

Algorithm Design Meets Information Design: Price Recommendation Algorithms on Online Platforms*

Nathaniel Hickok[†]

August 21, 2024

Abstract

Platforms often use price recommendation algorithms to suggest prices to firms based on the platform’s private information about consumer demand. I develop a theoretical model and algorithmic experiments to study the impact of platform price recommendations under three types of firm conduct: collusion, competition, and algorithmic pricing. When firms behave either collusively or competitively, I prove that the platform’s optimal price recommendation system is fully informative, and that this outcome is consumer-pessimal over the space of possible recommendation system designs. When firms instead use pricing algorithms, simulated algorithmic experiments show that the introduction of a price recommendation system reduces average consumer surplus by 31%.

JEL: D43, D82, D83, L13, L41

Keywords: Algorithmic Pricing, Information Design, Price Recommendation Algorithms

*I thank Nikhil Agarwal, Anne Carlstein, Stephen Morris, Aroon Narayanan, Vincent Rollet, Kartik Vira, and Mike Whinston, as well as participants at the MIT Industrial Organization Lunch for many helpful comments and suggestions. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant Nos. 1745302 and 2141064. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

[†]Department of Economics, Massachusetts Institute of Technology. Email: nhickok@mit.edu

1 Introduction

Modern online platforms such as Amazon, eBay, and Airbnb receive millions of visitors each day. At the same time, they host millions of sellers. The presence of such a massive number of firms and consumers interacting on these platforms, and the sheer quantity of data generated by these interactions, has led to a recent surge in firms using machine learning algorithms to set prices, hoping to gain an edge over their competitors. In turn, these algorithms have generated concern from regulators¹ and economists² alike because they can lead to *de facto* collusive outcomes without the need for any human intervention, thus skirting existing antitrust laws.

However, there is another, often overlooked, consequence of this environment: it introduces an acute informational asymmetry between platforms and the firms that sell on them.³ While firms are typically only able to receive at best coarse information about the set of potential consumers on the platform via monitoring their own transactions, platforms are able to generate detailed data about the characteristics of these consumers. In response to this informational asymmetry, a number of platforms have developed price recommendation algorithms similar to the pricing algorithms used by firms themselves. These algorithms then offer firms suggestions on how to price their goods based on the platform’s private information.

The consequences of these price recommendation algorithms for the firms and consumers who use online platforms are not well understood. Pavlov and Berman (2019) study their use in the context of whether or not a platform should centralize pricing and find that decentralized pricing combined with price recommendations can benefit consumers via increased competition compared to centralized pricing. However, they critically assume that firms interpret the platform’s recommendation as cheap talk. By contrast, in this paper I take the view that constructing a price recommendation algorithm solves both a *technological* problem—it would otherwise be infeasible for the platform to process its data and provide useful price recommendations to a potentially large number of firms—and also a *strategic* problem—it allows the platform to commit to a particular scheme of sharing its private information with firms, thus ensuring its recommendations are not solely cheap talk.

Because price recommendation algorithms allow platforms to commit to a given price recommendation scheme, platforms face an information design problem when constructing them. On the one hand, they can share more information with firms, allowing firms to set

¹See e.g. Delrahim (2018) and OECD (2017).

²See e.g. Assad et al. (2024), Johnson et al. (2023), or Calvano et al. (2020).

³The informational asymmetry between platforms and consumers is comparatively well-studied, see e.g. Anderson and Renault (2009) or Roesler and Szentes (2017).

prices more accurately in response to uncertain demand conditions. On the other hand, platform and firm incentives are not completely aligned. For instance, platforms typically take a cut of the revenue generated by each sale on their platform, making them total-revenue maximizers. Firms, on the other hand, seek to maximize their individual profits. Hence, the platform might wish to leverage its informational advantage to influence the behavior of firms so that their actions more closely align with the platform’s own objectives. Such behavior has consumer welfare impacts that are *ex ante* ambiguous: if platforms try to lower prices (reflecting the fact that they are revenue maximizing rather than profit maximizing) then consumers might benefit, but if platforms try to raise prices (reflecting the fact that they wish to maximize total revenue and price competition is typically a game of strategic complements) then consumers might suffer.

To study this problem, I take two approaches. First, I develop a theoretical model of platform price recommendation system design using tools from the burgeoning literature on information design (Rayo and Segal, 2010; Kamenica and Gentzkow, 2011).⁴ The model allows me to study the impact of price recommendations under two types of firm conduct: collusion and competition. However, the ubiquity of pricing algorithm usage by firms has led to a third type of conduct that is neither completely collusive nor completely competitive. Instead, firms typically set supra-competitive prices somewhere in the convex hull of the competitive and collusive prices.⁵ To study the impact of price recommendation algorithms when firms themselves are using pricing algorithms, as well as to assess the quantitative importance of this new information design channel, I develop simulated experiments in which machine learning algorithms set prices, taking as one of their inputs a platform-recommended price. Since the space of possible platform recommendation systems is large, I consider two benchmark systems that represent the possible extremes: a price recommendation system that is fully informative of the platform’s private information (always gives the firm-optimal price) and a price recommendation system that is completely uninformative (only gives a randomly generated recommendation).

In the context of my theoretical model, I find that the optimal price recommendation system for platforms is generically fully informative, despite the fact that they have different preferences for prices than firms. Intuitively, this result follows because platforms are constrained in the ways they can influence firm behavior through the use of their information. Under this constraint, the platform’s desire to influence firm behavior is dominated by its desire for prices to adjust to the realized state of demand on the platform. In other words, despite their differences, platform and firm incentives are “aligned enough” that a fully in-

⁴See Bergemann and Morris (2019) for a recent survey of this literature.

⁵Calvano et al. (2020) were among the first to document this feature of pricing algorithms.

formative price recommendation system is optimal. While this is the best possible outcome for firms, I also show that, by contrast, it is generically consumer-pessimal. Moreover, the constraints on how platforms can influence firms via information induces a monotonically decreasing relationship between consumer surplus and the (Blackwell) informativeness of the platform’s price recommendation system.

The results of my simulations lead to qualitatively similar conclusions. Under an informative price recommendation system, the platform is able to increase its revenue and consumer surplus declines. This decline in consumer surplus is substantial, with the average consumer surplus decreasing by more than 30%. The introduction of an informative price recommendation system also increases the extent to which algorithmic pricing leads to supra-competitive prices. This channel is quantitatively important, nearly doubling consumer surplus losses relative to a counterfactual in which an informative price recommendation system is introduced but the propensity of algorithmic pricing to lead to supra-competitive pricing is held fixed.

Related Literature A growing literature studies the use of pricing algorithms by firms and the associated potential for anti-competitive outcomes. These pricing algorithms have been shown to lead to supra-competitive pricing in both simulation (Calvano et al., 2020; Calvano et al., 2021; Johnson et al., 2023; Banchio and Skrzypacz, 2022) and empirical evidence (Assad et al., 2024). Recent work has also uncovered some important mechanisms and limiting factors, such as the exact learning protocol used, the richness of the set of prices that the algorithm considers, and the ability of algorithms to increase the speed at which firms can change their prices (Klein, 2021; Asker et al., 2022; Brown and MacKay, 2023). In this paper I study how pricing algorithms used by firms can interact with pricing recommendation algorithms used by platforms. I find that, when offered informative recommendations, algorithms are able to effectively learn how to adjust prices in response to those recommendations. Moreover, informative recommendations worsen the problem of algorithmic collusion, leading to higher prices and lower consumer welfare. My results uncover a novel mechanism—information—that enables pricing algorithms to learn to collude.

I also build on a large literature that studies the design of various aspects of platforms, particularly in data-rich environments. Recent empirical and theoretical work has studied the choice of pricing mechanism (Einav et al., 2018; Pavlov and Berman, 2019; Buchholz et al., 2020; Banchio and Skrzypacz, 2022), the design of consumer search (Dinerstein et al., 2018; Johnson et al., 2023), and how the use of fees that are per transaction or proportional to revenue can impact the platform’s other decisions (Teh, 2022). In this paper I consider a different aspect of platform design: the use of information as an instrument to influence

seller behavior. I find that allowing a platform to pass their private information to sellers through a price recommendation algorithm can increase platform revenue but only at the cost of consumer welfare.

Within this strand of the literature, Pavlov and Berman (2019) is most closely related to my work. They consider the impacts of platform price recommendation systems in the context of a platform deciding whether or not to centralize pricing. They find that recommendation systems can in fact help consumers because they make decentralized pricing more attractive, and the competition induced by decentralized pricing is generally consumer-welfare enhancing. By contrast, in this paper I take as given that firms set prices in a decentralized manner and explicitly take up the question of the design of the price recommendation system itself. I also critically assume that platforms can commit to the design of their price recommendation system, ensuring that their recommendations are not solely cheap talk. Under these assumptions, I reach a very different qualitative conclusion: the introduction of a price recommendation system is unambiguously bad for consumers.

