Buyer-Optimal Algorithmic Consumption*

Shota Ichihashi Alex Smolin

April 17, 2024

Abstract

We study a bilateral trade model in which a product is recommended to a buyer by an algorithm, based on the product's value to the buyer and its price. We fully characterize an algorithm that maximizes the buyer's ex ante payoff and show that it strategically biases consumption to incentivize lower prices. Under the optimal algorithmic consumption, informing the seller about the buyer's value does not change the buyer's ex ante payoff but leads to a more equitable distribution of interim payoffs.

*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. For valuable suggestions and discussions, we would like to thank Heski Bar-Isaac, David Martimort, Matthew Mitchell, Alessandro Pavan, as well as seminar participants at Carnegie Mellon University, Compass Lexecon (London), EODS Symposium UPenn 2023, Northwestern–TSE IO Workshop 2023, and Digital Economics Conference TSE 2024. The authors acknowledge support from the CTIC grant and the NET Institute grant. 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.

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). Algorithms negotiate with suppliers (Van Hoek and Lacity (2023)) and bid in online advertising auctions (Balseiro et al. (2021)). 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 financial securities, smart devices that control electricity use, and price-trackers that seek and pinpoint lower-priced products.

In this paper, we ask how algorithmic consumption may empower consumers. To answer this question, we characterize an algorithm that maximizes consumer surplus in a canonical bilateral trade setting and study how such an algorithm interacts with price discrimination. Our model captures three features of algorithmic consumption. First, an algorithm processes consumer and product data, facilitating the learning about trade value. Second, an algorithm can base recommendations on price.[1] Third, the product price is strategically chosen by a seller, who optimally responds to the algorithmic demand.

Specifically, we study the following model. A buyer and a seller can trade a single product, with both the trade cost and trade value being uncertain. The seller privately knows the cost, which constitutes her type. Initially, the buyer knows neither the value nor 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 Bayesian 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.

[1] The dependence of information on price can be programmed in directly, as seen in Amazon's search ranking algorithms (Lee and Musolff (2021), Farronato et al. (2023)) or it can arise indirectly through consumer feedback technology (Luca and Reshef (2021), Chakraborty et al. (2022)), wherein higher prices, all else being equal, lead to lower consumer satisfaction and ratings.

For any demand curve, the seller sets a price to maximize her profit.

We characterize a buyer-optimal algorithm, i.e., an algorithm that maximizes the buyer's expected payoff. Such an algorithm must balance a tradeoff. On the one hand, to incentivize the seller to lower the price, the algorithm should reward low prices by recommending the product more often, and it should punish high prices by recommending the product less often. On the other hand, the algorithm should strive to maximize the benefits of a trade at any give price and not forego beneficial trade opportunities.

We show that this trade-off is optimally resolved by an algorithm with a threshold structure, which recommends the product when its value reaches a threshold that varies with the price. Thus, we can recast the algorithm design as the design of a threshold function, which we approach and solve as a nonlinear screening problem.

We fully characterize the buyer-optimal algorithm and establish its two key features (Proposition 1). First, consistent with the strategic aspect of design, the algorithm is *biased* relative to the ex post optimal algorithm: at high prices, the buyer-optimal algorithm does not recommend the product even when the values exceeds the price; at low prices, the algorithm recommends the product even when the value is below the price. These ex post mistakes render the buyer's demand more price elastic and incentivize the seller to lower prices across different costs. This finding highlights the importance of strategic context for an algorithm assessment and thus contributes to the ongoing debate on AI regulation (e.g., European Commission (2021); Biden (2023)).

Second, the buyer-optimal algorithm attains the same outcome as if the buyer had full commitment power and designed a mechanism with arbitrary monetary transfers. That is, even though the algorithm serves only information, it effectively shifts market power to the buyer. This observation offers a novel perspective on the role of the buyer's information in bilateral trade: When the information policy can be contingent on prices, it can serve as an indirect optimal mechanism for the monopsony's mechanism design problem, thus delivering notably stronger benefits to the buyer than price-independent learning (cf. Roesler and Szentes (2017), and see further Section 3.1).

Finally, we show that such an algorithmic transfer of market power changes the welfare implications of third-degree price discrimination. Specifically, we assume that the seller can observe an informative signal about the buyer's valuation, representing a consumer segment. The seller then sets a price against an algorithm, buyer-optimal in that segment. We show that the buyer's ex ante expected payoff does not depend on the signal structure, i.e., on the granularity of market segmentation (Proposition 2). Moreover, within a natural class of signals, a more informative signal, i.e., a finer market segmentation, results in a mean-preserving contraction of surplus among buyers with different values. Intuitively, informing the seller about the buyer's value 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.

We conclude that algorithmic consumption may not only protect consumers against price discrimination but could also exploit such discrimination to realize the societal benefits of fairness and equality. This finding suggests that promoting algorithmic consumption may be a powerful consumer protection policy, complementary to the existing regulatory methods detailed, for example, by Scott Morton et al. (2019).

