare (also see Ali, Lewis and Vasserman, 2023).

Elliott, Galeotti, Koh and Li (2021) take an information design approach: they study what kind of data an intermediary would need to share with sellers to induce a desired outcome in equilibrium.

We study the value that the seller attaches to consumer data. A related question is how to sell the data to the monopolist and what is the resulting price. Taylor (2004), Bergemann and Bonatti (2015), Bergemann, Bonatti and Smolin (2018) and Segura-Rodriguez (2021) answer this question in a variety of settings.

3 THE MODEL

The market consists of a single seller and a mass of consumers indexed by $i \in C = [0, 1]$. A variety of goods can be produced and sold by the seller. Each type of good is characterized by a quality parameter $q \in \mathbb{R}_+$. There are two types of consumers who vary by their taste for quality: those with low marginal willingness to pay for quality, $t_{\ell} \in \mathbb{R}_+$, and those with a high willingness to pay, $t_h \in \mathbb{R}_+$. The difference between these is denoted by $\delta = t_h - t_{\ell} > 0$ and a consumer *i*'s marginal willingness to pay for quality is denoted by $\tau(i)$.

Each consumer demands at most one good. If consumer i buys a good of quality q, the following social surplus is realized:

$$s(i,q) = 2\tau(i)q - q^2$$

A transaction price *p* determines how the surplus is split between the seller and the buyer:

$$u(i,q,p) = s(i,q) - p$$
$$\pi(i,q,p) = p,$$

where *u* is the buyer's payoff and π is the seller's profit from this transaction. The consumers' outside option is valued at zero.

The seller can use consumer data to price the goods. We assume that the data is freely available, but the consumers can, at a cost, secretly manipulate their own records before the seller harvests the data.² To model this we assume the following. Each consumer *i* is endowed with *K* ex ante private attributes represented by the vector

$$\omega(i) \in \mathscr{A} = \{0, 1\}^K.$$

The vector of *ex post* public attributes—i.e., the attributes that can be observed by the seller—is *chosen* by consumer *i* and is denoted by

$$\alpha(i) \in \mathscr{A} = \{0, 1\}^K.$$

The cost of choosing this vector of attributes is

$$\frac{\|\alpha(i) - \omega(i)\|}{K}c,$$

where c > 0. This cost is linear in the *share* of attributes the consumer manipulates.

Consumer data generically contains information about demand. To quantify this, consider a collection of attribute vectors $S \subseteq \mathcal{A}$ and let λ be a Lebesgue measure on *C*. By $m(S) = \lambda(\{i \in C : \tau(i) = t_{\ell}, \alpha(i) \in S\})$ we denote the mass of consumers with quality valuation t_{ℓ} and a public vector of attributes from *S*. Similarly, we denote the mass of consumers with quality valuation t_h and a public vector of attributes from *S* by $n(A) = \lambda(\{i \in C : \tau(i) = t_h, \alpha(i) \in S\})$.³

All the seller needs to know about the demand to set prices optimally is the proportion of consumers with different willingness to pay for quality. This ratio conditional on attributes from *S* is

$$h(S) = \frac{n(S)}{m(S)}.$$

By \overline{m} and \overline{n} we denote the total masses of consumers with quality valuations t_{ℓ} and t_h respectively. The aggregate hazard ratio for the entire market is $\overline{h} = \frac{\overline{m}}{\overline{n}}$.

Finally, we assume that the consumers are anonymous and can be distinguished by the seller only through their ex post public attributes α . Thus, the offer made to a consumer *i* depends only on $\alpha(i)$ and cannot depend on *i* explicitly.

3.1 Optimal menu pricing

Suppose the seller offers the same profit-maximizing screening menu to a group of consumers with attribute vectors in $S \subseteq \mathcal{A}$. Let $\pi(i)$ be the profit that results from consumer *i* purchasing her favorite item from this menu. By $\rho(S)$ we denote the total profit over all consumers $i \in S$, normalized by the number of the consumers with low valuation:

$$\rho(S) = \frac{1}{m(S)} \int_{\alpha(i) \in S} \pi(i) di$$

As shown by Mussa and Rosen (1978), the profit-maximizing menu consists of two items. A premium item (p_h, q_h) which is designed for consumers with high valuation for quality,

 $^{^{2}}$ The consumers cannot deny the seller access to their data, but they can tamper with their records in an attempt to mislead the seller.

