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. We find that randomly restricting monopolist's access to consumer data may increase 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 the spoils of trade are 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 show the opposed interests. When some parts of the 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 in consumer surplus. In contrast with the seller's value of data, the consumer surplus loss may remain substantial when data becomes rich.

The hide and seek nature of monopolist-consumer relation 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.

[1]In 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.

This version: February 1, 2023.

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 addition 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.

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

The use of data by sellers is tightly related to consumer privacy. Bonatti and Cisternas (2020) investigate how the seller can condition the current prices on the past choices of the consumers via an aggregator that assigns a score to each consumer. Bhaskar and Roketskiy (2021) consider unrestricted use of past purchases for the equilibrium optimal pricing.

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 privacy-preserving pricing. They show that an optimal privacy-preserving menu always contains a finite number of alternatives.

We study markets where the seller uses data to segment the market and tailor prices to the demand in each segment. In a different environment, Hidir and Vellodi (2021) show that it is possible to find a segmentation that relies on consumer volunteering the private information on their valuations.

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. 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 anonymous consumers indexed by $i \in C$. A variety of goods can be produced and sold by the seller. Each type of good is characterized by a quality parameter $q \in \mathbb{R}_+$. A consumer i's marginal willingness to pay for quality $\tau(i)$ depends on her private type. There are consumers with the low willingness to pay for quality t_ℓ, and consumers with a high one $t_h > t_\ell$. We denote the difference between the two by $d = t_h - t_\ell$.

A market transaction between the seller and the consumer i generates social surplus

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

that depends on the quality of the good q and the consumer's type $t(i)$. 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. Consumers' outside option is valued at zero.

Consumer i is endowed with K *observable* and *manipulable* attributes represented by vector

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

The attributes can be secretly manipulated by the consumer at a cost. We assume that it costs c to change the entire vector and the cost is linear in the number of attributes that the consumer changes. For example, a consumer with an attribute vector $\hat{\mathbf{a}} = (0,0,1,0,1)$ can pay cost $\frac{2}{5}c$ to change the value of the second and fifth attributes. As a result, instead of the original $\hat{\mathbf{a}}$, this consumer would display his new attribute vector $\mathbf{a} = (0,1,1,0,0)$. From now on, we use a "hat" ($\hat{}$) to denote a value of a variable before the consumers made changes to their attributes. All variables without hats are ex post.

The consumer attributes may be correlated with the consumer type. By $m : 2^C \longrightarrow \mathbb{R}_+$ we will denote the measure of consumers with type L. Similarly, $n : 2^C \longrightarrow \mathbb{R}_+$ is the measure of consumers with type H. With some abuse of notation we denote sets of consumers by their characteristics: for example, $m(\mathbf{a}) = m(\{i \in C : \alpha(i) = \mathbf{a}\})$ is the mass of all consumers with type L and the attribute vector $\mathbf{a}$.

3.1 Optimal menu pricing

Suppose the seller offers a profit-maximizing screening menu to a group of consumers $S \subseteq C$. 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 the low willingness to pay:

$$\rho(S) = \frac{1}{m(S)} \int_S \pi(i) d[n(i) + m(i)],$$

As we know from Mussa and Rosen (1978), the profit-maximizing menu consists of two items. A premium item (p_h, q_h) is designed for consumers with high valuation for quality, and a basic item (p_ℓ, q_ℓ) is for everyone else. The seller's profit solves the following program:

$$\rho(S) = \max_{q_h \geq q_\ell \geq 0} \left\{ 2t_\ell q_\ell - q_\ell^2 - 2h(S) q_\ell d + h(S)\left[2t_h q_h - q_h^2\right] \right\},$$

where

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

is a hazard ratio for the group S. This ratio is a sufficient statistics of consumer heterogeneity and it plays a key role in our analysis. The solution to the seller's program is

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

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

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

The monopolist's profit is

$$\rho(S) = h(S)(t_\ell + d)^2 + [\max\{0, t_\ell - h(S)d\}]^2.$$

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.

