

Online Privacy and Information Disclosure by Consumers

JOB MARKET PAPER

Latest draft available online at stanford.edu/~shota2/

Shota Ichihashi*

November 16, 2017

Abstract

I study the question of what information consumers should disclose to firms. A consumer discloses information to a multi-product firm, which learns about his preferences, sets prices, and makes product recommendations. While the consumer benefits from disclosure as it enables the firm to make accurate recommendations, the firm may use the information to price discriminate. I show three main results. First, the firm prefers to commit not to price discriminate, which encourages the consumer to provide information that is useful for product recommendations. This explains why price discrimination by online sellers seems to be uncommon, which, empirically, has been puzzling. Second, nondiscriminatory pricing hurts the consumer, who would be better off by precommitting to withhold some information. This provides a rationale for privacy regulations that limit the amount of personal data that firms can possess. Third, equilibrium is often inefficient even if the consumer can potentially disclose any information about his preferences.

*Stanford University, Department of Economics, Email: shota2@stanford.edu. I am grateful to Paul Milgrom, Gabriel Carroll, Matthew Gentzkow, and Ilya Segal for their unfailing support and guidance. I would like to thank Mohammad Akbarpour, Luca Braghieri, Isaias Chaves, Weixin Chen, Piotr Dworzak, Takako Fujiwara-Greve, Pedro Gardete, Michihiro Kandori, Fuhito Kojima, Shengwu Li, Sitian Liu, Giorgio Martini, Guido Martirena, Shunya Noda, Takuo Sugaya, and participants in Paul Milgrom's group meeting, Stanford Theory Lunch, and Stanford Market Design Seminar. I am grateful for the financial support from the Yoshida Scholarship Foundation and the Stanford Graduate Fellowship.

1 Introduction

This paper studies the question of what information consumers should disclose to firms, which is a first-order issue in the Internet economy. Online firms can observe detailed information about consumers, such as their browsing histories, purchases, and characteristics; however, consumers can control whether and to what extent this information is revealed. For instance, they can delete cookies to hide their web-browsing histories, and they can create multiple accounts on shopping websites to obfuscate their purchasing behavior. For policymakers, information disclosure by consumers is an important consideration in formulating policies concerning online privacy.

In this paper, I focus on the following economic trade-off: The benefit for consumers to disclose information is that firms can recommend or advertise appropriate products. The cost is that firms may use this information to price discriminate. For instance, Amazon, Netflix, Spotify, and other e-commerce firms use consumers' personal data to offer product recommendations, which help consumers discover items that they might not have found otherwise, and Facebook, Twitter, and YouTube display online ads tailored to consumers' interests. However, these firms could potentially use such information to obtain estimates of consumers' willingness to pay for recommended or advertised products and, in turn, set prices on this basis.¹

I study a model that captures this trade-off. The model consists of a monopolistic firm that sells K products and a consumer with unit demand. The consumer has limited attention: He is initially uninformed of the valuations of the products, and it is prohibitively costly for him to assess all available products. At the beginning of the game, the consumer chooses a "disclosure policy," which determines what the firm learns about his willingness to pay for each product. For example, two disclosure policies could correspond to sharing his browsing history or not, where sharing it enables the firm to form a more accurate estimate of the valuations.

I consider two games that differ in the firm's ability to price discriminate. Under the *discriminatory pricing regime*, the firm sets prices after observing the information provided by the consumer. Under the *nondiscriminatory pricing regime*, the firm posts prices without observing the information. Comparing these two regimes is useful to illustrate how firms' ability to price discriminate affects consumers' incentive to disclose information. In either regime, after learning about the consumer's valuations, the firm selects one of the K products and recommends it. Finally, the consumer decides whether to buy the recommended product.

¹Section 5 shows that we can interpret my model as a game between a consumer and an online advertising platform that holds ad auctions, where advertisers set prices for their products and join the auctions.

I model information disclosure by the consumer as in the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011). In other words, without observing the valuations, the consumer chooses what information to be disclosed to the firm. The idea is that while it is difficult for consumers themselves to determine which item in a huge set of available products is most appropriate for them, firms can do this using consumers’ personal data. For instance, firms might analyze consumers’ browsing histories by using their knowledge of the products’ characteristics, the prior experiences of other consumers, and their computing power. These enable firms to map a given consumer’s history into estimates of the values of products, even though the consumer himself cannot learn the whole vector of valuations for all products in the market.

The model has a second interpretation, in which a large number of consumers disclose information to the firm. Formally, I can interpret the model as information disclosure by a continuum of consumers, where the different pricing regimes specify whether the firm can charge different prices to different consumers. Under nondiscriminatory pricing in which the firm commits to set a single price for each product, each consumer is non-pivotal in the sense that his disclosure does not affect prices. Thus, it is as if the firm sets prices without observing information disclosed. This alternative interpretation sheds light on the divergence between individual and collective incentives to disclose information.

I obtain three main findings, which have implications for understanding observed facts, designing privacy regulations, and the literature on information design. First, firms prefer to commit not to price discriminate. A rough intuition is that discriminatory pricing discourages consumers from providing information and lowers firms’ revenue. This result gives an economic explanation of an empirical puzzle: “The mystery about online price discrimination is why so little of it seems to be happening” (Narayanan, 2013). Namely, price discrimination by online sellers seems to be uncommon despite their potential ability to use consumers’ personal data to obtain estimates of consumers’ willingness to pay and, in turn, vary prices on this basis to capture more of the surplus.²

The second main finding is that consumers could be better off by committing to withhold some information. Under nondiscriminatory pricing, the consumer fully discloses information to obtain a better recommendation without worrying about how the firm uses the information to price discriminate. Expecting this full disclosure, the firm sets high (nondiscriminatory) prices. I formally prove this by showing that if the consumer reveals more information about which product is more valuable, the consumer draws the value of the recommended

²There have been several attempts by researchers to detect price discrimination by e-commerce websites. For instance, Iordanou et al. (2017) examine around two thousand e-commerce websites and they “conclude that the specific e-retailers do not perform PDI-PD (personal-data-induced price discrimination).”

product from a distribution associated with a lower price elasticity of demand, which gives the firm an incentive to set a higher price. As a result, the consumer could be better off by committing to disclose less information, which leads to lower prices. Furthermore, the second interpretation of the model reveals that consumers are collectively worse off under nondiscriminatory pricing because of a negative externality through prices associated with information sharing. This result gives a rationale for privacy regulations that limit consumer information disclosure, or equivalently, privacy regulations that limit the amount of personal data that firms can possess to a level below how much each consumer wants them to have.

The third main finding is that the equilibrium is often inefficient even if the consumer could potentially disclose enough information to achieve an efficient allocation. In particular, I show that there can be no efficient equilibrium under discriminatory pricing. This is in contrast to the single-product setting of [Bergemann, Brooks, and Morris \(2015\)](#), in which an efficient equilibrium always exists. My result highlights the difficulty of simultaneously achieving efficient price discrimination, efficient matching of products, and consumer-surplus maximization. I prove that an efficient equilibrium never exists under a mild condition without explicitly solving the consumer's equilibrium strategy, which is a solution of a multidimensional Bayesian persuasion problem. This approach could be useful to understand the qualitative nature of optimal information structures in complex information design problems.

The remainder of the paper is organized as follows. In [Section 2](#), after discussing related work, I present the baseline model of information disclosure, pricing, and product recommendation. I also provide a second interpretation of the model as information disclosure by a continuum of consumers. In [Section 3](#), I restrict the consumer to choosing from disclosure policies that reveal different amounts of information about his horizontal taste. I show that the firm prefers to commit to nondiscriminatory pricing, which hurts consumers. [Section 4](#) allows the consumer to choose any disclosure policy. I provide conditions under which equilibrium is inefficient and show how different pricing regimes lead to different kinds of inefficiency, then use these inefficiency results to show that the firm still prefers to adopt a nondiscriminatory pricing regime. [Section 5](#) extends the model to consider competition between multi-product firms. I show that disclosure benefits the consumer by intensifying price competition. This section also explains that the model could apply to the context of an online advertising platform that runs ad auctions. [Section 6](#) discusses the results and concludes.

1.1 Related Work

My work relates to two strands of literature: the literature on information design and that on the economics of privacy. In terms of modeling, the point of departure is [Bergemann, Brooks, and Morris \(2015\)](#). They consider a single-product monopoly pricing problem in which a monopolist has additional information about a consumer’s taste. I extend their framework by considering a multi-product firm and a consumer with limited attention, which renders information useful not only for pricing but also for product recommendation. These additional components yield different welfare consequences of information disclosure and discriminatory pricing. In contrast to [Bergemann et al. \(2015\)](#), who consider the entire set of attainable surplus, I restrict attention to the case in which the consumer discloses information to maximize his own payoff. Part of my analysis employs their “greedy algorithm,” which generates a consumer-optimal information disclosure policy given any prior valuation distribution.

My work also relates to the economics of privacy literature. As a growing number of transactions are conducted online and based on data about consumers’ behavior and personal traits, recent strands of literature have devoted considerable attention to the relationships between personal data and intertemporal price discrimination ([Acquisti and Varian, 2005](#); [Conitzer et al., 2012](#); [Fudenberg and Tirole, 2000](#); [Fudenberg and Villas-Boas, 2006](#); [Taylor, 2004](#); [Villas-Boas, 1999, 2004](#)). In these models, firms learn about a consumer’s preference from his purchase history. Thus, personal data endogenously arises as a history of a game. A consumer’s attempt to conceal personal information is often formulated as delaying purchase or erasing purchase history. In contrast, I assume that the consumer is endowed with his personal data at the outset.

Several papers, such as [Conitzer et al. \(2012\)](#) and [Montes et al. \(2017\)](#), examine consumers’ endogenous privacy choices when they are endowed with personal data at the outset. [Braghieri \(2017\)](#) studies a consumer search model in which a consumer can choose to be targeted by revealing his horizontal taste to firms instead of engaging in costly search. In the model, targeting plays a similar role to information disclosure in this paper: It enables consumers to find their favored products at low cost, but it can also hurt them because of discriminatory pricing. A similar trade-off can be observed in [De Corniere and De Nijs \(2016\)](#), who study a platform’s choice of disclosing consumers’ preferences to advertisers.

This paper differs from these works in three ways. Most importantly, my main focus is not on how and when consumers choose to protect their privacy or how consumers are harmed by discriminatory pricing. Instead, I aim to provide a potential explanation for why discriminatory pricing seems to be uncommon on the Internet, and why consumers seem to share their information so casually despite the potential negative effect of price

discrimination. The second difference is in the formulation. While these works assume that a consumer’s privacy choice is full or no disclosure, I assume that a consumer can choose what kind of information to disclose. This enables me to compare the amount and type of information disclosed under different environments. Third, and relatedly, I emphasize the difference between individual and group incentives to share personal data, which would be useful for understanding the nontrivial relationship between the actual market behavior of consumers and its welfare implications.

Beyond the context of online disclosure, my work relates to voluntary information disclosure in bilateral transactions (Glode et al., 2016). Also, as information disclosure with commitment can be interpreted as a combination of information gathering and truthful disclosure, my work also relates to information gathering by buyers before trade (Roesler, 2015; Roesler and Szentes, 2017).

2 Baseline Model

There is a monopolistic firm that sells multiple products $k = 1, \dots, K$ with the set of the products denoted by \mathcal{K} . There is a single consumer who has unit demand, in that he eventually consumes one of K products or nothing. The consumer’s value for product k , denoted by u_k , is drawn independently and identically across $k \in \mathcal{K}$ according to a probability distribution x_0 that has a density and is supported on $V \subset \mathbb{R}_+$.³ Let $u := (u_1, \dots, u_K)$ denote the consumer’s valuation vector. Appendix C contains the corresponding results for a more general x_0 , which may not have a density.

The consumer’s preferences are quasi-linear: If he purchases product k at price p , his ex post payoff is $u_k - p$. If he does not purchase any product, he obtains a payoff of zero. The firm’s payoff is its revenue. The consumer and the firm are risk-neutral.

At the beginning of the game, before observing u , the consumer chooses a *disclosure policy* (M, ϕ) from an exogenously given set \mathcal{D} . Each element of \mathcal{D} is a pair of a message space M and a function $\phi : V^K \rightarrow \Delta(M)$, where $\Delta(M)$ is the set of the probability distributions over M . After the consumer chooses a disclosure policy (M, ϕ) , Nature draws $u \in V^K$ and a message $m \in M$ according to $\phi(\cdot|u) \in \Delta(M)$. In this application, \mathcal{D} consists of consumers’ privacy choices, such as whether to share one’s browsing history or not. If \mathcal{D} consists of all disclosure policies, information disclosure takes the form of Bayesian persuasion studied by Kamenica and Gentzkow (2011). Hereafter, I sometimes write a disclosure policy as ϕ instead of (M, ϕ) .

³See Remark 4 for how the results extend to correlated values.

Next, I describe the firm’s pricing. I consider two games that differ in whether the firm can set prices contingent on the information provided by the consumer. Under the *nondiscriminatory pricing regime*, the firm sets the price of each product *simultaneously* with the consumer’s choice of a disclosure policy (M, ϕ) .⁴ Under the *discriminatory pricing regime*, the firm sets the price of each product *after* observing a disclosure policy (M, ϕ) and a realized message m . Note that under nondiscriminatory pricing, the firm not only does not base prices on a realized message m but also does not base prices on a disclosure policy ϕ . For example, if firms adopt this regime, they set prices based on neither browsing history nor whether consumers share their browsing history.

Under both pricing regimes, after observing a disclosure policy (M, ϕ) and a realized message m , the firm recommends one of the K products. The consumer observes the value and price of the recommended product and chooses whether to buy it. If he does not purchase the recommended product, he obtains a payoff of zero.

The timing of the game under each pricing regime, which is summarized in [Figure 1](#), is as follows. First, the consumer chooses a disclosure policy $(M, \phi) \in \mathcal{D}$. Under the nondiscriminatory pricing regime, the firm simultaneously sets the price of each product. Then Nature draws the consumer’s valuations u and a message $m \sim \phi(\cdot|u)$. After observing (M, ϕ) and m , the firm recommends a product. Under the discriminatory pricing regime, the firm sets the price of the recommended product at this point. Finally, the consumer decides whether to buy it.

My solution concept is subgame perfect equilibrium in which the firm breaks a tie in favor of the consumer whenever it is indifferent. Under nondiscriminatory pricing, I consider *symmetric equilibrium* in which each product has the same price.

Remark 1 (Discussion of Modeling Assumptions). I assume that the consumer commits to a disclosure policy before observing u . This would be suitable, for instance, if the consumer is not informed of the existence of products or products’ characteristics, but understands that his personal data enable the firm to learn about which product is valuable to him. In [Section 5](#), I provide a microfoundation for this idea in a model of two-sided private information in which the consumer is informed of his subjective taste and the firm

⁴I can alternatively assume that under nondiscriminatory pricing, the firm sets prices first and the consumer chooses a disclosure policy after observing the prices. This alternative assumption does not change the results in [Section 3](#). In contrast, this sequential move could affect the results in [Section 4](#) because the firm may set different prices for different products to induce the consumer to choose an asymmetric disclosure policy in equilibrium. This could give the firm greater revenue than some symmetric disclosure policy that would be sustained in equilibrium if the consumer and the firm move simultaneously. See the discussion after [Proposition 2](#).

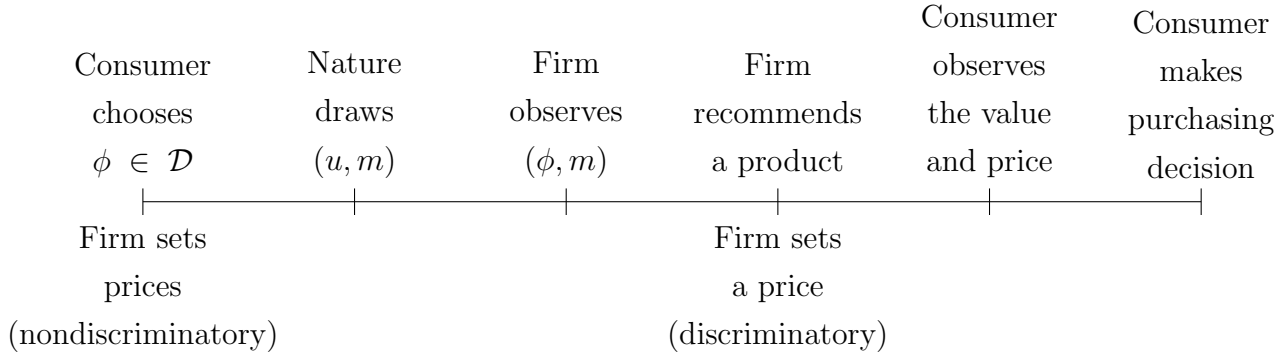


Figure 1: Timing of moves under each pricing regime

is informed of its products' characteristics. The commitment assumption is natural when studying information disclosure in online marketplaces, as consumers or regulators typically set disclosure rules up front and incentives to distort or misrepresent one's browsing history or characteristics seem to be less relevant.