Related literature.— Our paper contributes to a recent and rapidly growing literature on the economics of algorithmic decisions. The large focus of this literature has been on algorithmic pricing in competitive settings, either empirical (Calvano et al. (2020), Assad et al. (forth.)), theoretical (Salcedo (2015), Lamba and Zhuk (2023)), or both (Brown and MacKay (2023), Johnson et al. (2023)). This literature largely investigates whether and how algorithms can empower sellers by enlarging their collusion opportunities. We complement this literature by examining the other side of the market and asking whether and how algorithms can empower buyers.

Specifically, we show that algorithmic consumption can deliver a countervailing power in the spirit of Galbraith (1952) to buyers by giving them a stronger bargaining position vis-a-vis sellers.[2] In fact, out setting can be viewed as enabling a buyer from the classic setting of Myerson and Satterthwaite (1983) to commit to values and prices at which she

[2]Thus, algorithmic consumption can be viewed as an effective alternative to the 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).

would be purchasing a product. In this sense, we proceed in the opposite direction from the literature on limited commitment, which 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), Liu et al. (2019)).

Methodologically, our paper belongs to the recent strand of economic literature that examines methods of empowering buyers in monopolistic settings via information control. Roesler and Szentes (2017) analyze buyer-optimal learning in a bilateral trade setting. 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 and Bergemann et al. (2023) to auctions; Condorelli and Szentes (2020) analyze the buyer-optimal distribution of values within a given interval. We contribute to this literature by introducing seller heterogeneity[3] and by allowing the buyer's information, and thus posterior value distribution, depend on price, which are natural features of the online economy.[4]

Finally, our analysis provides a novel perspective on the classic question of monopolistic price discrimination based on consumer information studied in the literature on market segmentation (e.g., Bergemann et al. (2015), Yang (2022), Haghpanah and Siegel (2023), Ichihashi and Smolin (2023)). We show that the use of algorithms by consumers may introduce a new welfare implication whereby price discrimination attains a more equal distribution of consumer surplus without affecting average welfare outcomes. This finding also contributes to the recent literature that explores ways to promote equality and fairness through mechanism design (Kleinberg et al. (2018), Dworczak et al. (2021), Akbarpour et al. (forth.)) or information design (Doval and Smolin (forth.)), by putting forward algorithmic consumption as a means to mitigate the effects of price

[3]Thus, we combine the machinery of Bayesian persuasion (e.g., Kamenica and Gentzkow (2011)) with that of mechanism design (e.g., Baron and Myerson (1982)). Several other papers have combined these machineries in trade settings, typically to study revenue maximization, most recently including Lee (2021), Bergemann et al. (2022), Yang (2022), and Smolin (forth.).

[4]The dependence of information on price may also arise from a worst-case analysis, as in the work of Libgober and Mu (2021), in which the buyer's information is chosen to minimize the seller's profits.

discrimination on consumers and achieve a more equitable 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 her private *type*. The type distribution F has support $[0,1]$, positive density f, and continuous and increasing F/f. The value of the product to the buyer is $v \sim G$ and independent of the seller's type. The value distribution G has positive density g over its support $[0,1]$.

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. Specifically, an algorithm is a function $r : [0,1] \times \mathbb{R}_+ \to [0,1]$ such that for any pair (v,p) of a realized value $v \in [0,1]$ 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

$$\mathbb{E}[v \mid \text{recommended}, p] = \frac{\int_0^1 x r(x,p) g(x) \mathrm{d}x}{\int_0^1 r(x,p) g(x) \mathrm{d}x},$$

and then purchases the product whenever this value exceeds the price. A pair of an algorithm and a buyer's strategy induces 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 ex ante 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 seller types who trade and those who do not under a given algorithm and their posted 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 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. Thus, an optimal algorithm should maximize buyer surplus conditional on trade volume, prioritizing buyers with higher values. This observation enables us to identify a tractable class of algorithms.

Specifically, we say that an algorithm r is a *threshold algorithm* if there exists a *threshold function* $\hat{v} : \mathbb{R}_+ \to [0, 1]$ 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 exceeds a price-dependent threshold, and with probability 0 otherwise. If $\hat{v}(p) = 1$ at some price p, then the algorithm never recommends the product at price p.[5] If $\hat{v}(p) = 0$, then the algorithm always recommends the product at price p.

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 that threshold algorithms span a Pareto frontier in the space of buyer surplus and seller profit. In particular, if there exists a buyer-optimal algorithm, then it can be found in the class of threshold algorithms, and in what follows, we focus on threshold algorithms.