$$\bar{h} \in \left[\frac{c}{2d^2}, \frac{t_\ell}{d} - \frac{c}{2d^2}\right].$$

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 a basic item. In this sense, the third degree price discrimination is a more cost effective instrument than menu pricing. The seller could use consumer data to 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 sellers pricing decision, let us consider an arbitrary market segmentation, which is, by definition, a partition of the set of consumers: $\mathcal{S} = \left\{S_1, S_2, ... \middle| \bigcup_i S_i = C; \forall i \neq j : S_i \cap S_j = \emptyset \right\}$. If the seller's profit-maximizing menus serve both type of the consumers in every segment, the sellers' total profit under market segmentation $\mathcal{S}$ is

$$\pi_{\mathcal{S}} = \sum_{S \in \mathcal{S}} m(S)\rho(S) = \pi^* + d^2 \sum_{S \in \mathcal{S}} m(S)[h(S) - \bar{h}]^2,$$

where $\pi^* = m(C)\rho(C)$ is the profit of the monopolist without market segmentation. 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 variance of the hazard ratio across different market 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 is increasing in hazard ratio:

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

Therefore, 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 needs changing $A(S_1, S_2)/K$. In equilibrium, the following *no-arbitrage condition* must hold for any two segments S_1 and S_2:

$$|h(S_1) - h(S_2)| \leq \frac{c}{2d^2}\frac{A(S_1, S_2)}{K}. \tag{1}$$

To see that, recall that the hazard ratio is endogenous in the setup with manipulable data. Consider high valuation consumers in segment S_1 and assume that initially $\hat{h}(S_1) > \hat{h}(S_2)$. If the inequality (1) holds for $\hat{h}$, no consumer is tempted to manipulate their attributes. However, if it does not hold, a number of consumers would travel from S_1 to S_2, thus reducing the gap between $h(S_1)$ and $h(S_2)$. This number is determined in equilibrium. It is such that a marginal consumer is indifferent between the two options and, therefore, (1) would hold as equality for ex post hazard ratios h.

4 VALUE OF CONSUMER DATA

By segmenting the market according to $\mathcal{S}$, the seller increases his profit by

$$d^2 \sum_{S \in \mathcal{S}} m(S)[h(S) - \bar{h}]^2.$$

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, these ratios are not directly observable, which raises a question of how a third party could measure it using market data. To answer it, it is sufficient to examine the prices in the optimal menu and the quantity demanded by the consumers:

$$d^2 \sum_{S \in \mathcal{S}} m(S)[h(S) - \bar{h}]^2 = \frac{1}{d^2} \sum_{S \in \mathcal{S}} m(S) \left[p_h(S) - \frac{1}{\bar{m}} \sum_{\tilde{S} \in \mathcal{S}} m(\tilde{S}) p_h(\tilde{S}) \right]^2,$$

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

Because consumers are anonymous, not every segmentation can be implemented by the seller. The seller can either use observable attributes or random chance (or the combination of the two) to allocate consumers to market segments.

Definition 2. *Segmentation $\mathcal{S}$ is feasible if there exists a function $F : \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ such that for every market segment $S \in \mathcal{S}$ and every consumer $i \in C$: $\Pr\{i \in S\} = F(\alpha(i), S)$.*

Consider a segmentation that associates each possible value of attributes vector with a separate market segment:

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

Clearly this segmentation is feasible. Moreover, the next proposition shows that it is profit-maximizing among all feasible segmentations.

Proposition 3. *For any feasible market segmentation $\mathcal{S}$,*

$$\sum_{S \in \mathcal{S}} m(S)[h(S) - \bar{h}]^2 \leq \sum_{S \in \mathcal{S}^*} m(S)[h(S) - \bar{h}]^2.$$

Proof. The gain in profit is convex in hazard ratio and $\{h(S), m(S)\}_{S \in \mathcal{S}^*}$ weakly majorizes $\{h(S), m(S)\}_{S \in \mathcal{S}}$. □

This proposition shows that the seller uses segmentation $\mathscr{S}^*$ 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.

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.