Finally, this paper builds on a literature that studies the design of information in a platform context. These papers have largely focused on the design of targeted ads (Anderson and Renault, 2009; Bergemann et al., 2021; Bergemann and Bonatti, 2023) or the issue of third-degree price discrimination (Bergemann et al., 2015; Elliott et al., 2021). The work on third-degree price discrimination is most closely related to this paper, and similarly focuses on how a platform can use information to influence pricing decisions by firms. Bergemann et al. (2015) shows in the context of a monopolist seller that information can either help or harm consumers, and Elliott et al. (2021) extends that conclusion to an oligopoly setting. By contrast, since personalized pricing is virtually unused in practice, I instead shut down the price discrimination channel and assume firms use linear pricing. I show in both a monopolistic and oligopolistic setting that, under linear pricing, information is generically harmful for consumers. I also consider how the platform’s information design problem might interact with a third, and empirically relevant, form of conduct—algorithmic pricing. In this context, the use of information can help increase a platform’s revenue but worsens the problem of supra-competitive pricing by algorithms and decreases consumer surplus.

2 A Model of Platform Price Recommendations

There is a monopolistic retail platform that connects consumers and firms. Each firm sells a single product and faces a constant marginal cost $c \geq 0$. The platform takes a fraction of the revenue earned by each firm, leaving the remaining fraction $\phi > 0$ to the firm. The state of the world $\theta \in \Theta = [\underline{\theta}, \bar{\theta}]$ is initially unknown by both the platform and the firms,

and they have a common prior $\theta \sim \mu_0$.

The platform has the ability to (costlessly) learn about θ and can share its information with firms by recommending prices through a price recommendation algorithm.⁶ When constructing its price recommendation algorithm, the platform faces the choice of how much, if any, information the algorithm should communicate to the firms selling on the platform, an idea I will make formal below.

On the firm side, I assume firms compete in prices and consider two types of firm behavior: collusion and competition. In the case of collusion I abstract away from multi-good considerations and, for simplicity, assume that there is only a single homogeneous good being sold, or equivalently that there is only a single monopolistic firm selling on the platform. In the case of competition I assume that it takes the form of simple Hotelling competition. I also assume that the platform knows whether the firms are colluding or competing so that the two cases can be analyzed separately. This assumption is only for ease of exposition—the same price recommendation scheme will be optimal in both cases and so the results still hold even if the platform has no knowledge of firm conduct. All proofs are contained in Appendix A.

2.1 Collusive Firms

There are N firms selling a homogeneous good on the platform. They face demand $D(p; \theta)$, which gives the total quantity sold when the prices are $p = (p_1, \dots, p_N)$ and the state is θ . Since the goods are homogeneous and the firms are colluding it is without loss to consider only the case of a single firm, and so henceforth I will take $N = 1$. I also impose the following three conditions on the demand function:

Assumption 1 (Regularity). There is a non-degenerate interval $P = [0, \bar{p}]$ s.t. $D(p; \theta)$ is twice continuously differentiable and non-negative on $P \times \Theta$. Additionally, for any $\theta \in \Theta$, $D(p; \theta)$ is strictly decreasing and $D(p; \theta)(\phi p - c)$ is strictly concave in p on P and $D(p; \theta) = 0 \forall p \notin P$.

Assumption 2 (Non-triviality). $\forall \theta \in \Theta \exists p$ s.t. $D(p; \theta)(\phi p - c) > \max\{0, D(\bar{p}; \theta)(\phi \bar{p} - c)\}$

Assumption 3 (Demand Affinity). $D(p; \theta) = f_1(p) + \theta f_2(p)$ for some affine functions f_1, f_2 .

⁶Although the space of signals that the platform can use to communicate information with the firm is generically large, a revelation principle-style idea (see e.g. Bergemann and Morris 2019) shows that it is enough to limit attention to signals that take the form of obedient action recommendations. That is, the platform recommends a price to each firm, and conditional on receiving that price recommendation it is a best response for the firm to follow the recommendation.




Figure 1: Example demand curves

Notes: This figure plots demand curves $D(p; \theta)$ for $f_1(p) = 1 - p$, $f_2(p) = 2 - p$ across different values of the demand state.

The first assumption is a relatively standard regularity condition. It ensures that the firm’s problem always has a unique maximizer, and moreover that the maximizers are uniformly bounded. The content of the second assumption is to ensure that, regardless of the state, the firm’s problem is non-trivial: demand is always high enough that they can make positive profits, but demand is never so high that they want to price at the maximum price \bar{p} . It is also technically useful, as it ensures that the firm’s problem has a maximizer that is not just unique but also in the interior of P . Assumption 3 says that demand is affine in the both the price and the state. Since the distribution of θ is left arbitrary, I view the latter as a fairly weak restriction. While the additional assumption of price affinity is strong, the gain in tractability is large, since it reduces the platform’s information design problem to a linear form that is much better understood than general non-linear formulations (Dworczak and Martini 2019; Kleiner et al. 2021; Kolotilin 2018).

Figure 1 exemplifies how the state can affect demand under assumptions 3. In essence, it implies that the state can affect two objects. The first is the intercept of the demand curve, which can be interpreted as the total market size. The second is the slope of the demand curve, which will impact the elasticity of demand faced by the firm. Depending on the functions $f_1(\cdot)$ and $f_2(\cdot)$, the state may impact either of these objects, both, or neither.

The platform’s problem is to choose a price recommendation scheme. Since the platform is able to commit to its choice through the construction of a price recommendation algorithm, this choice can formally be modeled as an information design problem. In particular, given a price recommendation scheme, any price recommendation received by the firm will induce a

posterior belief μ over the distribution of θ . The firm will then choose a price $p^*(\mu)$ to solve

$$\max_p E_\mu[D(p; \theta)(\phi p - c)] \iff \max_p E_\mu[D(p; \theta)(p - C)]$$

where $C := \frac{c}{\phi}$. The choice of a price recommendation then amounts to choosing a distribution over posteriors subject to the Bayes plausibility constraint (Kamenica and Gentzkow, 2011). In particular, the platform solves

$$\max_{\tau \in \Delta(\Delta(\Theta))} \int_{\Delta(\Theta)} E_\mu[D(p^*(\mu); \theta)p^*(\mu)] d\tau(\mu)$$

subject to the Bayes plausibility constraint $\int_{\Delta(\Theta)} \mu d\tau(\mu) = \mu_0$. Assumption 3 ensures that the platform's expected revenue depends only on the expected state.

Lemma 1. *There exists $p(\cdot)$ s.t. $\forall \mu \in \Delta(\Theta)$, $p^*(\mu) = p(E_\mu[\theta])$. Moreover, $\forall \mu \in \Delta(\Theta)$, $E_\mu[D(p^*(\mu); \theta)p^*(\mu)] = \tilde{R}(E_\mu[\theta])$, where $\tilde{R}(E_\mu[\theta]) := D(p(E_\mu[\theta]); E_\mu[\theta])p(E_\mu[\theta])$.*

Since the platform's objective function therefore depends only on the induced posterior mean, rather than the full induced posterior distribution, the platform's choice of information structure reduces to choosing a distribution over posterior means. A distribution of posterior means μ can be induced by some information structure if and only if $\mu_0 \succ \mu$, where \succ is the mean-preserving spread relation, indicating that μ_0 is a mean-preserving spread of μ (Blackwell 1953; Gentzkow and Kamenica 2016).⁷ Intuitively, full information induces a degenerate posterior that puts full weight on the true state for every signal realization and so the distribution of posterior means is equal to the prior. By contrast, no information induces a degenerate distribution of posterior means, with full weight on the prior mean. Since any interior information structure can be represented as a garbling of full information (and analogously no information is a garbling of any information structure), the induced distribution of posterior means must also be a mean-preserving contraction of the induced distribution of posterior means from full information. The platform's problem can thus be written

$$\max_{\mu_0 \succ \mu} \int_{\Theta} \tilde{R}(x) d\mu(x)$$

The critical wedge between the preferences of the platform and the firm is that the firm faces a potentially non-zero marginal cost to production, so that the firm wants to maximize *profits* while the platform wants to maximize *revenue*. The existence of this wedge creates a central tension that the platform must confront when designing its price recommendation

⁷A distribution F is a mean-preserving spread of a distribution G if $E_F[\nu(x)] \geq E_G[\nu(x)]$ for all convex functions ν , where the inequality is strict for strictly convex functions if $F \neq G$.

system. On the one hand, providing more precise recommendations allows the firm to more carefully target its prices, raising prices in states with high demand and decreasing prices in states with low demand. Given free choice over prices, the revenue-maximizing platform would also prefer to have higher prices in states with high demand and likewise lower prices in states with low demand.

