Buyer-Optimal Algorithmic Consumption *

Shota Ichihashi Alex Smolin

September 21, 2023

Abstract

We analyze a bilateral trade model in which the buyer's value for the product and the seller's costs are uncertain, the seller chooses the product price, and the product is recommended by an algorithm based on its value and price. We characterize an algorithm that maximizes the buyer's expected payoff and show that the optimal algorithm underrecommends the product at high prices and overrecommends at low prices. Higher algorithm precision increases the maximal equilibrium price and may increase prices across all of the seller's costs, whereas informing the seller about the buyer's value results in a mean-preserving spread of equilibrium prices and a mean-preserving contraction of the buyer's payoff.

*Ichihashi: Queen's University, Department of Economics, shotaichihashi@gmail.com. Smolin: Toulouse School of Economics, University of Toulouse Capitole and CEPR, alexey.v.smolin@gmail.com. The authors acknowledge support from the CTIC grant and the NET Institute grants. Smolin acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program (grant ANR-17-EURE-0010). Javier Gonzalez-Morin provided excellent research assistance. Further acknowledgements will be added.

1 Introduction

Algorithmic decision-making is rapidly spreading in the modern economy, fueled by advancements in information technology and artificial intelligence. Algorithms make recommendations for bail (Angwin et al., 2016), health (Obermeyer et al., 2019), and lending (Jagtiani and Lemieux, 2019), among others. Furthermore, consistent with the predictions of Gal and Elkin-Koren (2016), algorithmic consumption is proliferating, as evidenced by chatbots that construct travel itineraries, robo-advisors that propose suitable financial securities, and price-trackers that automatically seek and pinpoint lower-priced products.

In this paper, we ask how algorithmic consumption affects the distribution of prices and welfare. To answer this question, we characterize algorithms that maximize consumer surplus in a monopolistic setting, and we study the consequences of higher algorithm precision and personalized pricing. Our approach captures two key features of algorithmic consumption. First, algorithms process consumer and product data, facilitating the assessment of trade value (cf. Roesler and Szentes (2017)). Second, with the capability to collect price information, algorithms can base recommendations on price, thereby exerting pressure on the seller's pricing decisions.

Specifically, we study a bilateral trade setting. A buyer and a seller can trade a single product, with both the trade costs and the trade value being uncertain. The seller privately knows the costs, which constitute her type. Initially, the buyer does not know the value or the existence of the product. However, an algorithm can discover the value and recommend the product based on the value and price posted by the seller. If recommended, the buyer forms a value estimate and decides whether to purchase the product at the posted price; otherwise, trade does not occur. Different algorithms are distinguished by their recommendation functions and the resulting demand curves. For any demand curve, the seller sets a price to maximize her profit. Our first goal is to characterize a buyer-optimal algorithm, i.e., an algorithm that maximizes the buyer's expected payoff, and the resulting pricing behavior.

A buyer-optimal algorithm must trade off price impact and efficiency. On the one hand, to impose pressure on the seller's pricing decisions, the algorithm should punish high prices by recommending the product less often. On the other hand, such a lack of recommendations may forego beneficial trade opportunities and damage consumer surplus.

We show that these tradeoffs are optimally resolved by a threshold algorithm, which recommends the product whenever its value reaches a threshold that varies with the price. The algorithm design can thus be recast as determining the range of acceptable prices and the trade volume, which in turn can be solved as a nonlinear screening problem similar to that studied by Baron and Myerson (1982). The optimal recommendations are strategic (Proposition 1): At high prices, the algorithm is overly selective, abstaining from recommending the product for some values above the price. At low prices, the algorithm may be indiscriminate, recommending the product even when its value is less than the price, but whether it happens depends on the distributions of costs and values. Overall, the buyer-optimal algorithmic consumption results in a range of equilibrium prices, with higher-cost sellers posting higher prices and experiencing lower trade volume.

Our characterization enables us to evaluate the price and welfare impact of two data interventions. First, we assess the impact of the algorithm having access to additional or higher-quality data, which results in a distribution of estimated buyer values undergoing a mean-preserving spread (see, e.g., Gentzkow and Kamenica (2016)). This change enhances the algorithm's performance and the buyer's expected payoff. Furthermore, as the algorithm becomes more accurate, it can identify more high-value buyers. As a result, the algorithm recommends the product more frequently at the top of the price range. This price leniency exerts the upward pressure on the equilibrium prices. In addition, the algorithm can also identify more low-value buyers and exclude them from trade, which tends to reduce trade volume and seller profits at the lower end of the price range. Overall, this increase in precision may lead to higher equilibrium prices across all seller types (Proposition 2). That is, enhancing algorithmic precision may result in both a universal price increase and greater consumer welfare.

Second, we assess the impact of information leakage to the seller about the buyer's value, e.g., via market segmentation, which enables the seller to engage in third-degree price discrimination as studied by Bergemann et al. (2015). We show that if the algorithm responds optimally to the leaked information, it can neutralize the effect of price discrimination on average, so that the information leakage has no impact on the expected payoffs of the buyer or seller. Moreover, such leakage enables and incentivizes the seller to set lower prices for low-value consumers and higher prices for high-value consumers, resulting in more dispersed prices and less dispersed consumer surplus (Proposition 3). Therefore, algorithmic consumption may not only protect consumers against price discrimination but could also exploit such discrimination to realize societal benefits by equalizing the distribution of welfare among consumers. This finding suggests that promoting algorithmic consumption may be a powerful consumer protection policy, complementary to the many existing regulatory methods detailed, for example, by Scott Morton et al. (2019).

Related literature.— First and foremost, our paper contributes to the recent strand of economic literature that examines the methods of empowering buyers in monopolistic settings. Roesler and Szentes (2017) analyze buyer-optimal learning in a bilateral trade setting with known seller costs. Similarly to us, they demonstrate that the buyer benefits from ex post imperfect decisions to influence the seller's pricing; that is, full learning about the value is not optimal. Deb and Roesler (2021) extend this analysis to the case of a multiproduct monopoly. Condorelli and Szentes (2020) analyze the buyer-optimal distribution of values within a given interval. We contribute to this literature by letting the buyer's information, and thus posterior value distribution, depend on price, a natural feature of algorithmic recommendations,[1] and show how this feature can provide additional leverage to consumers (see also Remark 1).[2]

Our setting can be viewed as enabling consumers to fully automate their consump-

[1]This feature is shared with the work of Libgober and Mu (2021), where the dependence of information on prices comes from their worst-case approach in which the "adversarial nature" chooses the buyer's information to minimize the seller's profits.

[2]In doing so, we build on the methodologies of Bayesian persuasion (e.g., Kamenica and Gentzkow (2011)) and mechanism design (e.g., Baron and Myerson (1982)). Several other papers have combined these methodologies in trade settings, typically to study revenue maximization, most recently including Lee (2021), Bergemann et al. (2022), and Yang (2022).

tion (Remark 2), i.e., to commit to at which values and prices they would be purchasing a product, marking the difference from the classic setting of Myerson and Satterthwaite (1983). A substantial body of literature on limited commitment investigates how the inability to commit, typically on the part of a seller or a mechanism designer, affects equilibrium trade outcomes (e.g., Mylovanov and Tröger (2014) and Liu et al. (2019)). We proceed in the opposite direction: we demonstrate that the algorithmic consumption empowers consumers with commitment and study how this power should be optimally exercised. Therefore, algorithmic consumption can be viewed as delivering a countervailing power (Galbraith (1952)) to buyers by giving them a greater bargaining position vis-a-vis sellers, and in doing so it can be viewed as an effective alternative to a joint use of an intermediary (see Decarolis and Rovigatti (2021) for an online advertising) or to a merger (see Loertscher and Marx (2022) for a multi-firm bargaining).

Our results in Section 4.1 contribute to the existing literature that explores the effects of disclosing information to buyers or bidders on welfare outcomes, especially focusing on the seller's revenue ((Milgrom and Weber, 1982; Ganuza and Penalva, 2010; Lewis and Sappington, 1994)). In contrast, we examine the impact of providing information to the algorithm, which optimally decides which information to pass on to the buyer.

Finally, the analysis in Section 4.2 speaks to the literature on market segmentation. Papers such as Bergemann et al. (2015), Yang (2022), and Haghpanah and Siegel (2023) study the welfare implications of price discrimination based on consumer information. We show that the buyer-optimal algorithm introduces a new implication, where price discrimination attains a more equal distribution of consumer surplus without affecting the average welfare outcomes. This finding also resonates with recent literature on inequality and fairness. In economics, Dworczak et al. (2021) and Akbarpour et al. (forth.) investigate the optimal welfare redistribution through markets and mechanisms, whereas Doval and Smolin (forth.) characterize the limits of welfare redistribution by means of information provision. In parallel, a body of literature studies algorithmic fairness (e.g., Kleinberg et al. (2018)). We contribute to this literature by adding algorithmic consumption and personalized pricing as a way to attain a fairer welfare outcome.