The optimal choice of a threshold function must balance the trade-off between maximizing trade surplus and incentivizing the seller to lower the price. One option is to set $\hat{v}(p) \equiv v_0$ so that regardless of the price, the product is recommended whenever the value is sufficiently high. This algorithm ignores the impact of algorithm design on the seller's pricing. At another extreme, one can set $\hat{v}(p) = \mathbb{1}(p > p_0)$ so that the product is recommended whenever the price is below some cutoff p_0. This algorithm overlooks the importance of values for buyer surplus. Another natural option is to set $\hat{v}(p) = p$ so that the product is recommended if and only if the value exceeds the price. This *ex post optimal algorithm* maximizes the buyer's payoff given fixed prices. However, the algorithm fails to maximize buyer surplus, because it underuses the opportunity to dampen equilibrium prices.

We can cast the designer's problem as a screening problem in which the optimal threshold responds to a price in a manner that depends on the value distribution and the virtual cost function. Let $\Gamma(c) \triangleq c + F(c)/f(c)$ denote the virtual cost function. Define $\bar{c}$ as $\Gamma(\bar{c}) = 1$.[6] Define the following pricing strategy:

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

[5] Precisely, the algorithm recommends the product only if $v = 1$, an event that occurs with zero probability.

[6] Type $\bar{c}$ exists and is unique because $\Gamma(\cdot)$ is strictly increasing and continuous on $[0, 1]$ and $\Gamma(0) = 0 < 1 \leq \Gamma(1)$.

where the conditional expectation is set to equal $\Gamma^{-1}(1)$ whenever the conditioning event occurs with probability zero, i.e., when $c \geq \overline{c}$. Because $\Gamma(c)$ is increasing and v has full support over $[0,1]$, $\tilde{p}(c)$ is a continuous function that equals $\underline{p} \triangleq \mathbb{E}_{v\sim G}[\Gamma^{-1}(v)]$ at $c = 0$, strictly increases on $(0, \overline{c})$, and equals $\overline{p} \triangleq \Gamma^{-1}(1)$ for $c \geq \overline{c}$. Define the inverse function of $\tilde{p}$ as $\tilde{p}^{-1}$ with the (nonstandard) convention that $\tilde{p}^{-1}(p) = \overline{c}$ for $p > \overline{p}$.

Proposition 1. (Buyer-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 $\hat{v}(\underline{p}) = 0$ and $\hat{v}(\overline{p}) > \overline{p}$.

2. In equilibrium, type $c \in [0, \overline{c})$ posts a price $\tilde{p}(c)$ and trades whenever $v \geq \Gamma(c)$. 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 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 simultaneously determines 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 cost. 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 reveals three notable features. First, less efficient types $c > \overline{c}$ 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 impact of the value and cost distributions can be decoupled. The optimal trade allocation in the space of costs and values, $v \geq \Gamma(c)$, 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 in Section 4.

Finally, the optimal algorithm makes two types of *ex post* errors: If the product price is high and close to $\overline{p}$, we have $\hat{v}(p) > p$, so the algorithm does not recommend the product even when the value is above the price; if the product price is low and close to $\underline{p}$, we have $\hat{v}(p) < p$, so the algorithm recommends the product even when the value is below the price. Thus compared to the ex post optimal algorithm, the buyer-optimal algorithm distorts purchasing decisions by rewarding low prices and punishing high prices, leading to an algorithmic demand that incentivizes the seller to set lower prices.

The following result strengthens the last observation—that the optimal algorithm differs from the ex post optimal algorithm at low and high prices. Within a flexible class of distributions, the buyer-optimal algorithm differs from the ex post optimal algorithm for almost all active prices, inducing over-consumption at low prices and under-consumption at high prices.

Corollary 1. *Assume that $F(c) = c^{\alpha}$ and $G(v) = v^{\beta}$ for some $\alpha, \beta > 0$. Then, under the buyer-optimal algorithm, a price $p^* \in (\underline{p}, \overline{p})$ exists such that $\hat{v}(p) < p$ for any $p \in [\underline{p}, p^*)$ and $\hat{v}(p) > p$ for any $p \in (p^*, \overline{p}]$.*

Example 1 (Uniform). We examine the buyer-optimal algorithm in a canonical case in which c and v are uniformly distributed on $[0, 1]$, which corresponds to setting $\alpha = \beta = 1$ in Corollary 1. The virtual cost is $\Gamma(c) = 2c$. By Proposition 1, under the optimal algorithm, the equilibrium pricing is $\tilde{p}(c) = \mathbb{E}[v/2 | v \geq 2c] = (1 + 2c)/4$ 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