³We assume that functions τ , α and ω are measurable.

and a basic item (p_l, q_l) which is for everyone else. The seller's profit solves the following program:

$$\rho(S) = \max_{q_h \ge q_\ell \ge 0} \left\{ 2t_\ell q_\ell - q_\ell^2 - 2h(S)q_\ell \delta + h(S) \left(2t_h q_h - q_h^2 \right) \right\}.$$

The solution to the seller's program is

$$p_{\ell} = t_{\ell} \max\{0, t_{\ell} - h(S)\delta\}$$
$$p_{h} = t_{h}^{2} - \delta \max\{0, t_{\ell} - h(S)\delta\}$$
$$q_{h} = t_{h}$$
$$q_{\ell} = \max\{0, t_{\ell} - h(S)\delta\}.$$

When shopping from this menu, consumers with the low willingness to pay receive no surplus. A consumer $i \in S$ with the high willingness to pay receives a surplus

$$u(i) = \max\{0, 2\delta(t_{\ell} - h(S)\delta)\}.$$

The monopolist's profit is

$$\rho(S) = h(S)(t_{\ell} + \delta)^{2} + (\max\{0, t_{\ell} - h(S)\delta\})^{2}$$

Because we study the value of consumer data for the seller, we focus on markets that are characterized by a high degree of consumer heterogeneity. Formally, this is reflected in the following assumption:

Assumption 1.

$$\overline{h} \in \left(\frac{c}{\delta^2}, \frac{t_\ell}{\delta} - \frac{c}{\delta^2}\right).$$

This assumption ensures that the seller's profit-maximizing menus deliver variety in terms of quality and prices.

3.2 Group pricing with data

Eliciting consumers' valuations via a screening menu is costly. The seller has to reduce the price difference between the premium and the basic items and distort the quality of the basic item. In this sense, third-degree price discrimination is a more cost effective instrument than menu pricing. The seller could use consumer data to estimate consumers' valuations and, thus, rely on menu pricing only to elicit residual uncertainty that is not explained by the available data.

To get a better understanding of the third-degree price discrimination component in the seller's pricing decision, let us consider a notion of market segmentation. Let $\mathscr{S} = \{S_1, S_2, ..., S_N\}$ be a partition of set \mathscr{A} . Every element of this partition corresponds to a market segment that can be identified by the seller using the consumer data.

The sellers' total profit under this market segmentation \mathcal{S} is

$$\pi_{\mathcal{S}} = \sum_{s \in \mathcal{S}} m(s)\rho(s) = \pi^* + \delta^2 \sum_{s \in \mathcal{S}} m(s) \left(h(s) - \overline{h}\right)^2,$$

where $\pi^* = \overline{m}\rho(\mathscr{A})$ is the profit of the monopolist without any market segmentation, or equivalently, with the coarsest segmentation possible (i.e., when every consumer gets the same menu). The second term of this expression represents the profit gain the seller could get by segmenting the market. Perhaps not surprisingly, this term is proportional to the

weighted variance of the hazard ratio across different market segments. The whole purpose of segmenting the market is that the seller is able to make *different* offers to consumers in different segments. The variance—a measure of spread—quantifies how different these offers are across the segments. This suggests that the seller would prefer the finest market segmentation he could achieve given the available consumer data.

3.3 Data manipulation

By the nature of third-degree price discrimination, the consumers expect differences in prices for the premium quality good across market segments. These differences motivate consumers to perform arbitrage—they manipulate their attributes to "travel" to a segment with lower prices. Note that consumers with low willingness to pay receive their reservation utility regardless of the market segment they find themselves in. Therefore, they have no incentives to manipulate their attributes.

Focus on high valuation consumers and recall that the price for the high quality product in segment *S* is increasing in hazard ratio in this segment:

$$p_h(S) = h(S)\delta^2 + t_h^2 - t_\ell\delta$$

Consider two segments, S_1 and S_2 such that $h(S_1) \ge h(S_2)$. A consumer's gain from "traveling" from segment S_1 to S_2 is proportional to the difference between the hazard ratios in these segments $h(S_1) - h(S_2)$. The cost on the other hand is proportional to the share of attributes that need to be changed to "travel" between the segments. Thus, in equilibrium, the following *no-arbitrage condition* must hold:

$$h(S_1) - h(S_2) \le \frac{c}{\delta^2 K} \min_{\substack{\mathbf{b} \in S_2, \\ \mathbf{a} \in S_1: m(\mathbf{a}) > 0}} \|\mathbf{a} - \mathbf{b}\|.$$
(1)

To see that, recall that the hazard ratio is endogenous in the setup with manipulable data. If the condition does not hold, there is a positive mass of high valuation consumers with attribute vector $a \in S_1$ who could set it to $b \in S_2$ instead at a cost that is lower than the gain from switching market segments. When they do so, they reduce the difference in the hazard ratios across the two segments.