On the other hand, the firm generically wants higher prices than the platform at any given demand level because, unlike the platform, they face a potentially non-zero marginal cost. Hence, in targeting its prices, the firm will raise them *too much* in the high demand states and decrease them *too little* in the low demand states. The platform can thus potentially benefit by providing recommendations for lower prices. However, the platform is constrained to inducing mean-preserving contractions. So, informally speaking, it cannot simply recommended lower prices across all demand states but instead must pool a recommendation for a low price across both lower and higher demand states. This type of pooling carries the cost of lower precision as prices will be less targeted toward the exact realization of the demand state.

Intuitively, the platform's optimal price recommendation strategy should balance these two forces. However, it turns out that the cost of imprecise prices is always greater than the cost of having prices larger than would be revenue-maximizing, and so the platform's precision motive dominates its lower-price motive. Hence, the optimal information structure for the platform is actually to always fully reveal the state. This result is immediate in the case when $c = 0$, since then the platform and the firm have the same preferences. Proposition 1 shows that indeed this full information result still holds independent of the marginal cost faced by the firm.

Proposition 1. *Full information is an optimal information structure, and generically⁸ it is the uniquely optimal information structure.*

There are two striking features of this result. The first is that it is independent of the prior beliefs held by the firm and the platform. Hence, regardless of the platform and the firm's initial beliefs, the platform will want to use fully informative price recommendations. Moreover, an outside analyst can likewise be sure of the platform's optimal price recommendation structure even without any knowledge of the distribution of the state. The second feature is that the platform's optimal price recommendation scheme is independent of the marginal cost faced by firms and the percentage of revenue taken by the platform.⁹ Thus, an

⁸I use generically here to mean except for on a measure zero set of parameter values, in particular the coefficients of the functions f_1 and f_2 .

⁹This independence holds only under assumptions 1 and 2, which do implicitly place some restrictions on the firm's marginal cost and the platform's revenue cut.

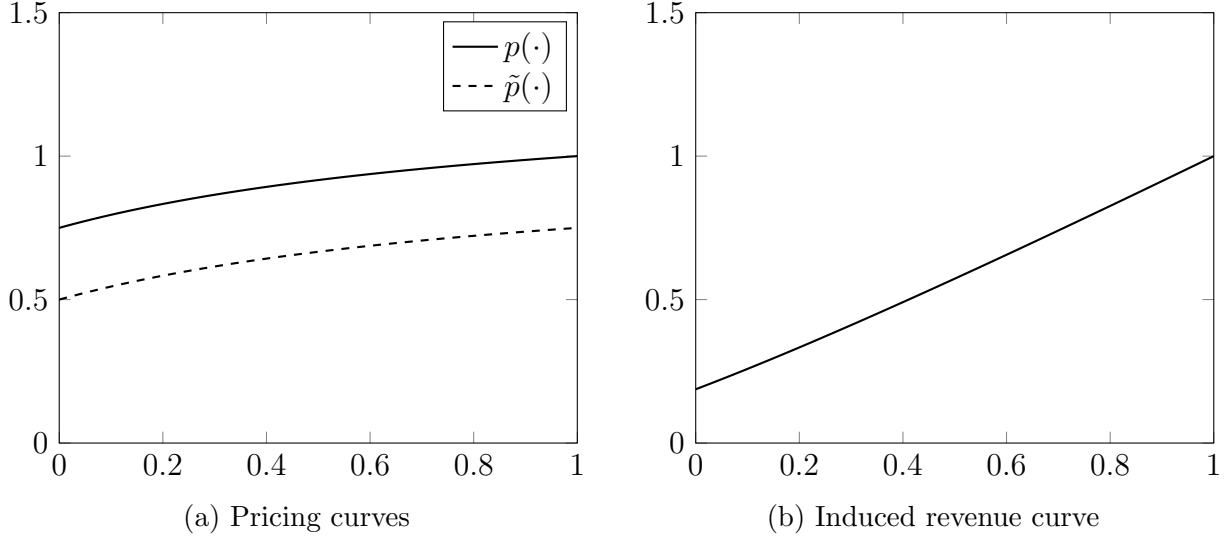


Figure 2: Pricing and induced revenue curves

Notes: In panel (a) this figure plots the firm pricing rule $p(\cdot)$ and the platform pricing rule $\tilde{p}(\cdot)$ as a function of the posterior expected demand state when $f_1(p) = 1 - p$, $f_2(p) = 2 - p$, $P = [0, 1]$, $\Theta = [0, 1]$, and $C = 1/2$. In panel (b) it plots the induced platform revenue curve $\tilde{R}(\cdot)$.

outside analyst can also know the platform's optimal price recommendation structure even without knowledge of these key institutional features.

We can build up some intuition for this result by considering a simple example with $f_1(p) = 1 - p$, $f_2(p) = 2 - p$ on $P = [0, 1]$, $C = 1/2$, and $\Theta = [0, 1]$. Figure 2 plots the firm price $p(\cdot)$ and the platform-optimal price $\tilde{p}(\cdot)$ as a function of the posterior expected state and likewise the platform's induced revenue curve $\tilde{R}(\cdot)$ as a function of the posterior expected state. In this parametrization of the problem, a higher state means demand is uniformly higher, and so both price and revenue are increasing in the posterior expected state. Moreover, the percentage difference between the firm price and the platform's desired price decreases as the state increases (the platform price curve is simply a downward shift of the firm price curve). Hence, revenue is increasing faster for higher posterior expected states and so the induced revenue curve is convex. This convexity is the key feature that drives the platform towards full information.

To see why, recall that the platform's choice among price recommendation schemes can be represented as choosing among distributions over posterior means subject to the restriction that the distribution is a mean-preserving contraction of the prior. But when the induced revenue curve is convex, such a contraction can only make the platform worse off, since by Jensen's inequality its value at a weighted average of points lies below the weighted average of the value of those points. Hence, full information is optimal for the platform.

2.1.1 Consumer Surplus Impacts

Having established that the platform will always want to use fully informative price recommendations, I now turn to considering the impact of the platform's price recommendation scheme on the consumers who use the platform. Intuitively, we would expect that the more information that the platform provides to firms, the worse off consumers become, as the platform is effectively reducing the information rents that consumers can obtain. As proposition 2 shows, this intuition can be formalized by ranking the platform's price recommendation schemes in terms of Blackwell informativeness. In this context, a price recommendation scheme R being Blackwell more informative than another price recommendation scheme R' reduces to the induced distribution of posterior means μ_R being a mean-preserving spread of the induced distribution of posterior means $\mu_{R'}$.

Proposition 2. *Consumer surplus is monotonically decreasing in the Blackwell informativeness of the platform's price recommendation scheme, and generically it is strictly decreasing in the Blackwell informativeness of the platform's price recommendation scheme.*

Corollary 1. *The platform-optimal information structure is generically the consumer-pessimal information structure.*

Given that proposition 1 shows that full information is generically the uniquely optimal information structure for the platform, proposition 2 yields the immediate corollary that generically (in particular whenever the platform-optimal information structure is unique), the platform-optimal information structure is in fact the consumer-pessimal information structure. In other words, any regulator with a consumer welfare standard should choose to ban platform price recommendations, as they can only serve to harm consumers.

To gain further intuition for this result, we can once again return to the simple parameterized example in which $f_1(p) = 1 - p$, $f_2(p) = 2 - p$, $P = [0, 1]$, $C = 1/2$, and $\Theta = [0, 1]$. Figure 3 plots consumer surplus in this case as a function of the posterior expected state. Unlike the induced revenue curve, which was convex, the consumer surplus curve is concave. Hence, it is precisely analogous logic to the case of the platform's optimal information structure that shows that consumers prefer less information. In particular, consider once again the representation of price recommendation structure as choosing among distributions over posterior means subject to the restriction that the distribution is a mean-preserving contraction of the prior. As the expected consumer surplus curve is concave, Jensen's inequality again implies that such a contraction can only make consumers better off.




Figure 3: Expected consumer surplus

Notes: This figure plots expected consumer surplus as a function of posterior expected demand state when $f_1(p) = 1 - p$, $f_2(p) = 2 - p$, $P = [0, 1]$, $\Theta = [0, 1]$, and $C = 1/2$.

2.2 Competitive Firms

I now turn to the case of firm competition. For simplicity, I consider only the duopoly case when $N = 2$, although it is straightforward to generalize the results to the N firm case. The firms compete in standard Hotelling competition where the unknown state θ represents the inverse of the strength of consumer preferences. That is, there is a mass M of consumers, which I normalize to 1, each with unit demand. Each consumer obtains utility $v - \theta^{-1}x - p$ for purchasing good 1 at price p and $v - \theta^{-1}(1 - x) - p$ for purchasing good 2 at price p , where $x \sim U[0, 1]$ in the population. To keep utilities well-defined, I impose the additional assumption that $\underline{\theta} > 0$. Then under the typical full coverage assumption,¹⁰ if the firms know θ then they have symmetric price $p = C + \frac{1}{\theta}$ in the unique pure strategy equilibrium.

