WORKING PAPERS

N° 1672

October 2025

"Buyer-Optimal Algorithmic Recommendations"

Shota Ichihashi and Alex Smolin

Toulouse School of Economics

Buyer-Optimal Algorithmic Recommendations*

Shota Ichihashi Alex Smolin

February 15, 2025

Abstract

In markets where algorithmic data processing is increasingly prevalent, recommendation algorithms can substantially affect trade and welfare. We consider a setting in which an algorithm recommends a product based on its value to the buyer and its price. We characterize an algorithm that maximizes the buyer's expected payoff and show that it strategically biases recommendations to induce lower prices. Revealing the buyer's value to the seller leaves overall payoffs unchanged while leading to more dispersed prices and a more equitable distribution of surplus across buyer types. These results extend to all Pareto-optimal algorithms and to multiseller markets, with implications for AI assistants and e-commerce ranking systems.

*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. We would like to thank Nemanja Antic, Heski Bar-Isaac, Dirk Bergemann, Alessandro Bonatti, Daniele Condorelli, Piero Gottardi, Emir Kamenica, Caio Lorecchio, Matthew Mitchell, Alessandro Pavan, Nicola Persico, and Udayan Vaidya, and audiences at ESSET Gerzensee 2024, Berkeley-Columbia-Duke-Northwestern-MIT IO Theory Conference 2024, and other venues, for their valuable suggestions and discussions. The authors acknowledge support from the Economics of Digital Services initiative at the University of Pennsylvania and from the NET Institute. 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) and through the Artificial and Natural Intelligence Toulouse Institute (ANITI). Javier Gonzalez-Morin provided excellent research assistance.

1 Introduction

In today's digital economy, algorithmic decision-making is transforming how consumers navigate markets—from price trackers that seek and pinpoint lower-priced products, to robo-advisors that propose financial securities, to e-commerce platforms where ranking algorithms shape which products capture buyers' attention, and to AI-powered personal assistants that advise individuals on their purchasing decisions.[1] Algorithms are not merely processing information but are actively influencing consumer and seller behavior.

In this paper, we formalize this dynamic by developing a model of algorithmic recommendations and examining how they can reshape market outcomes and redistribute welfare. Our model emphasizes three key features of recommendation algorithms. First, they operate according to preprogrammed rules. Second, they are capable of locating and processing critical information about a product's existence, inherent value, and price. Third, by influencing when and how recommendations are made, these algorithms can affect the purchasing decisions by the buyers and pricing decisions by the sellers.

We begin by developing a baseline model of bilateral trade. In this setting, a single product is traded between a buyer and a seller in the presence of uncertainty about both trade costs and values. The seller privately observes her production cost—her type—while the buyer is initially uninformed about both the product's existence and its value. An algorithm, however, can discover the product's value or, equivalently, produce its best estimate and issue a recommendation based on this information combined with the seller's posted price. Upon receiving a recommendation, the buyer updates her belief and decides whether to purchase at the offered price; if no recommendation is received, no trade occurs—that is, the algorithm can serve as a gatekeeper. The seller, fully aware of the algorithm's design, sets her price strategically to maximize profit.

The algorithm shapes equilibrium trade and welfare. We abstract away from the source of the algorithm and instead analyze all algorithms with desirable welfare properties. Motivated by the fact that currently most revenue for AI assistants comes from

[1] As of February 2025, the top three productivity apps on the Apple Store are AI assistants: ChatGPT, DeepSeek, and Gemini.

their users, we present most of the analysis with the objective being the buyer's expected surplus, thus characterizing a *buyer-optimal* algorithm. However, our characterization and key results extend, with minimal changes, to any Pareto-optimal algorithm in the space of buyer and seller surpluses. Thus, our findings speak to a variety of market structures that include, for example, two-sided platforms aiming to accommodate both of their sides.

A buyer-optimal algorithm must navigate a trade-off: It needs to reward sellers for offering lower prices by increasing the frequency of recommendations, while ensuring that beneficial trades are not foregone. We show that this balance is optimally achieved by an algorithm that uses a *pseudo value* threshold—recommending the product when this adjusted metric exceeds the price, even if the true value might not. The equilibrium we characterize (Proposition 1) reveals three key insights. First, although the algorithm's design and the equilibrium prices depend on both the cost and value distributions, the final product allocation is driven solely by the cost distribution. Second, by deliberately deviating from a simple ex post optimal rule—eschewing recommendations at high prices even when the true value is favorable and endorsing products at low prices even when the value falls short—the algorithm increases the buyer's price sensitivity, thereby pressuring the seller to offer lower prices.[2] Third, compared to standard monopoly pricing with ex post optimal trade, buyer-optimal algorithmic recommendations substitute high-value, high-cost trades with low-value, low-cost trades (Proposition 2).

The complete characterization of an optimal algorithm and the equilibrium strategies allows us to uncover major changes in the welfare implications of third-degree price discrimination compared to standard monopoly pricing. We show that if the seller can distinguish and price discriminate among different buyer segments, the algorithm adapts accordingly and *fully mitigates* the average effects of market segmentation. As a result, price discrimination does not affect the average price, buyer's total surplus, seller profits, or product allocation (Proposition 3). At the same time, finer segmentation enables the seller to correlate prices with values, and within a class of monotone segmentations, it

[2]This finding highlights the importance of the strategic context for an algorithm assessment and AI regulation (cf. White House (2023); European Commission (2024)).

leads to more dispersed equilibrium prices and a more concentrated—and thus more equitable—distribution of consumer surplus (Proposition 4). These results hold broadly and demonstrate that algorithmic recommendations can serve as a powerful consumer protection tool, complementing existing regulatory methods (e.g., Scott Morton et al. (2019)).

Our analysis extends in several dimensions. First, our characterization and market segmentation insights extend to any Pareto-optimal algorithm by simply incorporating Pareto weights into the definition of a pseudo-value (Proposition 5). Second, most of our key results and characterizations generalize to settings with multiple competing sellers (Proposition 6). Notably, despite strategic competition among sellers, the specifics of market segmentation do not affect the average price, total buyer surplus, or seller profits (Proposition 7), and finer segmentation leads to more dispersed prices (Proposition 8). Third, we expand our analysis by allowing the buyer to be informed in advance about either the product's existence or its value, showing how this information can potentially harm the buyer and discussing how algorithm design can mitigate this harm (Proposition 9 and Proposition 10).

Related literature.— Our paper is closest 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. Like us, they show 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. Unlike us, they require learning to occur before the price is set, which limits its impact (see further Section 3.1 and Section 5.3). Deb and Roesler (forth.) 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 allowing the buyer's information to depend on the price and by allowing the seller to have private information.[3]

[3]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.

In our study of algorithm design under sellers' private information, we integrate Bayesian persuasion (e.g., Kamenica and Gentzkow (2011)) with mechanism design (Baron and Myerson (1982)). Several works have combined these frameworks in trade settings. Among these, the most closely related studies are Yang (2022), Bergemann and Bonatti (2024), Xu and Yang (2024), and Bergemann et al. (forth.).[4] We share the perspective that new information technologies enable powerful third parties to process data, extract valuable insights, and shape market outcomes by serving as informational gatekeepers. However, our focus differs in that those studies allow the intermediary to charge transfers and study revenue maximization, whereas we focus on achieving Pareto efficiency. We further differ from Bergemann and Bonatti (2024) and Bergemann et al. (forth.) in the nature of the friction faced by the designer: Whereas they contend with buyers' off-site purchase opportunities, we focus on the private information held by sellers. In this respect, we are closer to Yang (2022) and Xu and Yang (2024), sharing some features of the solution with them (Remark 2).

Our analysis in Section 4 offers a novel perspective on the classic question of the impact of price discrimination based on consumer information, as studied in the market segmentation literature (e.g., Bergemann et al. (2015) and Haghpanah and Siegel (2023)). We show that consumer use of algorithms may introduce a new welfare implication whereby price discrimination results in a more equal distribution of consumer surplus without affecting average welfare outcomes. This finding 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. (2024)) or information design (Doval and Smolin (2024)).

Finally, our work contributes broadly to recent research strands in the economics of algorithmic decision-making. A significant body of work has examined algorithmic pricing, largely focusing on how such tools bolster seller power. Empirical studies—such as those by Calvano et al. (2020), Asker et al. (2022), and Assad et al. (2024)—document how algorithms facilitate collusion and dynamic pricing, while theoretical work (e.g., Sal-

[4]See also Lee (2021), Bergemann et al. (2022), and Smolin (2023).

cedo (2015), Lamba and Zhuk (2023)) and combined empirical–theoretical approaches (e.g., Brown and MacKay (2023), Johnson et al. (2023)) investigate their strategic impact. In contrast, we explore how algorithmic recommendations can empower buyers by mitigating information asymmetries and counteracting traditional price discrimination. This buyer-centric perspective resonates with early insights from Galbraith (1952), offering an alternative approach to balancing market power.[5,6]

2 Baseline Model

There is a buyer and a seller. The buyer may purchase a product but initially knows neither the existence nor the value of the product: The value is denoted by v and is distributed according to distribution G with positive density g over its support $[0, 1]$. A *recommendation algorithm* or simply *algorithm* provides the buyer with information about the product by means of recommendations.[7] The algorithm is characterized by 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 that the buyer purchase the product with probability $r(v, p)$.[8]

The seller can produce one unit of a product at cost c, which is her private *type*. The type distribution F has support $[0, 1]$ with positive density f.

[5]Thus, algorithmic recommendations can be viewed as an effective alternative to the joint use of an intermediary (see Decarolis and Rovigatti (2021) for online advertising) or to a merger (see Loertscher and Marx (2022) for multifirm bargaining).

[6]One can view our setting as enabling a buyer from the classic setting of Myerson and Satterthwaite (1983) to commit to values and prices at which she 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)).

[7]In this setting, providing extra information about v upon recommendation is unnecessary.

[8]The dependence of recommendations on price can be requested by the buyer in a conversation with an AI assistant, programmed directly, as seen in Amazon's search ranking algorithms (Lee and Musolff (2023), 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.

Timing For any given algorithm, which is commonly known to both players, the timing of the game is 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 and the game ends. With probability $r(v,p)$, the algorithm recommends the product to the buyer, in which case the buyer 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.

Given the algorithm, the solution concept is a 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 weakly 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 profit.

Objective We call the buyer's ex ante expected payoff *buyer surplus* and the ex ante seller's expected payoff *seller profit*. An algorithm *attains a given buyer surplus* if this buyer surplus arises in an equilibrium under this algorithm. For most of the analysis, we focus on the recommendation algorithms that maximize buyer surplus. However, in Section 5.1, we show that the same analysis readily applies to the whole class of Pareto-optimal algorithms.

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

In what follows, 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 a strictly positive probability and is *inactive* otherwise.

3 Buyer-Optimal Algorithm

In this section, we characterize the buyer-optimal algorithm. 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.

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

The proofs of this and all other results are in the Appendix. Lemma 1 shows that threshold algorithms span a Pareto frontier in the space of buyer surplus and seller profit. Intuitively, the buyer can be set to follow the recommendations because the algorithm can anticipate and mimic the buyer's response. In turn, the Pareto efficiency of threshold algorithms follows from the observation that each seller type is concerned solely with trade volume whereas the buyer surplus is maximized when the higher values are prioritized. Consequently, if a buyer-optimal algorithm exists, 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 the trade surplus and incentivizing the seller to lower the price. One 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.

3.1 Known Seller Cost

The simplest case to illustrate the power of algorithmic recommendations is to consider a limit case of our model in which the seller cost is commonly known.

Claim 1 (Known Cost). *Suppose the seller cost is commonly known to be $c \in [0, 1]$. Then an algorithm that employs a threshold function $\hat{v}$ defined by $\hat{v}(c) = c$ and $\hat{v}(p) > 1$ for $p \neq c$ is optimal and extracts the full surplus from the seller.*

The condition $\hat{v}(p) > 1$ for $p \neq c$ effectively forces the seller to set the price equal to the marginal cost, thereby eliminating any seller rents, while $\hat{v}(c) = c$ ensures that trade is efficient and generates maximal surplus. In equilibrium, the entire surplus is generated and appropriated by the buyer, which confirms the algorithm's optimality.

Claim 1 highlights a key distinction between our setting and Roesler and Szentes (2017). In their model, the cost is commonly known, but the algorithm cannot condition information on price nor exclude the seller from trade. Consequently, even though equilibrium trade is efficient, the seller earns positive rents. In contrast, in our model, the algorithm can condition information on price and exclude the seller from trade by not informing the buyer of the existence of the product. The combination of these properties enables the buyer to extract full surplus when the seller's cost is known.

In Section 5.3.1, we elaborate on the roles of price dependence and gatekeeping: When cost uncertainty is low, gatekeeping is particularly valuable, while price dependence is not, because the algorithm effectively provides information at only a single price. Conversely, when cost uncertainty is high—and we provide precise distributional conditions under which this holds—price dependence is as effective whether or not the buyer could purchase the product without the recommendation.

While the optimal algorithm in Claim 1 is simple, it relies on precise knowledge of seller cost and is fragile to small perturbations: Even an arbitrarily small increase in cost would lead to a complete collapse of trade. In practice, the algorithm may lack the information necessary to identify seller cost. Moreover, this asymmetry can be exacerbated by idiosyncratic cost shocks arising from supply disruptions or alternative sale opportunities. These observations motivate our subsequent analysis of private seller information.