4 VALUE OF CONSUMER DATA

By segmenting the market according to S, the seller increases his profit by

$$\delta^2 \sum_{s \in \mathcal{J}} m(s) \left(h(s) - \overline{h} \right)^2$$

We use this gain to define the value of information for the seller. For analytical purposes it is convenient to represent it in terms of hazard ratios within market segments. However, one could find it more intuitive to represent this gain in terms of dispersion of prices across market segments:

$$\delta^2 \sum_{s \in \mathcal{S}} m(s) \left(h(s) - \overline{h} \right)^2 = \frac{1}{\delta^2} \sum_{s \in \mathcal{S}} m(s) \left(p_h(s) - \frac{1}{\overline{m}} \sum_{\tilde{s} \in \mathcal{S}} m(\tilde{s}) p_h(\tilde{s}) \right)^2,$$

where $p_h(s)$ is the price of the premium item in market segment *s* and m(s) is the volume of basic items sold in the same segment.

A market segmentation is the sellers choice. Intuitively, the seller should prefer to use *all* available data in the most efficient way possible. To formalize this idea, consider the finest segmentation possible—i.e., the segmentation that associates each possible value of attributes vector with a separate market segment:

$$\mathcal{S}^* = \{\{i \in C \mid \alpha(i) = \mathbf{a}\} \mid \mathbf{a} \in \mathscr{A}\}.$$

PROPOSITION 2. For any market segmentation S,⁴

$$\sum_{s \in \mathcal{S}} m(s) \left(h(s) - \overline{h} \right)^2 \le \sum_{s \in \mathcal{S}^*} m(s) \left(h(s) - \overline{h} \right)^2.$$

PROOF. The gain in profit is convex in hazard ratio and $\{h(s), m(s)\}_{s \in S^*}$ weakly majorises $\{h(s), m(s)\}_{s \in S}$.

This proposition shows that the seller would always use segmentation S^* unless he can credibly promise to the consumers to restrict the use of data for price discrimination. The latter would require some commitment mechanism such as outsourcing data collection to an independent intermediary. We consider this possibility in Section 4.2.

Note that under segmentation S^* , if a consumer manipulates any of her attributes, she necessarily changes her market segment as well. Thus, the set of constraints (4) can be simplified:

$$\forall \mathbf{a}, \mathbf{b} \in \mathscr{A} : |h(\mathbf{a}) - h(\mathbf{b})| \le \frac{c}{\delta^2} \frac{\|\mathbf{a} - \mathbf{b}\|}{K}$$

The seller's gain from using data depends on the correlation between the consumer attributes and preferences, which is represented by variance of hazard ratios across different segments of the market. Thus, the gain can be small or even zero if the attributes are independent of the consumer valuations.

On the other hand, no matter how informative the initial attributes are, the gain from using them for price discrimination is bounded by consumers' data manipulation. The higher the informativeness of the attributes, the higher the incentives of the consumers to manipulate them. The very question that we study in this paper is how this tug of war between the seller and the manipulative consumers limits the value of consumer data for the seller.

In pursuit of the maximal potential of the consumer data for price discrimination, we define its value as the maximal gain the seller can obtain when consumers can manipulate their attributes. That is, we consider the value of data under the most favorable conditions for the seller.

DEFINITION 3. The value of consumer data D_K is the largest gain the seller can obtain by segmenting the market according to S^* across all possible correlations between the attributes

⁴One could extend the definition of market segmentation by allowing the seller to assign consumers to segments randomly. The proposition would still hold. We omit this generalization to make the exposition simple.



Fig. 1. An example of a constraints graph.

and valuations:

$$D_{K} = \delta^{2} \max_{h:\mathscr{A} \to \mathbb{R}_{+}} \sum_{\mathbf{a} \in \mathscr{A}} m(\mathbf{a}) \left(h(\mathbf{a}) - \overline{h} \right)^{2}$$
(2)

s.t.
$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a}) \left(h(\mathbf{a}) - \overline{h} \right) = 0,$$
 (3)

$$\forall \mathbf{a}, \mathbf{b} \in \mathscr{A} : |h(\mathbf{a}) - h(\mathbf{b})| \le \frac{c}{\delta^2} \frac{\|\mathbf{a} - \mathbf{b}\|}{K}.$$
(4)

The program that defines D_K highlights the main intuition about the interaction of the data-hungry seller and the data-cautious consumers. On the one hand, it is in the seller's interest to increase the share of preference variation explained by the observed attributes (see objective (2)). On the other hand, if explanatory power of the attributes increases beyond a certain point, the consumers erode it via attribute manipulation (see constraint (4)).

To characterize the program's solution h^* , we need to identify pairs of market segments which are involved in consumer arbitrage. In mathematical terms, we look for program constraints that bind at $h^{*,5}$ It is convenient to think about the binding constraints as edges of a graph. In particular, for a given h, we define a constraints graph G(h) in the following way.

DEFINITION 4. A graph G(h) with the set of nodes \mathcal{A} is called a constraints graph for h if for every pair of $\mathbf{a}, \mathbf{b} \in \mathcal{A}$ the following statements are equivalent:

(1) **a** and **b** are connected, (2) $|h(\mathbf{a}) - h(\mathbf{b})| = \frac{c}{\delta^2} \frac{\|\mathbf{a} - \mathbf{b}\|}{K}$.