Now suppose that firms face symmetric uncertainty over θ . In particular, both firms believe $\theta \sim \mu$. Since θ enters affinely into the firms' objective functions, a version of lemma 1 holds in this environment and the unique pure strategy equilibrium will depend only on the expected state $E_\mu[\theta]$. Specifically, we will have $p_i(E_\mu[\theta]) = C + \frac{1}{E_\mu[\theta]}$ for $i = 1, 2$.

Lemma 2. *Given beliefs μ , the unique pure strategy equilibrium is $p_1 = p_2 = C + \frac{1}{E_\mu[\theta]}$.*

While in principle the platform's price recommendation scheme could induce asymmetric posterior beliefs among the firms, I restrict the platform to use only symmetric signals, or

¹⁰Formally, I assume full coverage for every state θ , which in particular says that v is always large enough that the potential markets of the two firms overlap, and so in equilibrium every consumer always purchases a good. Without this assumption the firms simply price at their monopoly level and do not engage in any meaningful competition, which is the case covered in the previous section.

equivalently (since the environment is symmetric) to send only publicly observable signals. This restriction ensures that the firms always have the same posterior beliefs and so the induced game is always of the above form. It also matches closely with the empirical reality of most price recommendation systems used in practice, as these systems typically make good-specific rather than firm-specific recommendations, thereby allowing firms to see the price recommendation given to another firm simply by entering the relevant information of the good sold by that firm.

Under the assumptions of full coverage and public signals, and since the firm's prices are symmetric in the induced equilibrium, the platform's revenue is proportional to the equilibrium price $p(x) = C + \frac{1}{x}$ for any posterior expectation x of the state. Since the function $1/x$ is strictly convex (for positive x), we can immediately see that once again full information will be optimal for the platform, yielding the following proposition.

Proposition 3. *Full information is the uniquely optimal information structure for the platform.*

Analogous to proposition 1, full information is optimal for the platform regardless of the prior beliefs of the firm, the marginal costs faced by the firms, and the fraction of revenue that the platform takes. Why is this the case? Intuitively, the platform faces almost exactly the same trade-off as in the case of collusion. On the one hand, it has a precision motive and wants to make sure that the prices are set in accordance with the demand state. On the other hand, the platform, if it could, would want to set prices at a different level than the equilibrium prices of the firms. Although in this case, due to the full coverage assumption and the strategic complementarity of the Hotelling competition game, the platform wants higher prices than the firms, not lower prices. However, despite this distinction, the precision motive still dominates, and so the platform optimally uses fully informative price recommendations.

2.2.1 Consumer Surplus Impacts

The impact of the platform's price recommendation system on consumers is again analogous to the case of competition. Intuitively, by revealing the consumer's private information to firms, the platform is reducing the information rents that consumers can achieve and therefore lowering their welfare. While we might ex ante anticipate that the impact of price recommendations would be worse for consumers in the case of competition rather than collusion—since the platform wants higher prices than the firms instead of lower prices than the firms—this turns out not to be the case since the platform will use fully informative price recommendations in both cases.

Formally, the analysis of consumer surplus in the case of competition is complicated by the fact that expected consumer surplus given posterior beliefs μ depends not just on the first moment $E_\mu[\theta]$ but instead the entire distribution μ .

Lemma 3. *Given posterior beliefs μ , expected consumer surplus is $E_\mu[CS] = v - C - \frac{1}{E_\mu[\theta]} - \frac{1}{4}E_\mu\left[\frac{1}{\theta}\right]$.*

Hence, in order to formally make statements about the consumer welfare impacts of the platform's price recommendation scheme, we must take a stance on the exact distribution of posteriors that the platform induces rather than just the distribution of posterior means. Since the platform and the firm are both indifferent between any distribution of posteriors that induces the same distribution of posterior means, this amounts to an assumption on the tiebreaking behavior of the platform. I make the most favorable assumption possible for consumers:

Assumption 4 (Tiebreaking). If the platform chooses a price recommendation scheme R that induces a distribution μ_R over posterior means, then it does so by inducing the distribution of posterior beliefs τ that, conditional on inducing μ_R , yields the maximum possible expected consumer surplus.

It is not immediately obvious that this assumption is well-formed. That is, it might be that there is no unique τ that maximizes consumer surplus conditional on inducing a certain distribution of posterior means. However, the following lemma rules out this possibility, and indeed identifies the maximizing τ —it is the distribution that only puts positive weight on degenerate posteriors.

Lemma 4. *Fix a distribution of posterior means μ satisfying $\mu_0 \succ \mu$. Then under assumption 4, the platform will induce the distribution of posterior means μ by inducing the distribution of posteriors τ satisfying*

$$d\tau(\tilde{\mu}) = \begin{cases} d\mu(x) & \text{if } \tilde{\mu} = \delta_x \\ 0 & \text{otherwise} \end{cases}$$

Since under assumption 4 the platform will thus only induce distributions of degenerate posteriors, lemmas 3 and 4 yield the immediate corollary that expected consumer surplus for a given posterior belief does in fact only depend on the posterior mean, as long as we assume that the platform breaks ties in favor of consumers. Since this tiebreaking does not affect firms, such an assumption merely amounts to assuming that the platform will not discard consumer surplus “for free.”

Corollary 2. *Under assumption 4, for posterior belief μ induced by a platform recommendation system, expected consumer surplus is $v - C - \frac{3}{4E_\mu[\theta]}$.*

We can thus return to the framework of thinking about price recommendation systems only in terms of the induced distribution of posterior means. Moreover, since $-1/x$ is strictly concave (for positive x), we can immediately see that once again full information will be pessimal for consumers, yielding the following proposition.

Proposition 4. *Under assumption 4, the platform-optimal information structure is the consumer-pessimal information structure, and the unique consumer-optimal information structure is no information.*

Proposition 4 is the analogue of corollary 1 in an environment in which firms are competing rather than colluding. It says that a fully informative price recommendation system is in fact the worst possible recommendation system for consumers, and that an uninformative price recommendation system is the best possible recommendation system for consumers (among all possible recommendation systems that the platform could choose subject to assumption 4). It thereby formalizes the intuition that informative recommendation systems are bad for consumers because they lower their information rents. Moreover, it shows that, just as in the collusive environment, any regulator with a consumer welfare standard should choose to ban platform price recommendations.

2.3 Discussion of Model Results and Assumptions

A number of key messages emerge from the previous analysis of platform price recommendation system design in both collusive and competitive settings. The first is that, despite the existence of wedges between the preferences of the platform and those of the firms, there is a strong sense in which their incentives are aligned. In particular, the platform will always want to use a price recommendation system that is fully informative. In other words, it will always recommend the firm-optimal prices conditional on the realized state. This conclusion holds independent of institutional features such as the marginal cost faced by firms, the platform revenue cut, and the type of firm conduct, despite the fact that these features in some sense control the strength of the wedge between platform and firm preferences. It is also therefore robust to arbitrary uncertainty about those features.

The second key message regards consumer welfare. Ex ante, it is unclear how consumers would be impacted by platform price recommendations. On the one hand, they might benefit if the platform tries to drive down the price to raise revenue, or if prices become more aligned with their preferences conditional on the realized state (for instance having

lower prices when demand is lower and higher prices when demand is higher). On the other hand, any information that the platform provides the firms can in some sense be thought of as reducing the information rents that consumers are able to achieve, and thus lowering their welfare. This second intuition turns out to be the correct one. Price recommendation systems are unambiguously bad for consumers, in the strong sense that the platform-optimal price recommendation system is in fact consumer-pessimal. This conclusion again holds essentially independent of institutional features and so is robust to arbitrary uncertainty about those features.

One important limitation to these conclusions is that I have assumed throughout that demand is affine. This assumption is explicit in the model of collusion but it is also implicit in the model of competition, as Hotelling competition induces affine residual demand functions facing each of the firms. In the next section I relax this assumption by using model simulations to study the case of logit demand, which is a workhorse model of demand in empirical applications. I have also only analyzed two extreme forms of firm conduct: perfect competition and perfect collusion. On many modern online platforms firms do not purely collude or compete. Instead, they set prices through the use of machine learning algorithms. Typically, these algorithms do not converge to either the collusive or the competitive outcome but instead to some supra-competitive prices in the convex hull between the competitive and collusive prices (Calvano et al. 2020; Johnson et al. 2023).