2 Baseline Model

There is a buyer and a seller. The seller can produce one unit of a product at cost c, which is privately known by the seller and constitutes her *type*. The type distribution F has support $C = [0, 1]$, positive density f, and increasing F/f. The value of the product to the buyer is $v \sim G$ and independent of c. The value distribution G has a strictly positive density g over its support $V = [\underline{v}, \overline{v}]$. For expositional convenience, we assume that the support of costs includes the support of values, $0 \leq \underline{v} < \overline{v} \leq 1$.

The buyer initially knows neither the existence nor the value of the product. However, a *recommendation algorithm* or simply *algorithm* provides the buyer with this information based on the value and price of the product. Specifically, an algorithm is a function $r : [\underline{v}, \overline{v}] \times \mathbb{R}_+ \to [0, 1]$ such that for any pair of a realized value $v \in [\underline{v}, \overline{v}]$ and a product price $p \in \mathbb{R}_+$, the algorithm recommends the buyer to purchase the product with probability $r(v, p)$. The algorithm is commonly known to the buyer and seller.

Given an algorithm, the game unfolds as follows: First, nature draws the seller's type c and the buyer's value v. Second, the seller privately observes her type c but not value v, and posts a price, p. With probability $1 - r(v, p)$, the algorithm does not recommend the product, in which case trade does not occur. With probability $r(v, p)$, the algorithm recommends the product to the buyer, who observes the recommendation and the price, and then decides whether to buy the product. If trade occurs, the buyer and seller obtain ex post payoffs $v - p$ and $p - c$, respectively. Otherwise, both players obtain zero payoffs.

The solution concept is perfect Bayesian equilibrium. If the product is recommended, the buyer updates the expected value of the product to

$$\mathbb{E}[v \mid \text{recommended}] = \frac{\int_{\underline{v}}^{\overline{v}} x r(x, p) g(x) \mathrm{d}x}{\int_{\underline{v}}^{\overline{v}} r(x, p) g(x) \mathrm{d}x},$$

and then decides whether to purchase the product by comparing this value to the price. A pair of an algorithm and a buyer's strategy induce a demand curve, which maps each price to a probability of trade. In equilibrium, each seller type takes this demand curve as given and chooses a price that maximizes her expected payoff.

We call the buyer's expected payoff *buyer surplus* and the seller's expected payoff *seller profit*. An algorithm *attains a given buyer surplus* if this buyer surplus arises in an equilibrium under this algorithm. Our focus is on the recommendation algorithms that maximize buyer surplus:

Definition 1. A recommendation algorithm is *buyer-optimal* if it attains a greater buyer surplus than any other recommendation algorithm.

It will be useful to distinguish between types who trade and those who do not under a given algorithm and their respective prices. Given an algorithm and an equilibrium, we say that a price is *active* if it results in a strictly positive trade probability and is *inactive* otherwise. Similarly, we say that a seller type is *active* if she posts an active price with strictly positive probability and is *inactive* otherwise.

3 Buyer-Optimal Algorithm

In this section, we characterize the buyer-optimal algorithm. Our first observation is that it is without loss to assume that the buyer purchases the product whenever it is recommended. This is because the algorithm can anticipate and mimic the buyer's response. Our second observation is that the seller is concerned solely with trade volume. Generally, there are multiple ways to attain the same trade volume by rearranging which values are recommended for trade. However, an optimal algorithm should maximize the buyer surplus conditional on trade volume, thus prioritizing buyers with higher values. This observation allows us to identify a tractable class of algorithms that we can focus upon.

Specifically, we say that an algorithm r is a *threshold algorithm* if there exists a *threshold function* $\hat{v} : \mathbb{R}_+ \to [\underline{v}, \overline{v}]$ such that $r(v,p) = \mathbb{1}(v \geq \hat{v}(p))$, i.e., the algorithm recommends the product with probability 1 if the value reaches a price-dependent threshold, and with probability 0 otherwise. If $\hat{v}(p) = \overline{v}$ at some price p, then the algorithm never recommends the product at that price.[3] If $\hat{v}(p) = \underline{v}$, then the algorithm always recommends the product at that price.

Lemma 1. (Threshold Algorithms) *For any algorithm r, there exists a threshold algorithm under which the buyer follows the recommendations and that yields a greater buyer surplus than r and the same seller profit as r.*

Lemma 1 shows a general result that threshold algorithms span a Pareto frontier in the space of buyer surplus and seller profit. A direct consequence of this result is that if there exists a buyer-optimal algorithm, then it can be found in the class of threshold algorithms, and in what follows, we continue the search within that class.

The optimal choice of a threshold function must balance efficiency maximization and price minimization. One option, often employed in recommender systems, is to set $\hat{v}(p) = v_0$ for any $p \in \mathbb{R}_+$ so that regardless of the price, the product is recommended whenever the value is sufficiently high. This option overlooks the importance of prices on buyer's surplus. At another extreme, one can set $\hat{v}(p) = \underline{v} + (\overline{v} - \underline{v})\mathbb{1}(p > p_0)$ so that the product is recommended whenever the price is sufficiently low. This option overlooks the importance of value on buyer's surplus. Yet another natural option, which seemingly bypasses the shortcomings of the previous options is to set $\hat{v}(p) = p$ so that the product is recommended whenever the value is above the price. This choice maximizes ex post efficiency but is still not optimal because it underutilizes the opportunity to dampen equilibrium prices. Instead, we show that the designer's problem can be viewed as a screening problem and the optimal threshold responds to price in a manner that depends on the value distribution and the virtual cost function.

We need a few preliminary definitions. Let $\Gamma(c) \triangleq c + F(c)/f(c)$ denote the virtual cost function. Let $\underline{c}$ and $\bar{c}$ be the types such that $\underline{v} = \Gamma(\underline{c})$ and $\overline{v} = \Gamma(\bar{c})$.[4] Define the

[3] Specifically, the algorithm recommends the product only if $v = \overline{v}$, an event that occurs with zero probability.

[4] These types are well-defined because $\Gamma(0) = 0 \leq \underline{v} < \overline{v} \leq 1 \leq \Gamma(1)$.

following pricing strategy:

$$\tilde{p}(c) \triangleq \mathbb{E}_{v \sim G}[\Gamma^{-1}(v) \mid v \geq \Gamma(c)], \tag{1}$$

where the conditional expectation is set to equal $\Gamma^{-1}(\overline{v})$ whenever the conditioning event occurs with probability zero, i.e., when $c \geq \overline{c}$. Thus, because $\Gamma(c)$ is monotonically increasing and v has full support over $[\underline{v}, \overline{v}]$, $\tilde{p}(c)$ is a continuous function that equals $\underline{p} \triangleq \mathbb{E}_{v \sim G}[\Gamma^{-1}(v)]$ for $c < \underline{c}$, strictly increases on $(\underline{c}, \overline{c})$, and equals $\overline{p} \triangleq \Gamma^{-1}(\overline{v})$ for $c \geq \overline{c}$. Define the inverse function of $\tilde{p}$ as $\tilde{p}^{-1}$ with the (nonstandard) convention that $\tilde{p}^{-1}(p) = \underline{c}$ for $p < \underline{p}$ and $\tilde{p}^{-1}(p) = \overline{c}$ for $p > \overline{p}$.

Proposition 1. (Optimal Algorithm)

1. A buyer-optimal algorithm has a threshold function $\hat{v}(p) = \Gamma(\tilde{p}^{-1}(p))$, which strictly and continuously increases on $(\underline{p}, \overline{p})$ with $\hat{v}(\underline{p}) = \underline{v}$ and $\hat{v}(\overline{p}) > \overline{p}$.

2. In equilibrium, type c posts a price $\tilde{p}(c)$ and trades whenever $v \geq \Gamma(c)$. Types $c < \underline{c}$ always trade, whereas types $c \geq \overline{c}$ are inactive.

Proof Outline. We can solve for an optimal algorithm by building on mechanism-design machinery. Even though the algorithm does not administer monetary transfers, its recommendations can depend on the price, and as such, it can calculate and control the seller's expected revenue. The choice of the equilibrium price and threshold then gives the algorithm standard tools to screen different seller types, mirroring the seminal analysis of Baron and Myerson (1982).

The choice of the threshold determines, at the same time, the expected trade surplus, valued by the buyer, and the expected trade volume, valued by the seller. The threshold that optimally trades off efficiency and incentives, when viewed as a function of seller type $\hat{v}(p(c))$, equals virtual costs. This relationship pins down the equilibrium trade volume, and the equilibrium price function (1) guarantees an incentive-compatible profit distribution across types. In turn, this allows us to calculate the optimal threshold as a function of an equilibrium price $\hat{v}(p)$. Finally, the optimal algorithm generates positive buyer surplus at any price, so the buyer is always willing to purchase recommended products. ☐

Proposition 1 fully characterizes a buyer-optimal algorithm and reveals several notable features. First, less efficient types $c > \bar{c}$ are inactive and thus excluded from trade. This happens because the buyer value is bounded from above; if it were unbounded, some values would always be above virtual costs, and all interior types would be active. Second, the optimal algorithm can make two types of *ex post* errors. First, if the product price is high, the algorithm does not recommend the product even when the value is above the price. This error is always a feature of the optimal algorithm. Moreover, when $\underline{v} < \underline{p} = \mathbb{E}_{v \sim G}[\Gamma^{-1}(v)]$, if the product price is low, the algorithm recommends the product even when the value is below the price. In this way, the buyer-optimal algorithm distorts purchasing decisions by both rewarding low prices and punishing high prices, leading to an algorithmic demand that incentivizes the seller to set overall lower prices.