In this graph, every node is a market segment and edges represent arbitrage opportunities, or equivalently, instances of consumers manipulating their attributes. The graph representation of binding constraints allows us to find a simple necessary condition for optimality (2):

PROPOSITION 5. For every solution h^* of the problem (2) the corresponding constraints graph $G(h^*)$ is connected.

 $^{^{5}}$ Note that the objective function is convex, therefore the solution will be on the boundary of the convex permissible set.

PROOF. By contradiction, let *h* be a solution to the problem (2) for which G(h) is not connected. We can partition the nodes of this graph into two sets *A* and *B* in such a way that there are no edges between the two sets, and $h(A) \ge h(B)$.

The objective can be rewritten as

$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a}) \left(h(\mathbf{a}) - \overline{h}\right)^2 = \sum_{\mathbf{a}\in A} m(\mathbf{a}) \left(h(\mathbf{a}) - h(A)\right)^2 + \sum_{\mathbf{a}\in B} m(\mathbf{a}) \left(h(\mathbf{a}) - h(B)\right)^2 + \sum_{Z\in\{A,B\}} m(Z) \left(h(Z) - \overline{h}\right)^2 + \sum_{Z\in\{A,B\}} m(Z) \left(h(Z) -$$

Because there are no edges in G(h) between A and B, for any $\mathbf{a} \in A$ and $\mathbf{b} \in B$ it holds that $|h(\mathbf{a}) - h(\mathbf{b})| < \frac{c}{\delta^2} \frac{\|\mathbf{a} - \mathbf{b}\|}{K}$. Let

$$h_{\epsilon}(\mathbf{a}) = \begin{cases} h(\mathbf{a}) + \frac{\epsilon}{m(A)}, \text{ if } \mathbf{a} \in A\\ h(\mathbf{a}) - \frac{\epsilon}{m(B)}, \text{ if } \mathbf{a} \in B. \end{cases}$$

Note that for small enough $\epsilon > 0$, h_{ϵ} satisfies all the constraints of the problem (2) and increases the value of the objective compared to h:

$$\sum_{\mathbf{a}\in\mathcal{A}}m(\mathbf{a})\left(h_{\epsilon}(\mathbf{a})\right)^{2}-\sum_{\mathbf{a}\in\mathcal{A}}m(\mathbf{a})\left(h(\mathbf{a})\right)^{2}=\epsilon\left(2(h(A)-h(B))+\epsilon\left(m(A)^{-1}+m(B)^{-1}\right)\right)>0,$$

hence h cannot be a solution to (2).

We can restate this result using the hazard ratios in different market segments:

COROLLARY 6. If h^* is a solution to the program (2), then the difference in hazard ratios between any two segments is a multiple of $c/(\delta^2 K)$:

$$\forall \mathbf{a}, \mathbf{b} \in \mathscr{A}, \exists j \in \{0, 1, 2, ..., K\} : |h^*(\mathbf{a}) - h^*(\mathbf{b})| = \frac{c}{\delta^2} \frac{j}{K}$$

This corollary sheds light on why dimensionality of the data is important: even though there are 2^{K} possible realizations of attribute vectors, when data is most informative there are at most *K* truly distinct market segments. Put differently, many attribute vectors offer identical predictions of the within segment demand.

4.1 More data?

We are now prepared to investigate how the value of data depends on its richness. The parameter that describes the richness of the data is the number of attributes *K*. However, it is possible to increase the number of attributes and add little or no new information, by making the attributes correlated with each other. A stark example of this is duplicate attributes. To rule this possibility out, we assume that the attributes are conditionally independent. Note that in our setting, the attributes for consumers who have high willingness to pay are endogenous, therefore we only impose the independent condition on the consumers with low willingness to pay.

Assumption 7. There exist marginal probabilities $\mu_j : \{0, 1\} \rightarrow \mathbb{R}_+, j = 1, ..., K$, such that for any $\mathbf{a} \in \mathcal{A}$:

$$m(\mathbf{a}) = \overline{m} \prod_{j=1}^{K} \mu_j(\mathbf{a}_j).$$

Under this assumption, we can use the number of attributes K as a measure of the richness of consumer data because every attribute carries new information and, therefore, improves the estimate of consumer demand.

This assumption is strong and can be relaxed. In particular, we can allow for some attributes not only to be correlated, but to be identical to each other, as long as the proportion of these attributes does not does not grow too large when we increase the number of attributes. At the same time, independence allows for a very tractable closed-form characterization of the value of information.

It is important to clarify that we consider the richness of the data in terms of available variables and not observations. In our setup, the seller perfectly understand the data generating process for any collection of variables. In particular, he understands that different sets of predictors have different predictive power (both exogenously and because of consumers manipulating them).