3.2 General Case

To find an optimal algorithm in the general case of an unknown seller cost, we build on Lemma 1 and frame the designer's problem as a nonlinear screening problem in which the recommendation threshold responds to the price. The choice of a threshold at any given price simultaneously determines the expected trade surplus, which is valued by the buyer, and the expected trade volume, which is valued by the seller. We recover the optimal threshold function by adapting the seminal analysis of Baron and Myerson (1982) and confirm that, with this algorithm, the buyer indeed finds it optimal to follow the recommendations.

The optimal algorithm and equilibrium pricing are easier to describe not in terms of price-dependent thresholds for values but in terms of value-dependent thresholds for prices. Specifically, denote the *virtual cost function* by $\gamma(c) \triangleq c + F(c)/f(c)$, and assume that it is continuous and strictly increasing on $[0,1]$.[9] For each $v \in [0,1]$, define the buyer's *pseudo value* as

$$y(v) \triangleq \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\tilde{v} \geq v]. \tag{1}$$

The pseudo value $y(v)$ is an increasing function of v. When the true value is sufficiently low, close to 0, the pseudo value is higher than the true value, $y(v) > v$, because $y(0) = \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})] > 0$. When the true value is sufficiently high, close to 1, the pseudo value is below the true value, $y(v) < v$, because $y(1) = \gamma^{-1}(1) < 1$. Define $\bar{c} \triangleq \gamma^{-1}(1)$.[10]

[9] If $\gamma(c)$ were not everywhere increasing, then one would simply use an ironed version of it.

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

Proposition 1 (Buyer-Optimal Algorithm). *A buyer-optimal algorithm recommends the product if and only if $y(v) \geq p$. Under this algorithm, the seller of type $c \leq \bar{c}$ posts price $p^*(c) = y(\gamma(c))$, and the seller of type $c > \bar{c}$ is inactive. Under this algorithm and pricing, the trade occurs if and only if $v \geq \gamma(c)$.*

Proposition 1 reveals two notable features. First, the impact of the value and cost distributions can be decoupled: Even though the optimal algorithm and the equilibrium prices depend both on the cost and value distributions, the optimal product allocation depends only on the cost distribution because the product is traded if and only if the value exceeds the virtual cost. This feature is crucial for the market segmentation results in Section 4.

Second, the buyer-optimal algorithm makes two types of *ex post* errors: When the true value is sufficiently high, the pseudo value is below the true value, and the optimal algorithm never recommends the product when the value is below the price. At the same time, for seller types that satisfy $y(v) < y(\gamma(c)) < v$, the algorithm does not recommend the product even though the value exceeds the price. In contrast, when the true value is sufficiently low, the pseudo value is higher than the true value and the algorithm always recommends the product when the value exceeds the price. However, for seller types that satisfy $v < y(\gamma(c)) < y(v)$, the algorithm recommends the product even though the value is below the price. As a result, compared with the ex post optimal algorithm, the buyer-optimal algorithm overrecommends at low prices and underrecommends at high prices. These distortions benefit the buyer because they incentivize the seller to set lower prices.

Example 1 (Uniform). We illustrate the buyer-optimal algorithm when c and v are uniformly distributed on $[0, 1]$. In this case, the virtual cost is $\gamma(c) = 2c$; the pseudo value is $y(v) = \mathbb{E}_{\tilde{v} \sim U[0,1]} \left[\frac{\tilde{v}}{2} \mid \tilde{v} \geq v \right] = (1 + v)/4$; and the associated threshold function $\hat{v}(p)$ is equal to 0 for $p < 1/4$, to $4p - 1$ for $p \in [1/4, 1/2]$ and to 1 for $p > 1/2$. The equilibrium price posted by active type c is $p^*(c) = y(\gamma(c)) = (1 + 2c)/4$ for $c \in [0, 1/2]$. Types $c > 1/2$ are inactive and post, for example, $p^*(c) = 1/2$. A buyer who receives a recommendation to purchase at price $p \in [1/4, 1/2]$ infers that the product's expected

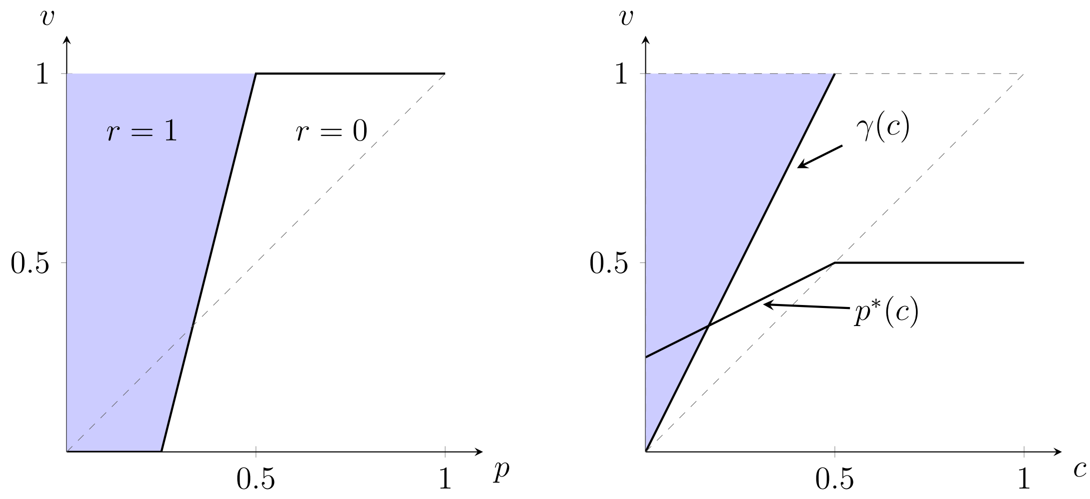

Figure 1: Optimal recommendation algorithm (left) and the resulting equilibrium pricing strategy and trade region (right). $v \sim U[0, 1]$, $c \sim U[0, 1]$.

value is $(4p - 1 + 1)/2 = 2p > p$ and is thus strictly 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 (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 below the price; 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 right side of Figure 1 depicts the resulting equilibrium pricing and trade: The price $p^*(c)$ 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 occurs whenever the buyer 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 $[1/4, 1/2]$.

We compare the equilibrium outcomes under the buyer-optimal algorithm and the ex post optimal algorithm. Under the ex post optimal algorithm, the seller's problem is a standard monopoly problem with the optimal price $p^m(c) = \frac{1+c}{2}$, and trade occurs if

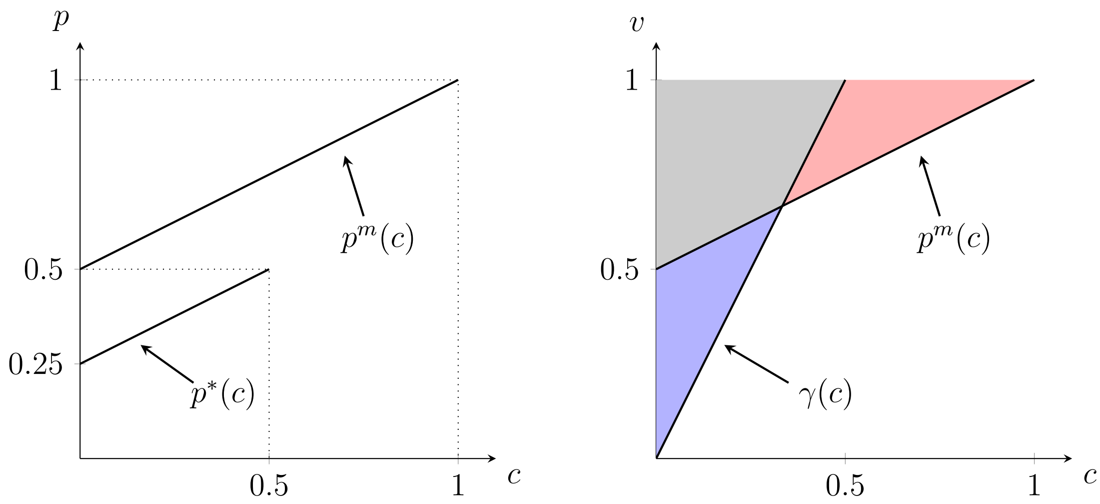

Figure 2: Comparison of the equilibrium outcomes under the ex ante optimal algorithm and the ex post optimal algorithm in terms of the seller's pricing strategies (left) and trade regions (right).

and only if $v \geq p^m(c)$. As indicated by the left panel of Figure 2, as we move from the ex post optimal algorithm to the buyer-optimal algorithm, each seller type $c \leq 0.5$—who transacts with a positive probability under both algorithms—decreases its price by 0.25; thus, the buyer-optimal algorithm induces a uniform downward shift of the equilibrium prices across *all* seller types. (In this example, the price change is associated with the uniform decrease of the seller profits: The interim profit for seller type c under the ex post optimal algorithm is $\frac{(1-c)^2}{4}$, whereas that under the buyer-optimal algorithm is $\frac{(1-2c)^2}{4}$ if $c \leq \frac{1}{2}$ and 0 otherwise. However, in general, some seller types may prefer the buyer-optimal algorithm due to the overrecommendations.)

Complementary, the right panel of Figure 2 depicts the sets of pairs of value v and cost c under which transaction occurs. Under the ex post optimal algorithm, transaction occurs when $v \geq p^m(c)$. The figure indicates that compared to the ex post optimal algorithm, the buyer-optimal algorithm reduces transactions between high-value buyers and high-cost sellers (the red region) and increases transactions between low-value buyers and low-cost sellers (the blue region).

$\Diamond$

The last observation of the previous example turns out to be more general. To describe the result, let Δ^* denote the set of (c, v) pairs such that trade occurs under the buyer-optimal algorithm, and let Δ^m denote the set of (c, v) pairs such that trade occurs under the ex post optimal algorithm.

Proposition 2 (Allocation Substitution). *If F has a strictly decreasing reversed hazard rate $\frac{f}{F}$ and G has a strictly increasing hazard rate $\frac{g}{1-G}$, then for any $(c, v) \in \Delta^* \setminus \Delta^m$ and $(c', v') \in \Delta^m \setminus \Delta^*$ such that $(c, v) \neq (c', v')$, we have $c < c'$ and $v < v'$.*

Proposition 2 shows that in regular environments, switching from monopoly pricing to algorithmic recommendations leads to a systematic shift in equilibrium trade: High-value, high-cost transactions are replaced by low-value, low-cost ones. Intuitively, algorithmic recommendations push prices down by under-recommending high-value products and over-recommending low-value ones. This, in turn, reduces trade with high-cost sellers while increasing trade with low-cost sellers.

Remark 1 (Comparison to Full Commitment Benchmark). As our proof reveals, the optimal algorithm in Proposition 1 attains the same outcome as when the buyer is aware of the product's existence, can observe the product's value, and has full commitment power over purchasing decisions and monetary transfers at different product values. As a result, even though the algorithm serves only information, it effectively transfers market power from the seller to the buyer. This observation has two consequences. First, the same algorithm remains optimal in the case of fully automated trade, i.e., if it could execute transactions without having the buyer in a loop, or if the algorithm could charge monetary transfers to the seller, e.g., referral or commission fees. Second, the same outcome would be optimal even if the seller could employ more general trade protocols than a posted price. In that case, a buyer-optimal algorithm would recommend products sold via posted prices according to the characterization in Proposition 1 and would never recommend products sold via alternative protocols.

Remark 2 (Optimal Product Allocation). The findings in Proposition 1 are related to those of Yang (2022) and Xu and Yang (2024). They study revenue-maximizing intermediation by a platform that can assess consumer value and can charge payments to both sides but is uncertain about seller costs. Our paper and their papers differ in the designer's objective, the proof techniques, and the optimal mechanisms. At the same time, the equilibrium product allocations coincide—the good is traded if and only if the value exceeds virtual costs. This outcome is not obvious *a priori* and suggests that a revenue-maximizing platform would be more aligned, in terms of the resulting product allocation, with the side about which it has more information.

4 Algorithm Design and Market Segmentation

Algorithmic recommendations have major consequences for third-degree price discrimination. To demonstrate this, we allow the seller to observe and base prices on signal $\mathcal{I} = (S, \pi)$ informative about the buyer's value. The signal consists of a set S of signal realizations s and a family of probability distributions $\{\pi(\cdot|v)\}_{v\in[0,1]}$ over S. We write $\hat{\pi}(s)$ for the marginal probability of signal realization $s \in S$ and G_s for the posterior value distribution conditional on s. Each signal can be viewed as a market segmentation, with $\hat{\pi}$ capturing the relative frequency of buyer segments and G_s capturing the distribution of buyer values within each segment (cf. Bergemann et al. (2015)). The signal is exogenous, and the signal realization is independent of the seller's type.

We assume that the algorithm can perfectly distinguish different market segments so that the cost information remains the only private information of the seller.[11] As the seller can set different prices in different segments, the optimal algorithm's recommendations should depend on the segment as well. In fact, the buyer-optimal segment-dependent algorithm must be buyer-optimal in each segment and is thus characterized in each segment s by Proposition 1, with the value distribution being G_s:[12] For each $s \in S$, the optimal algorithm recommends the product if and only if the corresponding

[11]As the algorithm knows the value, this assumption is trivially satisfied if the seller's signal is fully informative or, more generally, partitional.

[12]Formally, in Section 3, we assumed that the value distribution has full support on an interval. However, our derivation of Proposition 1 did not rely on this assumption, and thus the result applies to any market segment.

pseudo value exceeds the price, i.e.,

$$y_s(v) \triangleq \mathbb{E}_{\tilde{v}\sim G_s}[\gamma^{-1}(\tilde{v})|\tilde{v} \geq v] \geq p.$$

In equilibrium, the seller with type c posts a price of

$$p_s^*(c) \triangleq y_s(\gamma(c)) = \mathbb{E}_{\tilde{v}\sim G_s}[\gamma^{-1}(\tilde{v})|\tilde{v} \geq \gamma(c)] \tag{2}$$

and transacts whenever $v \geq \gamma(c)$.