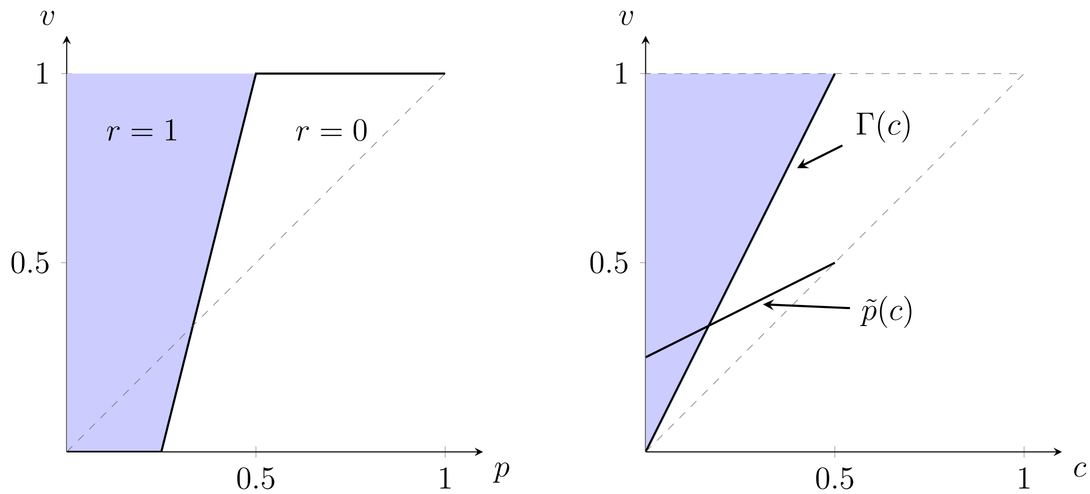

Figure 1: Optimal recommendation algorithm (left) and the resulting equilibrium pricing strategy and trade region (right). $v \sim U[0,1]$, $c \sim U[0,1]$.

value of the recommended product is $(4p-1+1)/2 = 2p > p$ 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 and formalized in Corollary 1, 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 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 piecewise 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 cost. 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 $[1/4, 1/2]$. $\diamondsuit$

3.1 Commitment and Information

The buyer-optimal algorithm combines commitment and information: By disclosing information about a product in a predetermined way, the algorithm enables the buyer to follow a certain demand schedule and obtain a higher surplus than in the standard monopoly setting. In this section, we further investigate the optimal algorithm as a tool to provide the buyer with commitment power and information.

First, the optimal algorithm in Proposition 1 attains the same outcome as when the buyer has full commitment power. Indeed, our proof reveals that the algorithm is an indirect implementation of an optimal mechanism in which the buyer can commit to a mechanism that determines the allocation of the product and monetary transfer based on the seller's reported cost and the buyer's true value. As a result, even though the algorithm serves only information, it effectively transfers market power from the seller to the buyer. Consequently, extending the framework by enabling the algorithm to charge a referral or a commission fee to the seller cannot increase buyer surplus.

This observation also implies that the same algorithm remains optimal in the case of a fully automated trade, e.g., if it could execute transactions without having the buyer in the loop. Formally, suppose now that, instead of giving a recommendation, the algorithm executes the transaction on behalf of the buyer with probability $r(v, p)$. In principle, the optimal algorithm in this setup may attain a higher buyer surplus than in our model, because it does not need to respect the buyer's incentive to follow the recommendation. However, the optimal algorithm in this scenario is the same as that described in Proposition 1, because that algorithm already achieves the buyer's full commitment outcome.

Finally, Proposition 1 remains relevant even in the *value disclosure setup*, which is a model where the buyer is initially uninformed about v but knows the existence of the product, so that the buyer can purchase the product even if the product is not recommended. In general, compared to our model, in the value disclosure setup the optimal algorithm may differ and the buyer may obtain a lower payoff. For example, suppose that $c = 0$ with probability 1. In our model, the buyer can extract the efficient surplus by the algorithm that recommends the product only at $p = 0$. In contrast, in the value disclosure setup, this algorithm would not be able to dissuade the buyer from purchasing when the seller sets a low positive price, so the seller would necessarily earn a positive profit. At the same time, in the uniform setting of Example 1 and in many other cases, the optimal algorithm and the resulting outcome in the value disclosure setup coincide with those in Proposition 1.

Claim 1. *If* $\int_0^{\Gamma(c)}[v - c]\mathrm{d}G(v) \leq 0$ *for each* $c \in [0, \bar{c}]$*, then the buyer-optimal algorithm in the value disclosure setup coincides with the algorithm in Proposition 1.*

The condition of Claim 1 ensures that whenever the product is not recommended in the value disclosure setup, the buyer infers that the expected value of the product is below the price and thus chooses not to buy it. To see the intuition, suppose that the seller posts a price of $\tilde{p}(c)$, and the algorithm recommends the buyer to not purchase the product, which by Proposition 1, reveals that $v \leq \Gamma(c)$. If the buyer purchases the product, his payoff must necessarily decrease, because the seller's profit increases but total surplus decreases because of $\int_0^{\Gamma(c)}[v - c]\mathrm{d}G(v) \leq 0$. Thus, the buyer is willing to follow the recommendation to not buy. Because the buyer is originally willing to follow the recommendation to buy under the buyer-optimal algorithm, we conclude that the buyer always follows the recommendation in the value disclosure setup.