Because of independence, the solution for program (2) for a large K can be constructed using solutions for the analogous problem for a smaller K. We can use induction on the number of attributes to characterize the solution.

PROPOSITION 8. If Assumption 7 holds, the value of information is

$$D_K = \frac{1}{K}\overline{m} \left(\frac{c}{\delta}\right)^2 \frac{\sum_{j=1}^K \mu_j(0)\mu_j(1)}{K}$$
(5)

PROOF. Recall that objective for the problem (2) for K + 1 attributes can be rewritten as

$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a}) \left(h(\mathbf{a}) - \overline{h} \right)^2 = \sum_{\mathbf{a}\in A} m(\mathbf{a}) \left(h(\mathbf{a}) - h(A) \right)^2 + \sum_{\mathbf{a}\in B} m(\mathbf{a}) \left(h(\mathbf{a}) - h(B) \right)^2 + \sum_{Z\in\{A,B\}} m(Z) \left(h(Z) - \overline{h} \right)^2, \tag{6}$$

where $A = \{\mathbf{a} \mid \mathbf{a} = (\mathbf{b}, 0), \mathbf{b} \in \mathcal{A}_K\}$ and $B = \mathcal{A}_{K+1} \setminus A$. Note that the first two components of this sum are the values of the objective for the problem (2) for *K* attributes. We can maximize the third component of the sum and check if the result violates any constraints of the original problem for K + 1 attributes. In particular, we solve for the contribution of the (K + 1)th attribute towards the overall objective

$$V_{K+1} = \delta^2 \max_{\{h(A), h(B)\}} \sum_{Z \in \{A, B\}} m(Z) \left(h(Z) - \overline{h} \right)^2$$

s.t.
$$\sum_{Z \in \{A, B\}} m(Z) \left(h(Z) - \overline{h} \right) = 0$$

$$|h(A) - h(B)| \le \frac{c}{\delta^2} \frac{1}{K+1},$$

The solution to this problem is

$$h(A) - \overline{h} = \frac{\mu_{K+1}(1)}{(K+1)} \frac{c}{\delta^2},$$

$$h(B) - \overline{h} = -\frac{\mu_{K+1}(0)}{(K+1)} \frac{c}{\delta^2}.$$

and, therefore, the maximal contribution of the (K+1)th attribute towards the overall objective is

$$W_{K+1} = \overline{m}\mu_{K+1}(0)\mu_{K+1}(1)\left(\frac{c}{\delta}\frac{1}{K+1}\right)^2$$

Let $D_K(c)$ be a value of information with K attributes and cost of manipulation c. Combining (6) with expression for V_{K+1} we get

$$D_{K+1}(c) = D_K\left(\frac{K}{K+1}c\right) + \overline{m}\mu_{K+1}(0)\mu_{K+1}(1)\left(\frac{c}{\delta}\frac{1}{K+1}\right)^2$$

The solution to this equation is

$$D_K(c) = \frac{1}{K}\overline{m}\left(\frac{c}{\delta}\right)^2 \frac{\sum_{j=1}^K \mu_j(0)\mu_j(1)}{K}$$

To establish if this value is feasible, we construct the solution h_{K+1}^* from a solution to problem with K attributes—i.e., h_{K+1}^* . We show that h_{K+1}^* also satisfies all relevant constraints.

Using h_K^* let us define

$$\tilde{h}_{K+1}(\mathbf{a}) = \overline{h} - \frac{K}{K+1} \left(\overline{h} - h_K^*(\mathbf{a}) \right)$$

Note that for any $a, b \in \mathcal{A}_K$ and $z \in \{0, 1\}$ the following constraint is satisfied because h^* satisfied all constraints in program (2):

$$\left|\tilde{h}_{K+1}(\mathbf{a}) - \tilde{h}_{K+1}(\mathbf{b})\right| \le \frac{c}{\delta^2} \frac{\|(\mathbf{a}, z) - (\mathbf{b}, z)\|}{K+1}$$

The value D_{K+1} is achieved at

$$h_{K+1}(\mathbf{a}) = \begin{cases} \tilde{h}_{K+1}(\mathbf{b}) - \frac{\mu_{K+1}(1)}{K+1} \frac{c}{\delta^2}, & \text{if } \mathbf{a} = (\mathbf{b}, 1) \\ \tilde{h}_{K+1}(\mathbf{b}) + \frac{\mu_{K+1}(0)}{K+1} \frac{c}{\delta^2}, & \text{if } \mathbf{a} = (\mathbf{b}, 0). \end{cases}$$

Note that the proposed solution h_{K+1} is obtained from h_K by applying the same transformation. It has the following features:

- (1) for any $\mathbf{a}, \mathbf{b} \in \mathscr{A}_{K+1}$ this transformation ensures that the constraint $|h_{K+1}(\mathbf{a}) h_{K+1}(\mathbf{b})| \le \frac{c}{\delta^2} \frac{\|\mathbf{a}-\mathbf{b}\|}{K+1}$ is satisfied. In particular, if
 - (a) $a_{K+1} = b_{K+1}$, this constraint is implied by the corresponding constraint for h_K ,
- (b) $a_{K+1} \neq b_{K+1}$, this constraint is satisfied because $\frac{\mu_{K+1}(1)}{K+1}\frac{c}{\delta^2} + \frac{\mu_{K+1}(0)}{K+1}\frac{c}{\delta^2} = \frac{c}{\delta^2}\frac{1}{K+1}$. (2) it maximizes the objective because it maximizes all three components of the sum in (6).



Fig. 2. The optimal market segmentation for K = 2 (on the left) and K = 4 (on the right) under symmetry.

There are two features present in expression (5) worthy of attention. First, the value of data is increasing in the average variance of the attributes for consumers with low valuation:

$$\frac{1}{K}\sum_{j=1}^K \mu_j(0)\mu_j(1).$$

There is a simple intuition for this. While consumers with high valuations manipulate their data, consumers with low valuation stay passive. Market segments with low concentrations of low valuation consumers are precisely the segments that consumers with high valuations are trying to avoid. Thus, the seller benefits from an as equal distribution of low valuation consumers across segments as possible. This attribute diversity among low valuation consumers is achieved by increasing the variance of their attributes.

Second, the value of data depends on the number of attributes in a stark way. Increasing the number of attributes has two conflicting effects: the predictive power of personal data increases if the attributes cannot be manipulated; while at the same time, the possibilities for manipulation grow which makes the data less reliable in aggregate. The following proposition proves that the second effect dominates.

PROPOSITION 9. If the data becomes arbitrarily rich—i.e., if the number of consumer attributes becomes large—the value of information becomes arbitrarily small:

$$\lim_{K\to\infty} D_K = 0$$

PROOF. First, since $\mu_j(1)\mu_j(0) \le \frac{1}{4}$ for any *j*, note that (Equation (5))

$$D_K = \frac{1}{K}\overline{m}\left(\frac{c}{\delta}\right)^2 \frac{\sum\limits_{j=1}^K \mu_j(0)\mu_j(1)}{K} \le \frac{1}{K}\frac{\overline{m}}{4}\left(\frac{c}{\delta}\right)^2$$

Thus, $\lim_{K \to \infty} D_K = 0.$

To gain more intuition about this result consider a symmetric case in which $\mu_j(0) = \mu_j(1)$ for every j = 1, ..., K. In this case every segment contains the same mass of consumers with low willingness to pay $m(\mathbf{a}) = \frac{\overline{m}}{2^K}$. When the number of attributes K increases, the share of segments with the hazard ratio close to \overline{h} grows whereas the share of segments with either large or small hazard ratios vanishes. This occurs because the mesh of binding no-arbitrage constraints becomes tighter. For the illustration of this effect see the two examples for K = 2 and K = 4 depicted on Figure 2. Because the data is used to predict deviations in demand from the average, the total value of this prediction becomes as K grows large.

To summarize, the larger the number of attributes becomes, the more combinations of the attributes correspond to approximately "average" market segments, i.e., the market segments in which the composition of consumers is similar to the average composition across the entire market. This is reminiscent of the Central Limit Theorem and makes it cheaper on average for the consumers to manipulate their data in pursuit of a better deal.

What are the welfare implications of richer consumer data? First, let us set aside the direct welfare costs associated with data manipulation and concentrate on the creation and division of surplus *after* the consumers changed their attributes. When the seller segments the market using consumer data, the total welfare is reduced:

$$\Delta W = -\delta^2 \sum_{\mathbf{a} \in \mathcal{A}} m(\mathbf{a}) \left(h_K(\mathbf{a}) - \overline{h} \right)^2 = -\frac{1}{K} \overline{m} \left(\frac{c}{\delta} \right)^2 \frac{\sum_{j=1}^K \mu_j(0) \mu_j(1)}{K}.$$

Therefore the consumer surplus is also reduced:

$$\Delta CS_{min} = -\frac{2}{K}\overline{m}\left(\frac{c}{\delta}\right)^2 \frac{\sum_{j=1}^{K} \mu_j(0)\mu_j(1)}{K}.$$

Just like the gain in the monopoly profit, these values become arbitrarily small when the number of attributes becomes arbitrarily large.