There are also three substantial assumptions on the firm's recommendation and the consumer's purchasing decision. First, the firm recommends one of the K products and cannot offer a more general mechanism. This reflects the idea that the consumer has limited attention and cannot examine complex contracts.

Second, the consumer observes his willingness to pay for the recommended product before he decides whether to buy it. This is reasonable in settings in which a consumer can learn the value after the purchase and return it for a refund whenever the price exceeds the value. It would be interesting to consider a setting in which firms have superior knowledge about how valuable products are to each consumer and can convey some information about the valuations through product recommendations.⁵

Third, the consumer can only buy the product the firm recommends to him. This is to simplify the purchasing decision so that I can focus on the consumer's incentives regarding information provision. Moreover, this assumption is in line with models of informative advertising in which a consumer can buy a product only if he receives an advertisement that informs him of the existence of the product (Butters, 1977; Grossman and Shapiro, 1984; Anderson and De Palma, 2012).

Finally, it is *not* without loss of generality to assume that production costs are equal

⁵This extension is similar to Bayesian persuasion but is different in that product recommendations affect both the consumer's learning and choice set. Rayo and Segal (2010) study a related problem in a single-product setting. Ichihashi (2017) studies the question of what information should be at firms' disposal in a single-product setting.

across products. (Assuming that they are equal, it is without loss to normalize them to zero.) For example, if the firm incurs low production cost for product 1, it has a greater incentive to recommend product 1 even if it is less valuable to the consumer than other products. Correspondingly, heterogeneous production costs are likely to affect the consumer’s incentive to disclose information. For instance, the consumer might be better off if the firm is less informed about the consumer’s values for high-margin products as this can make it less profitable for the firm to recommend those products. I leave this extension for future research.

Remark 2 (Other Applications). I use consumers’ privacy in online marketplaces as the main application. However, there are many other potential applications. For instance, the model could describe the following situation: An employer assigns his worker one of K tasks, the completion of which delivers a fixed value to the employer. The worker discloses information about cost c_k that he incurs to complete each task k . Two pricing regimes could correspond to whether the wage should be contingent on the information revealed.

2.1 Interpretation as a Model with a Continuum of Consumers

We can interpret the current setting as the reduced form of a model in which a *continuum of consumers* disclose information about their preferences. This interpretation enables us to draw insights about the divergence between individual and collective incentives to share information.

Formally, consider a continuum of consumers, each of whom discloses information in the same way as the consumer in the baseline model. I assume that the valuation vectors are independent across consumers.⁶ Under nondiscriminatory pricing, *after* observing the information provided, the firm sets a single price for each product. In other words, the firm cannot use individual-level information but can use aggregate information to set prices. Under discriminatory pricing, the firm can charge different prices to different consumers, which leads to exactly the same model as the baseline model. Under both regimes, the firm can recommend different products to different consumers. In this interpretation, the two pricing regimes do not differ in the timing of moves. Instead, they impose different restrictions on the firm’s strategy space.

To see that this alternative interpretation is equivalent to the original formulation, consider an equilibrium under nondiscriminatory pricing in which all consumers take the same

⁶The independence of valuation vectors across a continuum of consumers might raise a concern about the existence of a continuum of independent random variables. Sun (2006) formalizes the notion of a continuum of IID random variables for which the “law of large numbers” holds.

disclosure level δ^* . In the equilibrium, δ^* must be optimal for each consumer when he chooses a disclosure level taking prices as given, because an individual is nonpivotal and cannot influence the firm’s pricing.⁷ Moreover, the firm sets the price of each product to maximize its revenue given the commonly chosen disclosure level δ^* . This mutual best-response condition is the same as the equilibrium condition of the original formulation.

I derive the main results using the model of a single consumer because I can obtain clean statements without considering technical details specific to a game with a continuum of players. However, I refer to the alternative interpretation whenever it provides additional insights. [Appendix D](#) formally provides an appropriate equilibrium notion for a continuum of consumers under which I can identify the original formulation as the model with a continuum of consumers.

2.2 Horizontal and Vertical Information

My model contains two kinds of information. One is “horizontal information,” which is relevant to product recommendations. For instance, if the firm learns that product k is most valuable to the consumer, the firm would recommend it instead of other less valuable products. The other is “vertical information,” which is relevant to price discrimination. For instance, suppose the firm additionally learns that the consumer’s value of product k is high. Then, under discriminatory pricing, the firm would charge a higher price to extract more of the surplus. [Bergemann et al. \(2015\)](#) generally study how vertical information affects welfare in a single-product setting.

One might expect that the consumer prefers to reveal horizontal information to receive better recommendations and prefers to conceal vertical information to avoid price discrimination. However, as I show later, the consumer’s informational incentive is subtler than this intuition. In particular, the consumer may prefer not to reveal horizontal information, even if this could lower the chance that he is recommended the best product.

To isolate the novel features of the model, in [Section 3](#), I assume that \mathcal{D} consists of disclosure policies that reveal different amounts of information about the consumer’s horizontal taste. In [Section 4](#) I assume that \mathcal{D} consists of all disclosure policies.

3 Disclosing Horizontal Information

To study the disclosure of horizontal information, I assume that the firm sells two products ($K = 2$). Then, I identify \mathcal{D} with $[1/2, 1]$ and call each $\delta \in [1/2, 1]$ a *disclosure level*. Each

⁷[Appendix E](#) formalizes this.

disclosure level δ corresponds to disclosure policy $(\{1, 2\}, \phi_\delta)$ that sends message k with probability δ whenever $u_k = \max(u_1, u_2)$. Note that the greater δ is, the more informative ϕ_δ is in the sense of Blackwell. [Figure 2](#) shows the disclosure policy corresponding to δ .




Figure 2: Disclosure policy for $\delta \in [1/2, 1]$

Let F^{MAX} and F^{MIN} denote the cumulative distribution functions of $\max(u_1, u_2)$ and $\min(u_1, u_2)$, respectively. Conditional on a disclosure level δ and message $k \in \{1, 2\}$, the posterior distribution of u_k is $\delta F^{MAX} + (1 - \delta)F^{MIN}$. Thus, considering the stochastic relationship between F^{MAX} and F^{MIN} is useful. I show that F^{MAX} is greater than F^{MIN} in a sense stronger than the usual stochastic order.

Definition 1. Let G_0 and G_1 be two CDFs. G_1 is greater than G_0 in the *hazard rate order* if $\frac{1-G_1(z)}{1-G_0(z)}$ increases in $z \in (-\infty, \max(s_1, s_0))$. ($a/0$ is taken to be equal to $+\infty$ whenever $a > 0$.) Here, s_0 and s_1 are the right endpoints of the supports of G_0 and G_1 , respectively.

If G_0 and G_1 have densities g_0 and g_1 , the above definition is equivalent to

$$\frac{g_0(z)}{1 - G_0(z)} \geq \frac{g_1(z)}{1 - G_1(z)}, \forall z \in (-\infty, \max(s_1, s_0)).$$

The next lemma states that the valuation distribution of the preferred product is greater than that of the less preferred product in the hazard rate order. The proof is in [Section 5](#), where I prove the same result for a more general formulation of horizontal information.

Lemma 1. F^{MAX} is greater than F^{MIN} in the hazard rate order.

The intuition is as follows. Suppose that the firm recommends the most favored product to the consumer at price p , and he is willing to pay at least p for it. Then, suppose that the firm marginally increases the price from p to $p + \varepsilon$. After seeing this price increment, the consumer gives up buying only when both u_1 and u_2 are below $p + \varepsilon$ as the value is given by $\max(u_1, u_2)$. Alternatively, suppose that the firm recommends the least favored

product at price p , and the consumer is willing to pay at least p for it. If the firm marginally increases the price from p to $p + \varepsilon$, the consumer gives up buying whenever one of u_1 and u_2 is below $p + \varepsilon$ as the value is given by $\min(u_1, u_2)$. That is, the consumer is less likely to give up buying after observing marginal price increment if he is recommended his most favored product. This implies that the valuation distribution F^{MAX} of the most favored product has a lower hazard rate than F^{MIN} does.

Remark 3. Assuming $K = 2$ is to simplify the analysis. For a general K , for instance, I can conduct the same analysis by considering the following disclosure policy ϕ_δ corresponding to each disclosure level δ : ϕ_δ sends message k with probability δ if $k \in \arg \max_{k \in \mathcal{K}} u_k$, and it sends a message uniformly randomly from $\{1, \dots, K\}$ with probability $1 - \delta$.

Also, I can relax the assumption that the disclosure policy corresponding to $\delta = 1$ deterministically sends message k whenever $k \in \arg \max_k u_k$. The identical result holds, for instance, if I consider the disclosure policy that sends message k whenever $k \in \arg \max_k (u_k + \varepsilon_k)$ where ε_k is IID across k . [Section 5](#) more generally defines a disclosure policy that reveals horizontal information and shows that the identical results hold if the consumer can choose any garblings of such a disclosure policy.

3.1 Equilibrium Analysis

I solve the game backwards. The following lemma describes the firm's equilibrium recommendation strategy, which is common between the two pricing regimes. The result is quite intuitive: The firm can maximize its revenue by recommending the product that the consumer is more likely to prefer. For the proof, see [Appendix A](#).

Lemma 2. *Fix a pricing regime and take any equilibrium. Suppose that the consumer has chosen a disclosure level $\delta > 1/2$ and the firm follows the equilibrium strategy. Then, the firm recommends product $k \in \{1, 2\}$ if message $m = k$ realizes.*

Next, I show the key comparative statics: Under discriminatory pricing, information disclosure hurts the consumer through a higher price. To state the result generally, define $P(\delta) := \arg \max_{p \in \mathbb{R}} p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$. The following result holds for any prior distribution x_0 .

Lemma 3. *Suppose that the consumer has chosen a disclosure level δ under discriminatory pricing. Then, the firm sets price $p(\delta) = \min P(\delta)$ in equilibrium. Furthermore, $P(\delta)$ is increasing in δ in the strong set order, and thus $p(\delta)$ is increasing in δ .⁸*

⁸ $A \subset \mathbb{R}$ is greater than $B \subset \mathbb{R}$ in the strong set order if $a \in A$ and $b \in B$ imply $\max(a, b) \in A$ and $\min(a, b) \in B$.

Proof. By Lemma 2, the consumer is recommended his preferred product with probability δ . Then, at price p , the consumer buys the recommended product with probability $1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)$. Thus, the expected revenue is $p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$, which is maximized at $p \in P(\delta)$. Because the firm breaks a tie in favor of the consumer, it sets price $p(\delta)$ at equilibrium.

To show that $p(\delta)$ is increasing in δ , note that

$$\begin{aligned} & \log p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)] - \log p[1 - \delta' F^{MAX}(p) - (1 - \delta')F^{MIN}(p)] \\ &= \log \frac{1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)}{1 - \delta' F^{MAX}(p) - (1 - \delta')F^{MIN}(p)}. \end{aligned} \quad (1)$$

By Theorem 1.B.22 of Shaked and Shanthikumar (2007), if $\delta > \delta'$, $\delta F^{MAX} + (1 - \delta)F^{MIN}$ is greater than $\delta' F^{MAX} + (1 - \delta')F^{MIN}$ in the hazard rate order. Then, (1) is increasing in p . This implies that $\log p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$ has increasing differences in (p, δ) . By Topkis (1978), $P(\delta)$ is increasing in the strong set order. \square

The proof shows that the more information the consumer discloses, the less “elastic” demand the firm confronts. Indeed, the hazard rate order relates to the price elasticity of demand. Namely, for two CDFs F_1 and F_0 , F_1 is greater than F_0 in the hazard rate order if and only if the “demand curve” associated with F_1 has a lower price elasticity of demand than the one associated with F_0 . Here, as in Bulow and Roberts (1989), the demand curve for F is given by $D(p) = 1 - F(p)$. Then, the price elasticity of demand is given by $-\frac{d \log D(p)}{d \log p} = \frac{f(p)}{1 - F(p)}p$. Thus, the comparison of hazard rates is equivalent to that of the price elasticity of demands.

Note that first-order stochastic dominance would not be sufficient for the price to be increasing in δ . For example, suppose that F_0 puts equal probability on values 1 and 3 and that F_1 puts equal probability on 2 and 3. Though F_1 first-order stochastically dominates F_0 , the monopoly price under F_0 is 3, while the one under F_1 is 2.

The following result states that the firm prefers to commit to not price discriminate. While this encourages the consumer to provide information, he could be better off by committing to disclose less information, which is sustained in equilibrium under discriminatory pricing.

Theorem 1. *In the unique symmetric equilibrium under nondiscriminatory pricing, the consumer chooses the highest disclosure level 1. The firm’s equilibrium payoff is higher under nondiscriminatory pricing than under discriminatory pricing. The consumer is charged higher prices and obtains a lower equilibrium payoff under nondiscriminatory pricing than under discriminatory pricing.*

Proof. If the consumer is recommended his preferred and less preferred products at price p , his expected payoffs are $u^{MAX}(p) := \int_p^{+\infty} (v - p) dF^{MAX}(v)$ and $u^{MIN}(p) := \int_p^{+\infty} (v - p) dF^{MIN}(v)$, respectively.

Consider nondiscriminatory pricing. Let p^* denote the equilibrium price. If the consumer chooses δ , his expected payoff is

$$\delta u^{MAX}(p^*) + (1 - \delta) u^{MIN}(p^*),$$

which is uniquely maximized at $\delta = 1$. Thus, the firm sets price $p(1)$ at equilibrium. The consumer's resulting payoff is $u^{MAX}(p(1))$.

Consider discriminatory pricing. If the consumer chooses δ , his payoff is

$$\delta u^{MAX}(p(\delta)) + (1 - \delta) u^{MIN}(p(\delta)). \quad (2)$$

In equilibrium, the consumer obtains

$$\max_{\delta \in [1/2, 1]} \delta u^{MAX}(p(\delta)) + (1 - \delta) u^{MIN}(p(\delta)) \geq u^{MAX}(p(1)).$$

In contrast, the firm is better off under nondiscriminatory pricing. Note that $\delta F^{MAX} + (1 - \delta) F^{MIN}$ is stochastically increasing in δ and thus $p[1 - \delta F^{MAX}(p) - (1 - \delta) F^{MIN}(p)]$ is increasing in δ for any p . Then for any δ ,

$$\begin{aligned} p[1 - F^{MAX}(p)] &\geq p[1 - \delta F^{MAX}(p) - (1 - \delta) F^{MIN}(p)], \forall p \\ \Rightarrow p(1)[1 - F^{MAX}(p(1))] &\geq p[1 - \delta F^{MAX}(p) - (1 - \delta) F^{MIN}(p)], \forall p \\ \Rightarrow p(1)[1 - F^{MAX}(p(1))] &\geq p(\delta)[1 - \delta F^{MAX}(p(\delta)) - (1 - \delta) F^{MIN}(p(\delta))]. \end{aligned}$$

□

The intuition is as follows. Under nondiscriminatory pricing, the consumer fully reveals whether $u_1 > u_2$ to obtain a better recommendation without worrying about how the firm uses the information to set prices. Expecting that the consumer fully reveals information and the firm can make a perfect recommendation, it can set a high price for each product. In contrast, under discriminatory pricing, the consumer is the Stackelberg leader who balances the benefit of a more accurate recommendation and the cost from a higher price. As a result, the consumer typically withholds some information. This leads to lower prices and gives the consumer a greater payoff than under nondiscriminatory pricing.

The result gives an economic explanation of the observed puzzle—that is, online price discrimination based on personal data seems to be uncommon. My analysis suggests that if firms started price discriminating on the basis of past browsing or purchases, consumers

would have a strong incentive to hide or obfuscate this history. This would lower the quality of the match between consumers and products and hurt firms. In particular, the result emphasizes the importance of precommitment to nondiscriminatory pricing. This is in contrast to other rationales to avoid price discrimination, such as the existence of menu costs or customer antagonism (Anderson and Simester, 2010). In practice, a firm might commit to nondiscriminatory pricing by its privacy policy or a public statement. For instance, in 2000 Amazon CEO Jeff Bezos said, “We never have and we never will test prices based on customer demographics.”⁹

Theorem 1 also gives a rationale for privacy regulations that limit the amount of information consumers can share with firms. In my model, this is equivalent to restricting the amount of information firms can possess. Namely, if a regulator prohibits the firm from holding information exceeding the level that maximizes [equation \(2\)](#), then consumer welfare is greater than in the absence of the regulation.