Notably, Proposition 1 decouples the impact of the value and cost distributions. The optimal trade allocation in the space of costs and values depends only on the cost distribution via the virtual cost formula, whereas the optimal prices depend both on the cost and value distribution. This feature facilitates the analysis of the impact of data availability that we conduct in Section 4.

Example 1 (Uniform Distributions). Examine an optimal algorithm in a classic example in which c and v are uniformly distributed on $[0,1]$. In this case, the virtual cost is $\Gamma(c) = 2c$. By Proposition 1, under the optimal algorithm, the equilibrium pricing is $\tilde{p}(c) = \mathbb{E}\left[v/2 | v \geq 2c\right] = (1+2c)/4$ for $c \in [0, 1/2]$. For types $c > 1/2$, $\tilde{p}(c) = 1/2$, so those types are inactive. The recommendation threshold function $\hat{v}(p)$ equals 0 for $p < 1/4$, equals $4p - 1$ for $p \in [1/4, 1/2]$ and equals 1 for $p > 1/2$. A buyer who receives a recommendation at price p infers that the expected value of the recommended product is $(4p - 1 + 1)/2 = 2p$ and is thus willing to purchase it.

The left side of Figure 1 depicts the optimal recommendation threshold (solid line) along with the *ex post* optimal recommendation threshold $\hat{v} = p$ (dashed line). As we discussed above, the *ex ante* optimal algorithm is suboptimal *ex post* in two ways: if the product price is low, i.e., $p < 1/3$, it recommends the product even when the value is

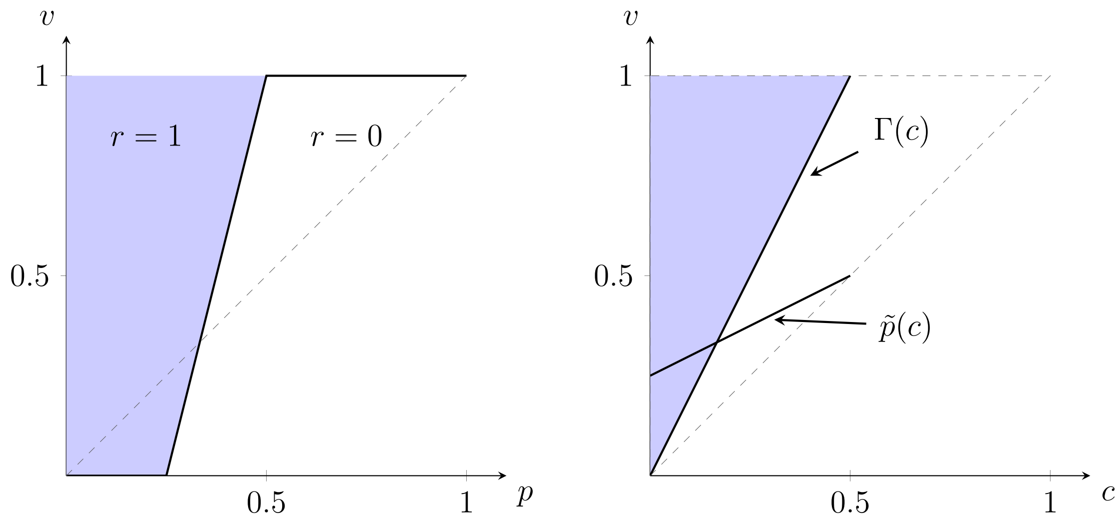

Figure 1: Optimal recommendation algorithm (left) and the resulting equilibrium pricing strategy and trade region (right). $v \sim U[0,1]$, $c \sim U[0,1]$.

below the price, leading to a negative ex post payoff to the buyer. Second, if the product price is high, i.e., $p > 1/3$, the algorithm does not recommend the product even when the value is above the price. The algorithm translates into a piece-wise linear demand curve, which each seller type considers as given when deciding which price to post.

The right side of Figure 1 depicts the resulting equilibrium pricing and trade: the price $\tilde{p}(c)$ posted by the seller of type c, the region of values and types in which the trade occurs (filled area), and the efficient trade region (area encircled by dashed lines). In accordance with Proposition 1, under an optimal algorithm, trade happens whenever the buyer's value is greater than the seller's virtual costs. Type $c = 0$ always trades. All higher types post progressively higher prices and serve progressively fewer buyers. Types $c > 1/2$ never trade. Equilibrium active prices span the interval $[1/4, 1/2]$. $\diamondsuit$

We conclude this section with two remarks on the nature of algorithmic consumption as revealed by our analysis.

Remark 1 (Commitment and Information). Our analysis advances the idea that the use of algorithms can confer a degree of commitment to consumers and shift the bargaining power in their favor. Perhaps the most illustrative case is when the seller's cost is commonly known. In this instance, a buyer-optimal algorithm recommends the product if and only if the price does not exceed the cost and the value is above the price. Such an algorithm induces efficient trade and allocates all trade surplus to the buyer.

When the seller's costs are uncertain, the optimal algorithm becomes more complex and leaves rents to the seller. However, even though the algorithm serves only information, it can achieve the same consumer surplus as would a mechanism with the ability to administer monetary transfers, e.g., one that could charge a referral or a commission fee.

This degree of commitment power is notably stronger than in the setting where the algorithm's recommendation cannot condition on the price and thus the algorithm design becomes pure information design, as studied by Roesler and Szentes (2017) in the case of known seller's cost. In that scenario, control of information grants only limited commitment power to the buyer: although the buyer's value estimate can be strategically manipulated, the buyer unambiguously purchases whenever this estimate is above the price, irrespective of what the price is. As a result, the seller obtains strictly positive rents even when her costs are known.

Remark 2 (Automation and Discovery). Motivated by existing real-world applications, we assumed that the buyer remains "in the loop" on purchasing decisions and makes the final decision based on the given recommendation. However, under the optimal algorithm, the buyer never wants to overturn the recommendation, and his incentive constraints are not binding. Intuitively, the algorithm is designed to maximize the buyer's surplus, and while it makes *ex post* imperfect decisions for some values, it generates nonnegative surplus for the buyer at any price on average. As such, the buyer would be willing to fully delegate the purchasing decisions to such an algorithm, and our analysis also applies to fully automated consumption.

Relatedly, we assumed that the recommendation algorithm not only provides information about the trade value but also aids in discovering the product, i.e., that the buyer cannot buy the product unless it is recommended. While this assumption may play a role in some settings, in others, it is innocuous. Consider, for instance, the setting of Example 1. It can be seen that the same algorithm remains optimal even if the buyer could always purchase the product. Indeed, under the algorithm, no seller type would set a price below 1/4. Furthermore, for any price above 1/4, the buyer would never purchase the product if it is not recommended, because its expected value would be below the price: $2p - 1/2$ for $p \in [1/4, 1/2]$ and $1/2$ for $p > 1/2$.

4 Impact of Data Availability

In this section, we examine how the availability of additional data to either the algorithm or the seller affects the equilibrium outcome under the buyer-optimal algorithm. Specifically, we assume that the seller or the algorithm gains access to a signal of the buyer's valuation and study how the equilibrium depends on the accuracy of the signal. Throughout, to obtain the sharpest results and unify the exposition, we focus on the truth-or-noise signals, although many of our results hold under more general signals:

Definition 2. The *truth-or-noise signal* (*with accuracy* α) is a random variable s such that with probability α, $s = v$, and with probability $1 - \alpha$, $s \sim G$ independent of v.

Truth-or-noise signals are frequently used in the economic literature (see, for example, Lewis and Sappington (1994) and Johnson and Myatt (2006)). The defining feature of a truth-or-noise signal is that it is *ex ante* distributed according to the prior value distribution, G, and the posterior expectations respond to the signal linearly and proportionally to its accuracy: $\mathbb{E}[v|s] = \alpha s + (1-\alpha)v_0$, where v_0 denotes the average value, $v_0 \triangleq \mathbb{E}[v]$.

4.1 More Data to the Algorithm

How does the equilibrium outcome of the buyer-optimal algorithm change as the algorithm gains access to more data? We begin with the observation that the characterization of the buyer-optimal algorithm in Proposition 1 applies equally to the setting in which the algorithm only observes a noisy signal of the buyer's valuation: because both the buyer and seller are risk-neutral, the algorithm's recommendations can, without loss of generality, depend solely on the posterior expectation of the value induced by each signal realization.