3 The Case of Algorithmic Pricing

In order to understand the impact of platform price recommendations in a setting that more closely matches the empirical reality on modern platforms, in this section I simulate the behavior of algorithmically-pricing firms in response to a price recommendation system implemented by the platform. I also assume logit, rather than affine, demand. Since the space of possible platform recommendation systems is large, I consider two benchmark systems: no information and full information.

In keeping with the theoretical model of the previous section, we should expect an uninformative price recommendation system to be consumer-optimal and a fully informative system to be platform-optimal if the firms are either competing or colluding. These conclusions are only suggestive, however, as the theory cannot speak directly to the case in which firms are neither colluding nor competing and consumer demand is not affine. Moreover, introducing a platform price recommendation system can (and indeed does) change the very nature of firm conduct in the algorithmic case, as it can make it easier or harder for firms to collude (O’Connor and Wilson, 2019). Intuitively, this duality stems from the fact that

additional information about demand both makes collusion more attractive (by increasing the payoffs to colluding) and at the same time raises incentives to cheat.

Following other recent work in the algorithmic pricing literature (e.g. Calvano et al. 2020; Johnson et al. 2023; Banchio and Skrzypacz 2022), I assume that firms price according to Q-learning algorithms. Q-learning algorithms are a popular class of reinforcement learning algorithms in both computer science and in practice (Sutton and Barto, 2018). In stationary, single-agent environments, Q-learning works by essentially learning the value function one state-action pair at a time, updating an estimate of the payoff and continuation value of playing a certain action in a certain state every time it plays that action in that state.

A typical method to ensure that the algorithm sufficiently explores its environment is called ϵ -greedy: in each period t the algorithm plays the action that yields the highest estimated value (inclusive of discounted continuation value) conditional on the state with probability $1 - \epsilon_t$, and with probability ϵ_t it plays a random action. With this learning strategy there are provable guarantees that, under certain conditions, Q-learning algorithms will converge to optimal play (Watkins and Dayan, 1992). In multi-agent settings such as firm competition, however, there are no provable guarantees that Q-learning algorithms will even converge, let alone converge to “optimal” play. The reason for this is that the environment is inherently non-stationary—it changes each time the other firms update their strategies. Nonetheless, in practice, convergence (in a sense defined below) is nearly always obtained (Calvano et al., 2020).

3.1 Simulation Environment

There are two different firms selling differentiated goods on the platform, as well as an outside option. The firms interact repeatedly in periods $t = 1, 2, \dots$ and, conditional on the demand state θ_t , face standard logit demand

$$s_{it} = \frac{\exp\left(\frac{a_i - \theta_t p_{it}}{\mu}\right)}{1 + \sum_{j=1}^2 \exp\left(\frac{a_j - \theta_t p_{j,t}}{\mu}\right)}$$

where a_i is an index of quality that captures potential vertical differentiation and μ is the scale of the logit error, capturing horizontal differentiation. In this context the demand state θ_t controls the price sensitivity of consumers. Each period it is drawn uniformly at random from the set Θ , independent of previous realizations. The firms face constant marginal cost C (interpreted as being adjusted for the platform revenue cut) and hence each period receive profits $\pi_{it} = (p_{it} - C)s_{it}$, where I normalize the mass of consumers to 1 so that $q_{it} = s_{it}$.

While the demand state is unknown to the firms, in each period t they receive a (public) signal ρ_t of the demand state from the platform (a “price recommendation”), which learns the state at the beginning of each period. The price recommendations are either fully informative ($\rho_t = \theta_t$) or completely uninformative ($\rho_t \sim U[\Theta]$). The firms also have bounded memory: they can remember prices from the previous m periods.¹¹ In making their period t decision, the firm can thus condition on ρ_t and $\{p_{1,t-i}, p_{2,t-i}\}_{i=1}^m$, which I will call the firm state, denoted f_{it} . Note that the assumption of bounded memory, which is necessary for computational tractability, is fairly weak as it still allows for complex dynamic interactions between the firms (Calvano et al., 2020).

As mentioned previously, I assume that the firms make pricing decisions according to Q-learning algorithms. Q-learning requires that both the state space and action space of each agent be finite, hence I restrict Θ to be finite and discretize the price space P as follows. I compute the Bertrand Nash and fully collusive prices for each $\theta \in \Theta$. Let \underline{p} be the minimum such price, rounded down to the nearest tenth, and \bar{p} be the maximum such price, rounded up to the nearest tenth. I take $P = \{\underline{p}, \underline{p} + .1, \underline{p} + .2, \dots, \bar{p}\}$.

The central object of Q-learning is the Q-matrix, which is (informally) defined as

$$Q(f, p) = E[\pi \mid f, p] + \delta E[\max_{p' \in P} Q(f', p') \mid f, p]$$

where $\delta < 1$ is the discount factor and assumed to be common across firms. In words, the Q-matrix gives the expected value (inclusive of discounted continuation value) of charging price p when in firm state f . Note that each firm has its own Q-matrix, but I suppress this dependence for notational simplicity. Of course, the Q-matrix is also not entirely well-defined in multi-agent settings such as this one, since it depends on the strategies of other agents. Nonetheless, there is still a sense in which a Q-learning algorithm, which attempts to learn the Q-matrix by exploring the payoffs to different prices in different firm states, can converge. I will define this convergence precisely shortly, but intuitively it can be thought of as convergence in *strategies*—once the firms have converged on strategies and cease exploring, the environment once again becomes stationary and the Q-matrix is well-defined.

A Q-learning algorithm has no knowledge of the Q-matrix, or indeed any of the model primitives, but instead attempts to learn the optimal strategy via exploration. In particular, starting from an arbitrary initial matrix Q_0 , in each period t the algorithm observes firm state f_t and chooses price p_t , observes payoff π_t and next period firm state f_{t+1} , and then

¹¹To keep the firm state space limited, I assume that firms cannot remember previous price recommendations. I view this as a weak restriction since firms do not engage in strategic interaction with the platform.

updates its guess of the Q-matrix by setting

$$Q_{t+1}(f_t, p_t) = (1 - \alpha)Q_t(f_t, p_t) + \alpha \left[\pi_t + \delta \max_{p \in P} Q_t(f_{t+1}, p) \right]$$

where I again suppress firm dependence for notational simplicity. The hyperparameter $\alpha \in (0, 1)$, the learning rate, determines how quickly the algorithm updates its guess of a particular element of the Q-matrix when receiving new information. Firm exploration is ϵ -greedy: in each period t , with probability $1 - \epsilon_t$ the firm chooses $p = \arg \max_{p'} Q_t(f_t, p')$ and with probability ϵ_t it chooses a random $p \in P$. In each period, a firm's strategy ψ_t gives the algorithm's guess of the optimal price for any firm state f . That is

$$\psi_t(f) = \arg \max_{p \in P} Q_t(f, p)$$

In keeping with Calvano et al. (2020), I define convergence by the following criterion: each firm's strategy has remained stable for the past 100,000 periods. I limit the maximum number of periods to one billion, but convergence is always achieved in my simulations and so this limit is never reached.

In my baseline simulations I take $\delta = .95, \mu = 1/4, \alpha = .15, a_1 = a_2 = 2, C = 1$, and $m = 1$. I also parametrize $\epsilon_t = \exp(-\beta t)$ and take $\beta = 5 \times 10^{-6}$. These parameter values align with those of Calvano et al. (2020) and Johnson et al. (2023), ensuring that my results have a comparable baseline in a setting without demand uncertainty or price recommendations. The one exception is that I take β to be smaller than their simulations, since the existence of demand uncertainty introduces additional noise into the simulations and so it is important to give the firms more opportunity to explore their environment. I take the set of possible demand state values to be the set $\Theta = \{.8, .9, 1., 1.1, 1.2\}$. Finally, in every simulation I initialize the environment in a random demand state and with Q_0 a matrix of zeros.

3.2 Simulation Results

I run the above simulations 1,000 times for both the case of fully informative price recommendations and the case of completely uninformative price recommendations. For each simulation run, I use the results of the final 100,000 stable periods to measure relevant outcomes such as prices and consumer surplus. Table 1 reports the average across simulation runs for a number of key metrics, including total revenue, prices, and consumer surplus. It also reports an average "collusion level." I define the collusion level to be a normalized

measure of how far prices are from the Bertrand Nash level. In particular, I take:

$$CL_t = \frac{p_t - p_{Nash,t}}{p_{Collusive,t} - p_{Nash,t}}$$

where, for simplicity since firms are symmetric, all prices are in terms of firm 1.¹² In simulations with fully informative price recommendations, the Bertrand Nash and collusive price is conditional on the demand state in that period, while in simulations with uninformative price recommendations it is the unconditional Bertrand Nash and collusive price.