That the firm prefers nondiscriminatory pricing might seem to be specific to the current setting, in which the optimal price depends on a disclosure level but not on a realized message itself. (Recall that the optimal price $p(\delta)$ does not depend on whether message 1 or 2 realizes.) In other words, the current setting excludes the potential benefit of discriminatory pricing by which the firm can set a price based on a realized message. For instance, if the consumer can disclose his willingness to pay for his favorite product, then the firm may want to adopt discriminatory pricing so that it can set the price based on the realized willingness to pay. However, in [Section 4](#), I show that the firm still prefers the nondiscriminatory pricing regime, even if the consumer can potentially disclose any information.

How much the firm benefits and the consumer loses under nondiscriminatory pricing depends on both the prior valuation distribution for each product and the number K of products. For instance, the outcome that is best for the firm and worst for the consumer arises if there are a large number of products. See [Remark 3](#) for the generalization to K products. The proof of the following result is in [Appendix B](#).

Proposition 1. *Suppose that the value of each product is independently drawn from a probability distribution that has a positive density and a compact interval support $[a, b]$. In the symmetric equilibrium under nondiscriminatory pricing, the firm’s payoff goes to b and the consumer’s payoff goes to 0 as $K \rightarrow +\infty$. Under discriminatory pricing, there exists $\underline{u} > 0$ such that the consumer’s equilibrium payoff is greater than \underline{u} for any K .*

⁹<http://www.e-commercetimes.com/story/4411.html>

3.2 Implication of [Theorem 1](#) for a Continuum of Consumers

[Theorem 1](#) is a classic tragedy of the commons in the alternative interpretation of the model, in which a continuum of consumers disclose information. First, as in [Theorem 1](#), each consumer in a continuum of population chooses the highest disclosure level under nondiscriminatory pricing, because the disclosure by an individual does not affect the price of each product but rather leads to a better recommendation for him. In contrast, under discriminatory pricing, each consumer chooses a disclosure level that balances the benefit and cost, as information directly affects the price offered to him.

Now, why are consumers worse off under nondiscriminatory pricing even though the firm sets prices after observing information? This is attributed to the presence of a negative externality associated with information sharing. Under nondiscriminatory pricing, if some consumers disclose more information, the firm sets a higher price for each product, which is shared by all consumers. Thus, information disclosure by some consumers lowers the welfare of other consumers through higher prices. As consumers do not internalize this effect, they prefer to fully reveal information in equilibrium. The full revelation results in higher prices for all consumers and lowers their joint welfare. [Appendix E](#) formalizes this observation.

3.3 Does Nondiscriminatory Pricing Increase Total Welfare?

Whether the firm's commitment to nondiscriminatory pricing increases efficiency depends on the prior valuation distribution x_0 , as shown in [Appendix F](#). nondiscriminatory pricing could increase total welfare because it encourages the consumer to disclose information and improves the quality of the match between the consumer and the product. At the same time, expecting this higher match quality, the firm prefers to set higher prices to extract more of the surplus, and this sometimes leads to the consumer's decision not to purchase, even though he has a positive valuation for the product. Intuitively, whether nondiscriminatory pricing leads to greater total welfare depends on whether the welfare gain from the better quality of the match exceeds the loss from monopoly price distortion.

The question posed here is equivalent to the question of whether an outward shift in a demand curve increases total surplus in the standard monopoly pricing problem. In this respect, the observation that total surplus can be either increasing or decreasing in disclosure level is consistent with the observation that an outward shift in demand can increase or decrease total surplus in the monopoly pricing problem.¹⁰

¹⁰[Cowan \(2004\)](#) provides a necessary and sufficient condition under which additive and multiplicative shifts of the demand and inverse demand function increase total surplus under monopoly. To the best of my knowledge, a more general condition has not been found in the industrial organization literature.

Remark 4. [Theorem 1](#) is robust to a variety of extensions.

Correlated Values: If K products are different versions of a similar product (such as economics textbooks with different levels of difficulty), then values might be correlated across products. In such a case, [Theorem 1](#) holds as long as the vector u of values is drawn from an exchangeable distribution whose multivariate hazard rate satisfies a condition described in Theorem 1.B.29 of [Shaked and Shanthikumar \(2007\)](#). Their condition ensures that $\max(u_1, u_2)$ dominates $\min(u_1, u_2)$ in the hazard rate order, which is sufficient to derive the theorem. I revisit correlated u_1 and u_2 in the model of competition in [Section 5](#).

Costly disclosure: The main insights continue to hold if the consumer incurs a “privacy cost” $c(\delta)$ from a disclosure level δ . While introducing the cost may change the equilibrium disclosure level, the consumer still discloses more information and is worse off under nondiscriminatory pricing. Finding an equilibrium under nondiscriminatory pricing requires a nontrivial fixed-point argument because the consumer may prefer different disclosure levels depending on the price he expects. The formal analysis is similar to the one presented in [Proposition 5](#) in [Section 5.1](#).

Informational Externality: In practice, firms may be able to learn about the preferences of some consumers from the information provided by other consumers. For instance, if two consumers share similar characteristics, the purchase history of one consumer might tell something about the preferences of the other consumer. I can capture this informational externality by using the model with a continuum of consumers. A simple way is to assume that a “true” disclosure level of consumer i is given by $\lambda\delta_i + (1 - \lambda)\bar{\delta}$ with an exogenous parameter $\lambda \in (0, 1)$, where $\bar{\delta} = \int_{i \in [0,1]} \delta_i di$ is the average disclosure level of the consumers. That is, the firm can learn about consumer i ’s preference not only from i ’s disclosure δ_i but also from other consumers’ $\bar{\delta}$. In this particular specification, a result identical to [Theorem 1](#) holds.

4 General Information Disclosure

In this section I assume that \mathcal{D} consists of all disclosure policies. Beyond theoretical curiosity, allowing the consumer to choose any disclosure policy is important for two reasons. First, it enables us to study whether consumer surplus maximization conflicts with efficiency. Note that in [Section 3](#), equilibrium is often inefficient, regardless of the pricing regime due to the standard monopoly price distortion. For instance, if the value of each product is distributed between 0 and 1, then the consumer does not purchase the recommended product whenever his value is sufficiently close to zero, as the firm always sets a positive price. However,

in the single-product setting of [Bergemann et al. \(2015\)](#), if the consumer can disclose any information about his willingness to pay and the firm can price discriminate, then inefficiency due to monopoly price distortion does not arise in equilibrium. The unrestricted \mathcal{D} is one neutral way to study whether the consumer's informational incentive leads to inefficiency in my multi-product setting.

The second reason is to study the robustness of the finding that firms prefer to commit not to price discriminate. The previous restriction on \mathcal{D} favored nondiscriminatory pricing, in that the two pricing regimes yield equal revenue for any *fixed* δ . This observation does not hold for a general disclosure policy: For a fixed disclosure policy, the firm typically achieves higher revenue under discriminatory pricing. For example, given a disclosure policy that fully discloses u , the firm can extract full surplus only by discriminatory pricing. Thus, the unrestricted \mathcal{D} enables us to study whether the firm's incentive to commit not to price discriminate persists in a more general environment.

In this section, I assume that the firm sells K products. For ease of exposition, I assume that the prior distribution x_0 of the value of each product is fully supported on a finite set $V = \{v_1, \dots, v_N\}$ with $0 < v_1 < \dots < v_N$ and $N \geq 2$. For each $x \in \Delta(V)$, $x(v)$ denotes the probability that x puts on $v \in V$. Abusing notation slightly, let $p(x)$ denote the lowest monopoly price when the value of each product is distributed according to $x \in \Delta(V)$:

$$p(x) := \min \left\{ p \in \mathbb{R}_+ : p \sum_{v \geq p} x(v) \geq p' \sum_{v \geq p'} x(v), \forall p' \in \mathbb{R} \right\}.$$

In particular, $p(x_0)$ is the price set by the firm if the consumer discloses no information. Note that $p(x)$ does not depend on K . Finally, let $X_{>v_1} \subsetneq \Delta(V)$ denote the set of the valuation distributions supported on V at which the optimal price is strictly above the lowest positive value:

$$X_{>v_1} := \{x \in \Delta(V) : p(x) > v_1\}.$$

4.1 Horizontal and Vertical Inefficiency

In the model of recommendation and pricing, there are two ways for an allocation to be inefficient. First, an allocation is inefficient if prices are so high that the consumer is sometimes excluded from the market. Second, an allocation is inefficient if the firm recommends a product other than the one he values the most. The following definition formalizes these two kinds of inefficiency.

Definition 2. An equilibrium is *vertically inefficient* if the consumer does not purchase any products with positive probability. An equilibrium is *horizontally inefficient* if the consumer purchases a product outside of $\arg \max_k u_k$ with positive probability.

I show that equilibrium is inefficient in both regimes whenever the firm sets a nontrivial monopoly price at the prior. Furthermore, I show that different pricing regimes lead to different kinds of inefficiency. First, an equilibrium under nondiscriminatory pricing is horizontally efficient: The consumer is recommended the most valuable product. However, trade may fail to occur because the firm cannot discount a price when the consumer has low values.

Proposition 2. *Under nondiscriminatory pricing, there exists a horizontally efficient equilibrium. If $x_0 \in X_{>v_1}$, this equilibrium is vertically inefficient.*

Proof. To show the first part, consider an equilibrium in which the consumer chooses (ϕ^*, M^*) defined as follows: $M^* = \mathcal{K}$ and $\phi^*(k|u) = \frac{1}{|\arg \max_{k \in \mathcal{K}} u_k|} \mathbf{1}_{\{k \in \arg \max_{k \in \mathcal{K}} u_k\}}$. Namely, ϕ^* discloses the name of the product with the highest value, breaking a tie uniformly if there is more than one such product. In equilibrium, if message k realizes, the firm recommends product k and sets an optimal price given the posterior induced by ϕ^* , which is the highest-order statistic of K samples drawn according to x_0 . Due to the uniform tie-breaking rule, the optimal price does not depend on k . Under nondiscriminatory pricing, the consumer chooses a disclosure policy taking prices as given. Because all products have the same price, it is optimal for the consumer to disclose which product he likes most so that the firm recommends product $k \in \arg \max_{k \in \mathcal{K}} u_k$. Thus, ϕ^* consists of an equilibrium.

To show the vertical inefficiency, let $x^{MAX} \in \Delta(V)$ denote the distribution of the highest-order statistic $\max_k u_k$. It follows that

$$p(x_0) \sum_{v \geq p(x_0)} x^{MAX}(v) \geq p(x_0) \sum_{v \geq p(x_0)} x_0(v) > v_1,$$

where the strict inequality follows from $x_0 \in X_{>v_1}$. Thus, the firm sets prices strictly greater than v_1 in the equilibrium, which implies vertical inefficiency. \square

Under nondiscriminatory pricing, the consumer discloses his favorite product without worrying how the firm uses the information to price discriminate. Such disclosure leads to better product match conditional on trade—but because the firm cannot tailor prices to the consumer’s willingness to pay, he is excluded from the market if he has low values for all products.

The equilibrium in [Proposition 2](#) is not necessarily unique, even up to consumer or total surplus. To see this, suppose that $V = \{L, H\}$ with $L < H$. [Table 1](#) and [2](#) show two equilibrium disclosure policies, where ϕ^S is the disclosure policy in [Proposition 2](#). Recall that $\phi(k|u_1, u_2)$ denotes the probability of sending message k conditional on values (u_1, u_2) .

	$\phi^A(1 u_1, u_2)$	$\phi^A(2 u_1, u_2)$
(H, H)	1	0
(H, L)	1	0
(L, H)	0	1
(L, L)	1	0

Table 1: Asymmetric equilibrium

	$\phi^S(1 u_1, u_2)$	$\phi^S(2 u_1, u_2)$
(H, H)	1/2	1/2
(H, L)	1	0
(L, H)	0	1
(L, L)	1/2	1/2

Table 2: Symmetric equilibrium

For any prior x_0 , there is an equilibrium in which the consumer chooses ϕ^A , the firm recommends product k given message k , and the firm sets price H for product 2. Furthermore, the equilibria associated with ϕ^A and ϕ^S can give different consumer and total surplus. For instance, if $x_0(L)$ is sufficiently large, the firm sets prices H and L for product 2 under ϕ^A and ϕ^S , respectively. In this case, both ϕ^A and ϕ^S achieve the efficient outcome but ϕ^A gives strictly lower consumer surplus. As another example, there can be x_0 such that the firm sets prices L and H for product 1 given ϕ^A and ϕ^S , respectively, while it always sets price H for product 2.¹¹ In this case, ϕ^A yields strictly greater total surplus than ϕ^S .¹²

This example illustrates the complementarity between information disclosure and pricing: If the firm sets a high price only for one product, then it could be a best response for the consumer to disclose information so that the firm can learn whether the consumer has a high value for the product or not, which in turn rationalizes the asymmetric pricing.

Does discriminatory pricing eliminate inefficiency? If $K = 1$, we can obtain the answer from [Bergemann et al. \(2015\)](#): There is a disclosure policy that achieves an efficient outcome without increasing the firm's payoff relative to no information disclosure. The consumer chooses such a disclosure policy in equilibrium, because no other disclosure policy can achieve strictly greater total surplus or strictly lower payoff of the firm.

Claim 1 ([Bergemann, Brooks, and Morris, 2015](#)). Suppose $K = 1$. Under discriminatory pricing, for any $x_0 \in \Delta(V)$, there is an efficient equilibrium. Furthermore, the firm is indifferent between the two pricing regimes whenever the efficient equilibrium is played under discriminatory pricing.

[Bergemann et al. \(2015\)](#) prove the efficiency of the equilibrium by directly constructing a disclosure policy that achieves an efficient outcome without increasing the firm's payoff.

¹¹Consider $(L, H) = (3, 4 + \varepsilon)$ and $x_0(L) = x_0(H) = 1/2$. For a small ε , the firm strictly prefers price 3 for product 1 at ϕ^A , while it strictly prefers price 4 under ϕ^S .

¹²The existence of multiple equilibria is not specific to the binary valuation in that I can construct a similar example for a general $V = \{v_1, \dots, v_N\}$ using priors that put sufficiently small weights on values other than v_1 and v_2 .

Here, I provide an alternative explanation for why there must be an efficient equilibrium if $K = 1$. This argument turns out to be useful for my multi-product setting. Suppose that an equilibrium is inefficient: At some realized message, the firm posts price p_0 , which the consumer rejects with positive probability. Under discriminatory pricing, the consumer can also disclose whether his value exceeds p_0 . This additional disclosure does not reduce his payoff, and thus the new disclosure policy continues to be an equilibrium.¹³ However, it strictly increases total surplus because if the value is below p_0 , the firm can set a lower price. Although this argument does not explain how to construct an equilibrium disclosure policy, it guarantees that there is at least one efficient equilibrium.

The next result shows that there could be *no* efficient equilibrium in the multi-product setting. In particular, horizontal inefficiency “generically” arises whenever the firm sets a nontrivial monopoly price at the prior. [Appendix G](#) contains the proof.

Proposition 3. *Under discriminatory pricing, there exists a vertically efficient equilibrium. For any $K \geq 2$, there is a measure-zero set $X_0 \subset \Delta(V)$ such that, for any $x_0 \in X_{>v_1} \setminus X_0$, any equilibrium is horizontally inefficient.*

As in the single-product case, there is a vertically efficient equilibrium because the consumer is willing to disclose whether his value of the recommended product exceeds the price whenever an equilibrium is vertically inefficient. The actual proof is subtler than $K = 1$, as providing additional information about one product could reveal information about another product, which could potentially affect the firm’s recommendation decision.

In contrast, there could be no equilibrium in which the firm can recommend the most valuable product to the consumer. Under discriminatory pricing, the consumer can affect prices by controlling what information to disclose. [Proposition 3](#) points out that this gives the consumer an incentive to obfuscate his horizontal taste whenever the monopoly price at the prior exceeds the lowest possible value.

Under discriminatory pricing, an equilibrium disclosure policy maximizes consumer surplus among all the disclosure policies. Thus, an equilibrium disclosure policy would also be chosen by a regulator or an Internet intermediary who cares about consumers and wants to release their information to firms in order to maximize consumer welfare. [Proposition 3](#) shows that such a regulator or an intermediary has to balance consumer welfare and total welfare, as the former cannot be maximized without sacrificing the latter.