As a result, we can find the buyer-optimal algorithm by applying Proposition 1 to the case of $G = G_\alpha$, where G_α is the distribution of posterior expectations induced by the signal (the baseline model corresponds to $\alpha = 1$). The question—how the outcome of the optimal algorithm depends on data—now reduces to comparative statics with respect to G_α. Under a truth-or-noise signal, distribution G_α has support $[\underline{v}_\alpha, \overline{v}_\alpha]$ with $\underline{v}_\alpha \triangleq (1-\alpha)v_0 + \alpha\underline{v}$ and $\overline{v}_\alpha \triangleq (1-\alpha)v_0 + \alpha\overline{v}$ and equals[5]

$$G_\alpha(v) = \Pr(\alpha s + (1-\alpha)v_0 \le v) = G\left(\frac{v}{\alpha} - \frac{(1-\alpha)v_0}{\alpha}\right), \tag{2}$$

so that the distribution of posterior expectations is a scaled version of the distribution of true values (cf. Persico (2000)). As such, the analysis applies equally to any source of increased accuracy in value estimation, be it additional data or improved algorithm architecture.

For a given algorithm and the seller's pricing strategy $p(\cdot)$, define the *individual profit for type* c, $\pi_\alpha(c)$, as the expected payoff of type c:

$$\pi_\alpha(c) \triangleq \mathbb{E}_{v\sim G_\alpha}[(p(c)-c)r(v,p(c)) \mid c]. \tag{3}$$

Proposition 2. (More Data to the Algorithm) *If the algorithm's accuracy α increases from α_L to $\alpha_H > \alpha_L$, then the buyer surplus increases, the individual profit increases for types $c > \hat{c}$ and decreases for types $c < \hat{c}$, the set of active seller types expands from $[0, \overline{v}_L)$ to $[0, \overline{v}_H)$, $\overline{v}_H > \overline{v}_L$, and the prices change as follows:*

1. *The highest active price increases.*
2. *If Γ is convex, then the lowest active price decreases.*
3. *If Γ is concave, then each seller type posts a higher price.*

[5] If $\alpha = 0$, then $G_\alpha(v) = \mathbb{1}(v \ge v_0)$.

Proposition 2 is easiest to understand by comparing the settings with full data (i.e., $\alpha_H = 1$) and no data (i.e., $\alpha_L = 0$). If no data are available, then the posterior value estimate is deterministically equal to prior mean v_0. In this case, the optimal algorithm recommends the product only if a price is below $p^* = \Gamma^{-1}(v_0)$, and in equilibrium, sellers with costs below p^* post and trade at price p^*, while all other types are inactive. Providing full data to the algorithm causes two changes. First, the algorithm can identify buyers with values above v_0 and then recommend the product of high-cost sellers—who would have been excluded in the no-data case—to these buyers. This effect unambiguously increases the highest active price, which by Proposition 1 is equal to $\Gamma^{-1}(\overline{v})$. In turn, this lenience for higher prices propagates through the incentive constraints and increases the information rents and equilibrium prices of lower-type sellers, driving prices upward. Second, symmetrically, the algorithm with full data can identify low-value buyers and reduce trade volume between these buyers and sellers with costs below $\Gamma^{-1}(v_0)$. This effect makes demand more sensitive and places downward pressure on prices.

If virtual cost Γ is convex, then the second effect dominates for sufficiently low types, resulting in a decrease in the lowest active price, which by Proposition 1 is equal to $\mathbb{E}_G[\Gamma^{-1}(v)]$ and thus the expansion of active price range. If the virtual cost is concave, we show that the first effect dominates across all seller types because, roughly, identifying high-value buyers allows the algorithm to accommodate many high-cost seller types.

Example 1 (continued). We revisit our leading example in which v and c are uniformly distributed on $[0, 1]$. By Equation 2, the posterior means induced by a truth-or-noise signal with accuracy α are uniformly distributed over $[\underline{v}_\alpha, \overline{v}_\alpha]$ with $\underline{v}_\alpha = (1 - \alpha)/2$ and $\overline{v} = (1 + \alpha)/2$, so that the support of posterior estimates is symmetric around $v_0 = 1/2$ and its size is equal to α.

By Proposition 1, the equilibrium pricing strategy is

$$p_\alpha(c) = \mathbb{E}_{v \sim U[\underline{v}_\alpha, \overline{v}_\alpha]}\left[\frac{v}{2} \mid v \geq 2c\right] = \begin{cases} \frac{1}{4}, & \text{if} \quad c \in [0, \frac{1-\alpha}{4}], \\ \frac{4c+\alpha+1}{8}, & \text{if} \quad c \in [\frac{1-\alpha}{4}, \frac{1+\alpha}{4}]. \end{cases}$$

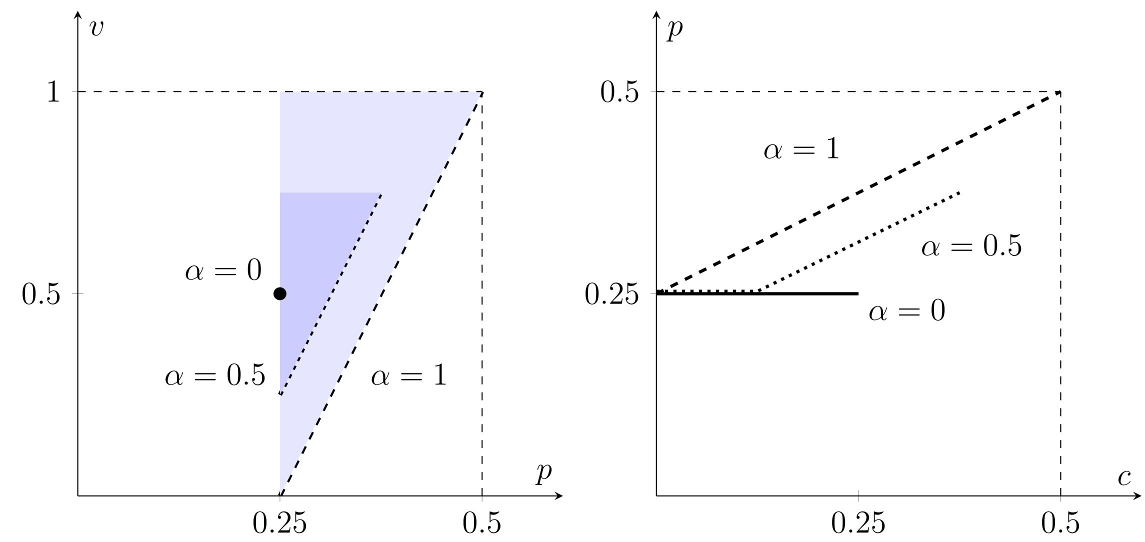

Figure 2: Optimal recommendation algorithm (left) and equilibrium pricing (right) at different algorithm's accuracies, shown at active prices. $v \sim U[0,1]$, $c \sim U[0,1]$.

Types $c > \frac{1+\alpha}{4}$ are inactive. Figure 2 illustrates the range of active prices and equilibrium pricing strategies for $\alpha = 0$, 0.5, and 1. In this example, the virtual costs are linear and thus both concave and convex. As a result, all parts of Proposition 2 apply, and as the algorithm gains access to more data, active prices become pointwise higher and cover a greater range, that is, the distribution of active prices shifts upward in the first-order stochastic sense. Nevertheless, the minimal active price $\mathbb{E}_{G_\alpha}[\Gamma^{-1}(v)]$ remains the same because Γ is linear. $\Diamond$

4.2 More Data to the Seller

We turn to the question of how the outcome of the buyer-optimal algorithm changes when the seller gains access to data about the buyer's value and can engage in third-degree price discrimination, i.e., in personalized pricing. This setting can be viewed as market segmentation, in which the seller has the ability to target different consumer segments with different prices.

To study this question, we extend the setting as follows. As in the baseline model, the algorithm has full information about the buyer's value. However, the seller now observes a public truth-or-noise signal about the value with accuracy α, independent of his type, and the algorithm can make recommendations based on both the realized signal and the value and price of the product.[6] Our baseline model corresponds to $\alpha = 0$.

Because the signal is public and exogenous, our characterization of the buyer-optimal algorithm straightforwardly extends to this setting, as applied to each signal realization. For any given $\alpha \in [0, 1]$, let $G_{\alpha,s}$ denote the posterior distribution of the buyer's value conditional on signal realization s. Under the truth-or-noise signal, $G_{\alpha,s}$ is a weighted average of a point mass $v = s$ and the value distribution G; consequently, the support of the posterior belief is the same for all s.[7] By Proposition 1, under the optimal algorithm, a seller of type c posts price

$$\tilde{p}_{\alpha,s}(c) = \mathbb{E}_{v\sim G_{\alpha,s}}[\Gamma^{-1}(v)|v \geq \Gamma(c)], \tag{4}$$

and the algorithm recommends the product if and only if $v \geq \Gamma(p^{-1}_{\alpha,s}(p))$, i.e., the value exceeds the virtual cost of the seller's type that posts price p.