Table 1: Simulation outcomes

	Informative	Uninformative
Average Revenue	1.39	1.36
Average Consumer Surplus	3.69	5.32
Average Collusion Level	0.50	0.29
Average Price (Firm 1)	1.69	1.60
Average Price (Firm 2)	1.69	1.60

Notes: This table shows average outcomes for simulations with fully informative and completely uninformative price recommendations. Averages are taken over the final 100,000 stable periods in each simulation run.

The results of these simulations are broadly consistent with the model developed in the previous section. When recommendations are informative, average revenue increases and consumer surplus decreases. The change in consumer surplus is substantial: going from uninformative to informative recommendations causes a 31% decrease in average consumer surplus. Moreover, firms are able to achieve much higher levels of collusion when recommendations are informative, which is also reflected in the fact that prices are higher on average. Finally, as expected given the symmetry of the environment, the prices of firm 1 and firm 2 are the same on average in both the informative and uninformative case.

3.2.1 Prices

When price recommendations are informative, average prices rise by 6%. However, the fact that average prices change does not imply that pricing algorithms are able to effectively use the information provided by the price recommendations to target prices to the demand state. Figure 4 shows the average (firm 1) price conditional on the demand state for the informative recommendation case and the uninformative recommendation case. Reassur-

¹²Calvano et al. (2020) use a similar measure except it is in terms of firm profits rather than prices. I prefer the measure in terms of prices since it allows for a type of welfare decomposition exercise (see section 3.2.3).

ingly, the algorithms are indeed able to learn to target prices when recommendations are informative. In the informative case, the average price when demand is highest is 1.8 while it is 1.6 when demand is lowest. By contrast, when recommendations are uninformative the price is constant at about 1.6 across all demand states. That the average price is uniformly higher across all demand states when recommendations are informative than when they are uninformative reflects the fact that the algorithms learn to collude much better in the informative case—if conduct were the same in both cases then we would still expect prices to rise in the high-demand states, but they should fall in the low-demand states.




Figure 4: Average price by demand state

Notes: This figure plots the average price charged by firm 1 in each demand state for fully informative price recommendations and completely uninformative price recommendations. Prices are averaged over the final 100,000 stable periods in each simulation run. The dashed and dot-dashed lines give the benchmark fully collusive and Bertrand Nash prices respectively.

3.2.2 Consumer Surplus

The average consumer surplus is lower in the informative case than in the uninformative case, a finding that is consistent with the theoretical results of section 2. Moreover, figure 5 shows that these losses stem from large declines in consumer surplus in high-demand states, while consumer surplus losses are much smaller in low-demand states (and consumer surplus in fact even increases slightly in the lowest-demand state despite the fact that the average price increased in that state). That the losses in the high-demand states are significantly larger than in the low-demand states reflects two factors. First, the average price increase is much higher in those states. And second, consumers in some sense care more about the products in the high-demand states than they do in the low-demand states, so that even if

the price increase was the same in all states the consumer surplus loss would still be higher in the high-demand states. These two factors compound, leading to large consumer surplus losses in the high-demand states.




Figure 5: Average consumer surplus by demand state

Notes: This figure plots consumer surplus conditional on demand state, averaged over the final 100,000 stable periods in each simulation. The solid line conditions on simulations with informative price recommendations, while the dashed line conditions on simulations with uninformative price recommendations.

3.2.3 Collusion Level

The average collusion level goes from 0.29 when recommendations are uninformative to 0.50 when they are informative, a 72% increase. To get a back-of-the-envelope sense of the importance of this difference in collusion levels, I perform the following exercise. First, I measure average consumer surplus assuming the price in each demand state is at the same collusion level as the average collusion level under informative recommendations: $p_\theta = 0.50p_{collusive,\theta} + 0.50p_{Nash,\theta}$. Next, I measure average consumer surplus assuming the price in each demand state is at the same collusion level as the average collusion level under




Figure 6: Welfare change decomposition

Notes: This figure decomposes the welfare change from uninformative to informative recommendations into the fraction due to changes in the collusion level and the fraction due to having prices adjust to the demand state.

uninformative recommendations: $p_\theta = .29p_{collusive,\theta} + .71p_{Nash,\theta}$. Finally, I measure average consumer surplus assuming the price in each demand state is at the same collusion level as the average collusion level under uninformative recommendations, but now relative to the unconditional Nash and collusive prices rather than the conditional Nash and collusive prices: $p_\theta = .29p_{collusive} + .71p_{Nash}$.

This exercise has the effect of decomposing the consumer surplus loss of informative recommendations into two pieces. The first piece is the loss from inducing better algorithmic collusion, holding constant the ability to condition the price on the demand state. The second piece is the loss from the ability to condition prices on the demand state, holding constant the collusion level. Figure 6 displays these two pieces visually. Going from a collusion level of .50 to .29 while keeping prices conditional on the demand state accounts for 45% of the overall welfare change between informative and uninformative recommendations. Going from conditional prices to a uniform price accounts for the other 55%. Hence, the increase in collusion level from introducing informative recommendations represents an important part of the change in consumer surplus.

3.2.4 Robustness

I assess the sensitivity of my simulation results to the chosen baseline parameters by computing the elasticity of each main outcome—consumer surplus, collusion level, revenue, and prices—to changes in parameter values. I do so by decreasing the focal parameter or hyperparameter by 10% while holding the other parameters fixed at their baseline values and then computing an implied elasticity from the change in outcomes. Using this procedure I assess sensitivity to changes in marginal cost (C), product quality (a), the scale of the logit error (μ), the discount factor (δ), the learning rate (α), and the exploration decay (β).

The elasticities are typically small (less than 1 in absolute value) and similar across both informative and uninformative recommendations, suggesting that both the quantitative and qualitative conclusions from the baseline simulations are fairly robust. The most sensitive outcome is the collusion level, for which the elasticity is typically between 2.5 and 7.5 in absolute value. However, the elasticities are similar for both informative and uninformative recommendations, suggesting that the relative comparison between the two is still robust. The outcomes are most sensitive to the quality parameter, although the elasticities are again similar for both informative and uninformative recommendations. The full set of robustness elasticities can be found in Appendix B.

4 Conclusion

In this paper I study the impact of platform price recommendations under three types of firm conduct: collusion, competition, and when firms use pricing algorithms. For the cases in which firms are either colluding or competing, I develop a theoretical model of platform price recommendation design. I show that, in the case of affine demand, the platform’s optimal price recommendation system is generically fully informative independent of prior beliefs and other model primitives. Moreover, the platform-optimal price recommendation system is generically consumer-pessimal.

To understand how platform price recommendations might interact with firms using pricing algorithms and more empirically realistic consumer demand, I develop simulated experiments of algorithmic pricing in an environment with demand uncertainty and price recommendations. I study two benchmark price recommendation systems: a fully informative system and a completely uninformative system. In my simulations, I find that the introduction of a price recommendation system reduces average consumer surplus by about 31%. Approximately 45% of this loss can be attributed to the price recommendation system making it easier for pricing algorithms to learn to collude while the other 55% can be attributed to firms learning to adjust prices in response to changes in the demand state.

There are a number of interesting questions left unanswered by this paper that would be fruitful directions for future work. For one, it is important to more fully understand the conditions under which platform and firm preferences are “aligned enough” that fully informative recommendation systems are optimal for the platform. It would also be useful to further develop our understanding of how price recommendation systems change the ability of pricing algorithms to learn to collude.

References

- S. P. Anderson and R. Renault. Comparative advertising: disclosing horizontal match information. *The RAND Journal of Economics*, 40(3):558–581, 2009.
- J. Asker, C. Fershtman, and A. Pakes. Artificial Intelligence, Algorithm Design, and Pricing. *AEA Papers and Proceedings*, 112:452–456, May 2022.
- S. Assad, R. Clark, D. Ershov, and L. Xu. Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. *Journal of Political Economy*, 132(3):723–771, Mar. 2024.
- M. Banchio and A. Skrzypacz. Artificial Intelligence and Auction Design, Feb. 2022. arXiv:2202.05947 [cs, econ].
- D. Bergemann and A. Bonatti. Markets for Information: An Introduction. *Annual Review of Economics*, 11(Volume 11, 2019):85–107, Aug. 2019. Publisher: Annual Reviews.
- D. Bergemann and A. Bonatti. Data, Competition, and Digital Platforms, 2023. Version Number: 1.
- D. Bergemann and S. Morris. Information Design: A Unified Perspective. *Journal of Economic Literature*, 57(1):44–95, Mar. 2019.
- D. Bergemann, B. Brooks, and S. Morris. The Limits of Price Discrimination. *American Economic Review*, 105(3):921–957, Mar. 2015.
- D. Bergemann, A. Bonatti, and A. Smolin. The Design and Price of Information. *American Economic Review*, 108(1):1–48, Jan. 2018.
- D. Bergemann, T. Heumann, and S. Morris. Selling Impressions: Efficiency vs. Competition, July 2021.
- D. Blackwell. Equivalent Comparisons of Experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. Publisher: Institute of Mathematical Statistics.
- Z. Y. Brown and A. MacKay. Competition in Pricing Algorithms. *American Economic Journal: Microeconomics*, 15(2):109–156, May 2023.
- N. Buchholz, L. Doval, J. Kastl, F. Matějka, and T. Salz. Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides. Technical Report w27087, National Bureau of Economic Research, Cambridge, MA, May 2020.
- E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello. Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10):3267–3297, Oct. 2020.
- E. Calvano, G. Calzolari, V. Denicoló, and S. Pastorello. Algorithmic collusion with imperfect monitoring. *International Journal of Industrial Organization*, 79:102712, Dec. 2021.