Importantly, in contrast to the recommendation function or the seller's pricing, the product allocation, when viewed as a function of value and cost, is the same across all segments. This enables us to derive sharp implications of finer market segmentation. For any active type c, we define the *distribution of prices of type c* as the distribution of $p_s^*(c)$ when we fix c and draw s from distribution $\hat{\pi}$. The corresponding *profit of type c* is

$$\pi(c) \triangleq \mathbb{E}_{s,v}[(p_s^*(c) - c)\mathbb{1}(y_s(v) \geq p_s^*(c)) \mid c\,], \tag{3}$$

where the expectation is taken with respect to $v \sim G$ and $s \sim \pi(\cdot|v)$.

Proposition 3 (Segmentation Neutrality). *For any signal $\mathcal{I}$ available to the seller, the buyer-optimal algorithm induces the same ex post product allocation, the same expected price and profit of each seller type, and the same ex ante buyer surplus.*

Proof Outline. Consider the buyer-optimal algorithms with and without the seller's signal $\mathcal{I}$. Under either algorithm, in equilibrium, the trade occurs if and only if $v \geq \gamma(c)$, and the highest type, $c = 1$, earns zero profit as she never trades. We can view the two algorithms as indirect mechanisms that induce the same allocation rule and yield the same profit for type $c = 1$ when the seller prices in an incentive-compatible way. A version of the revenue equivalence theorem (Lemma 2 in the appendix) implies that the individual profit of each seller type with and without the seller's signal $\mathcal{I}$ must coincide. Consequently, the buyer surplus, which is the total surplus minus the seller's ex ante profit, is also identical in the two settings. □

Proposition 3 establishes in a stark manner that no seller types benefit from having more information about the buyer value, and the buyer neither benefits from nor is harmed by the release of such information *on average*, as long as this release is accounted for in the algorithm design.

Despite the neutral aspects highlighted by Proposition 3, a change in market segmentation does affect the optimal algorithm, the equilibrium pricing, and the distribution of payoffs across buyers with different valuations. To analyze the redistribution effect, we define the *buyer surplus at value v and type c*, $w(v, c)$, as the equilibrium expected payoff of the buyer conditional on his value being v and the seller's type being c:

$$w(v, c) \triangleq \mathbb{E}_s[(v - p_s^*(c))\mathbb{1}(y_s(v)) \geq p_s^*(c)) \mid v\,], \tag{4}$$

where the expectation is taken with respect to signal realization $s \sim \pi(\cdot|v)$ conditional on value v. Similarly, we define the *distribution of buyer surplus at type c* as the distribution of $w(v, c)$ with $v \sim G$. Furthermore, we will obtain a cleaner characterization and stronger results for the natural class of *monotone partitional* signals.

Definition 2 (Monotone Partitional Signal). A signal $\mathcal{I}$ is *monotone partitional* if there exists a finite partition of $[0, 1]$ into intervals $\{I_1, ..., I_n\}$ such that for each interval I_k and each $v \in I_k$, $\pi(\cdot|v)$ assigns probability 1 to either (i) $s = v$ or (ii) $s = k$.

Under a monotone partitional signal, each market segment is either a singleton or an interval. Furthermore, different buyer values belong to different segments; thus, the segment dependency in the description of an algorithm, which already conditions on the value, is redundant. The buyer-optimal algorithm and equilibrium pricing can be described segment by segment. Consider a segment $[\underline{v}, \overline{v}]$, which either equals I_k for some k or satisfies $\underline{v} = \overline{v}$. By Proposition 1, for all $v \in [\underline{v}, \overline{v}]$, the buyer-optimal algorithm recommends the product if and only if the pseudo value $y(v) = \mathbb{E}_{\tilde{v} \sim G}[\gamma^{-1}(\tilde{v})|\overline{v} \geq \tilde{v} \geq v]$ is above the price. Given this algorithm, in the segment $[\underline{v}, \overline{v}]$, the seller of type c posts a price

$$p^*_{[\underline{v},\overline{v}]}(c) = \begin{cases} \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\overline{v} \geq \tilde{v} \geq \gamma(c)] & \text{if} \quad \underline{v} \leq \gamma(c) \leq \overline{v}, \\ \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\overline{v} \geq \tilde{v} \geq \underline{v}] & \text{if} \quad \gamma(c) < \underline{v}. \end{cases} \tag{5}$$

In particular, if the buyer value is revealed to be v (i.e., $\underline{v} = \overline{v} = v$), the algorithm recommends the product if and only if the price is below $\gamma^{-1}(v)$, and any seller with a cost below this value posts price $\gamma^{-1}(v)$.

Proposition 4 (Segmentation Redistribution). *If signal $\mathcal{I}_H$ is Blackwell more informative than $\mathcal{I}_L$, then the distribution of prices set by each seller type under $\mathcal{I}_H$ is a mean-preserving spread of that under $\mathcal{I}_L$. Furthermore, if signals $\mathcal{I}_H$ and $\mathcal{I}_L$ are monotone partitional, then the distribution of individual buyer surplus at any seller type under $\mathcal{I}_H$ is a mean-preserving contraction of that under $\mathcal{I}_L$.*

When the seller faces a more informative signal, the posterior beliefs on values conditional on signal realizations and the event $v \geq \gamma(c)$ become a mean-preserving spread of the posterior beliefs when the seller faces a less informative signal. By Equation 2, the equilibrium prices are linear in these beliefs; thus, the prices undergo a mean-preserving spread as the signal becomes more informative.

To gain intuition about the buyer surplus at different values, compare the seller who has no information and the seller who perfectly observes the value. When the seller has no information, the seller of type c posts a price of $\mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\tilde{v} \geq \gamma(c)]$ regardless of the value, and any buyer with value $v \geq \gamma(c)$ trades at that price. When the seller has full information, the price depends on the segment, and for the buyer with value v, the seller of type c posts a price of $\gamma^{-1}(v)$. The buyer trades at this price, which is increasing in value. As a result, the seller's information increases the prices set for buyers with higher values, whereas the average remains the same by the first part of the argument. Consequently, the seller's information 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.

The intuition behind the general monotone partitional signals is similar. In fact, the proof of Proposition 4 establishes an additional result: For any monotone partitional signal $\mathcal{I}$ and type c, a cutoff $v(c, \mathcal{I})$ exists such that the buyer surplus at value v and type c is greater with signal $\mathcal{I}$ than under no information if and only if $v \leq v(c, \mathcal{I})$.[13]

Example 1 (Continued). Let v and c be uniformly distributed on $[0, 1]$. Suppose that the seller has access to a monotone partitional signal with a uniform grid, i.e., each I_k is an interval between $\frac{k-1}{n}$ and $\frac{k}{n}$. All values in interval I_k are pooled into signal k. Let $I(v)$ denote the interval to which value v belongs.

We fix any type $c < \frac{1}{2}$ and examine how the equilibrium price and buyer surplus at v and c vary across $v \geq 2c$. Let $k = 1, \ldots, n$ satisfy $2c \in I_k$. If $v \in I_m$, the seller will observe signal realization m and infer that the value is uniformly distributed between $\frac{m-1}{n}$ and $\frac{m}{n}$. As a result, the price posted by type c against value v is

$$p(c, v) = \mathbb{E}_{\tilde{v} \sim G}\left[\frac{\tilde{v}}{2}\Big|\tilde{v} \geq 2c, \tilde{v} \in I(v)\right] = \begin{cases} \frac{2cn+k}{4n}, & \text{if} \quad v \in I_k, \\ \frac{2m-1}{4n}, & \text{if} \quad v \in I_m, m > k. \end{cases}$$

The buyer's surplus at $v \geq 2c$ is $w(v, c) = v - p(c, v)$.

Figure 3 depicts the optimal algorithm thresholds and the buyer surplus $\mathbb{E}_{c \sim F}[w(\cdot, c)]$ at different values for the uninformative signal, the signal represented by binary partition $\{[0, 0.5], (0.5, 1]\}$ and the fully informative signal. As the market segmentation becomes finer, the equilibrium price responds more to the buyer value, which redistributes the surplus from higher to lower values. As a result, the finer market segmentation, or more buyer information provided to the seller, makes the distribution of buyer surplus more equalized across different values.

Proposition 4 generalizes this observation for Blackwell-comparable market segmen-

[13] Monotone partitional signals are not the only class of signals under which the distribution of individual buyer surplus undergoes a mean-preserving contraction. For example, in the previous version of the draft, we established this result for the truth-or-noise signals (Lewis and Sappington, 1994). However, some signal restrictions are necessary: If the finer segmentation separates high-value buyers from medium-value buyers while pooling them with low-value buyers, the price charged to those buyers may decrease, exacerbating payoff inequality.

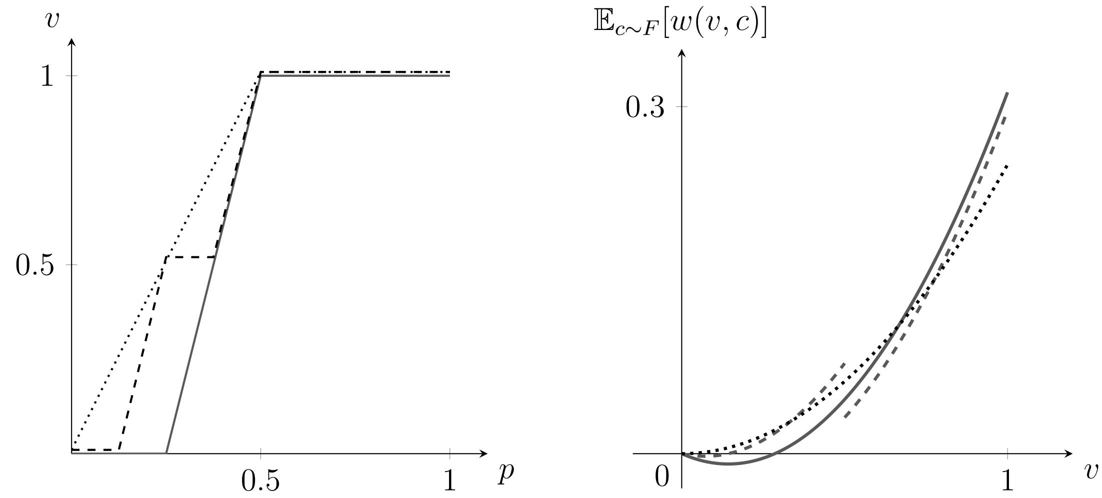

Figure 3: Buyer-optimal algorithm thresholds (left) and buyer surplus at different values (right). Computed at no segmentation (solid), binary partition (dashed), and full segmentation (dotted). $v \sim U[0,1]$, $c \sim U[0,1]$.

tations, which, in the case of uniform monotone partitions, correspond to partitions with n_H and n_L elements such that $n_H/n_L \in \mathbb{N}$. ◊

An important takeaway from our analysis in this section is that consumers can win the technological race against sellers. Advances in information technology not only drive the proliferation of algorithmic recommendations but also enable sellers to engage in third-degree price discrimination. Our neutrality results show that the benefits of buyer-optimal algorithmic consumption can fully offset the harm caused by price discrimination. However, the choice of the algorithm is crucial for this, e.g., an ex-post optimal algorithm could erode all buyer welfare under perfect price discrimination.

5 Extensions

5.1 Pareto-Optimal Algorithmic Recommendations

Thus far, we have focused on *buyer-optimal* algorithmic recommendations. A natural and broader question is what algorithmic recommendations are *Pareto optimal*, i.e., which buyer surplus and seller profit cannot be simultaneously improved upon. In Lemma 1, we showed that such recommendations are induced by threshold algorithms. In this section, we show that the characterization of Pareto-optimal algorithms is fully analogous to the characterization of a buyer-optimal algorithm.

To this end, assume that the designer's objective is a weighted average of buyer surplus and seller profit, with weights of α and $1-\alpha$, respectively. Call an algorithm *α-optimal* if it maximizes this objective for weight α. As α spans $[0,1]$, the α-optimal algorithms span the range from seller-optimal to socially-optimal to buyer-optimal. Define an α-virtual cost as follows:

$$\gamma_\alpha(c) \triangleq c + \max\left\{\frac{2\alpha-1}{\alpha}, 0\right\} \frac{F(c)}{f(c)}, \tag{6}$$

and assume that for all $\alpha \in [0,1]$, $\gamma_\alpha(c)$ is strictly increasing.[14] Define an α-pseudo value as:

$$y_\alpha(v) \triangleq \mathbb{E}_{\tilde{v}\sim G}[\gamma_\alpha^{-1}(\tilde{v})|\tilde{v} \geq v]. \tag{7}$$

Proposition 5 (α-Optimal Algorithm). *An α-optimal algorithm recommends the product if and only if $y_\alpha(v) \geq p$. Under this algorithm, type c posts $p^*(c) = y_\alpha(\gamma_\alpha(c))$. Under this algorithm and pricing, the trade occurs if and only if $v \geq \gamma_\alpha(c)$.*

Proposition 5 shows that the characterization of an α-optimal algorithm follows verbatim the characterization of a buyer-optimal algorithm, with the virtual cost and the pseudo value replaced by their α-analogs. Importantly, the product allocation continues to be independent of value distribution. Consequently, all the results on market segmentation in Section 4 apply to any Pareto-optimal algorithmic recommendations.