For a given algorithm and the seller's pricing strategy $p_s(c)$ at each (s, c), define the *individual surplus of buyer with value* v, $w_\alpha(v)$, as the expected payoff of the buyer conditional on his value being v:

$$w_\alpha(v) \triangleq \mathbb{E}_{\alpha,s,c}[(v - p_s(c))r(v, p_s(c)) \mid v], \tag{5}$$

where the expectation is taken with respect to (s, c). We continue to use the same notation $\pi_\alpha(c)$ for individual profit for type c as in the previous section, which is now written as

$$\pi_\alpha(c) = \mathbb{E}_{\alpha,s,v}[(p_s(c) - c)r(v, p_s(c)) \mid c], \tag{6}$$

[6] The assumption that the signal is public captures the idea that the seller has no information beyond that accessed by the algorithm. This assumption would be automatically satisfied if the seller's signal were a deterministic function of valuation, i.e., a partitional signal or full learning.

[7] In contrast to Section 4.1, the whole posterior belief distribution matters, not only its posterior expectation, because the algorithm has access to the true v.

where the expectation is taken with respect to (s, v). The following result examines how the seller's access to data affects the distributions of individual surpluses and profits, i.e., the distribution of $w_\alpha(v)$ with $v \sim G$ and that of $\pi_\alpha(c)$ with $c \sim F$.

Proposition 3 (More Data to the Seller). *Under the buyer-optimal algorithm, as the seller's data accuracy α increases, the individual profits π_α remain unchanged, the distribution of active prices p_α undergoes a mean-preserving spread, and the distribution of individual surpluses w_α undergoes a mean-preserving contraction.*

Proof Outline. Recall that in our baseline model, by Proposition 1, the equilibrium allocation of the product does not depend on the value distribution: regardless, the optimal algorithm executes trade if and only if $v \geq \Gamma(c)$. The same argument applies to each posterior induced by a signal. This means that regardless of α, the optimal algorithm attains the same total surplus, and if viewed as a mechanism, it results in the same mapping from each type to the trade volume. The revenue equivalence theorem then implies that the individual profit of each type does not change either.[8] As a result, buyer surplus also does not change with the seller's data accuracy.

At the same time, the seller's access to data redistributes the individual surplus. For example, when the seller has no information, the seller with cost c posts a price $p_0(c) = \mathbb{E}_{v \sim G}[\Gamma^{-1}(v) | v \geq \Gamma(c)]$ independent of v, and any buyer with value $v \geq \Gamma(c)$ trades at that price. Suppose now that the seller has full information about v. Under the buyer-optimal algorithm, type c posts price $p_1(c) = \Gamma^{-1}(v)$, and the buyer trades at this new price, which is now increasing in the value. As a result, aggregating over sellers with different costs, the seller's data decreases the individual surplus of buyers with high values and increases the individual surplus of buyers with low values, leading to a more equalized surplus distribution. An argument for noisy data has a similar intuition and builds on the analysis of stochastic orders, utilizing the structure of truth-or-noise signals. $\square$

Proposition 3 establishes, in a stark manner, that on average, the seller does not

[8]Note that the seller with the highest cost $c = 1$ never trades.

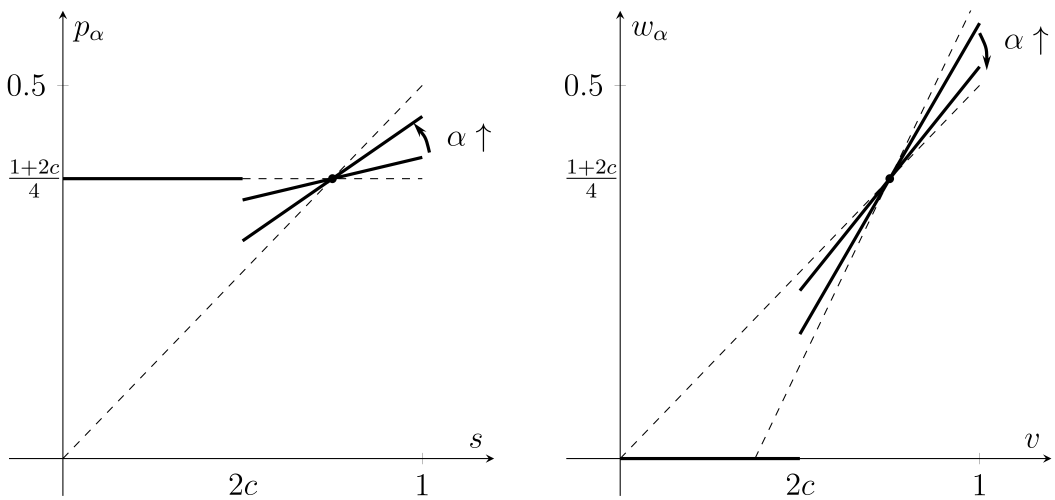

Figure 3: Equilibrium price posted by type c at different signal realizations and individual surplus at different values, shown at two levels of accuracy. $v \sim U[0,1]$, $c \sim U[0,1]$.

benefit from having more information about the buyer's value, and the buyer is not harmed by the release of such information, as long as this release is countered by the algorithm design. Moreover, such an information release may be considered beneficial if the designer prefers a more equal distribution of surplus across buyers.

Example 1 (continued). Consider the leading example in which v and c are uniformly distributed on $[0,1]$. Given a fixed α and a realized signal s, the posterior value distribution $G_{\alpha,s}$ places a mass of α on $v = s$ and a mass of $1-\alpha$ on $v \sim U[0,1]$. As a result, the price posted by active type $c < 1/2$, as presented in Equation 4, is

$$p_{\alpha,s}(c) = \mathbb{E}_{v\sim G_{\alpha,s}}\left[\frac{v}{2}\middle| v \geq 2c\right] = \begin{cases} \frac{1+2c}{4}, & \text{if} \quad s \leq 2c, \\ \alpha\frac{s}{2} + (1-\alpha)\frac{1+2c}{4}, & \text{if} \quad s \geq 2c. \end{cases}$$

Figure 3 depicts the equilibrium price of an active type c as a function of a realized signal s at two levels of accuracy. For a signal realization below $2c$, the seller does not base its price decision on the event that the signal is truth, because any value below $2c$ does not lead to trade. Thus, the seller charges a price based on the prior value distribution, leading to price $\frac{1+2c}{4}$. In contrast, the seller's price increases in a signal realization above $2c$. As the accuracy increases, the price responds more strongly to the realized signal, making the equilibrium price steeper as a function of the realized signal.

To see the implication of this price change on the buyer's individual surplus, let $w_\alpha(v,c)$ denote the buyer's equilibrium payoff conditional on having value v and facing seller type c. A buyer with $v < 2c$ does not trade and thus obtains a payoff of $w_\alpha(v,c) = 0$. A buyer with $v \geq 2c$ trades. Specifically, with probability α, the buyer faces signal realization $s = v \geq 2c$ and pays a price of $\alpha\frac{v}{2} + (1-\alpha)\frac{1+2c}{4}$. With the remaining probability, the signal is drawn from $U[0,1]$, leading to price $\frac{1+2c}{4}$ if $s \leq 2c$ or $\alpha\frac{s}{2} + (1-\alpha)\frac{1+2c}{4}$ if $s \geq 2c$. The resulting buyer's payoff is

$$w_\alpha(v,c) = v\left(1 - \frac{\alpha^2}{2}\right) - \frac{1+2c}{4}(1-\alpha^2).$$

Figure 3 depicts $w_\alpha(v,c)$ as a function of value v. As the accuracy of the signal increases, the equilibrium price is more strongly related to the buyer's value, which, in turn, makes the buyer's payoff less responsive to their willingness to pay. As a result, the seller's access to the signal makes the distribution of surplus $w_\alpha(v,c)$ more equalized across different values. The same property persists when we aggregate surpluses across different seller types, leading to the mean-preserving contraction property of the individual surpluses stated in Proposition 3. $\diamondsuit$

Before concluding, we would like to highlight that Proposition 3 assumes that the algorithm adopts a different recommendation rule for each realized signal. Specifically, while the resulting allocation of the good remains the same for any signal realization, the recommendation threshold function and the seller's equilibrium pricing depend on the posterior value distribution induced by each signal realization. Interpreted in the context of market segmentation, it means that the buyer-optimal recommendations should be personalized at the level of a market segment, so that the algorithm may optimally send different recommendations to buyers in different segments, even if the product price and the estimated trade values are the same.

5 Conclusion

We studied algorithmic decision-making by consumers in a bilateral trade setting. We showed that a buyer-optimal algorithm must strike a balance between increasing trade surplus by informing the buyer about the product and inducing low prices by withholding recommendations for products with high prices. We showed that an increase in algorithmic precision has a systematic impact on welfare outcomes and the distribution of product prices. Furthermore, the optimal algorithm can protect the total consumer surplus from personalized pricing and even use it to reduce surplus distribution inequalities.