The previous expression excludes the cost of data manipulations—it accounts only for quality distortions and higher prices. The aggregate cost of manipulation depends on how informative the *ex ante* data is (i.e., the private data represented by $\omega(i)$). For example, one extreme scenario is that the ex ante attributes are such that no consumer wants to change them. To capture the range of welfare losses, consider the upper bound on the total cost of manipulation:

$$\sum_{\mathbf{a}\in\mathscr{A}}m(\mathbf{a})h_K(\mathbf{a})\frac{\|\mathbf{a}-\mathbf{a}_m\|}{K}c,$$

where $\mathbf{a}_m = \arg \max_{\mathbf{a} \in \mathcal{A}} h_k(\mathbf{a})$. To understand how this upper bound depends on the richness of the data, consider a symmetric case in which $\mu_i(0) = \mu_i(1) = 1/2$ for all i = 1, ..., K. In this case,

$$h_K(\mathbf{a}) = \left(\frac{K}{2} - \sum_{i=1}^K a_i\right) \frac{c}{K\delta^2} + \overline{h},$$

and $\mathbf{a}_m = (0, \dots, 0)$. For maximal losses we get⁶

$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a})h_K(\mathbf{a})\frac{\|\mathbf{a}-\mathbf{a}_m\|}{K}c = \frac{\overline{m}c^2}{K\delta^2} \left(\sum_{k=1}^K \binom{K}{k} \frac{k}{K} \left(\frac{K}{2}-k+\frac{K\delta^2}{c}\overline{h}\right)\right) = \frac{\overline{n}c}{2} - \frac{1}{K}\frac{\overline{m}c^2}{2\delta^2}$$

If we combine this cost with the expression for consumer surplus we obtained earlier, we get the maximal reduction in consumer surplus:

$$\Delta CS_{max} = -\frac{\overline{n}c}{2}.$$

Thus, the presence of data reduces the consumer surplus, and in contrast to the effect on the seller's profit, the reduction in consumer surplus may not vanish when the number of consumer attributes becomes large.

4.2 Less data!

The value of information vanishes when consumer attributes become numerous because of the increase in opportunities of manipulation. There are ways for the seller to limit these opportunities. One possibility is to commit to the minimal use of data: the seller can promise the consumers to use a small fraction of the variables for pricing. However, this improve seller's ability to make personalized offers only if the seller is secretive about which attributes are used for pricing. In this section we consider an extreme example of such policy. The seller commits to use only one attribute without disclosing to the consumers which one exactly.

The result that we obtain in this section echoes similar observations about limiting the use of data made in other settings: Frankel and Kartik (2022) and Ball (2022) show that such a commitment, often implemented by introducing a data intermediary, improves endogenous precision of the data.

In our setting, the seller randomly chooses the attribute for the purpose of segmenting the market. If a particular attribute is not chosen in equilibrium, it becomes very informative, due to the consumers not manipulating it, and, therefore, using it would be a profitable deviation.

By γ_j denote the probability of the seller using attribute j = 1, ..., K for market segmentation. The no-arbitrage condition (4) in this case becomes

$$|h(a_j=0)-h(a_j=1)| \leq \frac{c}{\delta^2} \frac{1}{\gamma_j K},$$

and the gain from market segmentation based on attribute j is

$$\overline{m}\left(\frac{c}{\delta}\right)^2 \frac{\mu_j(0)\mu_j(1)}{\gamma_j^2 K^2}.$$

Because the firm chooses the attribute randomly, the gain must be the same for any two attributes. We can find the probabilities γ_j from this condition:

$$\gamma_j = \frac{\sqrt{\mu_j(0)\mu_j(1)}}{\sum\limits_{k=1}^{K} \sqrt{\mu_k(0)\mu_k(1)}}.$$

The likelihood of the seller using the attribute *j* for pricing is increasing in the variance of this attribute among consumers with low valuation $\mu_i(0)\mu_i(1)$. As we pointed out in the

⁶For binomial sums, see Boros and Moll (2004).

previous section, this variance measures how (un)attractive the particular attribute is for the purposes of manipulation. If the variance is low, the returns to manipulation are high.

PROPOSITION 10. If the seller commits to use only a single attribute for market segmentation without disclosing which attribute exactly, the value of consumer data is

$$D_K^r = \overline{m} \left(\frac{c}{\delta} \frac{\sum\limits_{j=1}^K \sqrt{\mu_j(0)\mu_j(1)}}{K} \right)^2.$$

If the seller adopts this data policy, the consumers' expected return to data manipulation becomes smaller. The reason is simple: when a consumer changes the value of attribute *i*, with probability $1 - \gamma_i$ she does not gain anything because the seller does not use this attribute for pricing. This implies that when the seller does use the attribute for pricing, it contains a great amount of information. This is true for every attribute, and therefore, the value of data is larger compared to the case when the seller uses all available attributes simultaneously.

Note that D_K^r is larger than D_K by a factor of $\sum_{j=1}^K \sqrt{\mu_j(0)\mu_j(1)}$ which increases linearly in K if $\mu_K(0)\mu_K(1)$ does not converge to zero. This implies that D_K^r does not vanish when K becomes large. Thus, using less data may result in higher overall value.