As the weight α attached to buyer surplus decreases, the difference between an α-virtual cost and a true cost decreases. By Proposition 5, this translates into the trade occurring over a broader range of costs and values, thus generating more total surplus.

[14] As before, if for some $\alpha \in [0,1]$, $\gamma_\alpha(c)$ were not everywhere increasing, then one would simply use an ironed version of it.

Moreover, for all $\alpha \leq 1/2$, the α-virtual cost coincides with the true cost. Therefore, a seller-optimal algorithm and a socially-efficient algorithm coincide and, as shown below, feature a simple recommendation structure:

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

The seller-optimal algorithm maximizes efficiency at the expense of the buyer. 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. Thus, the seller, regardless of cost, understands that she captures all the surplus generated, and her goal of maximizing profit aligns perfectly with efficiency. As a result, 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 product allocation is efficient, the seller obtains the maximal feasible surplus, and the buyer is left with no rent.[15]

Corollary 1 entails a notable feature: The persuasion constraint of the buyer, i.e., the requirement that the buyer is always willing to follow recommendations, does not constrain the designer at all points of the Pareto frontier except the seller-optimal one. Intuitively, by Corollary 1, α-optimal algorithms for $\alpha \in [0, 1/2]$ induce the same, efficient outcome, where the persuasion constraint holds with equality but does not constrain the designer. In turn, for $\alpha > 1/2$, as the weight given to buyer surplus increases, the recommendations become even more favorable to the buyer, and the persuasion constraint becomes slack.

[15] Gottardi and Mezzetti (2024) use a similar argument to construct an efficient one-shot mediation mechanism.

5.2 Competing Sellers

Thus far, we have focused on a monopoly setting in which a given product can be supplied only by a single seller. Applied seller-by-seller, this analysis also covers cases where multiple sellers offer noncompeting products. However, a natural alternative is a market in which sellers compete for the buyer so that the recommendation algorithm directs the buyer to one out of many alternatives (e.g., Hagiu and Jullien (2011), Hagiu et al. (2022), Elliott et al. (2022), Bar-Isaac and Shelegia (2022)). In this section, we show that our characterization of optimal algorithms and the main welfare implications extend to that setting.

Formally, we consider the following extension of the main setting. There is a single buyer with unit demand. There are J sellers indexed by $j = 1, \ldots, J$, each offering a single product. The buyer values for the products, $(v_1, ..., v_J) \in \mathbb{R}^J$, are drawn from a joint distribution $G \in \Delta([0,1]^J)$ and can be arbitrarily correlated. The cost of each seller j is drawn from F_j, independent of other costs or the value profile. For notational convenience, we introduce a *dummy seller* indexed by $j = 0$ with $v_0 = 0$ and $c_0 = 0$ that corresponds to the buyer's decision not to buy anything.

An algorithm is a function $r : [0,1]^J \times \mathbb{R}_+^J \to \Delta^{J+1}$, where Δ^{J+1} is a $J+1$ dimensional simplex, so that for any profiles of realized values $v = (v_1, ..., v_J)$ and prices $p = (p_1, ..., p_J)$, the algorithm recommends that the buyer purchase one of the products or none according to $r(v, p) \in \Delta^{J+1}$. The algorithm is commonly known to the buyer and sellers. Given an algorithm, nature draws the seller types c_j and the buyer values v_j. All sellers privately observe their types but not the buyer values or the types of other sellers and simultaneously post their prices p_j. The algorithm makes recommendations according to r. If no product is recommended, trade does not occur. If a product of seller j is recommended, the buyer observes the recommendation and the price and then decides whether to buy the product. If product j is purchased, the buyer and seller j obtain ex post payoffs $v_j - p_j$ and $p_j - c_j$, respectively. Otherwise, the players obtain zero payoffs. The solution concept is perfect Bayesian equilibrium.

Despite featuring strategic interaction between the sellers, this setting can be analyzed analogously to the single seller case. The main idea is that each seller's private information and thus the incentive constraints are similar in both cases: From the perspective of each seller, the value and cost uncertainty, as well as the strategic behavior of other sellers, matter only insofar as they affect her demand curve, which can be encoded in a single variable.

Specifically, for each j, denote by $\gamma_j(c_j)$ her virtual cost. For the dummy seller 0, set $\gamma_0(c_0) = 0$. Assume that for $j = 1, \ldots, J$, γ_j is strictly increasing and continuous in c_j. Define $\bar{c}_j \triangleq \gamma_j^{-1}(1)$. Define an auxiliary random variable

$$\theta_j = v_j - \max_{k \in \{0,1,\ldots,J\} \setminus j} \{v_k - \gamma_k(c_k)\},$$

i.e., θ_j is the value of seller j's product minus the highest virtual surplus among all other sellers as long as the latter is positive. Define

$$p_j^*(c_j) \triangleq \mathbb{E}_{\theta_j}[\gamma_j^{-1}\big(\theta_j\big) \mid \gamma_j^{-1}(\theta_j) \geq c_j], \tag{8}$$

and observe that $p_j^*(c_j)$ is a strictly increasing function. Define the inverse function of p_j^* as p_j^{*-1} with the (nonstandard) convention that $p_j^{*-1}(p) = 0$ for $p < p_j^*(0)$ and $p_j^{*-1}(p) = 1$ for $p > p_j^*(1)$.

Proposition 6 (Buyer-Optimal Algorithm with Competing Sellers). *A buyer-optimal algorithm recommends the product of seller $j^*(v, p)$ such that*

$$j^*(v, p) \in \operatorname*{argmax}_{j \in \{0,1,\ldots,J\}} \; v_j - \gamma_j\left(p_j^{*-1}(p_j)\right), \tag{9}$$

with ties being broken arbitrarily. Under this algorithm, seller j of type $c_j \leq \bar{c}_j$ posts price $p_j^*(c_j)$ and seller j of type $c_j > \bar{c}_j$ is inactive. Under this algorithm and pricing, for any realized profile of v and c, the buyer trades with seller $j^* \in \operatorname{argmax}_{j \in \{0,1,\ldots,J\}} v_j - \gamma_j(c_j)$.

Proposition 6 directly extends Proposition 1. To see this, observe that in the case of $J = 1$, the condition $v_1 - \gamma_1\left(p_1^{*-1}(p_1)\right) \geq v_0 - \gamma_0\left(p_0^{*-1}(p_0)\right) = 0$ is equivalent to the condition $y(v_1) \geq p_1$; thus, the two propositions describe the same algorithm albeit in different terms. In the case of many sellers, the condition $v_i - \gamma_i \left(p_i^{*-1}(p_i)\right) \geq v_j - \gamma_j \left(p_j^{*-1}(p_j)\right)$ for $i, j \neq 0$ cannot be easily translated into the language of pseudo values, so we present the buyer-optimal algorithm as in (9).

Importantly, as in the case of a single seller, Proposition 6 establishes that the equilibrium product allocation does not depend on the distribution of product values. Similarly, it allows us to succinctly analyze the impact of market segmentation. Formally, market segmentation is defined by an information structure $\mathcal{I} = (S, \pi)$ that consists of a set $S = \times_j S_j$ of signal realizations s_j privately observed by each seller, and a family of probability distributions $\{\pi(\cdot|v)\}_{v \in [0,1]^J}$ over S. The signal is commonly known and exogenous, the signal realizations are independent of the seller types but can be arbitrarily correlated across sellers, and the algorithm can base recommendations on the realized signals, the values, and the product prices.[16]

Proposition 7 (Segmentation Neutrality with Competing Sellers). *For any market segmentation, the buyer-optimal algorithm induces the same ex post product allocation, the same expected price and profit of each seller type, and the same ex ante buyer surplus.*

Proposition 7 implies the remarkable neutrality of market segmentation if the buyer uses an algorithm to guide consumption choices. As in Section 4, leaking buyer data not only does not harm the buyer but also cannot benefit him, despite competition among sellers. Intuitively, because the algorithm can assess the buyer's value and is designed prior to pricing decisions, it shifts the bargaining power to the buyer, and informing the sellers can only reduce the attainable buyer surplus. Furthermore, by adapting to the specifics of market segmentation, the optimal algorithm design can perfectly absorb the impact of information leakage on total buyer surplus, seller profits, and product allocation.

At the same time, market segmentation does affect equilibrium pricing and the re-

[16]As before, the assumption that the algorithm can base recommendations on the realized signals captures the idea that the sellers have no information beyond that accessed by the algorithm. This assumption is automatically satisfied if the signals are deterministic functions of the value profile, i.e., partitional or fully informative signals.

distribution of buyer surplus across different value profiles. Our analysis behind Proposition 7 reveals that equilibrium pricing can be decomposed across sellers, with each seller's pricing strategy depending only on her beliefs about the value profile, and being indifferent to information observed by other sellers. This immediately allows us to claim an impact of finer market segmentation on prices. Specifically, say that $\mathcal{I}_H$ is Blackwell more informative than $\mathcal{I}_L$ for seller j, if the corresponding marginal (S_j, π_j) with $\pi_j : [0,1]^J \to \Delta(S_j)$ derived from (S, π) is Blackwell more informative about v. Then, we have:

Proposition 8 (Segmentation Redistribution with Competing Sellers). *If signal $\mathcal{I}_H$ is Blackwell more informative than $\mathcal{I}_L$ for seller j, then for any given type of seller j, the distribution of prices set by seller j under $\mathcal{I}_H$ is a mean-preserving spread of that under $\mathcal{I}_L$.*

By Proposition 8, any finer market segmentation—regardless of how the additional information is correlated across sellers—results in a clear pattern of more dispersed prices, just as under a monopoly. However, the impact on surplus distribution is subtler than that under a monopoly. With multiple sellers, finer segmentation may group lower values for one seller's product with higher values for another seller's product, leading to a higher trade price and lower surplus for some low-value buyers, thus violating the mean-preserving contraction property.

5.3 Informed Buyer

Up to this point, we have deliberately assumed that the algorithm has full control over the buyer's information both about the product's existence and about the product's value. This assumption offered a clear benchmark for studying algorithmic recommendations, provided the algorithm with maximal information to control and thus established the upper bound on achievable buyer surplus. In this section, we relax this assumption and show how our analysis remains relevant even if the buyer is partially informed.

5.3.1 Information about Product Existence

We have assumed that the buyer cannot purchase the product if it is not recommended. This assumption is relevant in online settings where recommendation systems are used primarily to discover and bring products to the buyer's attention, and thus effectively serve as gatekeepers.[17] A natural alternative setting is one in which the buyer already knows the product exists and where to purchase it but may still be unsure about the match value. To address that setting, in this subsection, we allow the buyer to purchase the product even if it is not recommended, and impose a constraint that the buyer prefers not to purchase the nonrecommended product at each price.

First, consider the simplest case in which the seller's cost is commonly known to be c_0. This case is closest to the paper by Roesler and Szentes (2017) and differs only in the timing of the recommendations. In their setting, recommendations come before the price is posted and thus cannot condition on the price; in our setting, the recommendations can condition on the price.

When the costs are known to be c_0, a natural candidate for a buyer-optimal algorithm is to recommend the product if and only if $p = c_0$ and $v \geq c_0$. If this algorithm suffices to incentivize the seller to set $p = c_0$, then it is buyer-optimal because the outcome is efficient and leaves the seller with zero profit.

If $\mathbb{E}[v] \leq c_0$, an algorithm can attain this outcome by revealing no information whenever $p > c_0$ and thus dissuading the buyer from purchasing at a price above c_0. In contrast, if $\mathbb{E}[v] > c_0$, the seller can secure a positive profit by charging a price in $(c_0, \mathbb{E}[v])$ because no algorithm can make the buyer believe that the expected value of the product is always below $p < \mathbb{E}[v]$. In this case, the buyer-optimal algorithm deters the seller from setting a higher price via *adversarial persuasion*: At each price, the algorithm provides the buyer with information that minimizes the probability of trade.

[17]This assumption is consistent with the consideration set approach to model recommender systems. See, for example, Dinerstein et al. (2018) and Lee and Musolff (2023).

Specifically, the algorithm reveals whether $v < \hat{v}(p)$, where $\hat{v}(p)$ is such that

$$\mathbb{E}[\tilde{v}|\tilde{v} < \hat{v}(p)] = \min\{p, \mathbb{E}[v]\}, \tag{10}$$

and persuades the buyer to purchase only when $v \geq \hat{v}(p)$. The maximum profit the seller can guarantee against adversarial persuasion is

$$\underline{\pi} \triangleq \max_{p \geq c_0}(p - c_0)[1 - G(\hat{v}(p)]. \tag{11}$$

The buyer-optimal algorithm induces an efficient trade while leaving the seller with this profit:

Proposition 9 (Known Product, Known Cost). *Suppose that F is concentrated at $c_0 \in [0, 1)$ and that the buyer can purchase the product even when not recommended. The buyer-optimal algorithm recommends the product if $p^* = c_0 + \frac{\underline{\pi}}{1-G(c_0)}$ and $v \geq c_0$ and follows all other prices with adversarial persuasion. Under this algorithm, the seller posts a price p^*, and the equilibrium trade is efficient.*

Like in the setting of Roesler and Szentes (2017), a buyer-optimal algorithm leads to efficient trade. Unlike the setting of Roesler and Szentes (2017), the seller's rent is driven by adversarial persuasion price-by-price and thus is lower, reaching zero when $\mathbb{E}[v] < c_0$, e.g., when the product can be counterfeit or harmful with a high probability.