We view our work as a stepping stone toward a better understanding of optimal algorithmic design in strategic settings. Within the context of algorithmic consumption, we lay the groundwork for several future research possibilities. For example, we deliberately restricted the buyer's source of information such that the buyer does not learn anything beyond what is provided by the recommendation algorithm. In practice, however, buyers may be able to assess or search for the product on their own. Incorporating these considerations could further enrich the analysis of algorithm design. As another example, we confined our analysis to bilateral trade involving a single product. It could be instructive to explore the implications of algorithmic consumption in broader settings, such as those involving a multiproduct monopolist or competing sellers. The latter extension could complement the studies on strategic steering by online platforms, e.g., those by Hagiu and Jullien (2011), Hagiu et al. (2022), and Bar-Isaac and Shelegia (2022).

References

Akbarpour, M., P. Dworczak, and S. D. Kominers (forth.): "Redistributive Allocation Mechanisms," *Journal of Political Economy.*

Angwin, J., J. Larson, S. Mattu, and L. Kirchner (2016): "Machine Bias:

There's Software Used across the Country to Predict Future Criminals. And It's Biased Against Blacks." *ProPublica*, 23, 77–91.

Bar-Isaac, H. and S. Shelegia (2022): "Monetizing Steering," *Working paper*.

Baron, D. P. and R. B. Myerson (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 911–930.

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

Bergemann, D., T. Heumann, and S. Morris (2022): "Screening with Persuasion," *Working paper*.

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

Deb, R. and A.-K. Roesler (2021): "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Working paper*.

Decarolis, F. and G. Rovigatti (2021): "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising," *American Economic Review*, 111, 3299–3327.

Doval, L. and A. Smolin (forth.): "Persuasion and Welfare," *Journal of Political Economy*.

Dworczak, P., S. D. Kominers, and M. Akbarpour (2021): "Redistribution through Markets," *Econometrica*, 89, 1665–1698.

Gal, M. S. and N. Elkin-Koren (2016): "Algorithmic Consumers," *Harvard Journal of Law and Technology*, 30, 309–353.

Galbraith, J. K. (1952): *American Capitalism: The Concept of Countervailing Power*, Boston: Houghton Mifflin.

Ganuza, J.-J. and J. S. Penalva (2010): "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions," *Econometrica*, 78, 1007–1030.

Gentzkow, M. and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 106, 597–601.

Haghpanah, N. and R. Siegel (2023): "Pareto-Improving Segmentation of Multi-product Markets," *Journal of Political Economy*, 131, 000–000.

Hagiu, A. and B. Jullien (2011): "Why Do Intermediaries Divert Search?" *RAND Journal of Economics*, 42, 337–362.

Hagiu, A., T.-H. Teh, and J. Wright (2022): "Should Platforms Be Allowed to Sell on Their Own Marketplaces?" *RAND Journal of Economics*, 53, 297–327.

Jagtiani, J. and C. Lemieux (2019): "The Roles of Alternative Data and Machine Learning in Fintech Lending: Evidence from the LendingClub Consumer Platform," *Financial Management*, 48, 1009–1029.

Johnson, J. P. and D. P. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 756–784.

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

Kleinberg, J., J. Ludwig, S. Mullainathan, and A. Rambachan (2018): "Algorithmic Fairness," *AEA Papers and Proceedings*, 108, 22–27.

Krishna, V. (2009): *Auction Theory*, Academic Press.

Lee, C. (2021): "Optimal Recommender System Design," *Working paper.*

Lewis, T. R. and D. E. Sappington (1994): "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 309–327.

Libgober, J. and X. Mu (2021): "Informational Robustness in Intertemporal Pricing," *Review of Economic Studies*, 88, 1224–1252.

Liu, Q., K. Mierendorff, X. Shi, and W. Zhong (2019): "Auctions with Limited Commitment," *American Economic Review*, 109, 876–910.

Loertscher, S. and L. M. Marx (2022): "Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration," *American Economic Review*, 112, 616–649.

Milgrom, P. R. and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1089–1122.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Myerson, R. B. and M. A. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

Mylovanov, T. and T. Tröger (2014): "Mechanism Design by an Informed Principal: Private Values with Transferable Utility," *Review of Economic Studies*, 81, 1668–1707.

Obermeyer, Z., B. Powers, C. Vogeli, and S. Mullainathan (2019): "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," *Science*, 366, 447–453.

Persico, N. (2000): "Information Acquisition in Auctions," *Econometrica*, 68, 135–148.

Roesler, A.-K. and B. Szentes (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107, 2072–80.

Scott Morton, F., P. Bouvier, A. Ezrachi, B. Jullien, R. Katz, G. Kimmelman, A. D. Melamed, and J. Morgenstern (2019): "Committee for the Study of Digital Platforms: Market Structure and Antitrust Subcommittee Report," *Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business*, 36.

Shaked, M. and J. G. Shanthikumar (2007): *Stochastic Orders*, Springer.

Yang, K. H. (2022): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," *American Economic Review*, 112, 1364–93.

Appendix

A Omitted Formalism

Proof of Lemma 1 Take any algorithm r. For each $p \geq 0$, let $q_r(p) \triangleq \int_{\underline{v}}^{\overline{v}} r(v,p)dG(v)$ denote the probability with which the product is recommended, and thus purchased, under r. Define a new algorithm $\hat{r}$ as $\hat{r}(v,p) \triangleq \mathbb{1}(v > G^{-1}(1 - q_r(p)))$. At each price p, this algorithm recommends the product with the same probability as r, $1 - G(G^{-1}(1 - q_r(p))) = q_r(p)$. Moreover, the expected value of the product, conditional on recommendation, is greater under $\hat{r}$ than under r. As a result, the buyer will purchase the product whenever it is recommended by $\hat{r}$, and at each price p, the seller will earn the same profit under both r and $\hat{r}$. Therefore, $\hat{r}$ has an equilibrium that attains a greater buyer surplus than r with the same seller profit as r.

Proof of Proposition 1 By the revelation principle, we can study the algorithm design by analyzing direct mechanisms in which the seller reports the type to the designer, and the designer chooses which valuations to allocate to the seller and at which price. Furthermore, by Lemma 1, we can focus on threshold allocations. The designer's problem can thus be stated as:

$$\max_{\hat{v}:[0,1]\to[\underline{v},\overline{v}],\, p:[0,1]\to\mathbb{R}_+} \int_0^1 \int_{\hat{v}(c)}^{\overline{v}} (v - p(c))\, \mathrm{d}G\, \mathrm{d}F, \tag{7}$$

$$\text{s.t.} \int_{\hat{v}(c)}^{\overline{v}} (p(c) - c)\mathrm{d}G \geq \int_{\hat{v}(c')}^{\overline{v}} (p(c') - c)\mathrm{d}G \quad \forall\, c, c' \in [0,1],$$

$$\int_{\hat{v}(c)}^{\overline{v}} (p(c) - c)\mathrm{d}G \geq 0 \quad \forall\, c \in [0,1].$$

The easiest way to solve this problem is to reformulate it in familiar terms. Because the value is continuously distributed, the expected trade probability $q \triangleq \int_{\hat{v}}^{\overline{v}} \mathrm{d}G$ is strictly decreasing in $\hat{v}$, spanning $[0,1]$ as $\hat{v}$ spans $[\underline{v}, \overline{v}]$. Hence, q and v are in a one-to-one relationship, and instead of maximizing over $\hat{v}(c)$, we can maximize over $q(c)$. With a small abuse of notation, denote by $\hat{v}(q)$ the threshold that results in a given q and by $V(q) \triangleq \int_{\hat{v}(q)}^{\overline{v}} v\mathrm{d}G$ the corresponding trade surplus. The trade surplus is strictly increasing in q with $V(0) = 0$ and $V(1) = \mathbb{E}[v]$. Moreover,

$$\frac{\mathrm{d}V}{\mathrm{d}q} = \frac{\partial V/\partial \hat{v}}{\partial q/\partial \hat{v}} = \frac{-\hat{v}g(\hat{v})}{-g(\hat{v})} = \hat{v}(q), \tag{8}$$

and as such, $V(q)$ is a concave function with $V'(0) = \overline{v}$ and $V'(1) = \underline{v}$. Finally, denote the expected revenue by $t(c) \triangleq p(c) \int_{\hat{v}(c)}^{\overline{v}} \mathrm{d}G$. In these variables, we can restate problem (7) as:

$$\max_{q:[0,1]\to[0,1],\, t:[0,1]\to\mathbb{R}_+} \int_0^1 (V(q(c)) - t(c))\, \mathrm{d}F, \tag{9}$$

$$\text{s.t. } t(c) - cq(c) \geq t(c') - cq(c') \quad \forall\, c, c' \in [0,1],$$

$$t(c) - cq(c) \geq 0 \quad \forall\, c \in [0,1].$$

Problem (9) is precisely the problem analyzed by Baron and Myerson (1982) if q is interpreted as a quantity produced and V is interpreted as the welfare generated by producing quantity q. Its celebrated solution sets the optimal quantity to equalize marginal welfare benefits with virtual costs and the optimal transfer to guarantee the incentive-compatible profit distribution:

$$V'(q(c)) = c + \frac{F(c)}{f(c)},$$
$$t(c) - q(c)c = \int_c^1 q(x)\,\mathrm{d}x = \int_c^1 1 - G(\Gamma(x))\,\mathrm{d}x.$$