Definition 4. *The value of consumer data D_K is the largest gain the seller can obtain by segmenting the market according to $\mathscr{S}^*$ across all possible correlations between the attributes and valuations:*

$$D_K = d^2 \max_{h:\mathscr{A}\to\mathbb{R}_+} \sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a})[h(\mathbf{a}) - \bar{h}]^2 \tag{2}$$

$$s.t. \quad \sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a})[h(\mathbf{a}) - \bar{h}] = 0, \tag{3}$$

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

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^*.[2] 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 5. *A graph $G(h)$ with the set of nodes $\mathscr{A}$ is called a constraints graph for h if for every pair of $\mathbf{a}, \mathbf{b} \in \mathscr{A}$ the following statements are equivalent:*

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

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

[2]Note that the objective function is convex, therefore the solution will be on the boundary of the convex permissible set.

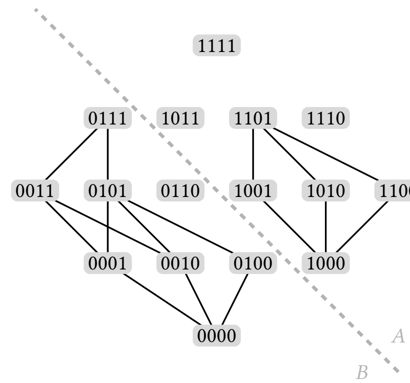

Fig. 1. An example of a constraints graph.

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 $h(A) \geq h(B)$.

The objective can be rewritten as

$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a})[h(\mathbf{a})-\bar{h}]^2 = \sum_{\mathbf{a}\in A} m(\mathbf{a})[h(\mathbf{a})-h(A)]^2 + \sum_{\mathbf{a}\in B} m(\mathbf{a})[h(\mathbf{a})-h(B)]^2 + \sum_{Z\in\{A,B\}} m(Z)[h(Z)-\bar{h}]^2.$$

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

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

hence h cannot be a solution to (2). □

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

COROLLARY 7. *If h^* is a solution to the program* (2), *then*

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

4.1 More data?

We are now prepared to investigate how this value depends on the richness of the data. The parameter that guides 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. 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 8. *There exist marginal probabilities* $\mu_i : \{0,1\} \to \mathbb{R}_+, i = 1, ..., K$, *such that for any* $\mathbf{a} \in \mathscr{A}$:

$$m(\mathbf{a}) = \bar{m} \prod_{i=1}^{K} \mu_i(\mathbf{a}_i).$$

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.

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 9. *If Assumption 8 holds, the value of information is*

$$D_K = \frac{1}{K} \bar{m} \left[\frac{c}{2d}\right]^2 \frac{\sum_{j=1}^{K} \mu_j(0)\mu_j(1)}{K}$$

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

$$\begin{aligned} \sum_{\mathbf{a} \in \mathscr{A}} m(\mathbf{a})[h(\mathbf{a}) - \bar{h}]^2 = & \sum_{\mathbf{a} \in A} m(\mathbf{a})[h(\mathbf{a}) - h(A)]^2 \\ & + \sum_{\mathbf{a} \in B} m(\mathbf{a})[h(\mathbf{a}) - h(B)]^2 \\ & + \sum_{Z \in \{A,B\}} m(Z)[h(Z) - \bar{h}]^2, \end{aligned} \tag{5}$$

where $A = \{\mathbf{a} \mid \mathbf{a} = (\mathbf{b}, 0), \mathbf{b} \in \mathscr{A}_K\}$ and $B = \mathscr{A}_{K+1} \setminus A$. Note that the first two components of this sum are the values of the objective for the problem (2) for K attributes. Since the solution to this problem is known, we can attempt to 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

$$\begin{aligned} V_{K+1} = d^2 \max_{\{h(A),h(B)\}} & \sum_{Z \in \{A,B\}} m(Z)[h(Z) - \bar{h}]^2 \\ s.t. & \sum_{Z \in \{A,B\}} m(Z)[h(Z) - \bar{h}] = 0 \\ & |h(A) - h(B)| \leq \frac{c}{2d^2} \frac{1}{K+1}, \end{aligned}$$

The solution to this problem is

$$\begin{aligned} [h(A) - \bar{h}] &= \frac{\mu_{K+1}(1)}{(K+1)} \frac{c}{2d^2}, \\ [h(B) - \bar{h}] &= -\frac{\mu_{K+1}(0)}{(K+1)} \frac{c}{2d^2}. \end{aligned}$$