While the horizontal inefficiency in [Proposition 3](#) is conceptually similar to the consumer’s choice of a low disclosure level in [Theorem 1](#), the actual proof is more complicated for three

¹³In particular, if the firm learns that the value exceeds p_0 , it continues to post price p_0 , as the firm’s new objective function is $p \cdot \frac{\mathbf{P}(u_1 \geq p)}{\mathbf{P}(u_1 \geq p_0)}$, which has the same maximizer as the original objective $p \cdot \mathbf{P}(u_1 \geq p)$.

reasons. First, a more informative disclosure policy does not necessarily lead to higher prices, and thus there is no analogue to a greater disclosure level that leads to higher prices. Second, it is challenging to directly characterize an equilibrium policy, which is a solution of a multidimensional Bayesian persuasion problem. Third, even though equilibrium under discriminatory pricing is unique up to the consumer’s payoff, there can be multiple equilibria that give different total surplus.

I prove the horizontal inefficiency in two steps. First, I solve a constrained Bayesian persuasion problem: I characterize disclosure policy ϕ^G that maximizes the consumer’s payoff among all the disclosure policies that lead to horizontally efficient outcomes, given the firm’s equilibrium behavior. Second, I modify ϕ^G in a horizontally inefficient way to strictly increase the consumer’s payoff. This completes the proof, as it shows that no horizontally efficient disclosure policy maximizes the consumer’s payoff.

To construct ϕ^G , I apply a procedure similar to the greedy algorithm in [Bergemann et al. \(2015\)](#) to each realized posterior of the disclosure policy that reveals $\arg \max_{k \in \mathcal{K}} u_k$. Then, I show that ϕ^G can be modified in a horizontally inefficient way to strictly increase consumer surplus. To modify ϕ^G in an inefficient way, I show that ϕ^G generically sends messages m^* and m_1 with the following properties. First, at m^* , the consumer obtains payoff zero, has the lowest value v_1 for all nonrecommended products, and has value $v > v_1$ for recommended product k . At m_1 , the firm is willing to recommend product $k' \neq k$ and is indifferent between any prices in V . I modify ϕ^G so that it sends message m_1 with a small probability whenever ϕ^G is supposed to send m^* . Note that this modification yields a horizontally inefficient disclosure policy. While such modification does not change the consumer’s payoff or the firm’s action, at the “new” m_1 , the firm strictly prefers to set the lowest price v_1 for product k' . Finally, I further modify the disclosure policy by pooling the new m_1 and another message m_2 at which the consumer obtains a payoff of zero but has a value greater than v_1 for product k' . In this way, the consumer can strictly increase his payoff. The following example illustrates this idea in a binary environment.

Example 1. Suppose that $K = 2$, $V = \{1, 2\}$, and $(x_0(1), x_0(2)) = (1/3, 2/3)$. Consider disclosure policy ϕ in [Table 3](#). ϕ only discloses which product the consumer weakly prefers. Because such information is necessary to achieve a horizontally efficient allocation, any horizontally efficient disclosure policy is weakly more informative than ϕ .¹⁴

I characterize ϕ^G as a disclosure policy that maximizes consumer surplus among those more informative than ϕ . To find ϕ^G , I apply a procedure similar to [Bergemann et al.’s](#)

¹⁴Precisely, I am restricting attention to “symmetric” disclosure policies, which are shown to be without loss of generality.

	$\phi(1 u_1, u_2)$	$\phi(2 u_1, u_2)$
(2, 2)	1/2	1/2
(2, 1)	1	0
(1, 2)	0	1
(1, 1)	1/2	1/2

Table 3: Disclosure policy ϕ revealing horizontal information

(2015) greedy algorithm to the posterior distribution conditional on each message $k \in \{1, 2\}$ drawn by ϕ . Note that the greedy algorithm generates a disclosure policy that maximizes consumer surplus given any prior. Thus, my procedure generates a disclosure policy that maximizes consumer surplus among the policies that induce the same valuation distribution (e.g., the distribution of $\max_k u_k$) for each recommended product as ϕ does.

Table 4 shows disclosure policy ϕ^G obtained by this procedure, which decomposes message k of ϕ into messages $k1$ and $k2$ of ϕ^G .¹⁵ For each message $kj \in \{11, 12, 21, 22\}$, the firm recommends product k at price j . Note that the firm is indifferent between prices 1 and 2 at message $k1$.

	$\phi^G(11 u_1, u_2)$	$\phi^G(12 u_1, u_2)$	$\phi^G(21 u_1, u_2)$	$\phi^G(22 u_1, u_2)$
(2, 2)	0	1/2	0	1/2
(2, 1)	1/4	3/4	0	0
(1, 2)	0	0	1/4	3/4
(1, 1)	1/2	0	1/2	0

Table 4: Efficient disclosure policy ϕ^G .

	$\phi^G(11 u_1, u_2)$	$\phi^G(12 u_1, u_2)$	$\phi^G(21 u_1, u_2)$	$\phi^G(22 u_1, u_2)$
(2, 2)	0	1/2	ε'	$1/2 - \varepsilon'$
(2, 1)	1/4	$3/4 - \varepsilon$	ε	0
(1, 2)	0	0	1/4	3/4
(1, 1)	1/2	0	1/2	0

Table 5: Horizontally inefficient disclosure policy ϕ^I .

Now, I show that the consumer can strictly increase his payoff by modifying ϕ^G to

¹⁵As I discuss in the proof, the greedy algorithm does not pin down the valuation distribution of the nonrecommended products. Table 4 is derived based on the procedure I define in the proof, which uniquely pins down the joint distribution of (u_1, \dots, u_K) .

create ϕ^I in Table 5. At $(u_1, u_2) = (2, 1)$, ϕ^I sends message 21 (corresponding to m_1) with probability ε instead of message 12 (corresponding to m^*). This does not change consumer surplus but gives the firm a strict incentive to set price 1 at message 21. At $(u_1, u_2) = (2, 2)$, ϕ^I sends message 21 with probability ε' instead of message 22 (corresponding to m_2). For a small ε' , this modification does not affect the firm's pricing and recommendation at message 21. However, this modification strictly increases consumer surplus because consumer surplus conditional on $(u_1, u_2) = (2, 2)$ increases from zero to a positive number.

4.2 The Firm (Still) Prefers Nondiscriminatory Pricing

In the model of the restricted disclosure set, Theorem 1 shows that firms prefer to commit to nondiscriminatory pricing, under which consumers choose the highest disclosure level. This argument does not apply here. For an exogenously given disclosure policy, the firm can often achieve strictly greater revenue under discriminatory pricing, because it can set prices based on realized messages. Thus, even if we proved that the consumer chooses a more informative disclosure policy under nondiscriminatory pricing, this does not necessarily imply that the firm prefers nondiscriminatory pricing.

The following result, however, shows that the firm still prefers to commit not to price discriminate, and this commitment hurts the consumer. The proof utilizes the previous inefficiency results. In contrast to Theorem 1, the result states that the firm *strictly* gains and the consumer *strictly* loses under nondiscriminatory pricing. To state the result, let R_{ND} and CS_{ND} be the payoffs for the firm and the consumer, respectively, in the horizontally efficient equilibrium under nondiscriminatory pricing in Proposition 2. Also, let R_D and CS_D denote the payoffs of the firm and the consumer, respectively, in any equilibrium under discriminatory pricing.

Theorem 2. *Suppose that the consumer can choose any disclosure policy. Then, the firm is better off and the consumer is worse off under nondiscriminatory pricing: $R_{ND} \geq R_D$ and $CS_{ND} \leq CS_D$ hold. Furthermore, for any $K \geq 2$, there is a measure-zero set $X_0 \subset \Delta(V)$ such that, for any $x_0 \in X_{>v_1} \setminus X_0$, $R_{ND} > R_D$ and $CS_{ND} < CS_D$ hold.*

Proof. Let ϕ^H denote the disclosure policy associated with the horizontally efficient equilibrium in Proposition 2. Also, let ϕ^G denote the symmetric disclosure policy that maximizes consumer surplus, subject to the constraint that it leads to a horizontally efficient outcome under discriminatory pricing. As the proof of Proposition 3 shows, ϕ^G achieves full efficiency. Because the consumer is recommended his most favored product under both ϕ^H and ϕ^G , conditional on the event that product k is recommended, the valuation distribution of product k is identical between ϕ^H and ϕ^G .

Let R_{ND} and p^* denote the firm’s expected revenue and the optimal price for each product in the equilibrium with ϕ^H . First, R_{ND} is given by $R_{ND} = p^* \sum_{v \geq p^*} x^{MAX}(v)$ with x^{MAX} being the distribution of $\max_k u_k$. Second, under ϕ^G , p^* is optimal for any posterior that can realize with positive probability. This observation follows from the fact that I construct ϕ^G using the [Bergemann et al.’s \(2015\)](#) greedy algorithm. Denoting the firm’s expected revenue from ϕ^G under discriminatory pricing by R_G , I obtain $R_{ND} = R_G$. Now, under discriminatory pricing, the equilibrium payoff of the consumer is weakly greater than the one from ϕ^G . Because ϕ^G leads to an efficient allocation, $R_{ND} \geq R_D$ must hold. As [Proposition 3](#) shows, for a generic prior $x_0 \in X_{>v_1}$, the consumer chooses an inefficient disclosure policy ϕ^I to obtain a payoff strictly greater than the payoff from ϕ^G , which implies $CS_D > CS_{ND}$. Because the allocation under ϕ^G is efficient, $R_{ND} > R_D$ holds. \square

A rough intuition is as follows. As [Proposition 3](#) shows, under discriminatory pricing, the consumer withholds information about his horizontal taste to induce the firm to set lower prices. In this model, the firm could still benefit from discriminatory pricing because it can tailor prices based on the consumer’s willingness to pay, which nondiscriminatory pricing does not allow. However, it turns out that the consumer can disclose vertical information without increasing the firm’s payoff. This implies that the entire gain created by the disclosure of vertical information accrues to the consumer, as in the single-product model. As a result, under discriminatory pricing, the firm loses due to horizontal inefficiency without benefiting from vertical efficiency.

Although consumers in reality might not be endowed with the unrestricted \mathcal{D} as their privacy choices, [Theorem 2](#) still has the following economic content. Suppose that a sophisticated third party, such as a regulator or an Internet intermediary, wants to reveal its information about consumers to firms in order to maximize consumer welfare. The result states that if the third party tries to achieve the best possible way to reveal information, it in turn gives firms an incentive to commit to nondiscriminatory pricing, which hurts consumers.

4.3 Characterizing Consumer-Optimal Disclosure Policy

While the general characterization of a “consumer-optimal” disclosure policy is beyond the scope of this paper, I provide full characterization assuming that the firm sells two products with binary valuation. Under discriminatory pricing, a consumer-optimal disclosure policy is an equilibrium disclosure policy. Under nondiscriminatory pricing, it is a disclosure policy that maximizes consumer surplus, assuming that the firm optimally sets (nondiscriminatory) prices against the disclosure policy. In the model of a continuum of consumers, the consumer-optimal policy under nondiscriminatory pricing is a collective deviation by consumers that

maximizes their joint surplus given the firm's best response.

The characterization provides two insights. First, I illustrate that disclosing correlated information is necessary to maximize consumer surplus. In particular, though valuations are independent across products, it is suboptimal to disclose information about the value of each product separately. Second, I show that the insight of [Theorem 1](#) continues to hold. That is, equilibrium disclosure under nondiscriminatory pricing could be consumer-suboptimal, in that the consumer is better off committing to withhold some information.

First, I describe the consumer-optimal disclosure policy under discriminatory pricing. Let $K = 2$ and $V = \{L, H\}$ with $0 < L < H$. To begin with, I define three distributions. First, given the prior $x_0 \in \Delta(V)$ for the value of each product, let x^{MAX} be the distribution of the highest-order statistic $\max(u_1, u_2)$ of two IID draws from x_0 . Second, let $x^{OPT} \in \Delta(V)$ satisfy $x^{OPT}(H) = \frac{L}{H}$. At x^{OPT} , the firm is indifferent between prices L and H . Thus, x^{OPT} puts the largest probability on H subject to the constraint that the firm is willing to set price L , which implies that x^{OPT} maximizes consumer surplus among $\Delta(V)$ given the firm's best response. Third, let $\bar{x} \in \Delta(V)$ satisfy $\bar{x}^{MAX} = x^{OPT}$. \bar{x} is the valuation distribution such that the highest-order statistic of two IID draws from \bar{x} coincides with x^{OPT} .

(Case 1: $x^{MAX}(H) \leq x^{OPT}(H)$) This occurs when the prior x_0 puts a sufficiently large weight on value L . Consider ϕ^* , which sends messages 1 and 2 according to [Table 6](#) with $\alpha = 1$. Note that the firm recommends product k following message k . In the present case, even at

	$\phi^*(1 u_1, u_2)$	$\phi^*(2 u_1, u_2)$
(H, H)	1/2	1/2
(H, L)	α	$1 - \alpha$
(L, H)	$1 - \alpha$	α
(L, L)	1/2	1/2

Table 6: Equilibrium disclosure for $x^{MAX}(H) \leq x^{OPT}(H)$.

$\alpha = 1$, the firm prefers to set price L at each posterior because the valuation distribution for the recommended product is given by x^{MAX} , and thus $Hx^{MAX}(H) \leq Hx^{OPT}(H) \leq L$. This is optimal for the consumer, because the allocation is efficient and the price is always L .

(Case 2: $x_0(H) \leq x^{OPT}(H) < x^{MAX}(H)$) This occurs when $x_0(L)$ is not too small but also not too large. The optimal disclosure policy is described in [Table 6](#) with $\alpha < 1$. In this case, the disclosure policy is horizontally inefficient. α is chosen so that at each message k , the valuation distribution x^k of the recommended product k is x^{OPT} . Such α exists because $x^k(H) = x_0(H) \leq x^{OPT}(H)$ at $\alpha = 1/2$, and $x^{OPT}(H) < x^{MAX}(H) = x^k(H)$ at $\alpha = 1$. This is an equilibrium, because the consumer obtains $(H - L)x^{OPT}(H)$, which is the maximum

payoff he can obtain among all of the distributions over V given the firm's best response. Note that equilibrium consumer surplus and total surplus are constant for priors such that $x_0(H) \leq x^{OPT}(H) < x^{MAX}(H)$, which is in contrast to the single-product case in which consumer and total surplus always vary across priors.

(Case 3: $x^{OPT}(H) < x_0(H)$) In this case, an equilibrium disclosure is horizontally inefficient. Indeed, the consumer reveals no information about which product is more valuable. [Appendix H](#) shows the specific form of the equilibrium disclosure policy.

It turns out that the equilibrium disclosure policy under discriminatory pricing is also consumer-optimal under nondiscriminatory pricing. Note that the disclosure policies shown above send only two messages, each of which corresponds to each product. Thus, the firm can set the price of each product without observing a particular message realization. This implies that if the consumer commits to ϕ^* , it maximizes consumer surplus even under nondiscriminatory pricing.

Comparing ϕ^* and the horizontally efficient equilibrium in [Proposition 2](#), I can conclude that the consumer-optimal disclosure policy cannot be sustained in equilibrium of nondiscriminatory pricing if the prior puts a high probability on value H , i.e., $x^{MAX}(H) > x^{OPT}(H)$. This is consistent with the idea of [Theorem 1](#), that the consumer could increase his payoff by committing to withhold some information.

5 Extensions

5.1 Competition between Firms

A natural question would be how the welfare consequences of information disclosure depend on market structure. In this section I study the model of competition assuming that, as in [Section 3](#), the consumer chooses a disclosure level from $[1/2, 1]$. I show that the consumer has a stronger incentive to provide information, because disclosure intensifies price competition.