By Equation 8, we can translate this solution back to problem (7) as

$$\begin{aligned}
\hat{v}(c) &= c + \frac{F(c)}{f(c)}, \\
p(c) &= c + \frac{\int_c^1 1 - G(\Gamma(x))\,\mathrm{d}x}{1 - G(\Gamma(c))} \\
&= c + \frac{\int_c^1 (x-c) g\left(\Gamma(x)\right)\Gamma'(x)\mathrm{d}x}{1 - G\left(\Gamma(c)\right)} \quad \text{(integration by parts)} \\
&= c + \frac{\int_{\Gamma(c)}^{\Gamma(1)} (\Gamma^{-1}(v) - c) g(v)\mathrm{d}v}{1 - G\left(\Gamma(c)\right)} \quad \text{(change of variable with } v = \Gamma(x)\text{)} \\
&= \frac{\int_{\Gamma(c)}^{\Gamma(1)} \Gamma^{-1}(x) g(x)\mathrm{d}x}{1 - G\left(\Gamma(c)\right)} = \mathbb{E}[\Gamma^{-1}(v) | v \geq \Gamma(c)].
\end{aligned}$$

The optimal algorithm must generate positive buyer surplus at any price, i.e., for all active prices $\mathbb{E}[v|v \geq \hat{v}(p)] \geq p$. If this were not the case for some positive measure of prices, all those prices could be excluded from trade by setting $\hat{v}(p) = \overline{v}$, and such a modification would strictly improve buyer surplus, leading to a contradiction. Therefore, the buyer is always willing to purchase the product whenever it is recommended.

Finally, note that we have $\hat{v}(\overline{p}) = \overline{v} > \Gamma^{-1}(\overline{v}) = \overline{p}$, and the seller's incentive compatibility requires that the threshold function strictly increases in the range of active prices. Additionally, we have $\hat{v}(\underline{p}) = \underline{v} < \underline{p}$ if and only if $\underline{v} < \mathbb{E}_v[\Gamma^{-1}(v)]$ by construction. This completes the proof.

Proof of Proposition 2

Step 1. Take $\alpha, \alpha' \in [0, 1]$ with $\alpha' > \alpha$. Under the buyer-optimal algorithm, the buyer surplus is greater at α' than α because the data are comparable in Blackwell's order, so that the algorithm can garble the input to replicate the value distribution G_α from $G_{\alpha'}$.

Step 2. In the proof of Proposition 1, we showed that the individual profit of type c is

$$\pi_\alpha(c) = \int_c^1 1 - G_\alpha(\Gamma(x))\mathrm{d}x. \tag{10}$$

Suppose that for some c, $\pi_{\alpha'}(c) \geq \pi_\alpha(c)$, i.e.,

$$\int_c^1 1 - G_{\alpha'}(\Gamma(x))\mathrm{d}x \geq \int_c^1 1 - G_\alpha(\Gamma(x))\mathrm{d}x. \tag{11}$$

We show that inequality (11) continues to hold for all $c' \geq c$, which implies the threshold structure of seller gains.

Note that by Equation 2, $G_{\alpha'}(x) \geq G_\alpha(x)$ if $x \leq v_0$ and $G_{\alpha'}(x) \leq G_\alpha(x)$ if $x \geq v_0$, i.e., G_α rotates clockwise around $(v_0, G_\alpha(v_0))$ as accuracy α increases. This implies that $G_{\alpha'}(\Gamma(c)) \geq G_\alpha(\Gamma(c))$ if $c \leq \Gamma^{-1}(v_0)$ and $G_{\alpha'}(\Gamma(c)) \leq G_\alpha(\Gamma(c))$ if $c \geq \Gamma^{-1}(v_0)$.

Functions $\pi_{\alpha'}$ and π_α have the following properties: First, for any $c \leq \Gamma^{-1}(v_0)$, we have $\mathrm{d}\pi_{\alpha'}/\mathrm{d}c = -1 + G_{\alpha'}(\Gamma(c)) \geq -1 + G_\alpha(\Gamma(c)) = \mathrm{d}\pi_\alpha/\mathrm{d}c$. Second, for any $c \geq \Gamma^{-1}(v_0)$, we have $\pi_{\alpha'}(c) \geq \pi_\alpha(c)$, because $G_{\alpha'}(\Gamma(x)) \leq G_\alpha(\Gamma(x))$ for all $x \geq c \geq \Gamma^{-1}(v_0)$. Thus, $\pi_{\alpha'}$ increases faster than π_α at $c \leq \Gamma^{-1}(v_0)$ and is greater than π_α at $c \geq \Gamma^{-1}(v_0)$. Therefore, $\pi_{\alpha'}(c) \geq \pi_\alpha(c)$ implies $\pi_{\alpha'}(c') \geq \pi_\alpha(c')$ for all $c' \geq c$.

Step 3. As $\Gamma(c)$ is increasing, the highest active price $p(\bar{c}) = \Gamma^{-1}(\bar{v}_\alpha)$ increases in α because $\bar{v}_\alpha = \alpha\bar{v} + (1-\alpha)v_0$ increases in α. Furthermore, if Γ is convex, then Γ^{-1} is concave and increasing. By concavity and Jensen's inequality, the lowest active price $p(\underline{c}) = \mathbb{E}_{v\sim G_\alpha}[\Gamma^{-1}(v)]$ decreases in α.

If Γ is concave, then Γ^{-1} is convex and increasing. By Proposition 1, an active type c posts a price $p_\alpha(c) = \mathbb{E}_{v\sim G_\alpha}[\Gamma^{-1}(v)|v \geq \Gamma(c)]$. To show that this price increases in α, we use the following definitions and results on stochastic orders, as presented by Shaked and Shanthikumar (2007). For a random variable X with CDF F_X define $\bar{t}_X \triangleq \sup\{t : F_X(t) < 1\}$ and for any $t < \bar{t}_X$, let $[X|X > t]$ denote a random variable that has as its distribution the conditional distribution of X given $X > t$. In addition, for any $t \in \mathbb{R}$ define a *mean residual life* $m_X(t)$ as

$$m_X(t) = \begin{cases} \mathbb{E}[X - t | X > t], & \text{for} \quad t < \bar{t}_X, \\ 0, & \text{o/w.} \end{cases}$$

Furthermore, for any two random variables X and Y, say that Y *is greater than* X *in the mean residual life order*, denoted as $Y \geq_{mrl} X$, if $m_Y(t) \geq m_X(t)$ for all t. Say that Y *is greater than* X *in the increasing convex order*, denoted as $Y \geq_{icx} X$, if $\int \phi \mathrm{d}F_Y \geq \int \phi \mathrm{d}F_X$ for any increasing convex function ϕ. Finally, for $\gamma \in [0, 1]$ denote by $\gamma Y + (1 - \gamma)X$ a random variable that equals Y with probability γ and equals X with probability $1 - \gamma$.

The following results are Theorem 2.A.18 and Theorem 4.A.24 of Shaked and Shanthikumar (2007).

Claim 1. *If* $Y \geq_{mrl} X$*, then for any* $\gamma \in [0, 1]$*,* $Y \geq_{mrl} \gamma Y + (1 - \gamma)X \geq_{mrl} X$*.*

Claim 2. $Y \geq_{mrl} X$ *if and only if* $[Y|Y > t] \geq_{icx} [X|X > t], \forall\, t < \min(\bar{t}_X, \bar{t}_Y)$*.*

Now, set $X = \alpha v + (1 - \alpha)v_0$, $Y = \alpha' v + (1 - \alpha')v_0$. We show that $Y \geq_{mrl} X$, i.e., that the value estimate under accuracy α' is greater than the value estimate under accuracy α in the mean residual life order. Indeed, X is a convex combination of Y and a constant random variable $Z \equiv v_0$. Moreover, $Y \geq_{mrl} Z$, because $m_Y(t) = \mathbb{E}[Y|Y > t] - t$ for $t < \alpha\bar{v} + (1 - \alpha)v_0$ and 0 otherwise for Y, and $m_Z(t) = v_0 - t$ for $t < v_0$ and 0 otherwise, so $m_Y(t) \geq m_Z(t)$ for all t. By Claim 1, $Y \geq_{mrl} X$.