Recall from Claim 1 that if the seller's cost were known to be c_0 in our original setup, the buyer-optimal algorithm would recommend the product if and only if $p = c_0$ and $v \geq c_0$. This algorithm always attains the efficient outcome and leaves the seller with zero profit. Therefore, when $c_0 < \mathbb{E}[v]$, the design and consequences of the optimal algorithm depend on whether the buyer can purchase the nonrecommended product, whereas if $c_0 > \mathbb{E}[v]$ they do not depend on it.

Furthermore, when the seller's costs are uncertain, the buyer-optimal algorithms in both cases can coincide even when the cost could fall below $\mathbb{E}[v]$. This happens, for example, in the uniform setting of Example 1: Under the buyer-optimal algorithm, the lack of recommendation is a sufficiently negative signal at any price to dissuade the buyer from the purchase. More generally:

Proposition 10 (Known Product. Unknown Cost). *If* $\int_0^{\gamma(c)}[v-c]\mathrm{d}G(v) \leq 0$ *for each* $c \in [0, \overline{c}]$*, then even if the buyer can purchase the product when not recommended, the algorithm in Proposition 1 is buyer-optimal.*

The condition of Proposition 10 ensures that whenever the product is not recommended under the algorithm of Proposition 1, the buyer infers that the expected value of the product is below the price. This holds for many classes of distributions, 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.[18] To see the intuition behind this sufficient condition, suppose that the seller posts a price of $p^*(c)$, and the algorithm recommends that the buyer not purchase the product, which by Proposition 1, reveals that $v \leq \gamma(c)$. If the buyer purchases the product, his payoff must decrease because the seller's profit increases but total surplus decreases because $\int_0^{\gamma(c)}[v-c]\mathrm{d}G(v) \leq 0$. Thus, the buyer is willing to follow the recommendation not to buy.

5.3.2 Information about Product Value

Thus far, we have assumed that the buyer does not obtain any product information beyond what is provided by the algorithm. This stylized assumption is intended to capture the context of experience goods, which can be difficult to judge based on appearance and for which individual taste shocks are sufficiently variable yet can be estimated by a well-trained recommendation system. In this section, we allow for the possibility that the buyer observes additional information about the value when a product is recommended while maintaining the original ignorance of the product's existence.[19]

We show how our previous analysis informs this setting. First, observe that the buyer's incentives in the buyer-optimal algorithm of Proposition 1 are generally slack.

[18] Point (i) follows from direct calculation. For Point (ii), note that the condition $(\frac{F(c)}{c})' \geq 0$ is written as $\frac{\gamma(c)}{c} \leq 2$, which is equivalent to $\int_0^{\gamma(c)}[v-c]\mathrm{d}G(v) \leq 0$ when $G = U[0,1]$.

[19] The remaining case of a buyer perfectly informed about the product's existence and value is a textbook monopoly setting.

That is, whenever a product is recommended, the buyer strictly prefers to follow the recommendation. Therefore, a small amount of extraneous information that does not significantly lower the posterior expectation will not interfere with the algorithm's design.

Second, our market segmentation analysis implies that even if the buyer can perfectly assess the value of the product upon seeing it, the algorithm can still achieve the same total buyer surplus, seller surplus, and product allocation, although this would require informing both the seller and the buyer. Specifically, suppose that the buyer observes the value v of the product whenever it is recommended. When the seller does not know v, the algorithm in Proposition 1 is not incentive compatible because the buyer will ignore the recommendation when p and v are such that $v < p < y(v)$, which occurs for low values. However, if the algorithm perfectly informs the seller about v, then the buyer-optimal algorithm, as characterized in Section 4, recommends the product if and only if $p \leq \gamma^{-1}(v) < v$. Under this algorithm, buyers with all values v are willing to follow the recommendations because, intuitively, informed sellers lower the prices offered to low-value buyers. By Proposition 3, this algorithm implements the same total buyer surplus, seller profits, and product allocation. Remarkably, in the case of consumption driven by algorithmic recommendations, third-degree price discrimination not only does not harm the total buyer surplus but also may be beneficial if the algorithm cannot fully control the buyer value information.

6 Conclusion

In this paper, we studied the question of optimal algorithmic recommendations in the presence of strategic pricing. We showed that optimal recommendations must strike a balance between increasing the trade surplus and inducing low prices. Algorithmic recommendations drastically change the predictions of third-degree price discrimination, whereby finer market segmentations by the sellers do not affect the total consumer surplus or seller profits but result in larger price spreads and a more equitable surplus distribution.

We view our work as a stepping stone toward a better understanding of algorithmic design in strategic settings, an area of growing importance at the intersection of economics and computer science (e.g., Goktas et al. (2025)). First, our model of algorithms is deliberately stylized to analyze strategic motives in a clear and tractable way. A practical implementation would ideally incorporate many engineering concerns from which we abstracted away, such as value estimation details, computational complexity, and robustness. Second, it would be interesting to study market structures for algorithm providers and understand which of the algorithms that we characterize are favored by one or another market structure. Third, the developed ideas of algorithmic decisions can be exported beyond consumption settings, such as to algorithmic matching or algorithmic negotiations. All this further research can be built upon the analytical framework proposed in this paper.

References

Akbarpour, M., P. Dworczak, and S. D. Kominers (2024): "Redistributive Allocation Mechanisms," *Journal of Political Economy*, 132, 1831–1875.

Asker, J., C. Fershtman, and A. Pakes (2022): "Artificial Intelligence, Algorithm Design, and Pricing," *AEA Papers and Proceedings*, 112, 452–456.

Assad, S., R. Clark, D. Ershov, and L. Xu (2024): "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy*, 132, 723–771.

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. and A. Bonatti (2024): "Data, Competition, and Digital Platforms," *American Economic Review*, 114, 2553–2595.

Bergemann, D., A. Bonatti, and N. Wu (forth.): "How Do Digital Advertising Auctions Impact Product Prices?" *Review of Economic Studies*.

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*.

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 (forth.): "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Review of Economic Studies*.

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.

Dinerstein, M., L. Einav, J. Levin, and N. Sundaresan (2018): "Consumer price search and platform design in internet commerce," *American Economic Review*, 108, 1820–1859.

Doval, L. and A. Smolin (2024): "Persuasion and Welfare," *Journal of Political Economy*, 132, 000–000.

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 (2024): "Artificial Intelligence Act: Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence," https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai, accessed: September 05, 2024.

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

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

Goktas, D., A. Greenwald, T. Osogami, R. Patel, K. Leyton-Brown, G. Schoenebeck, D. Cornelisse, C. Daskalakis, I. Gemp, J. Horton, D. C. Parkes, D. M. Pennock, A. Prakash, S. S. Ravindranath, M. O. Smith, G. Swamy, E. Vinitsky, M. Wasserkrug, Segev Wellman, J. Wu, H. Xu, J. Zhang, Y. Zhang, S. Zhao, and Q. Zhu (2025): "Strategic Foundation Models," *Working paper*.

Gottardi, P. and C. Mezzetti (2024): "Shuttle Diplomacy," *Journal of Economic Theory*, 216, 105794.

Haghpanah, N. and R. Siegel (2023): "Pareto-Improving Segmentation of Multiproduct 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.

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 (2023): "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.

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. (2023): "Disclosure and Pricing of Attributes," *Rand Journal of Economics*, 54, 570–597.

White House (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: September 05, 2024.

Xu, W. and K. H. Yang (2024): "Equivalent Mechanisms for Information Intermediation," *Working paper*.

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

Appendix: Ommited Proofs

A 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. We 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 the 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. □

B 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:

$$\max_{\hat{v}:[0,1]\to[0,1],\, p:[0,1]\to\mathbb{R}_+} \int_0^1 \int_{\hat{v}(c)}^1 (v - p(c))\,\mathrm{d}G\,\mathrm{d}F, \tag{12}$$

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

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

One 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,

$$\frac{\mathrm{d}V}{\mathrm{d}q} = \frac{\partial V/\partial \hat{v}}{\partial q/\partial \hat{v}} = \frac{-\hat{v}g(\hat{v})}{-g(\hat{v})} = \hat{v}(q). \tag{13}$$

As such, $V(q)$ is a concave function with $V'(0) = 1$ and $V'(1) = 0$. Finally, we denote the expected revenue by $t(c) \triangleq p(c) \int_{\hat{v}(c)}^{1} \mathrm{d}G$. In terms of these variables, we can restate problem (12) as follows:

$$\begin{aligned} \max_{q:[0,1]\to[0,1],\, t:[0,1]\to\mathbb{R}_+} \quad & \int_0^1 (V(q(c)) - t(c))\,\mathrm{d}F, \\ \text{s.t.}\ & t(c) - cq(c) \geq t(c') - cq(c') \quad \forall\, c, c' \in [0,1], \\ & t(c) - cq(c) \geq 0 \qquad\qquad\qquad \forall\, c \in [0,1]. \end{aligned} \tag{14}$$

Problem (14) 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:

$$\begin{aligned} V'(q(c)) &= \gamma(c), \\ t(c) - q(c)c &= \int_c^1 q(x)\,\mathrm{d}x = \int_c^1 1 - G(\gamma(x))\,\mathrm{d}x. \end{aligned}$$

By Equation 13, we can translate this solution back to problem (12) as

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

We now show that the algorithm and the equilibrium in the statement attain the same outcome as above. First, the buyer is willing to purchase the product when recommended because

$$
\mathbb{E}[v|y(v) \geq p(c)] = \mathbb{E}[v|y(v) \geq y(\gamma(c))] = \mathbb{E}[v|v \geq \gamma(c)] \geq \mathbb{E}[\gamma^{-1}(v)|v \geq \gamma(c)] = p(c).
$$

Thus, the expected value of the product conditional on each possible price exceeds the price.

Second, the seller with each type c is willing to set price $y(\gamma(c))$. Deviating to another price in $[y(\gamma(0)), y(\gamma(\bar{c}))]$ is not profitable because of the incentive compatibility constraints of the mechanism. Deviating to a price below $y(\gamma(0))$ or above $y(\gamma(\bar{c}))$ is not profitable either because it results in a lower profit than $p = y(\gamma(0))$ or no trade.

Finally, if the buyer follows the recommendation and each type c sets price $y(\gamma(c))$, the trade occurs if and only if $y(v) \geq y(\gamma(c))$, or equivalently, if $v \geq \hat{v}(c) = \gamma(c)$.

In summary, the algorithm described in the statement implements the solution to problem (12) in equilibrium. Therefore, it is a buyer-optimal algorithm. □

C Proof of Proposition 2

Under the buyer-optimal algorithm, trade occurs if and only if $v \geq \gamma(c)$. If F has a decreasing reversed hazard rate, then $\gamma(c) = c + \frac{F(c)}{f(c)}$ implies that $\gamma'(c) \geq 1$. Under the ex post optimal algorithm, trade occurs if and only if $v \geq r^{-1}(c)$, where $r(v) = v - \frac{1-G(v)}{g(v)}$ is the virtual valuation function. If G has an increasing hazard rate, then $r'(v) > 1$, which implies that

$$\frac{d}{dc}r^{-1}(c) = \frac{1}{r'(r^{-1}(c))} < 1.$$

Thus, $\gamma(c)$ is steeper than $r^{-1}(c)$, meaning that in the (c, v)-space, the curve $v = \gamma(c)$ crosses the curve $v = r^{-1}(c)$ at most once and from below.

Since $\gamma(c)$ starts at 0 when $c = 0$ and ends above 1 when $c = 1$, while $r^{-1}(c)$ starts above 0 at $c = 0$ and ends at 1 at $c = 1$, it follows that for low values of c, $\gamma(c) < r^{-1}(c)$, whereas for high values of c, $\gamma(c) > r^{-1}(c)$. Consequently, there exists a unique crossing point c^* such that $\gamma(c^*) = r^{-1}(c^*)$. We then obtain qualitatively the same picture as the right panel of Figure 2. This establishes the result. □

D Proofs of Proposition 3 and Proposition 4

To prove Proposition 3, we establish a version of the payoff equivalence theorem for our model (cf. Myerson (1981) and Krishna (2009)).

Lemma 2 (Payoff Equivalence). *For each $i \in \{1, 2\}$, take an algorithm r_i, market segmentation $\mathcal{I}_i$, and equilibrium $\mathcal{E}_i$. Suppose that $\mathcal{E}_1$ and $\mathcal{E}_2$ have the same allocation rule in terms of v and c and the same profit of the seller at type $c = 1$. Then, the seller's profit of any type and the buyer surplus are identical between $\mathcal{E}_1$ and $\mathcal{E}_2$.*

Proof. Let $q(v, c)$ denote the probability of a trade when the value is v and the type is c, and let $q(c) = \int_0^1 q(v, c)\, \mathrm{d}G(v)$ denote the expected probability of a trade for type c. Upon calculating these objects, we take expectation with respect to the possible segments. Let $\overline{\pi}$ denote the profit of the seller with the highest type, $c = 1$. By assumption, $q(\cdot, \cdot)$ and $\overline{\pi}$ are the same between $\mathcal{E}_1$ and $\mathcal{E}_2$. Additionally, let $\pi_i(c)$ and $t_i(c)$ denote the profit and the expected monetary transfer, respectively, at type c in equilibrium $\mathcal{E}_i$.