The game is a natural extension of [Section 3](#). There are two firms, Firm A and Firm B, both of which sell products 1 and 2. The game proceeds similarly as before (see [Figure 3](#)). First, the consumer chooses a disclosure level δ . Under nondiscriminatory pricing, Firms A and B simultaneously set prices at this point. Second, Nature draws a type u and messages (m^A, m^B) . m^j denotes the message sent from the consumer to Firm $j \in \{A, B\}$. I assume that, conditional on u , m^A and m^B are stochastically independent. Third, after observing (δ, m^j) , each firm $j \in \{A, B\}$ simultaneously recommends one of two products to the consumer. Under discriminatory pricing, each firm also sets the price of the recommended

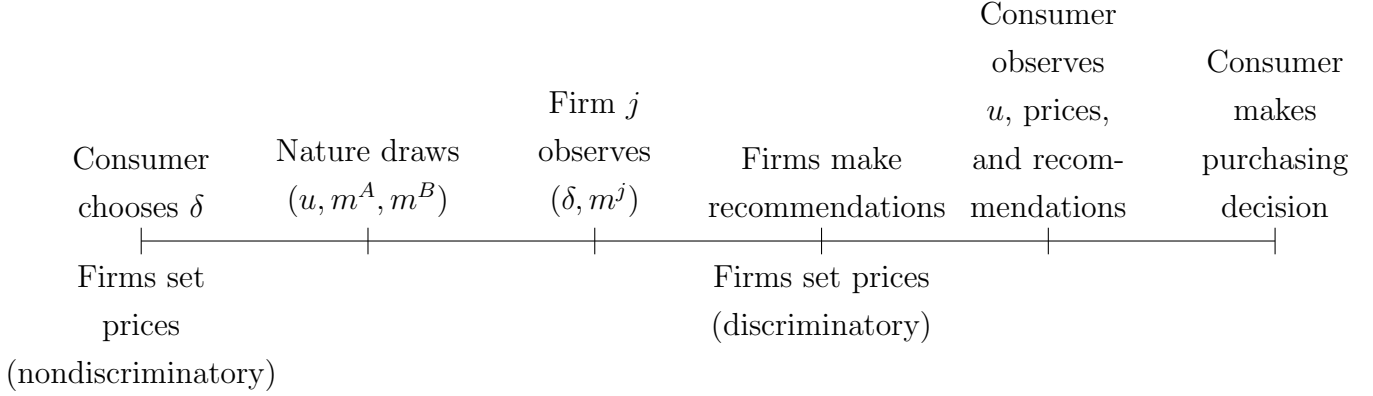


Figure 3: Timing of Moves: Competition

product. The consumer has three choices. He can choose to buy the recommended product from Firm A or B , or he can reject both.

To facilitate the analysis, I replace the distributional assumption on u_1 and u_2 . The following assumption implies that trade never occurs if the firm recommends a less preferred product. See [Remark 4](#) for how the results under a monopolist change if u_1 and u_2 are correlated.

Assumption 1. The joint distribution of (u_1, u_2) satisfies the following: The distribution F^{MIN} of $\min(u_1, u_2)$ puts probability 1 on 0, and the distribution F^{MAX} of $\max(u_1, u_2)$ has no probability mass and $p[1 - F^{MAX}(p)]$ is strictly concave for $p > 0$.

I solve the game backwards. Let $\Pi(p) := p[1 - F^{MAX}(p)]$, $p^M := p(1) = \arg \max_{p \geq 0} \Pi(p)$, and $\Pi^M := \Pi(p^M)$. Given δ , let $p^*(\delta)$ satisfy $\Pi(p^*(\delta)) = (1 - \delta)\Pi^M$.¹⁶ The following lemma describes an equilibrium price distribution given any disclosure levels. The proof is in [Appendix I](#).

Lemma 4. Define the following CDF parametrized by δ :

$$G_\delta^*(p) = \begin{cases} 0 & \text{if } p < p^*(\delta), \\ \frac{1}{\delta} - \frac{1 - \delta}{\delta} \cdot \frac{\Pi^M}{\Pi(p)} & \text{if } p \in [p^*(\delta), p^M], \\ 1 & \text{if } p > p^M. \end{cases}$$

Suppose that it is publicly known that the consumer has chosen a disclosure level δ . Under both pricing regimes, there is an equilibrium in which each firm j recommends product k whenever $m^j = k$. There is a unique equilibrium price distribution consistent with this recommendation rule: Each firm sets the price of each product according to G_δ^* .

¹⁶Such $p^*(\delta)$ exists, because $\Pi(p) = p[1 - F^{MAX}(p)]$ is concave and thus continuous.

Note that whenever $\delta > 1/2$, there is no equilibrium in which different firms specialize in recommending different products. For instance, suppose that Firm A and Firm B recommend only products 1 and 2, respectively. Then one firm, say Firm B , has an incentive to recommend product 1 at a price slightly lower than the price set by Firm A if the consumer is more likely to prefer product 1.

The next lemma states that information disclosure intensifies price competition.

Lemma 5. *Take any δ and δ' such that $\delta > \delta'$. $G_{\delta'}^*$ first-order stochastically dominates G_{δ}^* .*

Proof. By differentiating G_{δ}^* in δ , I obtain

$$\frac{\partial G_{\delta}^*(p)}{\partial \delta} = -\frac{1}{\delta^2} \left(1 - \frac{\Pi^M}{\Pi(p)} \right) \geq 0, \forall p,$$

where the inequality is strict for all $p \in (p^*(\delta), p^M)$. □

Intuitively, the more information the consumer discloses, the more likely that both firms will recommend the same product. This intensifies price competition and pushes prices downward. In the model with a continuum of consumers, the result highlights a “positive externality” associated with information sharing: Information disclosure by some consumers benefits other consumers through lower prices.

In the equilibrium of the entire game, the consumer fully discloses horizontal information and the first-best outcome is attained.

Proposition 4. *Suppose that the firms’ equilibrium strategies are given by Lemma 4. Then, under both pricing regimes, the consumer chooses $\delta = 1$ and the firms set the price of zero for each product.*

Proof. If the firms’ equilibrium strategies are given by Lemma 4, then under discriminatory pricing the highest disclosure level $\delta = 1$ gives the highest possible payoff for the consumer, because he is recommended his preferred product at price 0. Under nondiscriminatory pricing, $\delta = 1$ is a unique best response for the consumer because it uniquely maximizes the probability that he is recommended his preferred product without affecting prices. □

Under nondiscriminatory pricing, the consumer still fails to consider the impact of disclosure on prices. In this case, the impact is beneficial in the sense that disclosure not only improves the accuracy of recommendations but also lowers prices. However, because the consumer has an incentive to disclose information to obtain a better recommendation, there is no gap between equilibrium disclosure levels under nondiscriminatory and discriminatory pricing.

I show that this observation is specific to the assumption that information disclosure is costless. To see this, suppose the consumer incurs cost $c(\delta)$ from a disclosure level $\delta \in [1/2, 1]$. I assume that $c(\cdot)$ satisfies the following.

Assumption 2. $c(\cdot)$ is convex and differentiable with $c'(\frac{1}{2}) \leq 0 < c'(1)$.

$c(\cdot)$ can be nonmonotone, and if $c(\cdot)$ satisfies the assumption, $\gamma c(\cdot)$ also satisfies it for any $\gamma > 0$. The condition $c'(\frac{1}{2}) \leq 0 < c'(1)$ is rather a technical restriction, which ensures that there exists an equilibrium with an interior disclosure level. One interpretation for $c'(1) > 0$ might be that $c(\cdot)$ represents an intrinsic privacy cost, and a consumer prefers not to disclose full information about his taste given the potential cost of having too much information disclosed. Practically, the specific shape of disclosure cost could be influenced by other factors, such as a default privacy setting.

The following result is a converse of [Theorem 1](#): At any equilibrium of nondiscriminatory pricing regime, the consumer could be better off by precommitting to disclose *more* information. In contrast to the previous model, in which $\delta = 1$ is always a best response, this model requires a fixed-point argument to find an equilibrium under nondiscriminatory pricing. The proof is in [Appendix J](#).

Proposition 5. *Suppose information disclosure is costly. An equilibrium exists under each pricing regime, and the consumer is strictly worse off under nondiscriminatory pricing. Furthermore, for any equilibrium disclosure level δ^* under nondiscriminatory pricing, there is $\varepsilon > 0$ such that if the consumer chooses $\delta^* + \varepsilon$ under discriminatory pricing, he obtains a strictly greater payoff.*

The interpretation of the result is clear once we consider a continuum of consumers. Under nondiscriminatory pricing, each consumer ignores the fact that information disclosure collectively benefits consumers through intensifying price competition. Thus, in any equilibrium, if consumers collectively deviate to disclose more information, they could be better off.

Remark 5. The monopolist's model in [Section 3](#) and this competition model are special cases of the following mixed-population model. Suppose there is a unit mass of consumers. Fraction $\alpha \in (0, 1)$ of consumers are savvy and can compare recommendations from both firms as they do in this section. The remaining fraction $1 - \alpha$ of consumers are non-savvy and randomly choose one of the two firms. While price competition occurs with respect to mass α of consumers, each firm is a monopolist to mass $\frac{1-\alpha}{2}$ of consumers. Under nondiscriminatory pricing, information disclosure by savvy consumers benefits all other consumers through intensifying price competition. In contrast, information disclosure by non-savvy consumers

harms all other consumers because the firm, which is a monopolist for them, charges higher prices that are shared by everyone.

5.2 A Model of Two-Sided Private Information

I assume that the consumer does not observe the valuations of products when he chooses a disclosure policy. Focusing on the model of the restricted disclosure in [Section 3](#), I provide a microfoundation for this assumption. I consider a model in which the consumer is privately informed of his taste, the firm is privately informed of the products' characteristics, and the two pieces of information are necessary to determine the value of each product.

Formally, let $\theta \in \{1, 2\}$ denote the consumer's *taste*. $\theta = 1$ and $\theta = 2$ are equally likely and the consumer privately observes a realized θ . The firm is privately informed of the *products' characteristics* denoted by $\pi \in \{0, 1\}$. π and θ jointly determine which product is more valuable. If $\pi = 0$, the consumer with θ draws values of products θ and $-\theta$ from $\max\{u_1, u_2\}$ and $\min\{u_1, u_2\}$, respectively.¹⁷ If $\pi = 1$, the consumer with θ draws values of products $-\theta$ and θ from $\max\{u_1, u_2\}$ and $\min\{u_1, u_2\}$, respectively.

The game proceeds similarly as before. *After privately observing his taste* θ , the consumer (publicly) chooses a disclosure policy (M, ϕ) , which is now a pair of a message space and a probability distribution $\phi \in \Delta(M)$. As in the baseline model, the interpretation of a disclosure policy here is statistical information about his taste, such as browsing history. After observing (M, ϕ) , a realized $m \sim \phi$, and π , the firm recommends a product. The games under nondiscriminatory pricing and discriminatory pricing are defined in the same way as in the baseline model.

Note that being uninformed of the products' characteristics π , the consumer learns nothing about the values of products from θ . This model formalizes a natural setting in which combining consumers' subjective tastes and the products' characteristics is necessary for firms to give good product recommendations.

In this model, results identical with [Theorem 1](#) hold: Under nondiscriminatory pricing, the consumer with taste θ chooses a disclosure policy that sends message (say) θ with probability 1, because it maximizes the probability that he is recommended his preferred product. Under discriminatory pricing, the consumer chooses ϕ , which sends messages θ and $-\theta$ with probabilities δ^* and $1 - \delta^*$, respectively, where δ^* is an equilibrium disclosure level under discriminatory pricing in [Section 3](#).

¹⁷For $\theta \in \{1, 2\}$, $-\theta \in \{1, 2\} \setminus \{\theta\}$.

5.3 General Formulation of Horizontal Information

In [Section 3](#), I assume that the most informative disclosure policy available to the consumer discloses whether $u_1 > u_2$ or $u_1 < u_2$. Here, I provide a more general formulation of horizontal information under which all of the results in [Section 3](#) hold. I denote the prior CDF for the value of each product by x_0 .

Consider the following $(\bar{M}, \bar{\phi})$ with $\bar{M} = \{1, 2\}$. First, $\bar{\phi}(1|u_1, u_2) = \bar{\phi}(2|u_2, u_1)$ for any $(u_1, u_2) \in V^2$. Second, $\int_{u_2 \in V} \bar{\phi}(1|u_1, u_2) dx_0(u_2)$ is strictly increasing in $u_1 \in V$. One disclosure policy that satisfies this condition is the following: It sends message 1 with probability $h(u_1 - u_2)$, where $h(\cdot)$ is strictly increasing and $h(x) + h(-x) = 1$. (If h is a step function, I obtain the disclosure policy in [Section 3](#).) Intuitively, the higher value the consumer has for product $k \in \{1, 2\}$, the more likely message k is sent by $\bar{\phi}$.

Because of the symmetry, the posterior distribution for the value of product k conditional on message j depends only on whether $k = j$. Let F_1 and F_0 denote the posterior distributions for the value of product k conditional on messages k and $-k$, respectively. The next lemma extends [Lemma 1](#).

Lemma 6. F_1 is greater than F_0 in the hazard rate order.

Proof. Because $\int_{u_2} \bar{\phi}(1|u_1, u_2) dx_0(u_2)$ is increasing in u_1 , for any $u^+ \geq u^-$, I obtain

$$\begin{aligned} \frac{\int_{u_1 > u^+} \int_{u_2} \phi(1|u_1, u_2) dx_0(u_2) dx_0(u_1)}{1 - x_0(u^+)} &\geq \frac{\int_{u_1 > u^-} \int_{u_2} \phi(1|u_1, u_2) dx_0(u_2) dx_0(u_1)}{1 - x_0(u^-)} \\ \iff \frac{1 - F_1(u^+)}{1 - x_0(u^+)} &\geq \frac{1 - F_1(u^-)}{1 - x_0(u^-)}. \end{aligned}$$

Replacing $\phi(1|u_1, u_2)$ by $\phi(2|u_1, u_2)$, I obtain

$$\frac{1 - F_0(u^+)}{1 - x_0(u^+)} \leq \frac{1 - F_0(u^-)}{1 - x_0(u^-)}.$$

These inequalities imply

$$\frac{1 - F_1(u^+)}{1 - F_0(u^+)} \geq \frac{1 - F_1(u^-)}{1 - F_0(u^-)}$$

whenever the fractions are well-defined. Therefore, F_1 dominates F_0 in the hazard rate order. \square

Note that the derivation of [Theorem 1](#) only uses the property that the value of the consumer's preferred product dominates that of the less preferred product in the hazard rate order. Thus, I can conduct the same analysis under the assumption that consumers choose any disclosure policy less informative than $(\bar{M}, \bar{\phi})$.

5.4 Alternative Interpretation of the Model: Online Advertising Platform

I reinterpret the model of [Section 3](#) as a game between the consumer, an online advertising platform (such as Google or Facebook), and two advertisers. Advertisers 1 and 2 sell products 1 and 2, respectively. The platform auctions off the consumer’s impression using a second price auction with an optimal reserve price.

In this interpretation, first, the consumer chooses his disclosure level δ (e.g., whether to accept a cookie) and visits the platform. Each advertiser $k \in \{1, 2\}$ sets a price of product k and a bidding rule $b_k : \{1, 2\} \rightarrow \mathbb{R}$. Here, $b_k(j)$ is the bid that Advertiser k places for the impression of consumers with a realized message $j \in \{1, 2\}$.¹⁸ If Advertiser $k \in \{1, 2\}$ wins the auction for the impression, the ad of product k is shown to the consumer. The consumer sees the ad of the winning advertiser, learns the value and the price of the product, and decides whether to buy it.

The ad auction works as the product recommendation. Assume that each advertiser can set a price after observing a disclosure level. Suppose that the consumer chooses a disclosure level δ . In this subgame, the following equilibrium exists. Each advertiser sets price $p(\delta)$ and a bidding rule such that $b_x(x) = p(\delta)[1 - \delta F^{MAX}(p(\delta)) - (1 - \delta)F^{MIN}(p(\delta))]$ and $b_x(-x) < b_x(x)$. The platform sets a reserve price $p(\delta)[1 - \delta F^{MAX}(p(\delta)) - (1 - \delta)F^{MIN}(p(\delta))]$ to extract full surplus from the winning advertiser. If the advertisers and the platform adopt these strategies, the consumer is shown the ad of a product he prefers more with probability δ at price $p(\delta)$. In contrast, if each advertiser has to set a price before observing δ , then there is an equilibrium in which the consumer chooses disclosure level 1 and each advertiser sets price $p(1)$. Thus, I obtain the same result as [Theorem 1](#).

Now, if an advertiser cannot credibly commit to nondiscriminatory pricing, how can the platform prevent discriminatory pricing to promote information disclosure? One way is to adopt the following “privacy policy”: The platform commits to disclose advertisers which product the consumer is more likely to prefer, but without disclosing the exact likelihood. In other words, the platform only discloses the message realization and not δ itself. This privacy policy implements nondiscriminatory pricing: Even if Advertiser k can set a product price after learning about the consumer’s preference, it cannot set a price based on δ , which is not disclosed by the platform. This in turn implies that the consumer chooses the highest disclosure level without worrying about how the information is used by the advertisers. Expecting this, each advertiser sets a price of $p(1)$ in equilibrium. [Theorem 1](#) states that

¹⁸The analysis does not depend on whether an advertiser can place a bid contingent on δ . For simplicity, I assume that an advertiser cannot do so.

such a privacy policy benefits the platform and hurts the consumer.

6 Concluding Discussion

The question of what information consumers should disclose to firms is a first-order issue in today's economy. In this paper, I focus on the following core economic trade-off: Disclosing information could benefit consumers, because firms can provide personalized offerings such as product recommendations, online ads, and customized products. However, disclosure could hurt consumers if firms use the information to price discriminate. I capture this trade-off in a model in which a consumer discloses information about his preferences to a multi-product firm. The consumer has limited attention, in that he does not yet know his vector of values for the products and can evaluate only a small number of products relative to the huge variety of products available. As a result, the firm can use the provided information not only to extract surplus through pricing, but also to create surplus through product recommendations. This new modeling feature is important for studying the role of consumers' information in online marketplaces, in which firms often use consumers' personal data to make direct product recommendations and deliver ads to consumers, who are unable to view and assess all available products.