In a certain sense, the condition $\int_0^{\Gamma(c)}[v - c]\mathrm{d}G(v) \leq 0, \forall c \in [0, \bar{c}]$ requires that the cost distribution dominates the value distribution. This condition holds, for example, when (i) $F(c) = c^\alpha$ and $G(v) = v^\beta$ with $0 < \alpha \leq \beta$ or (ii) G is uniform and $F(c)/c$ is increasing. Either condition implies that the cost distribution dominates the value distribution in terms of first-order stochastic dominance.

3.2 Seller-Optimal Algorithmic Consumption

Thus far, we have focused on the question of buyer-optimal algorithmic consumption. Natural alternative questions are as follows: what is the seller-optimal algorithmic consumption, and what is the efficient algorithmic consumption, i.e., what recommendation algorithms maximize seller profit and total surplus, respectively? We show that the answers to these questions coincide and feature a simple recommendation structure.

Claim 2 (Seller-Optimal Algorithm). *A threshold algorithm with $\hat{v}(p)$ such that for all p in $[\mathbb{E}[v], 1]$, $\mathbb{E}[v \mid v \geq \hat{v}(p)] = p$ simultaneously maximizes the seller's profit and total surplus, and moreover, achieves an efficient trade.*

The algorithm presented in Claim 2 maximizes efficiency at the expense of the buyer. Indeed, for any price, the algorithm maximally pools products of different values to the extent that the buyer is still willing to purchase when recommended. This results in a threshold recommendation, and given the full support assumption on G, a threshold is uniquely defined for all $p \in [\mathbb{E}[v], 1]$. Under this algorithm, the buyer is guaranteed a zero expected payoff irrespective of the posted price. Hence, the seller of any cost understands that she appropriates all generated surplus and thus her objective of profit maximization is perfectly aligned with efficiency. As a result, because the algorithm allows any threshold allocation, the seller of type c will post a price $p(c)$ that leads to an efficient trade, i.e., $p(c) = \mathbb{E}[v|v \geq c]$. The resulting overall allocation is efficient, the seller obtains the maximal feasible surplus, and the buyer is left with no rent.

We note that the discussion of Section 3.1 applies with greater force to the seller-optimal (and efficient) algorithm. As the algorithm achieves the first-best efficient benchmark, it cannot be improved upon by the use of monetary transfers, and as the algorithm maximizes the seller's profit, it remains optimal in the value disclosure setting for any cost and value distributions.

4 Algorithm Design and Market Segmentation

We now turn to the question of how the design and the outcome of the buyer-optimal algorithm adapt to the setting in which the seller gains access to data about the buyer's value and can engage in third-degree price discrimination. This setting can be viewed as market segmentation, in which the seller has the ability to target different consumer segments with different prices (e.g., Bergemann et al. (2015), Elliott et al. (2022)).

To study this question, we assume that the seller observes a public signal about the buyer's value. Formally, a signal $\mathcal{I} = (S, \pi)$ consists of a set S of signal realizations s and a family of distributions $\{\pi(\cdot|v)\}_{v \in [0,1]}$ over S. The signal is independent of the seller's type, and the algorithm can make recommendations based on both the realized signal and the value and price of the product.[7] To provide a sharper characterization, a part of our analysis considers a specific class of signals (see, e.g., Lewis and Sappington (1994) and Johnson and Myatt (2006)):

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

Because the signal is public and exogenous, our characterization of the buyer-optimal algorithm extends to this setting, as applied to each signal realization. Take any signal. Let G_s denote the posterior distribution of the buyer's value conditional on signal realization s. By Proposition 1, under the optimal algorithm, a seller of type c posts price

$$\tilde{p}_s(c) = \mathbb{E}_{v \sim G_s}[\Gamma^{-1}(v) | v \geq \Gamma(c)], \tag{2}$$

and the algorithm recommends the product if and only if $v \geq \Gamma(\tilde{p}_s^{-1}(p))$, 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 $p_s(c)$, define the *individual*

[7] 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 or fully informative signal.

buyer surplus at value v, $w(v)$, as the expected payoff of the buyer conditional on his value being v:

$$w(v) \triangleq \mathbb{E}_{s,c}[(v - p_s(c))r(v, p_s(c)) \mid v], \tag{3}$$

where the expectation is taken with respect to (s, c). Analogously, define the *individual seller profit at cost* c, $\pi(c)$, as follows:

$$\pi(c) = \mathbb{E}_{s,v}[(p_s(c) - c)r(v, p_s(c)) \mid c], \tag{4}$$

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 $w(v)$ with $v \sim G$ and that of $\pi(c)$ with $c \sim F$.

Proposition 2 (Market Segmentation). *Under the buyer-optimal algorithm, the ex ante buyer surplus and the individual seller profit at each* c *are the same for any signal* $\mathcal{I}$*. If* $\mathcal{I}$ *is a truth-or-noise signal, then as accuracy* α *increases, the distribution of prices undertakes a mean-preserving spread, and the distribution of individual buyer surpluses 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 the signal, the 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 depend on the signal.