In equilibrium $\mathcal{E}_i$, type $c = 1$ cannot earn a strictly higher profit by imitating the pricing strategy of type c' in every segment in $\mathcal{I}_i$. This incentive compatibility constraint is

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

The envelope theorem implies that

$$\pi_i(c) = \overline{\pi} + \int_c^1 q(x)\, \mathrm{d}x.$$

The right-hand side does not depend on i. Thus, the seller's profit is the same between $\mathcal{E}_1$ and $\mathcal{E}_2$ for every seller type. The buyer surplus is the same between $\mathcal{E}_1$ and $\mathcal{E}_2$ because it is the total surplus from allocation rule $q(\cdot, \cdot)$ minus the seller's profit, neither of which depends on i. □

Proof of Proposition 3

Take any signal, $\mathcal{I}$. For any signal realization, the optimal algorithm induces a trade if and only if $v \geq \gamma(c)$. Hence, the ex post allocation of the product is independent of the signal, as is the total surplus. Furthermore, the highest seller type $c = 1$ always earns zero profits. Lemma 2 then implies that the seller profit of all types and the buyer surplus are independent of the signal. □

Proof of Proposition 4

First, we show that as the signal becomes more informative, the distribution of prices set by each active seller type undergoes a mean-preserving spread. To see this, consider any signals $\mathcal{I}_H$ and $\mathcal{I}_L$ such that $\mathcal{I}_H$ is more informative than $\mathcal{I}_L$. Recall that $\hat{\pi}_H$ and $\hat{\pi}_L$ denote the respective ex ante distributions of the signal realizations.

Fix any active type c. For each $\alpha \in \{L, H\}$, let $G^c_{\alpha,s} \in \Delta[0, 1]$ denote the posterior distribution of value v conditional on (i) signal s being realized under signal $\mathcal{I}_\alpha$ and (ii) $v \geq \gamma(c)$. The equilibrium price of type c after observing signal realization s under signal $\mathcal{I}_\alpha$ is

$$p(c|s,\alpha) \triangleq \int_0^1 \gamma^{-1}(v)\mathrm{d}G^c_{\alpha,s}(v). \tag{15}$$

Let $\mathcal{G}^c_\alpha \in \Delta\Delta[0,1]$ denote the distribution of posteriors $G^c_{\alpha,s}$ for a fixed c. Specifically, $\mathcal{G}^c_\alpha$ is the distribution of random variable $G^c_{\alpha,s}$ with $s \sim \hat{\pi}_\alpha$. Because signal $\mathcal{I}_H$ is more informative than signal $\mathcal{I}_L$, $\mathcal{G}^c_H$ is a mean-preserving spread of $\mathcal{G}^c_L$.[20] As the price is linear in posterior $G^c_{\alpha,s}$, the mean-preserving spread relation between the distributions of posteriors, $\mathcal{G}^c_H$ and $\mathcal{G}^c_L$, imply the mean-preserving spread relation between real-valued random variables $p(c|s,H)$ and $p(c|s,L)$. Therefore, we conclude that $p(c|s,H)$ with $s \sim \hat{\pi}_H$ is a mean-preserving spread of $p(c|s,L)$ with $s \sim \hat{\pi}_L$. Therefore, the distribution of prices set by each active type under signal $\mathcal{I}_H$ is a mean-preserving spread of the price distribution under signal $\mathcal{I}_L$.

Second, we establish the results on monotone partitions. In what follows, we view a monotone partitional signal as a partition of $[0,1]$ and use an "interval" to mean an interval with a positive length, excluding a singleton set.

Take any monotone partitional signals, $\mathcal{I}_H$ and $\mathcal{I}_L$, such that $\mathcal{I}_H$ is finer than $\mathcal{I}_L$. We can create partition $\mathcal{I}_H$ by applying the following operations finitely many times to partition $\mathcal{I}_L$: (i) taking an interval from $\mathcal{I}_L$ and dividing it into two subintervals or (ii) taking an interval from $\mathcal{I}_L$ and fully revealing the values within it. The latter operation means partitioning interval $[a,b]$ into $\{\{v\}\}_{v\in[a,b]}$. To obtain our result, it suffices to show that applying (i) or (ii) to any given monotone partitional signal leads to a mean-preserving contraction of the buyer surplus at any seller type. We consider these two operations in turn.

Operation (i). Fix any monotone partitional signal $\mathcal{I}$ that is different from the fully informative signal. Suppose that we take interval $[v_i, v_{i+1}]$ from $\mathcal{I}$ and split it into $[v_i, \hat{v}]$ and $[\hat{v}, v_{i+1}]$ for some $\hat{v} \in (v_i, v_{i+1})$. We show that after this partitioning, the

[20] Distribution $\mathcal{G}_H$ of posteriors being a mean-preserving spread of $\mathcal{G}_L$ means that there exist $\Delta[0,1]$-valued random variables Z_H and Z_L such that $Z_H \sim \mathcal{G}_H, Z_L \sim \mathcal{G}_L$ and $\mathbb{E}(Z_H \mid Z_L) = Z_L$.

buyer surplus $w(v,c)$, when we fix c but draw v from G, undergoes a mean-preserving contraction.

First, we consider the values and types that are affected by the operation, i.e., (v,c) such that $\gamma(c) \le v < v_{i+1}$. Before the operation, the buyer's ex post payoff is

$$w_0(v,c) \triangleq v - \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\tilde{v} \in [\gamma(c), v_{i+1}]], \forall v \in [\gamma(c), v_{i+1}].$$

After the operation, the buyer's ex post payoff is

$$w_1(v,c) \triangleq \begin{cases} v - \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\tilde{v} \in [\gamma(c), \hat{v}]] & \text{if} \quad v \in [\gamma(c), \hat{v}], \\ v - \mathbb{E}_{\tilde{v}\sim G}[\gamma^{-1}(\tilde{v})|\tilde{v} \in [\hat{v}, v_{i+1}]] & \text{if} \quad v \in [\hat{v}, v_{i+1}]. \end{cases}$$

Note that by applying Operation (i), the ex post payoff of the buyer with value $v \in [\gamma(c), \hat{v}]$ increases because of the lower price and that of $v \in [\hat{v}, v_{i+1}]$ decreases because of the higher price.

For each $k \in \{0,1\}$, consider the distribution of $w_k(v,c)$ when $v \sim G(\cdot|\tilde{v} \in [\gamma(c), v_{i+1}])$. First, they have the same mean because the expected price remains the same before and after the operation. Second, because $w_1(\cdot,c)$ crosses $w_0(\cdot,c)$ once from above, the CDF of $w_1(v,c)$ crosses the CDF of $w_0(v,c)$ once from above. The equal mean property and the single-crossing property imply, by Theorem 3.A.44 (Condition 3.A.59) of Shaked and Shanthikumar (2007), that $w_0(v,c)$ is a mean-preserving spread of $w_1(v,c)$ when $v \sim G(\cdot|\tilde{v} \in [\gamma(c), v_{i+1}])$.

Therefore, for a fixed c, the buyer's ex post surplus conditional on $v \in [\gamma(c), v_{i+1}]$ under signal $\mathcal{I}_L$ is a mean-preserving spread of that under signal $\mathcal{I}_H$. The same relationship trivially holds for the ex post surpluses of value $v < \gamma(c)$ or $v > v_{i+1}$ because those types do not trade or continue to face the same price. In summary, for any fixed c, the buyer's ex post surplus under signal $\mathcal{I}_L$ is a mean-preserving spread of that under signal $\mathcal{I}_H$ conditional on each of the three cases, $v \in [\gamma(c), v_{i+1}]$, $v < \gamma(c)$, and $v > v_{i+1}$. The mean-preserving spread relationship is closed under mixtures (e.g., Theorem 3.A.12(b) of Shaked and Shanthikumar (2007)). Thus, for any fixed c, the distribution of the buyer's surplus under signal $\mathcal{I}_L$ is a mean-preserving spread of the distribution of the buyer's surplus under signal $\mathcal{I}_H$.

Operation (ii). We can apply the same logic as the case of Operation (i) by defining $w_1(v,c)$ as

$$w_1(v,c) \triangleq v - \gamma^{-1}(v), \forall v \in [v_i, v_{i+1}].$$

□

E Proof of Proposition 5

The proof follows that of Proposition 1. The designer's problem is stated as follows:

$$\max_{\hat{v}:[0,1]\to[0,1],\, p:[0,1]\to\mathbb{R}_+} \int_0^1 \int_{\hat{v}(c)}^1 \alpha(v - p(c)) + (1-\alpha)(p(c) - c)\,\mathrm{d}G\,\mathrm{d}F, \tag{16}$$

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

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

For $\alpha > 0$, the designer's objective can also be written as

$$\int_0^1 \int_{\hat{v}(c)}^1 v - p(c) + \frac{1-\alpha}{\alpha}(p(c) - c)\,\mathrm{d}G\,\mathrm{d}F.$$

Repeat the same step to rewrite $\int_{\hat{v}(c)}^1 v\,\mathrm{d}G$ as in the proof of Proposition 1. We can then write the designer's problem as

$$\max_{q:[0,1]\to[0,1],\, t:[0,1]\to\mathbb{R}_+} \int_0^1 \left[V(q(c)) - t(c) + \frac{1-\alpha}{\alpha}(t(c) - cq(c))\right] \mathrm{d}F, \tag{17}$$

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

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

We consider two cases.

Case 1: $\alpha \geq \frac{1}{2}$ When the designer places a weakly higher weight on buyer surplus than seller profit (i.e., when $\alpha \geq \frac{1}{2}$), Problem (17) is analogous to the problem analyzed by Baron and Myerson (1982), where the weight "α" in their Lemma 2 is replaced by $\frac{1-\alpha}{\alpha}$ for our proof. Thus, its solution sets the optimal quantity to equalize marginal welfare benefits with α-virtual costs and the optimal transfer to guarantee the incentive-compatible profit distribution:

$$
\begin{aligned}
V'(q(c)) &= \gamma_\alpha(c), \\
t(c) - q(c)c &= \int_c^1 q(x)\,\mathrm{d}x = \int_c^1 1 - G(\gamma_\alpha(x))\,\mathrm{d}x.
\end{aligned}
$$

The rest of of the proof is identical with that of Proposition 1, where we replace $\gamma(c)$ in that proof with $\gamma_\alpha(c)$.

We use the following observation for the next case: At $\alpha = \frac{1}{2}$, the α-virtual costs equals the seller's true cost. Consequently, under the $\frac{1}{2}$-optimal algorithm, type c sets a price of

$$y_{\frac{1}{2}}(c) \triangleq \mathbb{E}_{\tilde{v}\sim G}[\tilde{v}|\tilde{v} \geq c], \tag{18}$$

and the algorithm recommends the product if and only if $v \geq c$. As a result, the product allocation is efficient and buyer surplus is 0.

Case 2: $\alpha < \frac{1}{2}$ In this case, the designer places a strictly higher weight on seller profit than buyer surplus. We show that the solution to the designer's problem with weight $\alpha' < \frac{1}{2}$ is equal to the solution with weight $\alpha = \frac{1}{2}$. Suppose to the contrary that between the solution with weight α' and the solution with weight α, either buyer surplus or seller profit is different. Then, both buyer surplus and seller profit must be different. For example, if the two solutions have the same buyer surplus but the solution with weight α' has a strictly higher seller profit, then the designer with $\alpha = \frac{1}{2}$ would mimic the designer with weight α'.

If both buyer surplus and seller profit are different, then it must be the case that under the solution with a higher weight on seller profit, α', the seller profit is strictly higher and the buyer surplus is strictly lower than under $\alpha = \frac{1}{2}$. However, it would mean that buyer surplus at α' would be negative because buyer surplus at $\alpha = \frac{1}{2}$ is 0. This is a contradiction, because in equilibrium, the buyer can secure zero payoffs by not purchasing anything. Therefore, the solution to the designer's problem with weight $\alpha' < \frac{1}{2}$ is equal to the solution at weight $\alpha = \frac{1}{2}$. In terms of the α-virtual costs, this means that we can use (6) for any $\alpha \in [0, 1]$. $\square$

F Proofs of Proposition 6 and Proposition 7

We prove a result that implies both Proposition 6 and Proposition 7 as corollaries. Assume that the sellers face an information structure $\mathcal{I} = (S, \pi)$ that consists of a set $S = \times_j S_j$ of signal realizations $s_j \in S_j$ privately observed by each seller and a family of probability distributions $\{\pi(\cdot|v)\}_{v \in [0,1]^J}$ over S. We write $\tilde{s}_j$ for seller j's signal as a random variable and $s_j \in S_j$ for a generic realization.

Denote the set of real sellers by $\mathcal{J} \triangleq \{1, ..., J\}$ and the set of all sellers, together with a dummy seller, by $\mathcal{J}_0 \triangleq \{0, 1, ..., J\}$. When we say that the buyer purchases from (or transacts with) seller 0, it means that the buyer does not purchase from any seller $j = 1, ..., J$. The profiles of the signal realizations, values, types, and prices are denoted as $\mathbf{s}, \mathbf{v}, \mathbf{c}$, and $\mathbf{p}$, respectively. When we refer to a profile that excludes seller j, we use notations such as $\mathbf{s}_{-j}$ and $\mathbf{v}_{-j}$. With a slight abuse of notation, we write F, G, F_{-j}, and G_{-j}, for the distributions of $\mathbf{c}$, $\mathbf{v}$, $\mathbf{c}_{-j}$, and $\mathbf{v}_{-j}$, respectively. Unless otherwise stated, these vectors and joint distributions exclude the dummy seller.