- M. Delrahim. “Life in the Fast Lane”: Antitrust in a Changing Telecommunications Landscape, Nov. 2018.
- M. Dinerstein, L. Einav, J. Levin, and N. Sundaresan. Consumer Price Search and Platform Design in Internet Commerce. *American Economic Review*, 108(7):1820–1859, July 2018.
- P. Dworczak and G. Martini. The Simple Economics of Optimal Persuasion. *Journal of Political Economy*, 127(5):1993–2048, Oct. 2019. Publisher: The University of Chicago Press.
- A. S. Edlin and C. Shannon. Strict Monotonicity in Comparative Statics. *Journal of Economic Theory*, 81(1):201–219, July 1998.
- L. Einav, C. Farronato, J. Levin, and N. Sundaresan. Auctions versus Posted Prices in Online Markets. *Journal of Political Economy*, 126(1):178–215, Feb. 2018.
- M. Elliott, A. Galeotti, A. Koh, and W. Li. Market Segmentation Through Information, Dec. 2021.
- M. Gentzkow and E. Kamenica. A Rothschild-Stiglitz Approach to Bayesian Persuasion. *American Economic Review*, 106(5):597–601, May 2016.
- J. P. Johnson, A. Rhodes, and M. Wildenbeest. Platform Design When Sellers Use Pricing Algorithms. *Econometrica*, 91(5):1841–1879, 2023.
- B. Jullien and A. Pavan. Information Management and Pricing in Platform Markets. *The Review of Economic Studies*, 86(4):1666–1703, July 2019.
- E. Kamenica and M. Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615, Oct. 2011.
- T. Klein. Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3):538–558, Sept. 2021.
- A. Kleiner, B. Moldovanu, and P. Strack. Extreme Points and Majorization: Economic Applications. *Econometrica*, 89(4):1557–1593, 2021.
- A. Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.
- A. Kolotilin and A. Wolitzky. Assortative Information Disclosure, June 2020.
- A. Kolotilin, T. Mylovanov, A. Zapechelnyuk, and M. Li. Persuasion of a Privately Informed Receiver. *Econometrica*, 85(6):1949–1964, 2017. Publisher: [Wiley, The Econometric Society].
- J. Levin and P. Milgrom. Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review*, 100(2):603–607, May 2010.
- J. Mensch. Monotone persuasion. *Games and Economic Behavior*, 130:521–542, Nov. 2021.

- J. O'Connor and N. E. Wilson. Reduced Demand Uncertainty and the Sustainability of Collusion: How AI Could Affect Competition. Working Paper 341, FTC Bureau of Economics, June 2019.
- OECD. Algorithms and Collusion: Competition policy in the digital age, 2017.
- V. Pavlov and R. Berman. Price Manipulation in Peer-to-Peer Markets and the Sharing Economy, Sept. 2019.
- L. Rayo and I. Segal. Optimal Information Disclosure. *Journal of Political Economy*, 118(5):949–987, 2010. Publisher: The University of Chicago Press.
- A.-K. Roesler and B. Szentes. Buyer-Optimal Learning and Monopoly Pricing. *American Economic Review*, 107(7):2072–2080, July 2017.
- R. Schmalensee. Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination. *The American Economic Review*, 71(1):242–247, 1981. Publisher: American Economic Association.
- M. Spann, R. Zeithammer, M. Bertini, E. Haruvy, S. D. Jap, O. Koenigsberg, V. Mak, P. Popkowski Leszczyc, B. Skiera, and M. Thomas. Beyond Posted Prices: the Past, Present, and Future of Participative Pricing Mechanisms. *Customer Needs and Solutions*, 5(1):121–136, Mar. 2018.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, Massachusetts, and London, England, 2nd edition, 2018.
- T.-H. Teh. Platform Governance. *American Economic Journal: Microeconomics*, 14(3): 213–254, Aug. 2022.
- C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8:279–292, 1992.
- P. Ye, J. Qian, J. Chen, C.-h. Wu, Y. Zhou, S. De Mars, F. Yang, and L. Zhang. Customized Regression Model for Airbnb Dynamic Pricing. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 932–940, London, UK, Aug. 2018. Association for Computing Machinery.

Appendix

A Proofs

Proof of Lemma 1. For any induced posterior μ , the firm solves $\max_p E_\mu[D(p, \theta)(p - C)]$. By state affinity, this can be written $\max_p D(p, E_\mu[\theta])(p - C)$ and so the firm's choice depends only on $E_\mu[\theta]$. We can thus write this choice as $p(E_\mu[\theta])$. The platform's objective function is then $E_\mu[D(p(E_\mu[\theta]); \theta)]p(E_\mu[\theta]) = D(p(E_\mu[\theta]); E_\mu[\theta])p(E_\mu[\theta])$, where the final equality again follows from state affinity. \square

Proof of Proposition 1. I solve for the optimal information structure directly. Since f_1, f_2 are affine, they can be written as $f_1(p) = a_1 + b_1p, f_2(p) = a_2 + b_2p$. Now fix the posterior expected state x . By assumptions 1 and 2, the solution to the firm's problem is given by the first-order condition:

$$\begin{aligned} (f_1'(p) + xf_2'(p))(p - C) + f_1(p) + xf_2(p) &= 0 \\ 2(b_1 + b_2x)p &= (b_1 + b_2x)C - (a_1 + a_2x) \\ p &= \frac{1}{2} \left(C - \frac{a_1 + a_2x}{b_1 + b_2x} \right) \end{aligned}$$

Note that $b_1 + b_2x < 0$ since demand is decreasing. Hence, for any posterior expectation x the firm will choose price $p(x) = \frac{1}{2} \left(C - \frac{a_1 + a_2x}{b_1 + b_2x} \right)$. This gives

$$\begin{aligned} \tilde{R}(x) &= D(p(x); x)p(x) \\ &= (a_1 + a_2x + (b_1 + b_2x)p(x))p(x) \\ &= \frac{1}{2}(a_1 + a_2x + (b_1 + b_2x)C)p(x) \\ &= \frac{1}{4} \left((b_1 + b_2x)C^2 - \frac{(a_1 + a_2x)^2}{b_1 + b_2x} \right) \end{aligned}$$

It is straightforward to verify that $\tilde{R}(x)$ is convex. For instance, its second derivative is

$$\frac{-2(b_1a_2 - b_2a_1)^2}{(b_1 + b_2x)^3} \leq 0$$

where the inequality follows since $b_1 + b_2x < 0$. Moreover, it is generically strictly convex, since the numerator is non-zero unless $b_1a_2 = b_2a_1$. It is then immediate from the definition of a mean-preserving spread that full information is always an optimal information structure (since this corresponds with the most spread distribution of posterior means) and that generically it is the uniquely optimal information structure. \square

Proof of Proposition 2. Suppose that the price is p and the state is θ . Then since demand

is affine on $[0, \bar{p}]$ and 0 afterward consumer surplus is given by

$$CS = \frac{1}{2}(\bar{p} - p)(D(p; \theta) + D(\bar{p}; \theta))$$

which is just the area of the consumer surplus trapezoid.