It follows by Claim 2 that if type c is active under both α and α', then it charges a greater price under α' because $\mathbb{E}_{v \sim G_{\alpha'}}[\Gamma^{-1}(v)|v \geq \Gamma(c)] \geq \mathbb{E}_{v \sim G_\alpha}[\Gamma^{-1}(v)|v \geq \Gamma(c)]$. The only other alternative is that type c is active under α' but not under α, i.e., $\alpha\bar{v} + (1 - \alpha)v_0 < \Gamma(c)$ and $\alpha' v + (1 - \alpha')v_0 > \Gamma(c)$. However, any price that is active under α' and not active under α is greater than any price that is active under both α and α'. Therefore, according to Equation 1, type c charges a higher price as well in that case.

Proof of Proposition 3

Step 1. We show that the buyer surplus and the individual profits do not depend on the accuracy α of the seller's signal. Regardless of the realized signal, the optimal algorithm recommends trade if and only if $v \geq \Gamma(c)$. Hence, the total surplus is independent of α. This also means that from an ex ante perspective, for any α, the optimal algorithm results in the same mapping from each type to the trade volume (i.e., type c produces $1 - G(\Gamma(c))$). Additionally, the highest cost type $c = 1$ is always inactive. Thus for any α, the buyer-optimal algorithm, as an indirect mechanism, attains the same allocation rule and the profit of the highest-cost seller. The revenue equivalence theorem then implies that the seller's individual profit is independent of α (Myerson, 1981; Krishna, 2009).

Step 2. We now show that the seller's information makes the distribution of active prices undertake a mean-preserving spread. Fix any active type c, and take two accuracies, $\alpha_H, \alpha_L \in [0, 1]$, with $\alpha_H > \alpha_L$. Let $G^c_{\alpha,s} \in \Delta[\underline{v}, \overline{v}]$ denote the posterior distribution of the buyer's value conditional on (i) signal s being realized under accuracy α and (ii) $v \geq \Gamma(c)$. The equilibrium price of type c after observing signal s with accuracy α is

$$p(c|s, \alpha) = \mathbb{E}_{v \sim G_{\alpha,s}}[\Gamma^{-1}(v)|v \geq \Gamma(c)] = \int_{\underline{v}}^{\overline{v}} \Gamma^{-1}(v) \mathrm{d}G^c_{\alpha,s}(v). \tag{12}$$

Let $\mathcal{G}^c_\alpha \in \Delta\Delta[\underline{v}, \overline{v}]$ denote the distribution of these posteriors. The (truth-or-noise) signal with accuracy α_H is Blackwell more informative than the signal with accuracy α_L. Therefore, $\mathcal{G}^c_{\alpha_H}$ is a mean-preserving spread of $\mathcal{G}^c_{\alpha_L}$, and in turn, $p(c|s, \alpha_H)$ is a mean-preserving spread of $p(c|s, \alpha_L)$ if viewed as random variables generated by s. This relationship holds for any given type c, and the mean-preserving spread relationship is closed under mixtures (e.g., Theorem 3.A.12(b) of Shaked and Shanthikumar (2007)). As a result, the ex ante prices under accuracy α_H are also a mean-preserving spread of the ex ante prices under accuracy α_L.

Step 3. We turn to the distribution of individual surpluses. First, we calculate the buyer's surplus for a given pair of (v, c). Because signal s is drawn from the truth-or-noise signal with accuracy α, posterior $G_{\alpha,s}$ places probability α on $v = s$ and probability

$1-\alpha$ on the event that $v \sim G$. Thus, Equation 12 can be expanded as:

$$p(c|s,\alpha) = \begin{cases} \mathbb{E}_{v\sim G}[\Gamma^{-1}(v)|v \geq \Gamma(c)], & \text{if} \quad s < \Gamma(c), \\ \alpha\Gamma^{-1}(s) + (1-\alpha)\mathbb{E}_{v\sim G}[\Gamma^{-1}(v)|v \geq \Gamma(c)], & \text{if} \quad s \geq \Gamma(c). \end{cases}$$

If the buyer has value $v < \Gamma(c)$, then in equilibrium the algorithm never recommends the product, and the buyer's expected payoff is nil.

If the buyer has value $v > \Gamma(c)$, then the realized signal is v with probability α and is an independent draw from G with probability $1-\alpha$, leading to the buyer's expected payoff:

$$\begin{aligned} w(v,c,\alpha) &= v - \alpha\left(\alpha\Gamma^{-1}(v) + (1-\alpha)\mathbb{E}_{\tilde{v}\sim G}[\Gamma^{-1}(\tilde{v})|\tilde{v} \geq \Gamma(c)]\right) - (1-\alpha)\mathbb{E}_{s\sim G}[p(c|s,\alpha)] \\ &= v - \alpha^2\Gamma^{-1}(v) - (1-\alpha^2)\mathbb{E}_{\tilde{v}\sim G}[\Gamma^{-1}(\tilde{v})|\tilde{v} \geq \Gamma(c)]. \end{aligned}$$

Only the first two terms in this expression depend on v. We can now view $w(v,c,\alpha_H)$ and $w(v,c,\alpha_L)$ as transformations of random variable $v \sim G(\cdot|v \geq \Gamma(c))$ with two properties. First, $w(v,c,\alpha_H)$ and $w(v,c,\alpha_L)$ have the same mean under $G(\cdot|v \geq \Gamma(c))$, i.e., $\mathbb{E}_{v\sim G}[w(v,c,\alpha_H)|v \geq \Gamma(c)] = \mathbb{E}_{v\sim G}[w(v,c,\alpha_L)|v \geq \Gamma(c)]$. The reason is as follows: Between accuracies α_H and α_L, the interim profit of type c and the allocation of the product remain the same (as shown above). Moreover, the buyer with any value $v < \Gamma(c)$ obtains zero payoffs. Hence, the buyer surplus conditional on $v \geq \Gamma(c)$ must be equal between α_H and α_L.

Second, the cumulative distribution function of $w(v,c,\alpha_H)$ crosses that of $w(v,c,\alpha_L)$ once and from below. To see this, note that because $\alpha_H > \alpha_L$ and Γ is monotonically increasing, we have

$$\frac{\partial}{\partial v}w(v,c,\alpha_L) = 1 - \alpha_L^2\frac{\partial}{\partial v}\Gamma^{-1}(v) > 1 - \alpha_H^2\frac{\partial}{\partial v}\Gamma^{-1}(v) = \frac{\partial}{\partial v}w(v,c,\alpha_H) \geq 0, \qquad (13)$$

where the last inequality holds due to the monotonicity of the hazard rate:

$$1-\alpha^2\frac{\partial}{\partial v}\Gamma^{-1}(v)\geq 1-\frac{\partial}{\partial v}\Gamma^{-1}(v)=1-\frac{1}{\Gamma'\left(\Gamma^{-1}(v)\right)}=1-\frac{1}{1+\left(\frac{F}{f}\right)'}\geq 0.$$

Inequality (13) and the equal mean property imply that $w(\underline{v},c,\alpha_L)<w(\underline{v},c,\alpha_H)$ and $w(\overline{v},c,\alpha_L)>w(\overline{v},c,\alpha_H)$. For $i\in\{H,L\}$, let J_i denote the CDF of a random variable $w(v,c,\alpha_i)$ (recall that c is fixed here). Let $[\underline{w}_i,\overline{w}_i]$ be the range of $w(v,c,\alpha_i)$, and define $w^{-1}(x,\alpha_i)\triangleq\max\{v\in[\underline{v},\bar{v}]:w(v,c,\alpha)=x\}$, which is well-defined on $[\underline{w}_i,\overline{w}_i]$. Then, we have

$$J_i(x)=\begin{cases}0, & \text{if}\quad x<\underline{w}_i,\\ G(w^{-1}(x,\alpha_i)|x\geq\Gamma(c)), & \text{if}\quad \underline{w}_i\leq x\leq\overline{w}_i,\\ 1, & \text{if}\quad \overline{w}_i<x.\end{cases}$$

Because $w(v,c,\alpha_L)$ crosses $w(v,c,\alpha_H)$ once and from below as a function of v, $J_H(x)$ crosses $J_L(x)$ once and from below.

In turn, these equal mean and single-crossing properties imply, by Theorem 3.A.44 (Condition 3.A.59) of Shaked and Shanthikumar (2007), that $w(v,c,\alpha_H)$ is a mean-preserving spread of $w(v,c,\alpha_L)$.

We showed that for a fixed (v,c), the individual surpluses of values $v\geq\Gamma(c)$ under accuracy α_H are a mean-preserving spread of those individual surpluses under accuracy $\alpha_L<\alpha_H$. The same relationship trivially holds for individual surpluses of values $v<\Gamma(c)$, because those types do not trade. In other words, for any fixed c, the buyer surplus $w(v,c,\alpha_H)$ is a mean-preserving spread of $w(v,c,\alpha_L)$, both conditional on $v<\Gamma(c)$ and $v\geq\Gamma(c)$. Because those two scenarios are mutually exhaustive, it follows that for any fixed c, the distribution of the buyer's surplus under accuracy α_H is a mean-preserving spread of distribution of the buyer's surplus under accuracy α_L. Because this relationship holds for each c and the mean-preserving spread relationship is closed under mixtures, the same property holds *ex ante*.