Recall that we defined an auxiliary random variable

$$\theta_j = v_j - \max_{k \in \mathcal{J}_0 \setminus \{j\}} \{v_k - \gamma_k(c_k)\}.$$

For each j, let $\overline{v}_j(s_j)$ be the supremum of the support of the posterior distribution of v_j conditional on $\tilde{s}_j = s_j$. Define $\overline{c}_j(s_j) = \gamma_j^{-1}(\overline{v}_j(s_j))$. For each $j \in \mathcal{J}$, $s_j \in S_j$, and

$c_j \in [0, \bar{c}_j(s_j)]$, define

$$p_j^*(c_j, s_j) \triangleq \mathbb{E}_{\theta_j}[\gamma_j^{-1}\left(\theta_j\right) \mid \theta_j \geq \gamma_j(c_j), \tilde{s}_j = s_j]. \tag{19}$$

The conditional expectation is well-defined for any $c_j \leq \bar{c}_j(s_j)$ or equivalently $\gamma_j(c_j) \leq \bar{v}_j(s_j)$ because $\theta_j = \bar{v}_j(s_j)$ is in the support of the posterior distribution of v_j conditional on $\tilde{s}_j = s_j$. This is because $v_j = \bar{v}_j(s_j)$ is in the support, and $v_k \leq \gamma_k(c_k)$ for all $k \neq j$ could occur with a positive probability. Given signal realizations $\mathbf{s}$, we say that seller j's type c_j is active if $c_j \leq \bar{c}_j(s_j)$. Otherwise, the type is inactive. Seller j's price p_j is said to be active if p_j is in the range of $p_j^*(\cdot, s_j)$; otherwise, the price is called inactive. Note that any active type sets an active price.

In this appendix, for simplicity, we focus on the case in which for each seller j and active price p_j, there exists a unique type, denoted by $p_j^{*-1}(p_j, s_j)$, that solves $p_j^*(c_j, s_j) = p_j$. This is the case, for example, if value v_j has a full support on $[0, 1]$ conditional on each signal realization s_j. The proof for the case in which multiple types may set the same price is relegated to the Supplementary Material.[21] For any inactive price, we set $p_j^{*-1}(p_j, s_j) = 1$; for the dummy seller $j = 0$, we set $p_j^{*-1}(p_j, s_j) = 0$.

We now define an algorithm that we will prove to be optimal for the buyer.

Definition 3. Define the *candidate algorithm* as follows: At each profile of signal realizations $\mathbf{s} = (s_1, ..., s_J)$, values $\mathbf{v} = (v_1, ..., v_J)$, and prices $\mathbf{p} = (p_1, ..., p_J)$, the candidate algorithm recommends trading with seller $j^*(\mathbf{v}, \mathbf{p}, \mathbf{s})$ such that

$$j^*(\mathbf{v}, \mathbf{p}, \mathbf{s}) \in \underset{j \in \mathcal{J}_0}{\operatorname{argmax}}\, v_j - \gamma_j\left(p_j^{*-1}(p_j, s_j)\right). \tag{20}$$

If multiple sellers attain the maximized value in Equation 20, the algorithm breaks ties in favor of sellers such that $p_j^{*-1}(p_j, s_j) > 0$. Other than this restriction, ties are broken arbitrarily.

The following result characterizes the buyer-optimal algorithm and equilibrium under

[21] Multiple active types may set the same price if, for example, the signal is such that the support of the posterior distribution for v_j is non-convex for some signal realizations.

any information structure.

Proposition 11 (Market Segmentation with Competing Sellers). *For any information structure $\mathcal{I}$, the corresponding candidate algorithm is a buyer-optimal algorithm. In equilibrium, seller j of type $c_j \leq \overline{c}_j(s_j)$ posts price $p_j^*(c_j, s_j)$, and any type $c_j > \overline{c}_j(s_j)$ sets some inactive price above 1. Under this algorithm and pricing, for any realized profile of values and costs, the buyer trades with seller $j^* \in \operatorname{argmax}_{j \in \mathcal{J}_0} v_j - \gamma_j(c_j)$. Moreover, the profit of any seller of any type and the total buyer surplus are independent of $\mathcal{I}$.*

Proof. The proof consists of three steps.

Step 1: Characterizing a buyer-optimal direct mechanism. First, we derive a buyer-optimal direct mechanism, where the direct mechanism is in the sense of Myerson (1981), i.e., a mechanism that fully controls product allocation and transfers across all players. By the revelation principle, since any information structure combined with an algorithm can be viewed as an indirect mechanism, a buyer-optimal mechanism must achieve a weakly higher buyer surplus than a buyer-optimal algorithm under any information structure.

Given a profile of values $\mathbf{v}$ and reported types $\mathbf{c}$, let $q_j(\mathbf{v}, \mathbf{c})$ be the probability of allocating seller j's product to the buyer, and let $t_j(\mathbf{v}, \mathbf{c})$ be the monetary transfer from the buyer to seller j. The direct mechanism design problem can be written as:

$$\max_{q:[0,1]^{2J} \to [0,1],\, t:[0,1]^{2J} \to \mathbb{R}} \sum_{j=1}^{J} \int_{[0,1]^J} \int_{[0,1]^J} \left(v_j q_j(\mathbf{v}, \mathbf{c}) - t_j(\mathbf{v}, \mathbf{c})\right) \mathrm{d}F \mathrm{d}G \tag{21}$$

$$\begin{aligned}
\text{s.t. } & T_j(c_j) - c_j Q_j(c_j) \geq T_j(c_j') - c_j Q_j(c_j'), && \forall\, j \in \mathcal{J}, c_j, c_j' \in [0,1],\\
& T_j(c_j) - c_j Q_j(c_j) \geq 0, && \forall\, j \in \mathcal{J}, c_j, c_j' \in [0,1],\\
& Q_j(c_j) = \int_{[0,1]^J} \int_{[0,1]^{J-1}} q_j(\mathbf{v}, c_j, \mathbf{c}_{-j})\, \mathrm{d}F_{-j} \mathrm{d}G, && \forall\, j \in \mathcal{J}, c_j \in [0,1],\\
& T_j(c_j) = \int_{[0,1]^J} \int_{[0,1]^{J-1}} t_j(\mathbf{v}, c_j, \mathbf{c}_{-j})\, \mathrm{d}F_{-j} \mathrm{d}G, && \forall\, j \in \mathcal{J}, c_j \in [0,1],\\
& \sum_{j \in \mathcal{J}} q_j(\mathbf{v}, \mathbf{c}) \leq 1, && \forall\, \mathbf{v}, \mathbf{c} \in [0,1]^J.
\end{aligned}$$

The standard mechanism design arguments imply that only the participation constraint of $c_j = 1$ for each seller j binds at the optimum, and that the IC constraints are equivalent to the local IC constraints (see Equation 22 below) with $Q_j(\cdot)$ being weakly decreasing for each j (cf. Baron and Myerson (1982)). Using the local IC constraints, we can rewrite the expected transfer as follows:

$$\begin{aligned}
\sum_{j=1}^{J} \int_{[0,1]^J} \int_{[0,1]^J} t_j(\mathbf{v}, \mathbf{c}) \, \mathrm{d}F \mathrm{d}G &= \sum_{j=1}^{J} \int_0^1 T_j(x) f_j(x) \, \mathrm{d}x \\
&= \sum_{j=1}^{J} \int_0^1 \left(x + \frac{F_j(x)}{f_j(x)} \right) Q_j(x) f(x) \, \mathrm{d}x \\
&= \sum_{j=1}^{J} \int_0^1 \gamma_j(x) Q_j(x) f(x) \, \mathrm{d}x.
\end{aligned}$$

Plugging this into the objective and using $Q_j(x) = \int_{[0,1]^J} \int_{[0,1]^{J-1}} q_j(\mathbf{v}, x, \mathbf{c}_{-j}) \, \mathrm{d}F_{-j} \mathrm{d}G$, we can rewrite the designer's problem as the choice of a product allocation rule to maximize virtual surplus:

$$\int_{[0,1]^J} \int_{[0,1]^J} \sum_{j=1}^{J} (v_j - \gamma_j(c_j)) \, q_j(\mathbf{v}, \mathbf{c}) f(\mathbf{c}) \, \mathrm{d}\mathbf{c} \, \mathrm{d}G.$$

We can maximize the virtual surplus by choosing $\{q_j(\mathbf{v}, \mathbf{c})\}_{j \in \mathcal{J}}$ to maximize the integrand for each $(\mathbf{v}, \mathbf{c})$. The optimal mechanism allocates seller j's product to the buyer, $q_j(\mathbf{v}, \mathbf{c}) = 1$, if seller j has the highest virtual surplus $v_j - \gamma_j(c_j)$ and it is nonnegative; otherwise, $q_j(\mathbf{v}, \mathbf{c}) = 0$. Let q^D be this optimal product allocation rule and $Q_j^D(c_j)$ be the interim allocation probability for seller j with type c_j. Under the optimal mechanism, the monetary transfer T^D must satisfy

$$T_j^D(c_j) = Q^D(c_j) c_j + \int_{c_j}^1 Q_j^D(x) \, \mathrm{d}x, \forall c_j \in [0,1]. \tag{22}$$

Under the optimal mechanism, the participation constraints for the highest types bind and thus each seller j with the type $c_j = 1$ earns zero profit.

Step 2: Connecting with the candidate algorithm. In this step, we show that the candidate algorithm has an equilibrium in which the product allocation rule and the profits of the highest types are the same as those in the optimal direct mechanism.

First, suppose that each seller follows the pricing strategy described in the proposition. Take any profile of signal realizations $\mathbf{s}$, values $\mathbf{v}$, and types $(\hat{c}_1, ..., \hat{c}_J)$. Let $\mathbf{p}$ be the resulting price profile posted by the sellers. For each seller that posts an active price p_j, the candidate algorithm calculates the unique type that sets price p_j according to Equation 19 and recommends a seller that maximizes virtual surplus. Also, the candidate algorithm never recommends a seller that sets an inactive price, because their corresponding virtual surplus is always negative. Thus if all sellers use the pricing rule in Equation 19 and the buyer always follows the recommendations, then for any profile of prices that can arise, the candidate algorithm recommends the product of the seller with the highest virtual surplus and thus induces the same product allocation as the buyer-optimal mechanism. In particular, the buyer never purchases the product from a seller who has a negative virtual surplus, which means that seller j with $c_j = 1$ earns zero profit.

We now show that the pricing rule in Equation 19 is indeed an equilibrium if the buyer always purchases the recommended product. Take any active seller $j \in \mathcal{J}$ with type c_j. The seller cannot profit from setting an inactive price, because the candidate algorithm never recommends a seller at an inactive price. Alternatively, suppose that the seller has type c_j but deviates to an active price which would be chosen by type c_j'. Let H be the distribution of $\gamma_j^{-1}(\theta_j)$ conditional on $\tilde{s}_j = s_j$. We can compare the profits without and with the deviation as follows:

$$
\begin{aligned}
&\Pr(\gamma_j^{-1}(\theta_j) - c_j \geq 0 | \tilde{s}_j = s_j) \cdot \mathbb{E}_{\theta_j}[\gamma_j^{-1}\Big(\theta_j\Big) - c_j \mid \gamma_j^{-1}(\theta_j) - c_j \geq 0, \tilde{s}_j = s_j] \\
= &\int_{c_j}^{1} (x - c_j) dH(x) \\
\geq &\int_{c_j'}^{1} (x - c_j) dH(x) \\
= &\Pr(\gamma_j^{-1}(\theta_j) - c_j' \geq 0 | \tilde{s}_j = s_j) \cdot \mathbb{E}_{\theta_j}[\gamma_j^{-1}\Big(\theta_j\Big) - c_j \mid \gamma_j^{-1}(\theta_j) - c_j' \geq 0, \tilde{s}_j = s_j]
\end{aligned}
$$

Here, the first line is the profit from following the candidate strategy, and the last line is the profit from deviation.

The other case is when a deviating seller j has an inactive type c_j. In this case, conditional on s_j, any possible realization of θ_j satisfies $\gamma_j(c_j) > \theta_j$. Thus, the profit from the deviation to active type c'_j, which is given by the last line of the above inequalities, will be negative. We conclude that each seller has no profitable deviation.

The last part of this step is to show that the buyer is willing to purchase the product whenever recommended, i.e., conditional on knowing the identity and the price of the recommended product. We present a substantially stronger statement: The buyer is willing to follow recommendations even if she additionally observes the realized signal s_j and the type c_j of the recommended seller j. For each c_j, we have:

$$\begin{aligned}
&\mathbb{E}\left[v_j \middle| v_j - \gamma_j(c_j) \geq \max_{k \in \mathcal{J}_0 \setminus \{j\}} v_k - \gamma_k(c_k), \tilde{s}_j = s_j\right] \\
=&\mathbb{E}\left[v_j \middle| \gamma^{-1}\left(\theta_j\right) \geq c_j, \tilde{s}_j = s_j\right] \\
\geq&\mathbb{E}\left[\gamma_j^{-1}(\theta_j) \middle| \gamma^{-1}\left(\theta_j\right) \geq c_j, \tilde{s}_j = s_j\right] \\
=&p_j^*(c_j, s_j),
\end{aligned}$$

where the inequality holds because:

$$v_j \geq \theta_j = v_j - \max_{k \in \mathcal{J}_0 \setminus \{j\}} \{v_k - \gamma_k(c_k)\} \geq \gamma_j^{-1}(\theta_j).$$

Step 3: Establishing the "payoff equivalence." Let $\{(Q_j(c_j), T_j(c_j)\}_{j \in \mathcal{J}, c_j \in [0,1]}$ be the interim allocation probability $Q_j(c_j)$ and expected revenue $T_j(c_j)$ for each seller j and type c_j under the candidate algorithm. Recall that $\{(Q_j^D(c_j), T_j^D(c_j)\}_{j \in \mathcal{J}, c_j \in [0,1]}$ denote the corresponding objects in the optimal direct mechanism.