Since the firm price depends only on the expected state, it is easy to see that a version of lemma 1 still holds: for posterior beliefs μ , expected consumer surplus is given by

$$E_\mu[CS] = \frac{1}{2}(\bar{p} - p(E_\mu[\theta]))(D(p(E_\mu[\theta]); E_\mu[\theta]) + D(\bar{p}; E_\mu[\theta]))$$

which depends only on the posterior mean. Therefore, consider expected consumer surplus given posterior mean x , written $CS(x)$:

$$\begin{aligned} CS(x) &= \frac{1}{2}(\bar{p} - p(x))(D(p(x); x) + D(\bar{p}; x)) \\ &= \frac{1}{2} \left(\bar{p}D(p(x); x) + \underbrace{\bar{p}D(\bar{p}; x)}_{\text{affine}} - p(x)D(\bar{p}; x) - \underbrace{p(x)D(p(x); x)}_{\tilde{R}(x)} \right) \end{aligned}$$

For the other two terms, recall that $p(x) = \frac{1}{2} \left(C - \frac{a_1 + a_2 x}{b_1 + b_2 x} \right)$. Hence,

$$\begin{aligned} \bar{p}D(p(x); x) &= \bar{p} \left((a_1 + b_1 p(x)) + (a_2 + b_2 p(x))x \right) \\ &= \bar{p} \left((a_1 + a_2 x) + \frac{1}{2}(b_1 + b_2 x) \left(C - \frac{a_1 + a_2 x}{b_1 + b_2 x} \right) \right) \\ &= \frac{\bar{p}}{2} \left((a_1 + a_2 x) + (b_1 + b_2 x)C \right) \\ p(x)D(\bar{p}; x) &= \frac{1}{2} \left(C - \frac{a_1 + a_2 x}{b_1 + b_2 x} \right) \left((a_1 + a_2 x) + (b_1 + b_2 x)\bar{p} \right) \\ &= \frac{1}{2} \left(C \left((a_1 + a_2 x) + (b_1 + b_2 x)\bar{p} \right) - (a_1 + a_2 x)\bar{p} - \frac{(a_1 + a_2 x)^2}{b_1 + b_2 x} \right) \end{aligned}$$

We can then see that the former term is also affine and the latter term has the same second derivative as $\tilde{R}(x)$. The terms with the same second derivative as $\tilde{R}(x)$ enter with a negative, and so $CS(x)$ is therefore always concave, and generically it will be strictly concave. The full result then follows from analogous logic to the proof of proposition 1. \square

Proof of Lemma 2. Note that firms face effective marginal cost $C = \frac{c}{\phi}$. So for state θ , firm i has the standard Hotelling best response function

$$BR_i(p_j) = \arg \max_p \frac{1}{2}(p - C)(1 + \theta(p_j - p))$$

Hence, when facing uncertainty over θ in the form of believing $\theta \sim \mu$, their best response

function is

$$BR_i(p_j) = \arg \max_p \frac{1}{2}(p - C)(1 + E_\mu[\theta](p_j - p))$$

These best response functions are of the same form as a standard Hotelling game with strength of consumer preferences $t = \frac{1}{E_\mu[\theta]}$. Thus, the unique pure strategy equilibrium has $p_1 = p_2 = C + \frac{1}{E_\mu[\theta]}$, as desired. \square

Proof of Proposition 3. The result follows immediately from the fact that $\frac{1}{x}$ is strictly convex for positive x . \square

Proof of Lemma 3. Fix a posterior distribution of beliefs μ . Consumers will face price $p = C + \frac{1}{E_\mu[\theta]}$. Moreover, regardless of the realized state, consumers will always buy the closest product. Hence, expected consumer surplus given posterior beliefs μ is

$$\begin{aligned} E_\mu[CS] &= E_\mu \left[v - C - \frac{1}{E_\mu[\theta]} - \frac{1}{\theta} \left(\int_0^{1/2} x dx + \int_{1/2}^1 x dx \right) \right] \\ &= v - C - \frac{1}{E_\mu[\theta]} - \frac{1}{4} E_\mu \left[\frac{1}{\theta} \right] \end{aligned}$$

\square

Proof of Lemma 4. Recall that expected consumer welfare conditional on posterior belief $\tilde{\mu}$ is given by $v - C - \frac{1}{E_{\tilde{\mu}}[\theta]} - \frac{1}{4} E_{\tilde{\mu}} \left[\frac{1}{\theta} \right]$. By Jensen's inequality, $E_{\tilde{\mu}} \left[\frac{1}{\theta} \right] > \frac{1}{E_{\tilde{\mu}}[\theta]}$ except when $\tilde{\mu} = \delta_{E_{\tilde{\mu}}[\theta]}$, in which case the two are equal (recall that $\Theta \subseteq \mathbb{R}_{>0}$). Now consider some distribution of posteriors $\tilde{\tau}$ that induces the distribution of posterior means μ and suppose it puts positive weight on a non-degenerate distribution $\tilde{\mu}$. From above we can see that expected consumer surplus would be strictly increased by replacing $\tilde{\mu}$ with the degenerate distribution $\delta_{E_{\tilde{\mu}}[\theta]}$, and moreover that by construction this does not change the induced distribution of posterior means. By proceeding with transformations of this kind, we can transform $\tilde{\tau}$ into τ since τ is the unique distribution that induces the distribution of posterior means μ and only puts positive weight on degenerate posteriors. As each transformation strictly increases expected consumer welfare, the result follows. \square

Proof of Proposition 4. The result follows immediately from the fact that $\frac{-1}{x}$ is strictly concave for positive x . \square

B Robustness

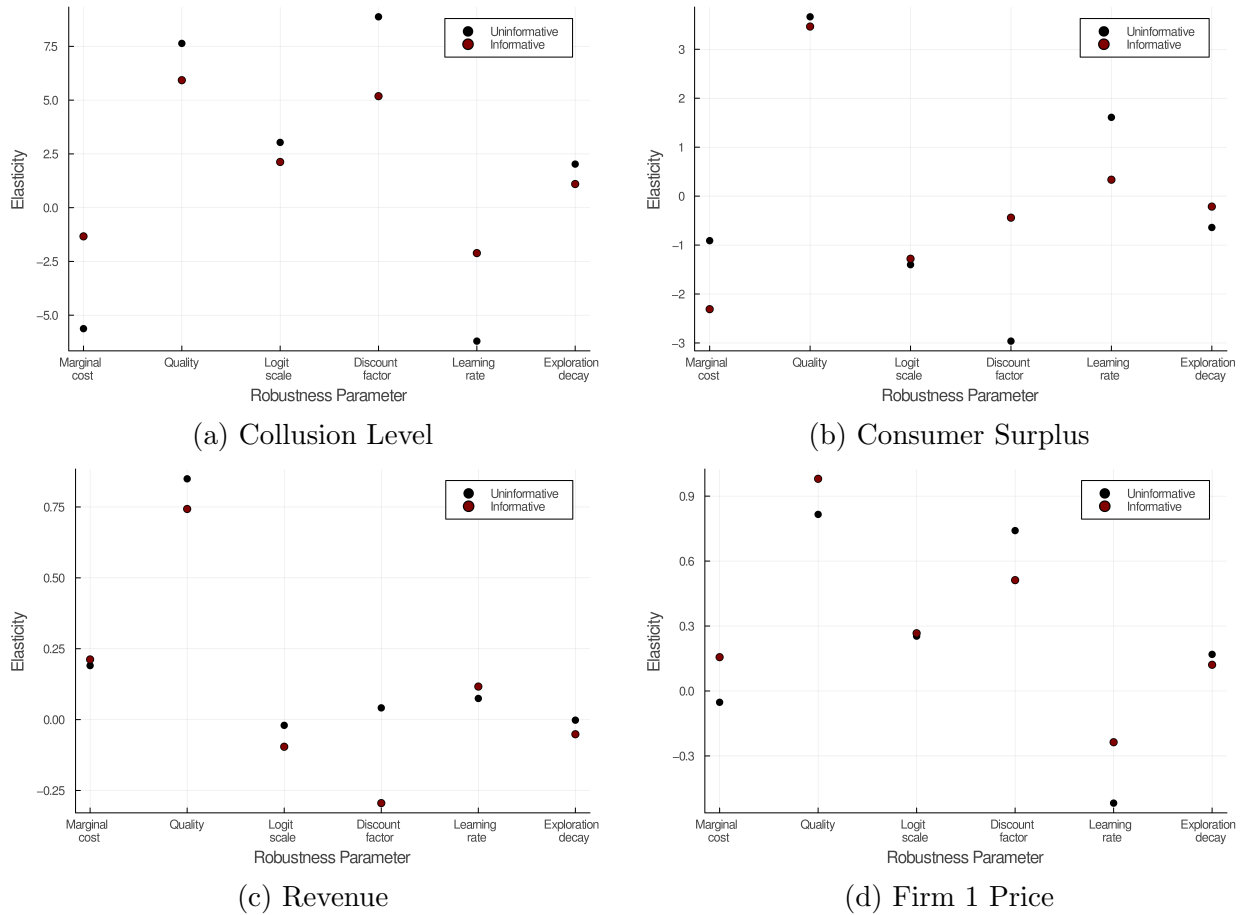


Figure 7: Robustness

Notes: These figures plot the elasticity of each outcome to a change in a particular parameter, separately for uninformative and informative recommendations. Panel (a) plots the elasticity of the collusion level, panel (b) plots the elasticity of the average consumer surplus, panel (c) plots the elasticity of the average revenue, and panel (d) plots the elasticity of the price for firm 1. Elasticities were computed by decreasing each parameter 10% from its baseline value while holding the other parameters fixed.