The paper's contributions are threefold. One is to give an economic explanation of a somewhat puzzling observation in the Internet economy: Firms seem to not use individual data to price discriminate, and consumers seem to casually share their information with online sellers. The model explains this phenomenon as firms' strategic commitment and consumers' best response. I show that this outcome robustly arises in two settings that differ in the information-disclosure technologies available to consumers.

The second contribution is to provide a framework for use in the design of privacy regulations. For instance, the model shows that nondiscriminatory pricing and the resulting full information revelation are consumer-suboptimal. Restricting the amount of information firms can possess could benefit consumers, even if consumers are rational and can decide on their own what information to disclose.

The third contribution is to expand the theory of information disclosure by consumers. The model of general disclosure reveals that even with very fine-grained control of information, consumers or a regulator cannot simultaneously achieve efficient price discrimination and efficient matching of products without sacrificing consumer welfare. I also extend the model to study the impact of competition on information disclosure. This extension demonstrates that the market structure and consumer attentiveness are two factors that affect what information consumers should disclose.

There are various interesting directions for future research. For example, the models could be extended to consider information sharing between firms or the presence of data brokers, which is likely to add new policy implications. Moreover, this paper highlights the value of the equilibrium analysis to study consumers' privacy choices in the Internet economy. It would also be fruitful to study how consumers' information disclosure collectively affects welfare in other aspects of online privacy.

References

- Acquisti, Alessandro and Hal R Varian (2005), "Conditioning prices on purchase history." *Marketing Science*, 24, 367–381.
- Anderson, Eric T and Duncan I Simester (2010), "Price stickiness and customer antagonism." *The Quarterly Journal of Economics*, 125, 729–765.
- Anderson, Simon P and André De Palma (2012), "Competition for attention in the information (overload) age." *The RAND Journal of Economics*, 43, 1–25.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015), "The limits of price discrimination." *The American Economic Review*, 105, 921–957.
- Braghieri, Luca (2017), "Targeted advertising and price discrimination online."
- Bulow, Jeremy and John Roberts (1989), "The simple economics of optimal auctions." *Journal of Political Economy*, 97, 1060–1090.
- Butters, Gerard R (1977), "Equilibrium distributions of sales and advertising prices." *The Review of Economic Studies*, 465–491.
- Conitzer, Vincent, Curtis R Taylor, and Liad Wagman (2012), "Hide and seek: Costly consumer privacy in a market with repeat purchases." *Marketing Science*, 31, 277–292.
- Cowan, Simon (2004), "Demand shifts and imperfect competition."
- De Corniere, Alexandre and Romain De Nijs (2016), "Online advertising and privacy." *The RAND Journal of Economics*, 47, 48–72.
- Fudenberg, Drew and Jean Tirole (2000), "Customer poaching and brand switching." *RAND Journal of Economics*, 634–657.

- Fudenberg, Drew and J Miguel Villas-Boas (2006), “Behavior-based price discrimination and customer recognition.” *Handbook on Economics and Information Systems*, 1, 377–436.
- Glode, Vincent, Christian C Opp, and Xingtang Zhang (2016), “Voluntary disclosure in bilateral transactions.”
- Grossman, Gene M and Carl Shapiro (1984), “Informative advertising with differentiated products.” *The Review of Economic Studies*, 51, 63–81.
- Ichihashi, Shota (2017), “Limiting sender’s information in bayesian persuasion.”
- Jordanou, Costas, Claudio Soriente, Michael Sirivianos, and Nikolaos Laoutaris (2017), “Who is fiddling with prices?”
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615.
- Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso M Valletti (2017), “The value of personal information in markets with endogenous privacy.”
- Narayanan, Arvind (2013), “Online price discrimination: conspicuous by its absence.” <http://33bits.org/2013/01/08/online-price-discrimination-conspicuous-by-its-absence>. Accessed: 2017-08-19.
- Ponomarev, Stanislav P (1987), “Submersions and preimages of sets of measure zero.” *Siberian Mathematical Journal*, 28, 153–163.
- Rayo, Luis and Ilya Segal (2010), “Optimal information disclosure.” *Journal of political Economy*, 118, 949–987.
- Roesler, Anne-Katrin (2015), “Is ignorance bliss? rational inattention and optimal pricing.”
- Roesler, Anne-Katrin and Balázs Szentes (2017), “Buyer-optimal learning and monopoly pricing.” *American Economic Review*, 107, 2072–2080.
- Shaked, Moshe and George Shanthikumar (2007), *Stochastic orders*. Springer Science & Business Media.
- Sun, Yeneng (2006), “The exact law of large numbers via fubini extension and characterization of insurable risks.” *Journal of Economic Theory*, 126, 31–69.

Taylor, Curtis R (2004), “Consumer privacy and the market for customer information.” *RAND Journal of Economics*, 631–650.

Topkis, Donald M (1978), “Minimizing a submodular function on a lattice.” *Operations Research*, 26, 305–321.

Villas-Boas, J Miguel (1999), “Dynamic competition with customer recognition.” *The Rand Journal of Economics*, 604–631.

Villas-Boas, J Miguel (2004), “Price cycles in markets with customer recognition.” *RAND Journal of Economics*, 486–501.

7 Appendix

A Proof of Lemma 2

Without loss of generality, suppose that message 1 realizes. If the firm recommends products 1 and 2 to the consumer, he draws values from distributions $\delta F^{MAX} + (1 - \delta)F^{MIN}$ and $\delta F^{MIN} + (1 - \delta)F^{MAX}$, respectively. Because $\delta > 1/2$ and F^{MAX} strictly first-order stochastically dominates F^{MIN} , $\delta F^{MAX} + (1 - \delta)F^{MIN}$ strictly first-order stochastically dominates $\delta F^{MIN} + (1 - \delta)F^{MAX}$. This implies that it is a unique best response for the firm to recommend products 1 in both pricing regime: otherwise, the firm can change recommendation and set the price optimally to strictly increase its revenue.

B Proof of Proposition 1

For each K , the consumer chooses $\delta = 1$ in the unique symmetric equilibrium under nondiscriminatory pricing because disclosure does not affect prices and increases his payoff through a better recommendation. Let F denote the valuation distribution for each product. Take any $\varepsilon > 0$. Suppose that the firm sets a nondiscriminatory price of $b - \varepsilon/2$ for each product. For a sufficiently large K , the probability $1 - F(p)^K$ that the consumer buys the recommended product goes to 1. Thus, there is \underline{K} such that the firm’s revenue is at least $b - \varepsilon$ if $K \geq \underline{K}$. This implies that the consumer’s payoff is at most ε for any such K . This completes the proof of the first part.

To see that the consumer can always guarantee some positive payoff \underline{u} under discriminatory pricing, observe that the consumer can choose to disclose no information and obtain a payoff of $\int_{p_0}^b v - p_0 dF(v)$ where $p_0 < b$ is the optimal price given no disclosure, which is independent of K .

C A Prior Distribution without a Density Function

I show the analogue of [Theorem 1](#) when x_0 is a general non-degenerate distribution. In this case, I can no longer say that the consumer chooses disclosure level 1 in the unique equilibrium. This is because the consumer may get a payoff of zero for any disclosure level. (Consider x_0 that puts probability almost 1 on $u_k = 100$ and the remaining probability on $u_k = 1$. The firm sets price 100 for any δ and thus the consumer is indifferent between any disclosure levels.)

I show that the consumer still discloses more information under nondiscriminatory pricing in the strong set order. Let $\Delta_N \subset [1/2, 1]$ denote the set of disclosure levels δ_N in symmetric equilibrium under nondiscriminatory pricing. Similarly, let $\Delta_D \subset [1/2, 1]$ denote the set of disclosure levels sustained in equilibrium under discriminatory pricing.

Theorem 3. *Δ_N is greater than Δ_D in the strong set order. Take the most informative equilibrium under each pricing regime. Then, the firm is better off and the consumer is worse off under nondiscriminatory pricing than under discriminatory pricing. Similarly, take the least informative equilibrium under each pricing regime. Then, the firm is better off and the consumer is worse off under nondiscriminatory pricing.*

Proof. Take any $(\delta_N, \delta_D) \in \Delta_N \times \Delta_D$ with $\delta_D > \delta_N$. By [Lemma 2](#), under the nondiscriminatory pricing, the payoff of the consumer is

$$\delta_N u^{MAX}(p(\delta_N)) + (1 - \delta_N) u^{MIN}(p(\delta_N)).$$

If he deviates to $\delta_D > \delta_N$, he obtains

$$\delta_D u^{MAX}(p(\delta_N)) + (1 - \delta_D) u^{MIN}(p(\delta_N)) \leq \delta_N u^{MAX}(p(\delta_N)) + (1 - \delta_N) u^{MIN}(p(\delta_N)).$$

Then, I obtain $u^{MAX}(p(\delta_N)) - u^{MIN}(p(\delta_N)) = 0$. Because $u^{MAX}(p) - u^{MIN}(p)$ is non-negative and non-increasing in p , I also obtain $u^{MAX}(p(\delta_D)) - u^{MIN}(p(\delta_D)) = 0$. Thus,

$$\delta_D u^{MAX}(p(\delta_D)) + (1 - \delta_D) u^{MIN}(p(\delta_D)) \geq \delta u^{MAX}(p(\delta_D)) + (1 - \delta) u^{MIN}(p(\delta_D))$$

for any δ . Thus, $\delta_D \in \Delta_N$. Also,

$$\begin{aligned} \delta_N u^{MAX}(p(\delta_N)) + (1 - \delta_N) u^{MIN}(p(\delta_N)) &\geq \delta_D u^{MAX}(p(\delta_N)) + (1 - \delta_D) u^{MIN}(p(\delta_N)) \\ &\geq \delta_D u^{MAX}(p(\delta_D)) + (1 - \delta_D) u^{MIN}(p(\delta_D)). \end{aligned}$$

Thus, $\delta_N \in \Delta_D$.

□

D Equivalence to the Model of a Continuum of Consumers

This subsection consists of two parts. First, I formally describe the model with a continuum of consumers with an appropriate equilibrium notion. Second, I show that for *any* set of disclosure policies \mathcal{D} , the set of equilibria with a single consumer is equal to the set of symmetric equilibria with a continuum of consumers. For ease of exposition I call the model with a continuum of consumers as model (C). Similarly, I call the original model with a single consumer as model (S).

First, I describe the model (C). There is a continuum of consumers uniformly distributed on $[0, 1]$. Each consumer is endowed with the identical preference and the strategy space as the consumer in model (S). The firm has enough units of each of K products to serve all the consumers. At the game beginning of the game, each consumer $i \in [0, 1]$ chooses a disclosure policy $\phi_i \in \mathcal{D}$. Then, the valuation vector u^i and a message $m_i \sim \phi_i(\cdot|u^i)$ realize for each i . Then, the firm observes $(\phi_i, m_i)_{i \in [0, 1]}$, sets prices, and recommends one of the K products to each consumer. Under nondiscriminatory pricing, the firm sets a single price on each product. Under discriminatory pricing, the firm can charge different prices to different consumers.

My solution concept for model (C) is subgame perfect equilibrium (hereafter, SPE) which satisfies the following four properties. First, at each information set, the firm recommends products to maximize not only her total revenue but also the expected payment of each consumer, who has measure zero. Furthermore, under discriminatory pricing, I impose the same requirement on the firm's pricing. Without this restriction, any disclosure policies are sustained in some SPE because the firm can "punish" the unilateral deviation of a consumer, who has measure zero, by offering a completely random recommendation. The second property requires that the firm chooses prices and recommendations to maximize consumers' payoffs whenever the firm is indifferent given the first requirement. This excludes the multiplicity of equilibrium due to the firm's tie-breaking rule. The third property requires that the firm recommends each of K products to a positive mass of consumers on the equilibrium path. The fourth property requires the symmetry: each consumer takes the same disclosure policy on the equilibrium path. I use equilibrium (C) to mean a SPE with these four properties.

Note that under discriminatory pricing, model (S) and (C) differ only in the number of consumers. Thus, the equivalence is trivial. The following result shows the equivalence under nondiscriminatory pricing.

Proposition 6. *Take any \mathcal{D} . Consider nondiscriminatory pricing. Take any equilibrium disclosure policy ϕ^* and product prices $p^* = (p_1^*, \dots, p_K^*)$ of model (S) in which each product*

is recommended with positive probability. Then, (ϕ^*, p^*) consists of an equilibrium (C) in model (C). Similarly, if (ϕ^*, p^*) consists of an equilibrium (C) of model (C), then it also consists of an equilibrium of model (S).

Proof. Take any (ϕ^*, p^*) consisting of an equilibrium (C) in model (C). I show that (ϕ^*, p^*) is an equilibrium under model (S). First, if everyone chooses ϕ^* in models (C) and (S), then the firm’s revenue as a function of prices and recommendation strategy is the same between model (C) and model (S) with the only difference being whether the expectation is taken based on mass of consumers or on probability. Thus, p^* and the associated recommendation strategy in model (C) continues to be optimal in model (S). Next, each consumer in model (C) is choosing a disclosure policy from \mathcal{D} taking prices as given, which is the same best response condition as in model (S). Note that the fact that each consumer is a price taker relies on the first, second, and third requirements I impose in equilibrium (C). The equivalence of the other direction proceeds in the same way. \square

E Presence of “Negative Externality” with a Continuum of Consumers

I show that in the second interpretation of the model, information disclosure by a positive mass of consumers lowers the welfare of other consumers. To see this, note that if each consumer i chooses a disclosure level δ_i and the firm sets price p for each product, then the total revenue is given by

$$\begin{aligned} & \int_{i \in [0,1]} p[1 - \delta_i F^{MAX}(p) - (1 - \delta_i) F^{MIN}(p)] di \\ & = p[1 - \bar{\delta} F^{MAX}(p) - (1 - \bar{\delta}) F^{MIN}(p)] \end{aligned}$$

where $\bar{\delta} = \int_{i \in [0,1]} \delta_i di$ is the average disclosure level. This implies that the optimal price under nondiscriminatory pricing is given by $p(\bar{\delta})$. If a positive mass of consumers disclose more information, $\bar{\delta}$ increases. This increases $p(\bar{\delta})$ and decreases the payoffs of other consumers who have not changed disclosure levels.

F Does nondiscriminatory Pricing Increase Total Welfare?: Example

Example 2. nondiscriminatory pricing can be welfare-detrimental. Consider the prior x_0 and the resulting probability mass function f_1 of $\max(u_1, u_2)$ in Table 7. If the firm commits to nondiscriminatory pricing, then in equilibrium, the consumer chooses disclosure level 1

	$v = 2$	$v = 3$
x_0	1/2	1/2
f_1	1/4	3/4

Table 7: nondiscriminatory pricing reduces welfare

	$v = 1$	$v = 2$	$v = 3$
x^*	1/4	1/4	1/2
f_1	1/16	3/16	3/4

Table 8: nondiscriminatory pricing increases welfare

and the value of the recommended product is drawn from f_1 (F_1). In this case, the monopoly price is 3, which gives a total surplus of $9/4$. If the consumer discloses no information, the value is drawn from x^* . In this case, the monopoly price is 2, which gives a total surplus of $5/2 > 9/4$. At equilibrium of the discriminatory pricing regime, the consumer chooses the highest δ at which the firm is willing to set price 2. This gives the total surplus greater than no disclosure because both the consumer and the firm obtain greater payoffs. Thus, the equilibrium total surplus is higher under discriminatory pricing in which the consumer discloses less information than under nondiscriminatory pricing.

Example 3. Conversely, nondiscriminatory pricing can be welfare-enhancing. Consider the prior x^* and the resulting probability mass function f_1 of $\max(u_1, u_2)$ in Table 8. If the firm commits to nondiscriminatory pricing, then in equilibrium, the consumer chooses disclosure level 1 and the value of the recommended product is drawn from f_1 . In this case, the monopoly price is 3, which gives total surplus of $9/4$. Under discriminatory pricing, the consumer discloses no information. Indeed, the prior x^* is such that the firm is indifferent between prices 2 and 3, which it strictly prefers to price 1. Thus, any disclosure level $\delta > 1/2$ makes the firm prefer to set price 3. Given no disclosure, the monopoly price is 2, which gives total surplus of $2 < 9/4$. Thus, the equilibrium total surplus is higher under nondiscriminatory pricing in which the consumer discloses more information. The result is not a knife-edge case in the sense that qualitatively the same result holds if the prior is stochastically smaller but close to x^* .