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

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

of $\mathbb{E}_{v\sim G}[\Gamma^{-1}(v)|v \geq \Gamma(c)]$ 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 a price of $\Gamma^{-1}(v)$ for the buyer with value 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 decrease 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. □

Proposition 2 establishes, in a stark manner, that on average, the seller does not 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). Assume again that v and c are uniformly distributed on $[0, 1]$. Suppose that the seller has access to a public truth-or-noise signal with accuracy α. Given 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 2, is

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

Figure 2 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

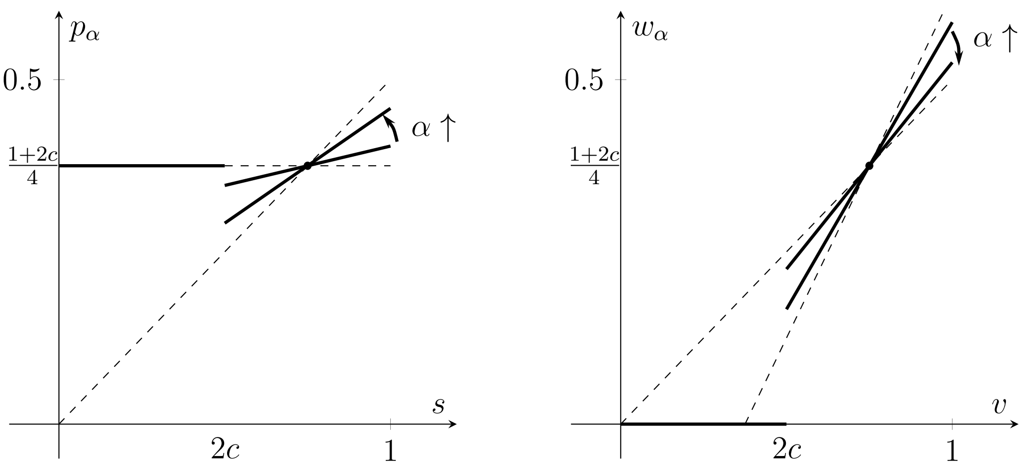

Figure 2: Equilibrium price posted by type c at different signal realizations and individual surplus at different values, shown at two levels of accuracy. $v \sim U[0,1]$, $c \sim U[0,1]$.

realized signal, making the equilibrium price steeper as a function of the realized signal.

To see the implication of this price change for 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 2 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 2. $\Diamond$

Proposition 2 requires the algorithm to adopt 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. The optimal algorithm can protect 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 with incomplete information. 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, which could be incorporated into optimal algorithm design. As another example, we confined our analysis to bilateral trade involving a single product. One could 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.

Assad, S., R. Clark, D. Ershov, and L. Xu (forth.): "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy*.

Balseiro, S., Y. Deng, J. Mao, V. Mirrokni, and S. Zuo (2021): "Robust Auction Design in the Auto-bidding World," *Advances in Neural Information Processing Systems*, 34, 17777–17788.

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*.

——— (2023): "Bidder-Optimal Information Structures in Auctions," *Working paper*.

Biden, J. R. (2023): "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/, accessed: January 2, 2024.

Brown, Z. Y. and A. MacKay (2023): "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics*, 15, 109–156.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020): "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review*, 110, 3267–97.

Chakraborty, I., J. Deb, and A. Oery (2022): "When Do Consumers Talk?" *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.

Elliott, M., A. Galeotti, A. Koh, and W. Li (2022): "Market Segmentation through Information," *Working paper*.

European Commission (2021): "Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts," https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206, accessed: January 2, 2024.

Farronato, C., A. Fradkin, and A. MacKay (2023): "Self-Preferencing at Amazon: Evidence from Search Rankings," *AEA Papers and Proceedings*, 113, 239–43.

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.

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.

Ichihashi, S. and A. Smolin (2023): "Data Provision to an Informed Seller," *Working paper.*

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.

Jameson, G. J. (2006): "Counting Zeros of Generalised Polynomials: Descartes' Rule of Signs and Laguerre's Extensions," *The Mathematical Gazette*, 90, 223–234.

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

Johnson, J. P., A. Rhodes, and M. Wildenbeest (2023): "Platform Design When Sellers Use Pricing Algorithms," *Econometrica*, 91, 1841–1879.

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.

Lamba, R. and S. Zhuk (2023): "Pricing with Algorithms," *Working paper*.

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

Lee, K. H. and L. Musolff (2021): "Entry into Two-Sided Markets Shaped by Platform-Guided Search," *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.

Luca, M. and O. Reshef (2021): "The Effect of Price on Firm Reputation," *Management Science*, 67, 4408–4419.

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.

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

Salcedo, B. (2015): "Pricing Algorithms and Tacit Collusion," *Working paper*.