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

$$V_{K+1} = \bar{m}\mu_{K+1}(0)\mu_{K+1}(1)\left[\frac{c}{2d}\frac{1}{K+1}\right]^2$$

Let $D_K(c)$ be a value of information with K attributes and cost of manipulation c. Combining (5) with expression for V_{K+1} we get

$$D_{K+1}(c) = D_K\left(\frac{K}{K+1}c\right) + \bar{m}\mu_{K+1}(0)\mu_{K+1}(1)\left[\frac{c}{2d}\frac{1}{K+1}\right]^2$$

The solution to this equation is

$$D_K(c) = \frac{1}{K}\bar{m}\left[\frac{c}{2d}\right]^2 \frac{\sum_{j=1}^{K}\mu_j(0)\mu_j(1)}{K}$$

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

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

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

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

The value D_{K+1} is achieved at

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

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 $|h_{K+1}(\mathbf{a}) - h_{K+1}(\mathbf{b})| \le \frac{c}{2d^2}\frac{\|\mathbf{a}-\mathbf{b}\|}{K+1}$ 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 $\frac{\mu_{K+1}(1)}{K+1}\frac{c}{2d^2} + \frac{\mu_{K+1}(0)}{K+1}\frac{c}{2d^2} = \frac{c}{2d^2}\frac{1}{K+1}$.
(2) it maximizes the objective because it maximizes all three components of the sum in (5).

□

Proposition 10. *When the data becomes arbitrarily rich—i.e., when the number of consumer attributes becomes large—the value of information becomes arbitrarily small:*

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

Proof. Because $\mu_K(1)\mu_K(0) \le \frac{1}{4}$ for any K, we note that

$$0 \le \lim_{K\to\infty} \bar{m}\left[\frac{c}{2d}\right]^2 \frac{\sum_{j=1}^{K}\mu_j(0)\mu_j(1)}{K} \le \frac{\bar{m}}{4}\left[\frac{c}{2d}\right]^2.$$

Thus, $\lim_{K\to\infty} D_K = 0$. □

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.

What are the welfare implications of better 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 by

$$\Delta W = -d^2 \sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a})[h_K(\mathbf{a}) - \bar{h}]^2 = -\frac{1}{K}\bar{m}\left[\frac{c}{2d}\right]^2 \frac{\sum_{j=1}^{K} \mu_j(0)\mu_j(1)}{K}.$$

Therefore the consumer surplus is reduced by

$$\Delta CS = -\frac{2}{K}\bar{m}\left[\frac{c}{2d}\right]^2 \frac{\sum_{j=1}^{K} \mu_j(0)\mu_j(1)}{K}.$$

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

These calculations do not take into account the costs of data manipulations. It is natural to assume that manipulating attributes carries no explicit value for the consumers and therefore it reduces welfare. The reduction in welfare via this channel depends on how informative the data is *before* consumers make changes. It is possible 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 $\mathbf{a}_m = \arg\max_{\mathbf{a}\in\mathscr{A}} h_k(\mathbf{a})$. To understand how this upper bound depends on the richness of the data, consider a case in which $\mu_i(0) = \mu_i(1) = 1/2$ for all $i = 1, \ldots, K$. In this case,

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

and $\mathbf{a}_m = (0, \ldots, 0)$. For maximal losses we get[3]

$$\sum_{\mathbf{a}\in\mathscr{A}} m(\mathbf{a}) h_K(\mathbf{a}) \frac{\|\mathbf{a} - \mathbf{a}_m\|}{K} c = \frac{\bar{m}c^2}{2Kd^2}\left[\sum_{k=1}^{K} \binom{K}{k} \frac{k}{K} \left(\frac{K}{2} - k + \bar{h}\frac{2Kd^2}{c}\right)\right] = \frac{\bar{n}c}{2} - \frac{1}{K}\frac{\bar{m}c^2}{8d^2}$$

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

$$\overline{\Delta CS} = -\frac{\bar{n}c}{2}.$$

We can see from this expression that the presence of the data reduces the consumer surplus. In contrast to the effect on the seller's profit, the reduction in consumer surplus does not vanish when the number of consumer attributes becomes large.