We have shown that (i) $\{(Q_j(c_j), T_j(c_j)\}_{j \in \mathcal{J}, c_j \in [0,1]}$ is an equilibrium object and thus satisfies the first two constraints of Equation 21, i.e., the incentive compatibility and participation constraints; (ii) in the candidate algorithm, the profits of the highest seller types are 0; and (iii) $Q_j = Q_j^D$ for each seller j because they come from the same ex post product allocation rule. Thus, the interim expected revenue of each seller j under the candidate algorithm must satisfy

$$\begin{aligned} T_j(c_j) &= Q(c_j)c_j + \int_{c_j}^{1} Q_j(x)dx \\ &= Q^D(c_j)c_j + \int_{c_j}^{1} Q_j^D(x)dx \\ &= T_j^D(c_j), \end{aligned}$$

where the first equality comes from (i) and (ii), the second from (iii), and the third from Equation 22. Therefore, the interim profit of each seller, $T_j(c_j) - Q(c_j)c_j$, is the same between the candidate algorithm and the optimal mechanism. As a result, the buyer surplus, which is the total surplus (uniquely determined by q^D) minus the seller profit, is the same between the algorithm described in the statement and the optimal mechanism in Step 1. □

Proofs of Proposition 6 and Proposition 7 Proposition 6 holds by setting information structure $\mathcal{I}$ to the uninformative structure, e.g., $S_j = \{\emptyset\}$ for each seller j. Proposition 7 is a direct corollary of Proposition 11. □

G Proof of Proposition 8

The proof follows that of Proposition 4. Consider any seller j with type c_j. Consider a fictitious situation in which seller j's prior belief over $(\mathbf{v}, \mathbf{c}_{-j})$ is given by the conditional distribution of $(\mathbf{v}, \mathbf{c}_{-j})$ given $\theta_j \geq c_j$ (calculated starting from the true prior). We call seller j's prior in this situation as the fictitious prior and any posterior belief updated from the fictitious prior a fictitious posterior.

Under any information structure, the price posted by seller j with type c_j, which is (19), is linear in seller j's fictitious posterior belief over $(\mathbf{v}, \mathbf{c}_{-j})$ (which pins down θ_j) conditional on signal realization s_j. If seller j obtains more information about $\mathbf{v}$ in the sense stated in the proposition, the seller also gains more information about $(\mathbf{v}, \mathbf{c}_{-j})$. Thus, the distribution of seller j's fictitious posterior belief over $(\mathbf{v}, \mathbf{c}_{-j})$ undergoes a mean-preserving spread. The same logic as in the proof of Proposition 4 implies that seller j's price undergoes a mean-preserving spread. $\square$

H Proof of Proposition 10

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. 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 $p \in [0, p(\overline{c}))$, the condition is written as

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

or

$$V(q(c)) - cq(c) - \int_c^1 q(x)dx \geq \mu - c - \frac{\int_c^1 q(x)\mathrm{d}x}{q(c)}. \tag{23}$$

Because $q(c) \leq 1$, a sufficient condition for inequality (23) 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, equivalently,

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

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

Supplementary Material for Appendix F

In the appendix, we assumed that each active price is posted by a unique type. In this Supplementary Material, we drop this assumption and prove that the candidate algorithm continues to maximize virtual surplus.

Recall the pricing equation (Equation 19). For each $j \in \mathcal{J}$, $s_j \in S_j$, and active price $p_j \in \mathbb{R}$, define

$$C_j(p_j, s_j) \triangleq \{c_j \in [0,1] : p_j^*(c_j, s_j) = p_j\}$$

as the set of the types of seller j that choose price p_j. For each p_j, we define function $p_j^{*-1}(p_j, s_j)$ as follows: For each $j \in \mathcal{J}$,

$$p_j^{*-1}(p_j, s_j) = \begin{cases} \max C_j(p_j, s_j) & \text{if} \quad C_j(p_j, s_j) \neq \emptyset \\ 1 & \text{if} \quad C_j(p_j, s_j) = \emptyset. \end{cases}$$

For the dummy seller $j = 0$, we set $p_j^{*-1}(p_j, s_j) = 0$.

Suppose that each seller follows the pricing strategy described in the proposition. Take any profile of signal realizations $\mathbf{s}$, values $\mathbf{v}$, and types $(\hat{c}_1, ..., \hat{c}_J)$. Let $\mathbf{p}$ be the resulting price profile posted by the sellers. Suppose that the candidate algorithm recommends seller $j^* \in \mathcal{J}$. Without loss, assume $j^* = 1$. We show that seller 1 has the highest, nonnegative virtual surplus. For each seller $j \in \mathcal{J}$, there exists some c_j such that

$$v_j - \gamma_j\left(p_j^{*-1}(p_j, s_j)\right) = v_j - \gamma_j(c_j). \tag{24}$$

If p_j is an active price, $c_j = \max C_j(p_j, s_j)$. If p_j is an inactive price, $c_j = 1$ by construction. Let $c_j(p_j)$ be the type that satisfies Equation 24. Then Equation 20 implies that

$$v_1 - \gamma_1\left(c_1(p_1)\right) \geq \max_{k \in \mathcal{J}_0 \setminus \{1\}} v_k - \gamma_k\left(c_k(p_k)\right), \tag{25}$$

so that seller 1 has the largest virtual surplus under type profile $(c_1(p_1), ..., c_J(p_J))$. The rest of the proof is devoted to showing that seller 1 has the largest virtual surplus under type profile $(\hat{c}_1, ..., \hat{c}_J)$ as well, i.e.,

$$v_1 - \gamma_1 (\hat{c}_1) \geq \max_{k \in \mathcal{J}_0 \setminus \{1\}} v_k - \gamma_k (\hat{c}_k) . \tag{26}$$

Note that if p_k is inactive for some non-recommended seller k, any type $\hat{c}_k$ that sets p_k satisfies $\hat{c}_k > \bar{c}_k(s_k)$, which implies $v_k - \gamma_k (\hat{c}_k) < 0$. Thus, replacing one inactive type $c_k(p_k)$ with another inactive type $\hat{c}_k$ in the RHS of Equation 26 does not affect the inequality (and it does not affect the argument below).[22] Thus, to simplify exposition, we assume that all sellers set active prices, or equivalently, $c_k(p_k) \leq \bar{c}_k(s_k)$ for each $k \in \mathcal{J}$. We change the type of each seller from $c_j(p_j)$ to $\hat{c}_j$ one by one and show that seller 1 continues to maximize virtual surplus at each step.

First, suppose that we change the type of seller 1 from $c_1(p_1)$ to $\hat{c}_1$. Because $c_1(p_1) = \max C_1(p_1, s_1)$, we have $c_1(p_1) \geq \hat{c}_1$. Therefore, we obtain

$$v_1 - \gamma_1 (\hat{c}_1) \geq \max_{k \in \mathcal{J}_0 \setminus \{1\}} v_k - \gamma_k (c_k(p_k)) . \tag{27}$$

Next, for each seller $\ell = 2, ..., J$, we replace $c_\ell(p_\ell)$ in the RHS of Equation 27 with another active type $\hat{c}_\ell \neq c_\ell(p_\ell)$ that sets the same price p_ℓ, and show that the inequality is preserved. To begin with, for a given $\ell \geq 2$, we consider the following inequalities:

$$v_1 - \gamma_1 (\hat{c}_1) \geq \max \left\{ \max_{k \in \{2,3...,\ell-1\}} v_k - \gamma_k (\hat{c}_k) , v_\ell - \gamma_\ell (c_\ell(p_\ell)) , \max_{k \in \{\ell+1,...,J,0\}} v_k - \gamma_k (c_k(p_k)) \right\} , \tag{28}$$

where we ignore $\max_{k \in \{2,3...,\ell-1\}} v_k - \gamma_k (\hat{c}_k)$ when $\ell = 2$, and

$$\max \left\{ \max_{k \in \{1,...,\ell-1\}} v_k - \gamma_k (\hat{c}_k) , \max_{k \in \{\ell+1,...,J,0\}} v_k - \gamma_k (c_k(p_k)) \right\} < v_\ell - \gamma_\ell (\hat{c}_\ell) . \tag{29}$$

Equation 28 means that seller 1 continues to maximize virtual surplus after we change the type of each seller $k \leq \ell - 1$ from $c_k(p_k)$ to $\hat{c}_k$. Equation 29 means that the inequality is

[22] This argument also implies that if the candidate algorithm does not recommend any seller at a given price profile, then under any type profile that is consistent with the price profile, all the sellers have non-positive virtual surplus.

reversed after we replace the type of seller ℓ. When we consider a version of Equation 28 in which we replace ℓ with k', we refer to the inequality as Equation 28(k'). Note that Equation 28(ℓ) is Equation 28, and we have already shown that Equation 28(2) holds.

We also consider the following inequalities:

$$\max_{k \in \mathcal{J}_0 \setminus \{\ell\}} v_k - \gamma_k (c_k) > v_\ell - \gamma_\ell (c_\ell(p_\ell)) \tag{30}$$

and

$$\max_{k \in \mathcal{J}_0 \setminus \{\ell\}} v_k - \gamma_k (c_k) < v_\ell - \gamma_\ell (\hat{c}_\ell) . \tag{31}$$

We will show that if Equation 28 and Equation 29 hold, then Equation 30 and Equation 31 hold for a positive measure of type profiles $c_{-\ell}$, which, as we show, leads to a contradiction.

For each $\ell \geq 2$, assume that Equation 28(k') for $k' = 2, ..., \ell$ and Equation 29 hold. Note that Equation 29 implies $c_\ell(p_\ell) > \hat{c}_\ell \geq 0$. Thus, we have $c_\ell(p_\ell) > 0$. We consider four cases.

Case 1. First, suppose that (i) Equation 28 holds with strict inequality or (ii) $\hat{c}_1 > 0$. In either case, we can find a positive measure of type profiles $c_{-\ell}$ such that Equation 30 and Equation 31 hold. For example, if $\hat{c}_1 > 0$, then replacing $\hat{c}_1$ with a slightly lower $\hat{c}_1 - \epsilon$ satisfies Equation 28 and Equation 29 with strict inequalities, because the virtual cost functions are assumed to be continuous and strictly increasing. We can then find a positive measure of type profiles $c_{-\ell}$ that maintain these strict inequalities, leading to Equation 30 and Equation 31 for a positive measure of $c_{-\ell}$ (conditional on s_ℓ). This is a contradiction, because the existence of such $c_{-\ell}$'s implies that there is a positive measure of θ_ℓ's such that $\theta_\ell > \gamma_\ell(c_\ell(p_\ell))$ and $\theta_\ell < \gamma_\ell(\hat{c}_\ell)$, i.e., types $c_\ell(p_\ell)$ and $\hat{c}_\ell$ set different prices (see Equation 19).

Case 2. Suppose that (i) Equation 28(k') holds with equality at every step $k' \leq \ell$, (ii) $\hat{c}_1 = 0$, (iii) and the RHS of Equation 28 is positive. Note that Points (i) and (ii) imply that $c_1(p_1) = \hat{c}_1 = 0$, so seller 1 sets a price for which there is a unique active type (to see this, note that if $c_1(p_1) > \hat{c}_1$ and Equation 28(2) holds with equality, then

Equation 25 would fail). This means that seller ℓ does not attain the maximized value of the RHS in Equation 28, because if both sellers 1 and ℓ maximize virtual surplus and $\hat{c}_1 = 0 < c_\ell(p_\ell)$, the candidate algorithm would recommend seller ℓ instead of seller 1 (see the tie-breaking rule described in Definition 3). Because the RHS of Equation 28 is not determined by seller ℓ but both sides are positive, we can slightly increase the cost of each seller $k \neq 1, \ell$ to make Equation 28 and Equation 29 strict. We can then find a positive measure of type profiles $c_{-\ell}$ such that Equation 30 and Equation 31 hold, which leads to a contradiction by the same argument as in Case 1.

Case 3. Suppose that (i) Equation 28 holds with equality but there is some step $k' < \ell$ at which Equation 28(k') is strict, (ii) $\hat{c}_1 = 0$, (iii) and the RHS of Equation 28 is positive. Point (i) implies that there is some $k' \in \{2, ..., \ell - 1\}$ such that Equation 28(k') is strict but Equation 28($k' + 1$) holds with equality, which occurs only when the RHS of the inequality increases as we move from Equation 28(k') to Equation 28($k' + 1$). It means that the RHS of Equation 28 is not determined by seller ℓ, whose type did not change in earlier steps. By the same argument as Case 2, we can find type profiles $c_{-\ell}$ where Equation 28 and Equation 29 hold with strict inequalities, which leads to a contradiction.

Case 4. Suppose that both sides of Equation 28 are 0, and $\hat{c}_1 = 0$. Both sides of Equation 28(k') are 0 for every step $k' = 2, ..., \ell$, because after each replacement, the LHS of Equation 28(k') remains 0 and the RHS is weakly greater than 0. By the same argument as in Case 2, we conclude that that seller ℓ does not attain the maximized value of the RHS in Equation 28 (otherwise, seller ℓ would be recommended according to the tie-breaking rule of the candidate algorithm). But it means that seller ℓ has a negative virtual surplus, which is a contradiction.

In summary, we have shown that for any information structure, if all sellers use the pricing rule in Equation 19 and the buyer always follows the recommendations, for any profile of prices that can arise, the candidate algorithm recommends the product of the seller with the highest virtual surplus. The rest of the proof follows the same argument as in the appendix.