G Proof of Proposition 3

First, I show that there is a vertically efficient equilibrium. Take any (M^*, ϕ^*) which leads to a vertically inefficient allocation given the firm's best response. Let $x \in \Delta(V^K)$ denote a

realized posterior at which trade may not occur.¹⁹ Without loss of generality, suppose that product 1 is recommended at price v_ℓ at x . I show that there is another disclosure policy ϕ^{**} which gives a weakly greater payoff than ϕ^* to the consumer and achieves a strictly greater total surplus. Suppose that ϕ^{**} discloses whether the value for product 1 is weakly greater than v_ℓ or not whenever posterior x realizes, in addition to the information disclosed by ϕ^* . Let x^+ and $x^- \in \Delta(V^K)$ denote the posterior beliefs of the firm after the consumer discloses that the value for product 1 is weakly above and strictly below v_ℓ , respectively. Then, $x = \alpha x^+ + (1 - \alpha)x^-$ holds for some $\alpha \in (0, 1)$.

I show that ϕ^{**} weakly increases the payoff of the consumer. First, conditional on the event that the value is below v_ℓ , the consumer gets a greater payoff under ϕ^{**} than under ϕ^* because the consumer obtains a payoff of zero under ϕ^* . Second, I show that, conditional on the event that the value is weakly above v_ℓ , the firm continues to recommend product 1 at price v_ℓ . To show this, suppose to the contrary that the firm strictly prefers to recommend another product m at price v_k . If $m = 1$, v_k is different from v_ℓ . Let $x_1^+ \in \Delta(V)$ and $x_m^+ \in \Delta(V)$ be the marginal distributions of u_1^i and u_m^i given x^+ , respectively. Because the firm prefers recommending a product m at price v_k to recommending a product 1 at price v_ℓ , I obtain

$$v_k \sum_{j=k}^K x_m^+(v_j) > v_\ell \sum_{j=\ell}^K x_1^+(v_j),$$

which implies

$$v_k \sum_{j=k}^K \alpha x_m^+(v_j) + (1 - \alpha)x_m^-(v_j) \geq v_k \sum_{j=k}^K \alpha x_m^+(v_j) > v_\ell \sum_{j=\ell}^K \alpha x_1^+(v_j) = v_\ell \sum_{j=\ell}^K \alpha x_1^+(v_j) + (1 - \alpha)x_1^-(v_j).$$

The last equality follows from $x_1^-(v) = 0$ for any $v \geq v_\ell$. This contradicts that the firm prefers to recommend product 1 at price v_ℓ at x .

Consider the following mapping $\Phi : \mathcal{D} \rightarrow \mathcal{D}$: given any disclosure policy $\phi \in \mathcal{D}$, Φ chooses a posterior belief x induced by ϕ at which trade fails to occur with a positive probability. If there are more than one such belief, Φ chooses the posterior belief corresponding to the lowest price and the smallest index $k \in \mathcal{K}$.²⁰ $\Phi(\phi)$ is a disclosure policy which discloses whether the value for the recommended product is weakly greater than the price or not whenever posterior x realizes, in addition to the information disclosed by ϕ .

¹⁹Because $|V^K| < +\infty$, without loss of generality, I can assume $|M^*| < +\infty$. Then, each message realizes with a positive probability from the ex-ante perspective. Thus, there must be a posterior $x \in \Delta(V^K)$ which realizes with a positive probability and trade may fail to occur given x .

²⁰If this does not pin down a posterior uniquely, then I define Φ so that it first modifies ϕ by merging multiple beliefs at which the same product is recommended at the same price.

To show that there exists a vertically efficient equilibrium, take any equilibrium disclosure policy ϕ_0 . Define $\Phi^1(\phi_0) = \Phi(\phi_0)$ and $\Phi^{n+1}(\phi_0) = \Phi(\Phi^n(\phi_0))$ for each $n \geq 1$. Because $|V^K| < +\infty$, there exists n^* such that $\Phi^{n^*} = \Phi^{n^*+1}$. Define $\phi^* := \Phi^{n^*}(\phi_0)$. By construction, ϕ^* gives a weakly greater payoff to a consumer than ϕ_0 . Thus, it is an equilibrium under discriminatory pricing. Moreover, at each realized posterior, trade occurs with probability 1. Therefore, ϕ^* is a vertically efficient equilibrium.

Next, I show that for a generic prior, any equilibrium is horizontally inefficient whenever $x_0 \in X_{>v_1}$. While the consumer has private type u drawn from $x_0 \times \cdots \times x_0 \in \Delta(V^K)$, for the ease of exposition, I interpret the model as having the total mass one of consumers with mass $\prod_{k=1}^K x^*(u_k)$ having a valuation vector $u = (u_1, \dots, u_K) \in V^K$.

Let $\mathcal{E} \subset \mathcal{D}$ denote the set of disclosure policies which lead to an efficient allocation for some best response of the firm. Take any disclosure policy $\phi^E \in \mathcal{E}$. Under ϕ^E , if the firm prefers to recommend product k , then $k \in \arg \max_{\ell \in \mathcal{K}} u_\ell$. Thus, if both ϕ^E and $\hat{\phi}^E$ achieve an efficient allocation, they only differ in terms of which product is recommended to consumers who have the same valuation for more than one product. I show that without loss of generality, I can focus on disclosure policies that recommend each product in $\arg \max u_k$ with equal probability whenever $|\arg \max u_k| \geq 2$.

To show this, take any $(M, \phi) \in \mathcal{E}$. Let $P \subset \mathcal{K}^{\mathcal{K}}$ be the set of the permutations of $\{1, \dots, K\}$. Define ϕ^E as the following disclosure policy. First, ϕ^E publicly draws a permutation $\tau \in P$ uniformly randomly. Second, ϕ^E discloses information according to $\phi(u_{\tau(1)}, \dots, u_{\tau(K)}) \in \Delta(M)$ for each realization (u_1, \dots, u_K) . Then, from the ex-ante perspective, the consumer is recommended a product $k \in \arg \max u_j$ with probability $\frac{1}{|\arg \max u_j|}$.

I further modify ϕ^E to obtain $\phi^G \in \mathcal{E}$ which maximizes the consumer's payoff among the disclosure policies which lead to an efficient allocation. First, ϕ^G decomposes the prior x^* into K segments so that $x^* = x_1 + \cdots + x_K$. (As a disclosure policy, $\frac{1}{\sum_{n=1}^N x_k(v_n)} x_k \in \Delta(V^K)$ is a posterior belief that ϕ^G draws) Each x_k consists of $\frac{1}{|\arg \max u_j|} \cdot \prod_{j=1}^K x^*(u_j) \cdot 1_{\{k \in \arg \max u_j\}}$ mass of consumers with $u \in V^K$. Now, I apply the following procedure to each segment x_k . Without loss of generality, I explain the procedure for x_1 . I apply the ‘‘greedy algorithm’’ in [Bergemann et al. \(2015\)](#) to x_1 with respect to the value for product 1 so that I can decompose x_1 into $x_1 = \alpha_1 x_1^{S_1} + \cdots + \alpha_{N_1} x_1^{S_{N_1}}$. Here, $S_1 = V$ and $S_{n+1} \supset S_n$ for $n = 1, \dots, N_1 - 1$. Moreover, the marginal distribution of each $x_1^{S_n}$ with respect to u_1 is supported on $S_n \subset V$, and the firm is indifferent between charging any price for product 1 inside the set S_n if the value for product 1 is distributed according to $x_1^{S_n}$. In contrast to [Bergemann et al. \(2015\)](#), the consumer's type is K -dimensional. Thus, directly applying the algorithm does not pin down the valuation distribution for product $k \neq 1$ in each segment $x_1^{S_n}$. To pin down the distribution of values for product $k \neq 1$ in each $x_1^{S_n}$, I assume the following: whenever the

algorithm picks consumers from x_1 to construct $x_1^{S_n}$, it picks consumers whose value for product 2 is lower. If this does not uniquely pin down the valuation vector to pick, it picks consumers whose value for product 3 is lower, and so on. In this way, the algorithm pins down a unique segmentation.

Consumer surplus under ϕ^G is weakly greater than under ϕ^E . This is because the segmentation created by the greedy algorithm maximizes consumer surplus and that the valuation distribution of each recommended product is identical between ϕ^E and ϕ^G . Also, under ϕ^G , the firm is willing to recommend product k to consumers in $x_k^{S_n}$ because $x_k^{S_n}$ only contains consumers such that $u_k \geq u_{k'}$ for any $k' \in \mathcal{K}$.

Next, I show the following: there exists a set $D \subset \Delta(V)$ satisfying the following: D has Lebesgue measure zero in \mathbb{R}^N , and for any prior $x_0 \in \Delta(V) \setminus D$, all consumers in $x_1^{S_{N_1}}$ constructed by the last step of the algorithm have the same value for product k . The proof of this part consists of two steps.

In the first step, take any subsets of V as $S_1 \supset S_2 \supset \dots \supset S_{N_1}$ such that $|S_{N_1}| \geq 2$. Then, define

$$Y(S_1, \dots, S_{N_1}) := \left\{ y \in \mathbb{R} : y = \sum_{n=1}^{N_1} \alpha_n x_1^{S_n}, \exists (\alpha_1, \dots, \alpha_{N_1}) \in \Delta^{N_1-1} \right\}$$

where Δ^{N_1-1} is the $(N_1 - 1)$ -dimensional unit simplex. Because $|S_{N_1}| \geq 2$, $N_1 \leq N - 1$. Thus, $Y(S_1, \dots, S_{N_1})$ is a subset of at most $N - 1$ dimensional subspace, which has Lebesgue measure zero in $\Delta(V) \subset \mathbb{R}^N$. Define \mathcal{S} as

$$\mathcal{S} = \{(S_1, \dots, S_{N_1}) : \exists N_1 \in \mathbb{N}, V \supset S_1 \supset S_2 \supset \dots \supset S_{N_1}, |S_{N_1}| \geq 2\}.$$

Let \mathcal{Q} be the set of $x \in \Delta(V)$ such that consumers in $x_k^{S_n}$ constructed in the last step of the algorithm have different values for product k . I can write it as $\mathcal{Q} = \cup_{(S_1, \dots, S_{N_1}) \in \mathcal{S}} Y(S_1, \dots, S_{N_1})$. Because $|\mathcal{S}| < +\infty$ and each $Y(S_1, \dots, S_{N_1})$ has measure zero, \mathcal{Q} has Lebesgue measure zero as well.

In the second step, to show that there exists D with the desired property, consider a function φ which maps any prior $x \in \Delta(V)$ to the valuation distribution of product k conditional on the event product k is recommended under ϕ^E . Because the distribution does

not depend on k , I consider $k = 1$ without loss of generality. φ is written as follows.

$$\varphi(x) = K \cdot \begin{pmatrix} \frac{1}{K} x_1^K \\ x_2 \sum_{\ell=0}^{K-1} x_1^{K-1-\ell} x_2^\ell \cdot \frac{1}{\ell+1} \binom{K-1}{\ell} \\ x_3 \sum_{\ell=0}^{K-1} (x_1 + x_2)^{K-1-\ell} x_3^\ell \cdot \frac{1}{\ell+1} \binom{K-1}{\ell} \\ \vdots \\ x_N \sum_{\ell=0}^{K-1} (x_1 + \dots + x_{N-1})^{K-1-\ell} x_N^\ell \cdot \frac{1}{\ell+1} \binom{K-1}{\ell} \end{pmatrix}.$$

φ is infinitely differentiable and its Jacobian matrix J_φ is a triangular matrix with the diagonal elements being positive as long as $x_n > 0$ for each $n = 1, \dots, N$. Thus, $J_\varphi(x)$ has full rank if x is *not* in a measure-zero set

$$\{(x_1, \dots, x_N) \in \Delta(V) : \exists n, x_n = 0\}. \quad (3)$$

By Theorem 1 of [Ponomarev \(1987\)](#), $\varphi : \mathbb{R}^N \rightarrow \mathbb{R}^N$ has the “0-property”: the inverse image of measure-zero set by φ has measure zero. In particular, $D := \varphi^{-1}(\mathcal{Q})$ has measure zero. Thus, there exists a measure-zero set D such that for any $x \in \Delta(V) \setminus D$, all consumers in $x_k^{S_n}$ constructed in the last step of the algorithm have the same value for product k .

Consider the algorithm applied to product k . Recall that $x_k^{N_k}$ is the segment created at the last step. As I have shown, generically, all consumers in $x_k^{N_k}$ have the same value for product k . Let v^* denote the value. In equilibrium, consumers in $x_k^{N_k}$ obtain a payoff of zero given the firm’s optimal price v^* . Moreover, if the optimal price at the prior is strictly greater than v_1 (i.e., $p(x_0) \geq v_2$), then $v^* > v_1$. Indeed, if $v^* = v_1$, then $v_1 \in S_n$ for $n = 1, \dots, N_1$. This implies that v_1 is an optimal price for each $x_1^{S_n}$ and thus for $x_1 = \sum_{n=1}^{N_1} \alpha_n x_1^{S_n}$, which is a contradiction. To sum up, except for a Lebesgue measure zero set of priors, if the optimal price is strictly greater than v_1 under the prior, then consumers in $x_k^{N_k}$ obtain a payoff of zero given the firm’s optimal price strictly above v_1 .

Now, I modify ϕ^G to create a horizontally inefficient ϕ^I that yields consumer surplus strictly greater than ϕ^G , which completes the proof. To simplify the exposition, for any $S \subset \mathcal{K}$, let $v_S^* \in V^K$ denote a vector whose coordinate for each $k \in S$ is v^* and other coordinates are v_1 . First, I replace ε mass of $v_{\{2\}}^*$ in the segment $x_2^{S_1}$ for product 2 created by the first step of the algorithm (applied for product 2) by the same probability mass of $v_{\mathcal{K}}^*$ in the segment $x_2^{S_{N_2}}$. Now, this does not affect consumer surplus generated from product 2. However, I now have ε mass of $v_{\{2\}}^*$ remaining. I pool this ε mass of $v^*(2)$ with segment $x_1^{S_1}$. Let \hat{x}_1^1 denote the segment created in this way. First, under \hat{x}_1^1 , price v_1 is uniquely optimal because I add a positive mass of consumers having value v_1 to x_1^1 , and price v_1 is optimal for

x_1^1 . Second, the firm is willing to recommend product 1 for \hat{x}_1^1 , as long as ε is small. This follows from the fact that the firm strictly prefers to set price v_1 if the firm recommended product $k \neq 1$ for x_1^1 . Indeed, at $x_1^{S_1}$, the firm's optimal price is v_1 no matter which product it recommends. While the firm is indifferent between recommending any prices of product 1, consumers who have value v_1 for all products but product 1 reject any price strictly greater than v_1 . (Note that such consumers must be in the segment in $x_1^{S_1}$.) Thus, for product 2, price v_1 is uniquely optimal at x_1 .

Because the firm strictly prefers to recommend product 1 at price v_1 compared to any other choices, for some $\delta > 0$, I can bring mass δ of v_K^* from $x_1^{S_{N_1}}$ who originally receives zero payoff. Let \tilde{x}_1^1 denote the segment created in this way. As long as δ is small, at \tilde{x}_1^1 , the firm still recommends product 1 at price v_1 . This strictly increases consumer surplus because consumers who obtain zero payoff at segment x_1^1 now obtain a positive payoff at \tilde{x}_1^1 without changing surplus accruing to other consumers. However, the resulting allocation is inefficient.

Therefore, for any disclosure policy which leads to an efficient allocation, there exists a horizontally inefficient disclosure policy which gives a strictly greater consumer surplus. This completes the proof.

H Disclosure Policy for Case 3 of Section 4.3

I consider the third case: $x^{OPT}(H) < x_0(H)$. One representation of an equilibrium disclosure policy is shown in Table 9. β and γ are determined as follows. Let $x'(H)$ denote a threshold

	$\phi^*(Low u_1, u_2)$	$\phi^*(High u_1, u_2)$
(H, H)	β	$1 - \beta$
(H, L)	γ	$1 - \gamma$
(L, H)	γ	$1 - \gamma$
(L, L)	1	0

Table 9: Equilibrium disclosure for $x_H^* > x^{OPT}(H)$.

value of $x_0(H)$ such that, if $x_0(H) = x'(H)$, then the firm is indifferent between charging prices L and H given message *Low* with $(\beta, \gamma) = (0, 1)$. Note that $x^{OPT}(H) < x'(H)$. If $x_0(H) \leq x'(H)$, the optimal disclosure sets $\gamma = 1$, and $\beta < 1$ is such that the firm is indifferent between prices L and H conditional on message *Low*. If $x_0(H) > x'(H)$, then $\beta = 0$ and γ is such that the firm is indifferent between prices L and H conditional on message *Low*. This is an equilibrium because the consumer obtains the first-best payoff $(H - L)x_H^{OPT}$ with the maximum probability. The following is the formal proof.