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.

Smolin, A. (forth.): "Disclosure and Pricing of Attributes," *RAND Journal of Economics*.

Van Hoek, R. and M. Lacity (2023): "Procurement in the Age of Automation," *MIT Sloan Management Review*, 64, 1–10.

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 $q_r(p) \triangleq \int_0^1 r(v,p)dG(v)$ 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 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 follows:

$$\begin{aligned} \max_{\hat{v}:[0,1]\to[0,1],\, p:[0,1]\to\mathbb{R}_+} \quad & \int_0^1 \int_{\hat{v}(c)}^1 (v - p(c))\, \mathrm{d}G\, \mathrm{d}F, && (5) \\ \text{s.t.} \quad & \int_{\hat{v}(c)}^1 (p(c) - c)\mathrm{d}G \geq \int_{\hat{v}(c')}^1 (p(c') - c)\mathrm{d}G \quad && \forall\, c, c' \in [0,1], \\ & \int_{\hat{v}(c)}^1 (p(c) - c)\mathrm{d}G \geq 0 && \forall\, c \in [0,1]. \end{aligned}$$

The easiest way to solve this problem is to reformulate it in familiar terms. Because the value is continuously distributed, the expected trade probability $q \triangleq \int_{\hat{v}}^1 \mathrm{d}G$ is strictly decreasing in $\hat{v}$, spanning $[0,1]$ as $\hat{v}$ spans $[0,1]$. 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 $V(q) \triangleq \int_{\hat{v}(q)}^1 v\mathrm{d}G$ the corresponding trade surplus. The trade surplus is strictly increasing in q with $V(0) = 0$ and $V(1) = \mathbb{E}[v]$. Moreover,

$$\frac{\mathrm{d}V}{\mathrm{d}q} = \frac{\partial V / \partial \hat{v}}{\partial q / \partial \hat{v}} = \frac{-\hat{v} g(\hat{v})}{-g(\hat{v})} = \hat{v}(q), \tag{6}$$

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

(5) as follows:

$$\max_{q:[0,1]\to[0,1],\, t:[0,1]\to\mathbb{R}_+} \int_0^1 (V(q(c)) - t(c))\,\mathrm{d}F, \tag{7}$$
$$\text{s.t. } t(c) - cq(c) \geq t(c') - cq(c') \quad \forall\, c, c' \in [0,1],$$
$$t(c) - cq(c) \geq 0 \qquad \forall\, c \in [0,1].$$

Problem (7) is analogous to 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 6, we can translate this solution back to problem (5) as

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

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 $\hat{v}(p) = 1$, 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 $\hat{v}(\overline{p}) = 1 > \Gamma^{-1}(1) = \overline{p}$, and the seller's incentive compatibility requires that the threshold function strictly increases in the range of active prices. Additionally, we have $\hat{v}(\underline{p}) = 0 < \underline{p}$ if and only if $0 < \mathbb{E}_v[\Gamma^{-1}(v)]$ by construction. This completes the proof.

Proof of Corollary 1 Let $A \triangleq 1 + \frac{1}{\alpha}$. We have

$$
\begin{aligned}
\Gamma(c) &= Ac, \\
\tilde{p}(c) &= \mathbb{E}_{v\sim G}[\Gamma^{-1}(v) \mid v \geq \Gamma(c)] \\
&= A^{-1}\frac{\int_{Ac}^{1} v\beta v^{\beta-1}\mathrm{d}v}{1-(Ac)^{\beta}} \\
&= A^{-1}\frac{\beta}{1+\beta}\left(\frac{1-(Ac)^{1+\beta}}{1-(Ac)^{\beta}}\right).
\end{aligned}
$$

Thus we have

$$
\underline{p} = \frac{A^{-1}\beta}{1+\beta} > 0 \quad \text{and} \quad \overline{p} = A^{-1} < 1.
$$

We show that a unique $p^* \in (\underline{p}, \overline{p})$ exists such that $\hat{v}(p) = \Gamma(\tilde{p}^{-1}(p)) < p$ for $p < p^*$ and $\hat{v}(p) = \Gamma(\tilde{p}^{-1}(p)) > p$ for $p > p^*$. The optimal price $\tilde{p}(c)$ is continuous and strictly increasing in c, and it varies from $\underline{p}$ to $\overline{p}$ as we vary c from 0 to A^{-1}. Thus finding p^* is equivalent to finding a unique $c^* \in (0,1)$ such that $\Gamma(c) < \tilde{p}(c)$ for $c < c^*$ and $\Gamma(c) > \tilde{p}(c)$ for $c > c^*$.