REFERENCES

- Acquisti, A., Taylor, C., Wagman, L. 2016. The economics of privacy. Journal of economic Literature. 54 (2), 442–492.
- Ali, S. N., Lewis, G., Vasserman, S. 2022. Voluntary Disclosure and Personalized Pricing. The Review of Economic Studies. 90 (2), 538–571.
- Ali, S. N., Lewis, G., Vasserman, S. 2023. Consumer control and privacy policies. AEA Papers and Proceedings. 113, 204–09.
- Ball, I. 2022. Scoring strategic agents.
- Bergemann, D., Bonatti, A. 2015. Selling cookies. American Economic Journal: Microeconomics. 7 (3), 259-294.
- Bergemann, D., Bonatti, A., Smolin, A. 2018. The design and price of information. American economic review. 108 (1), 1–48.
- Bhaskar, V., Roketskiy, N. 2021. Consumer privacy and serial monopoly. The RAND Journal of Economics. 52 (4), 917–944.
- Bonatti, A., Cisternas, G. 2020. Consumer scores and price discrimination. The Review of Economic Studies. 87 (2), 750–791.
- Bonatti, A., Huang, Y., Villas-Boas, J. M. 2023. A theory of the effects of privacy.
- Boros, G., Moll, V. 2004. Irresistible integrals: symbolics, analysis and experiments in the evaluation of integrals. Cambridge University Press.
- Caner, M., Eliaz, K. 2021. Non-manipulable machine learning: The incentive compatibility of lasso. arXiv preprint arXiv:2101.01144.
- Carroll, G., Egorov, G. 2019. Strategic communication with minimal verification. Econometrica. 87 (6), 1867–1892.
- Conitzer, V., Taylor, C. R., Wagman, L. 2012. Hide and seek: Costly consumer privacy in a market with repeat purchases. Marketing Science. 31 (2), 277–292.
- Cunningham, T., Moreno De Barreda, I. 2022. Effective signal-jamming.
- Dana, J., Larsen, B., Moshary, S. 2023. Fake ids and arbitrage: Price discrimination with misreporting costs.
- Deneckere, R., Severinov, S. 2022. Signalling, screening and costly misrepresentation. Canadian Journal of Economics. 55 (3), 1334–1370.
- Eilat, R., Eliaz, K., Mu, X. 2021. Bayesian privacy. Theoretical Economics. 16 (4), 1557-1603.
- Eliaz, K., Spiegler, R. 2022. On incentive-compatible estimators. Games and Economic Behavior. 132, 204-220.

Elliott, M., Galeotti, A., Koh, A., Li, W. 2021. Market segmentation through information.

Frankel, A., Kartik, N. 2019. Muddled information. Journal of Political Economy. 127 (4), 1739-1776.

Frankel, A., Kartik, N. 2022. Improving information from manipulable data. Journal of the European Economic Association. 20 (1), 79–115.

Hardt, M., Megiddo, N., Papadimitriou, C., Wootters, M. 2016. Strategic classification. In Proceedings of the 2016 ACM conference on innovations in theoretical computer science, 111–122.

- Hidir, S., Vellodi, N. 2021. Privacy, personalization, and price discrimination. Journal of the European Economic Association. 19 (2), 1342–1363.
- Hu, L., Immorlica, N., Vaughan, J. W. 2019. The disparate effects of strategic manipulation. In Proceedings of the Conference on Fairness, Accountability, and Transparency, 259–268.
- Ichihashi, S. 2020. Online privacy and information disclosure by consumers. American Economic Review. 110 (2), 569–595.
- Kartik, N. 2009. Strategic communication with lying costs. The Review of Economic Studies. 76 (4), 1359–1395.

Liang, A., Madsen, E. forthcoming. Data and incentives. Theoretical Economics.

- Milli, S., Miller, J., Dragan, A. D., Hardt, M. 2019. The social cost of strategic classification. In Proceedings of the Conference on Fairness, Accountability, and Transparency, 230–239.
- Mussa, M., Rosen, S. 1978. Monopoly and product quality. Journal of Economic theory. 18 (2), 301-317.
- Perez-Richet, E., Skreta, V. 2022. Test design under falsification. Econometrica. 90 (3), 1109-1142.

Perez-Richet, E., Skreta, V. 2023. Fraud-proof non-market allocation mechanisms.

Segura-Rodriguez, C. 2021. Selling data.

Severinov, S., Tam, Y. C. 2018. Screening under fixed cost of misrepresentation.

Tan, T. Y. 2023. Price discrimination with manipulable observables.

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

Yang, K. H. 2022. Selling consumer data for profit: Optimal market-segmentation design and its consequences. American Economic Review. 112 (4), 1364–1393.