Proof. Take any disclosure policy ϕ that maximizes consumer surplus. I show that ϕ is modified to be the form of Table 9 without decreasing consumer surplus. By Proposition 1 of Kamenica and Gentzkow (2011), without loss of generality, I assume that ϕ sends at most four messages $1L$, $1H$, $2L$, and $2H$, where a firm recommends product $k \in \{1, 2\}$ at price $j \in \{L, H\}$ conditional on message kj . Let x_{kj} be the segment that corresponds to kj . Namely, x_{kj} specifies the ex-ante probability mass of each type (u_1, u_2) which sends message kj . Thus, $x_{1L} + x_{1H} + x_{2L} + x_{2H} = x^*$. Now, without loss of generality, I can assume that x_{kH} puts no weight on $(u_1, u_2) = (L, L)$ and that x_{1L} and x_{2L} contain equal mass of (L, L) . Next, the firm must be indifferent between prices L and H at any x_{kL} . Indeed, if the firm strictly prefers price L , say at x_{1L} , then I can take a positive mass of any of (H, H) and (H, L) from $x_{1H} + x_{2H}$ and pool it to x_{1L} to strictly increase consumer surplus. Note that $x_{1H} + x_{2H}$ must contain (H, H) or (H, L) : if they do not, which means $x_{1L} + x_{2L}$ contains all the mass of (H, H) and (H, L) , it contradicts that $x^{OPT}(H) < x_0(H)$ and that the firm is willing to set price L at the original x_{1L} and x_{2L} .

Furthermore, x_{1L} contains a weakly greater mass of (H, L) than that of (L, H) . If any of x_{1H} and x_{2H} contains a positive mass of (L, H) , then I can put this (L, H) into x_{1L} . I consider three cases:

(Case I) Suppose that after putting all the mass of (L, H) in $x_{1H} + x_{2H}$ into x_{1L} , the new x_{1L} still contains a strictly greater mass of (H, L) than that of (L, H) and the firm is willing to set price L at the new x_{1L} . This implies that I can also put the remaining mass of (H, L) and (L, H) into x_{2L} . At the resulting x_{1L} and x_{2L} , segment x_{2L} contains a strictly greater mass of (L, H) than that of (H, L) because the total probability masses of (H, L) and (L, H) are equal. I show that this leads to a contradiction. To show this, I can exchange mass $\varepsilon/2$ of (H, L) in x_{1L} and mass $\varepsilon/2$ of (L, H) in x_{2L} , where ε is the difference between masses of (H, L) and (L, H) in x_{1L} . Then, I put a probability mass slightly greater than $\varepsilon/2$ of (H, H) from $x_{1H} + x_{2H}$ to each of x_{1L} and x_{2L} . Note that $x_{1H} + x_{2H}$ must be containing a mass strictly greater than ε of (H, H) . Indeed, if the total mass of (H, H) in $x_{1H} + x_{2H}$ is $\varepsilon' \leq \varepsilon$, then this implies that the firm is willing to set price L at x_{1L} and x_{2L} after we take all mass ε' of (H, H) from $x_{1H} + x_{2H}$ to each of x_{1L} and x_{2L} . However, this means that the firm prefers price H at x_{1L} because it prefers price H at the prior. This is a contradiction. This exchange strictly increases consumer surplus, which is a contradiction.

(Case II) Suppose that before putting all the mass of (L, H) in $x_{1H} + x_{2H}$ into x_{1L} , the new x_{1L} contains an equal mass of (H, L) and (L, H) . Then, I apply the same procedure to x_{2L} . After putting all the mass of (L, H) in $x_{1H} + x_{2H}$ into x_{2L} , if the new x_{2L} still contains a strictly greater mass of (H, L) than that of (L, H) , then this is a contradiction, because this implies that the total mass of (L, H) is strictly greater than that of (H, L) .

(Case III) Suppose before putting all the mass of (L, H) in $x_{1H} + x_{2H}$ into x_{1L} , the new x_{1L} contains an equal mass of (H, L) and (L, H) . Then, I apply the same procedure to x_{2L} . Suppose that before all the mass of (L, H) in $x_{1H} + x_{2H}$ into x_{2L} , the new x_{2L} also contains an equal mass of (H, L) and (L, H) . Because a consumers and the firm are now indifferent between two products at x_{1L} and x_{2L} , I can combine them to create a new segment *Low*.

If $x_0(H) \leq x'(H)$, then I can take all the remaining masses of (H, L) and (L, H) from $x_{1H} + x_{2H}$ to *Low*. If the firm strictly prefers price L after this procedure, I take a positive mass of (H, H) from $x_{1H} + x_{2H}$ to make the firm indifferent between prices L and H . This corresponds to [Table 9](#) with $\gamma = 1$.

In contrast, if $x_0(H) > x'(H)$, then, I take equal masses of (H, L) and (L, H) from $x_{1H} + x_{2H}$ to *Low* up to the point that the firm is indifferent between prices L and H . This corresponds to [Table 9](#) with $\gamma < 1$ and $\beta = 0$. \square

I Proof of [Lemma 4](#)

Proof. First, I show that there is an equilibrium in which each firm recommends product k conditional on message k . Because an identical argument applies to both pricing regimes, I consider discriminatory pricing. Take any equilibrium and consumer $i \in I$ who has chosen a disclosure level δ . Let CDF G denote the equilibrium price distribution of Firm B . Suppose that Firm B recommends product k with probability β conditional on message k . I calculate the payoff of Firm A conditional on the event that the consumer prefers a product 1.

Suppose that Firm A 's recommends product k with probability α_k conditional on message k . Note that conditional on the consumer's preferring product 1, the probability that Firm B recommends each product does not depend on (α_1, α_2) , because Firms A and B receive stochastically independent messages. With probability $b := \delta\beta + (1 - \delta)(1 - \beta)$, Firm B recommends a product 1. In this case, the consumer buys product 1 from Firm A as long as his value is above p and Firm B sets a price above p (which occurs with probability $[1 - G(p)]$), and Firm A recommends product 1. With the remaining probability $1 - b$, Firm B recommends product 2. In this case, the consumer buys product 1 from Firm A as long as his value is above p and Firm A recommends product 1, which happens with probability $\delta\alpha_1 + (1 - \delta)(1 - \alpha_2)$.

To sum up, Firm A 's payoff is

$$[\delta\alpha_1 + (1 - \delta)(1 - \alpha_2)] \cdot \{(1 - b)\Pi(p) + b[1 - G(p)]\Pi(p)\}.$$

Applying the same calculation for the event that the consumer prefers product 2, I obtain the payoff

$$[\delta\alpha_2 + (1 - \delta)(1 - \alpha_1)] \cdot \{(1 - b)\Pi(p) + b[1 - G(p)]\Pi(p)\}.$$

Thus, the ex-ante expected payoff is

$$\begin{aligned}\Pi_A &= \frac{1}{2} [\delta\alpha_1 + (1 - \delta)(1 - \alpha_2)] \cdot \{(1 - b)\Pi(p) + b[1 - G(p)]\Pi(p)\} \\ &\quad + \frac{1}{2} [\delta\alpha_2 + (1 - \delta)(1 - \alpha_1)] \cdot \{(1 - b)\Pi(p) + b[1 - G(p)]\Pi(p)\}.\end{aligned}$$

This is maximized at $\alpha_1 = \alpha_2 = 1$.

Next, I show that (G_δ^*, G_δ^*) is an equilibrium given the above recommendation strategy. First, I show that (G_δ^*, G_δ^*) is an equilibrium under both pricing regime. First, at price p^M , consumer i purchase from a firm only if the other firm recommends the products with value zero. Thus, a firm obtains expected payoff of $(1 - \delta)\Pi^M$ from price p^M conditional on the event that the firm recommends the product that the consumer values. The indifference condition of Firm j requires that any price $p \in [p^*(\delta), p^M]$ gives a payoff of $(1 - \delta)\Pi^M$.

The indifference condition is as follows:

$$(1 - \delta)\Pi(p) + \delta[1 - G_\delta^*(p)]\Pi(p) = (1 - \delta)\Pi^M, \quad \forall p \in [p^*(\delta), p^M].$$

The direct substitution shows that G_δ^* satisfies these equations. Finally, any price $p > p^M$ cannot be a profitable deviation because it gives a firm a payoff of $(1 - \delta)\Pi(p) \leq (1 - \delta)\Pi^M$. Any price $p < p^*(\delta)$ cannot be a profitable deviation either, because a firm obtains a revenue strictly lower than $(1 - \delta)\Pi^M$. Thus, (G_δ^*, G_δ^*) consists of an equilibrium.

Next, I show the uniqueness of the equilibrium. Take any equilibrium price distributions (G_A, G_B) . First, G_A and G_B^* do not have a positive mass at the same price: otherwise, one firm can profitably undercut the price of the other. Second, let Γ_j denote the support of G_j . Define $\bar{\Gamma}_j = \sup \Gamma_j$ for each $j \in \{1, 2\}$. Then, $\bar{\Gamma}_1 = \bar{\Gamma}_2 = p^M$ must hold. Suppose to the contrary that at least one of $\bar{\Gamma}_1 < p^M$ and $\bar{\Gamma}_2 < p^M$ holds. First, suppose that $\bar{\Gamma}_j > \bar{\Gamma}_{-j}$. Then, Firm j uses p^M instead of prices in $(\bar{\Gamma}_j, \bar{\Gamma}_{-j})$ in the equilibrium. Then, Firm $-j$ prefers to set prices arbitrarily close to p^M instead of those around $\bar{\Gamma}_{-j}$. This contradicts $\bar{\Gamma}_j > \bar{\Gamma}_{-j}$. Second, suppose that $\bar{\Gamma}_j = \bar{\Gamma}_{-j} < p^M$. Then, if one firm, say j , does not have a mass at $\bar{\Gamma}_j$, then the other firm prefers to deviate and set p^M , which contradicts $\bar{\Gamma}_{-j} < p^M$.

Third, $\underline{\Gamma}_1 := \inf \Gamma_1 = \inf \Gamma_2 =: \underline{\Gamma}_2$ hold. Otherwise, firm j with $\underline{\Gamma}_j < \underline{\Gamma}_{-j}$ has no incentive to charge below $\underline{\Gamma}_{-j}$. Now, $\underline{\Gamma}_1 \geq (1 - \delta)\Pi^M$ has to hold, because Firm B can always guarantee at least $(1 - \delta)\Pi^M$ whenever it recommends correct products at price p^M . Thus, $\underline{\Gamma}_2 \geq (1 - \delta)\Pi^M$. This implies that Firm A can get a profit arbitrarily close to $(1 - \delta)\Pi^M$, which implies $\underline{\Gamma}_1 \geq (1 - \delta)\Pi^M$.

Fourth, I show $\underline{\Gamma}_1 = \underline{\Gamma}_2 = (1 - \delta)\Pi^M$ for each j . Suppose to the contrary that $\underline{\Gamma}_j > (1 - \delta)\Pi^M$ for $j = 1, 2$. Then, Firm B has to win by setting p^M when Firm A recommends correct products at the same price. Otherwise, Firm B obtains a strictly lower payoff from

$\underline{\Gamma}_2$ than from p^M , which contradicts that p^M is in the support of the equilibrium strategy. For Firm B to win with positive probability at price p^M , G_A must have a mass at p^M . If G_A has a mass at p^M , Firm B undercuts p^M unless consumers accept the recommendations of Firm B with probability 1 when both firms recommend correct products at p^M . However, if Firm B wins for sure whenever both firms recommend correct products at price p^M , then Firm A obtains a payoff strictly lower than $(1 - \delta)\Pi^M$, which is a contradiction.

Next, I show that the supports of Γ_1 and Γ_2 are convex. Suppose to the contrary that Γ_j has a hole $(a, b) \subset (\underline{\Gamma}_j, \bar{\Gamma}_j)$. Then, Firm $-j$ does not set prices in (a, b) either. Define $p'_j := \inf \{p \in \Gamma_j : p \leq a\}$. Firm j with $p'_j > p'_{-j}$ has an incentive to set a price $b - \varepsilon$ instead of setting prices around p'_j . If $p'_j = p'_{-j}$, then a firm who has no mass at this price has an incentive to deviate in the same way.

Finally, I show that neither G_A nor G_B has a probability mass. Indeed, if G_j has a mass at x in Γ_j , then Firm $-j$ prefers to set a price $x - \varepsilon$ instead of prices in $(x, x + \varepsilon)$ for some $\varepsilon > 0$. This contradicts that Γ_1 and Γ_2 are convex.

To sum up, G_A and G_B are supported on $[p^*(\delta), p^M]$, and neither G_A nor G_B^* has a probability mass. The indifference conditions uniquely determine the price distributions. \square

J Proof of Proposition 5

Proof. Take any equilibrium under nondiscriminatory pricing. Suppose that the consumer chooses a disclosure level δ . Then, each firm sets a price according to G_δ^* in Lemma 4. Let $p_A(\delta)$ and $p_B(\delta)$ be independent random variables drawn from CDF G_δ^* .

If the consumer unilaterally deviates to δ_i , the expected payoff of the consumer $X(\delta_i, \delta)$ is as follows.

$$X(\delta_i, \delta) := \mathbb{E} [\delta_i^2 (u(1, \min \{p_A(\delta), p_B(\delta)\})) + 2\delta_i(1 - \delta_i)u(1, p_A(\delta))] - c(\delta_i). \quad (4)$$

Here, the expectation is taken with respect to the value of a recommended product and the realizations of prices $(p_A(\delta), p_B(\delta))$. The first and the second terms correspond to the events that the consumer is recommended a product he prefers from both firms and only one firm, respectively.

I show that $X(\delta_i, \delta)$ is strictly concave in δ_i . Define $X_1 := \frac{\partial X}{\partial \delta_i}$. It holds that

$$\frac{X_1}{2} = \mathbb{E} [\delta_i (u(1, \min \{p_A(\delta), p_B(\delta)\})) + (1 - 2\delta_i)u(1, p_A(\delta))] - c'(\delta_i). \quad (5)$$

X_1 is decreasing in δ_i . In particular, $\mathbb{E} [\delta_i (u(1, \min \{p_A(\delta), p_B(\delta)\})) + (1 - 2\delta_i)u(1, p_A(\delta))]$ is decreasing in δ_i because differentiating it in δ_i gives

$$\mathbb{E} [(u(1, \min \{p_1(\delta), p_B(\delta)\})) - 2u(1, p_A(\delta))] \leq 0,$$

where the inequality follows from

$$\begin{aligned} & \mathbb{E} [(u(1, \min \{p_A(\delta), p_B(\delta)\})) - 2u(1, p_A(\delta))] \\ &= \mathbb{E} [\max \{v - \min \{p_A(\delta), p_B(\delta)\}, 0\} - \max \{v - p_A(\delta), 0\} - \max \{v - p_B(\delta), 0\}] \leq 0. \end{aligned}$$

For each δ , define the correspondence $BR(\delta) := \arg \max_{\delta_i \in [1/2, 1]} X(\delta_i, \delta) - c(\delta_i)$. δ^* consists of a symmetric equilibrium if and only if $\delta^* \in BR(\delta^*)$. The compactness of $[1/2, 1]$ and the continuity of $X(\delta_i, \delta)$ in (δ_i, δ) guarantees that $BR_i(\cdot)$ is non-empty valued and has a closed graph. Because the consumer's payoff is concave in δ_i , $BR(\cdot)$ is convex-valued. By Kakutani's fixed point theorem, there exists δ^* such that $\delta^* \in BR(\delta^*)$. Now, $1 \notin BR_i(1)$, because $X_1(1, 1) - c'(1) = -c'(1) < 0$. Thus, there is no equilibrium in which the consumer choose disclosure level 1. Also, $1/2 \notin BR_i(1/2)$, because $X_1(1/2, 1/2) - c'(1/2) \geq X_1(1/2, 1/2) > 0$.

Take any equilibrium disclosure level δ^* under nondiscriminatory pricing. By the previous argument, $1/2 < \delta^* < 1$. Because $X_1(\delta^*, \delta^*) = c'(\delta^*)$, $X_1(\delta^*, \delta^*) + X_2(\delta^*, \delta^*) > c'(\delta^*)$. Thus, there is $\varepsilon > 0$ such that, if the consumer pre-commit to $\delta^* + \varepsilon$, he obtains a strictly greater payoff. This is equivalent to saying that the consumer obtains a strictly greater payoff if he was in discriminatory pricing and chooses $\delta^* + \varepsilon$. \square