[3]For binomial sums, see Boros and Moll (2004).

4.2 Less data!

The reason why value of information vanishes when consumer attributes become numerous has to do with increasing opportunities to manipulate the data. There are ways for the seller to limit these opportunities. One possibility is to credibly promise the consumers to use a small fraction of the available data. However, this would work only if the seller is secretive about which attributes are used for pricing. In this section we consider an extreme example of such policy—i.e., the seller committing to use only one attribute without disclosing to the consumers which one exactly.

Bound by the commitment, the seller will randomly choose 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_i denote the probability of the seller using attribute $i = 1, \ldots, K$ for market segmentation. The no-arbitrage condition (4) in this case becomes

$$|h(a_i = 0) - h(a_i = 1)| \le \frac{c}{2d^2} \frac{1}{\gamma_i K},$$

and the gain from market segmentation is

$$\bar{m} \left[\frac{c}{2d} \right]^2 \frac{\mu_i(0)\mu_i(1)}{\gamma_i^2 K^2}.$$

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

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

Proposition 11. *If the seller commits to use only a single attribute for market segmentation without disclosing which attribute exactly, the value of consumer data is*

$$D_K^r = \bar{m} \left[\frac{c}{2d} \frac{\sum_{j=1}^{K} \sqrt{\mu_j(0)\mu_j(1)}}{K} \right]^2.$$

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 deal 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.

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 $\mu_j(0)\mu_j(1)$ does not converge to zero. This implies that D_K^r does not vanish when K becomes large.

REFERENCES

Ian Ball. Scoring strategic agents. 2022.

Dirk Bergemann and Alessandro Bonatti. Selling cookies. *American Economic Journal: Microeconomics*, 7(3):259–294, 2015.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American economic review*, 108 (1):1–48, 2018.

V. Bhaskar and Nikita Roketskiy. Consumer privacy and serial monopoly. *The RAND Journal of Economics*, 52(4):917–944, 2021.

Alessandro Bonatti and Gonzalo Cisternas. Consumer scores and price discrimination. *The Review of Economic Studies*, 87 (2):750–791, 2020.

George Boros and Victor Moll. *Irresistible integrals: symbolics, analysis and experiments in the evaluation of integrals*. Cambridge University Press, 2004.

Mehmet Caner and Kfir Eliaz. Non-manipulable machine learning: The incentive compatibility of lasso. *arXiv preprint arXiv:2101.01144*, 2021.

Gabriel Carroll and Georgy Egorov. Strategic communication with minimal verification. *Econometrica*, 87(6):1867–1892, 2019.

Tom Cunningham and Inés Moreno De Barreda. Effective signal-jamming. 2022.

Raymond Deneckere and Sergei Severinov. Signalling, screening and costly misrepresentation. *Canadian Journal of Economics*, 55(3):1334–1370, 2022.

Ran Eilat, Kfir Eliaz, and Xiaosheng Mu. Bayesian privacy. *Theoretical Economics*, 16(4):1557–1603, 2021.

Kfir Eliaz and Ran Spiegler. On incentive-compatible estimators. *Games and Economic Behavior*, 132:204–220, 2022.

Alex Frankel and Navin Kartik. Muddled information. *Journal of Political Economy*, 127(4):1739–1776, 2019.

Alex Frankel and Navin Kartik. Improving information from manipulable data. *Journal of the European Economic Association*, 20(1):79–115, 2022.

Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122, 2016.

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

Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 259–268, 2019.

Navin Kartik. Strategic communication with lying costs. *The Review of Economic Studies*, 76(4):1359–1395, 2009.

Annie Liang and Erik Madsen. Data and incentives. *Theoretical Economics*, forthcoming.

Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 230–239, 2019.

Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317, 1978.

Eduardo Perez-Richet and Vasiliki Skreta. Test design under falsification. *Econometrica*, 90(3):1109–1142, 2022.

Carlos Segura-Rodriguez. Selling data. 2021.

Sergei Severinov and Yin Chi Tam. Screening under fixed cost of misrepresentation. 2018.