To show the single-crossing property, define $H(x)$ and $I(x)$ as follows:

$$
\begin{aligned}
H(x) &= x - A^{-1}\frac{\beta}{1+\beta}\left(\frac{1-x^{1+\beta}}{1-x^{\beta}}\right), \forall x \in [0,1), \\
I(x) &= x(1-x^{\beta}) - A^{-1}\frac{\beta}{1+\beta}\left(1-x^{1+\beta}\right), \forall x \in [0,1],
\end{aligned}
$$

and $H(1) \triangleq \lim_{x\to 1} H(x) = 1 - A^{-1}$. Because $H(\cdot)$ is continuous and $H(0) < 0 < H(1)$, there is at least one x^* such that $H(\cdot)$ crosses 0 at x^* from below. Suppose to the contrary that there are $y, z \in (0,1)$ such that $y \neq z$ and $H(y) = H(z) = 0$. We then have $I(y) = I(z) = I(1) = 0$, which means that the generalized polynomial $I(\cdot)$ has three real positive zeros. However, by Descartes' rule of signs for generalized polynomials (see, e.g., Jameson (2006)), $I(\cdot)$ can have at most two real positive zeros, which is a contradiction. Thus $I(\cdot)$ has a unique positive zero in $(0,1)$, which means that there is a unique x^* with $H(x^*) = 0$. We conclude that for type $c^* = \frac{x^*}{A}$, we have $\Gamma(c) < \tilde{p}(c)$ for $c < c^*$ and $\Gamma(c) > \tilde{p}(c)$ for $c > c^*$.

Proof of Claim 1

We borrow the notation from the proof of Proposition 1 and let $\mu = \mathbb{E}_{v\sim G}[v]$. Suppose that the buyer faces the optimal algorithm of Proposition 1 in the value disclosure setup. Because the buyer is willing to follow the algorithm's recommendation to purchase, it suffices to show that the buyer is also willing to follow the recommendation to not purchase. This constraint is equivalent to the condition that the buyer's ex ante payoff from following the recommendation weakly exceeds the payoff from always buying the product regardless of the recommendation. For any active price $p \in [0, p(\bar{c}))$ (i.e., a price that some type below $\bar{c}$ chooses), the condition is written as

$$V(q(c)) - t(c) \geq \mu - p(c)$$

or

$$V(q(c)) - cq(c) - \int_c^1 q(x)dx \geq \mu - c - \frac{\int_c^1 q(x)\mathrm{d}x}{q(c)}. \tag{8}$$

Because $q(c) \leq 1$, a sufficient condition for inequality (8) is

$$V(q(c)) - cq(c) \geq \mu - c.$$

We can rewrite this inequality as

$$\int_{\Gamma(c)}^1 v\mathrm{d}G(v) - c\int_{\Gamma(c)}^1 1\mathrm{d}G(v) \geq \int_0^1 v\mathrm{d}G(v) - c\int_0^1 1\mathrm{d}G(v),$$

or $\int_0^{\Gamma(c)}[v-c]\mathrm{d}G(v) \leq 0$.

Finally, the buyer follows the recommendation to not buy the product at any price p that is not active, i.e., $p \geq \tilde{p}(\bar{c})$. Recall that the buyer-optimal algorithm provides no information about v at price $p > \tilde{p}(\bar{c})$. Plugging $c = \bar{c}$ into $\int_0^{\Gamma(c)}[v-c]\mathrm{d}G(v) \leq 0$, we obtain $\mu - \bar{c} \leq 0$. Thus if $p > \tilde{p}(\bar{c})$, we have $\mu - p \leq \mu - \tilde{p}(\bar{c}) = \mu - \bar{c} \leq 0$.

Proof of Proposition 2

Step 1. We show that the buyer surplus and individual profits do not depend on the public signal, $\mathcal{I}$. 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 $\mathcal{I}$. This also means that from an ex ante perspective, for any $\mathcal{I}$, 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 $\mathcal{I}$, 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 $\mathcal{I}$ (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 $G^c_{\alpha,s} \in \Delta[0,1]$ 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

$$p(c|s,\alpha) = \mathbb{E}_{v \sim G_{\alpha,s}}[\Gamma^{-1}(v)|v \geq \Gamma(c)] = \int_0^1 \Gamma^{-1}(v)\mathrm{d}G^c_{\alpha,s}(v). \tag{9}$$

Let $\mathcal{G}^c_\alpha \in \Delta\Delta[0,1]$ 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 9 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

$$\frac{\partial}{\partial v}w(v,c,\alpha_L) = 1-\alpha_L^2\frac{\partial}{\partial v}\Gamma^{-1}(v) > 1-\alpha_H^2\frac{\partial}{\partial v}\Gamma^{-1}(v) = \frac{\partial}{\partial v}w(v,c,\alpha_H) \geq 0, \tag{10}$$

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 (10) and the equal mean property imply that $w(0,c,\alpha_L) < w(0,c,\alpha_H)$ and $w(1,c,\alpha_L) > w(1,c,\alpha_H)$. 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 $w^{-1}(x,\alpha_i) \triangleq \max\{v \in [0,\bar{v}] : w(v,c,\alpha) = x\}$, 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*.