

Ichihashi, Shota

Working Paper

Online privacy and information disclosure by consumers

Bank of Canada Staff Working Paper, No. 2019-22

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Ichihashi, Shota (2019) : Online privacy and information disclosure by consumers, Bank of Canada Staff Working Paper, No. 2019-22, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2019-22>

This Version is available at:

<https://hdl.handle.net/10419/210763>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Working Paper/Document de travail du personnel 2019-22

Online Privacy and Information Disclosure by Consumers



by Shota Ichihashi

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2019-22

June 2019

Online Privacy and Information Disclosure by Consumers

by

Shota Ichihashi

Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
sichihashi@bankofcanada.ca

Acknowledgements

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, Jason Allen, Susan Athey, Luca Braghieri, Isaias N. Chaves, Weixin Chen, Piotr Dworczak, Takako Fujiwara-Greve, Pedro Gardete, Michihiro Kandori, Fuhito Kojima, Shengwu Li, Sitian Liu, Qihong Liu, Giorgio Martini, Guido Martirena, Shunya Noda, Takuo Sugaya, Ben Tomlin, Erhao Xie, and participants in various seminars and conferences. I am grateful for the financial support from the Yoshida Scholarship Foundation and the Stanford Graduate Fellowship. The opinions expressed in this article are the author's own and do not reflect the views of the Bank of Canada. The author declares that he has no relevant material or financial interests that relate to the research described in this paper.

Abstract

I study the welfare and price implications of consumer privacy. A consumer discloses information to a multi-product seller, which learns about the consumer's preferences, sets prices, and makes product recommendations. While the consumer benefits from accurate product recommendations, the seller may use the information to price discriminate. I show that the seller prefers to commit to not using consumer information for pricing to encourage information disclosure. However, this commitment hurts the consumer, who could be better off by pre-committing to withhold some information. In contrast to single-product models, total surplus may be lower if the seller can base prices on information.

Bank topics: Economic models

JEL codes: D82, D83

Résumé

J'étudie les incidences sur le bien-être économique et sur les prix de l'usage des renseignements sur les consommateurs. Un consommateur divulgue des renseignements à un vendeur multiproduit, qui découvre ainsi ses préférences, établit le prix de ses produits et en recommande certains. Alors que le consommateur bénéficie de recommandations pertinentes, le vendeur peut exploiter ses renseignements pour exercer une discrimination de prix. Je démontre que le vendeur préfère s'engager à ne pas utiliser les renseignements sur le consommateur pour fixer ses prix afin d'encourager ce dernier à les divulguer. Cependant, cet engagement pénalise le consommateur, pour lequel il pourrait être plus avantageux de s'engager, au préalable, à ne pas divulguer certains renseignements. Contrairement aux résultats obtenus au moyen d'un modèle à produit unique, le surplus total pourrait être moins élevé si le vendeur utilisait ses renseignements pour fixer ses prix.

Sujets : Modèles économiques

Codes JEL : D82, D83

Non-technical Summary

E-commerce firms such as Amazon, eBay, and Netflix collect consumer information like purchase histories and personal characteristics. On the one hand, this information enables sellers to offer product recommendations and personalized services, which benefit consumers. On the other hand, sellers could potentially use such information for price discrimination. Given the benefit of personalization and the potential cost of price discrimination, how much information should consumers provide to sellers? Given consumers' incentive, how should a seller use their information?

To answer these questions, I study a simple game-theoretic model where a consumer discloses information to a multi-product seller, which learns about the consumer's preferences, sets prices, and makes product recommendations. For example, the consumer might choose to reveal information by sharing browsing activities or try to obfuscate information by disabling cookies.

I obtain two main findings. First, the seller is better off by committing to not use consumer information to set prices. This commitment encourages the consumer to disclose information that is useful for providing accurate recommendations. I show that under a mild condition, the seller's gain from accurate recommendations can exceed the loss from not being able to price discriminate. The result contrasts with the classical theory of third-degree price discrimination.

Second, the seller's commitment makes the consumer worse off. When the seller can price discriminate, the consumer withholds some information, taking into account how information affects prices. In contrast, under nondiscriminatory pricing, the consumer discloses a lot of information without worrying about price discrimination. However, expecting this greater level of disclosure and resulting accurate recommendations, the seller prefers to commit to a high price for each product up front, which lowers consumer welfare. I show that the consumer's loss from high prices can exceed the gain from better recommendations.

The result has three implications. First, it gives sellers a new rationale for committing to not price discriminate. Second, the result gives a potential explanation for why consumers do not seem to actively protect their privacy from online sellers in reality. The model justifies such behavior as a best response to sellers' commitment to not price discriminate. Third, the result suggests that such a situation could be undesirable for consumers, in the sense that consumers could be better off if they could pre-commit to withhold information.

1 Introduction

I study the welfare and price implications of consumers' privacy in online marketplaces. Online sellers can observe detailed information about consumers, such as their browsing histories, purchases, and characteristics; however, consumers can often affect whether and to what extent this information is revealed. For example, they can disable cookies to hide their web-browsing activities, or they can use their social networking accounts to log in to online shopping websites. For policymakers, information revelation by consumers is an important consideration in formulating policies concerning online privacy.

The paper considers the following economic trade-off: The benefit for consumers of disclosing information is that sellers can recommend or advertise products that are directly relevant. The cost is that sellers may use this information to price discriminate. For example, Amazon, Netflix, Spotify, and other e-commerce sellers use consumers' personal data to offer product recommendations, which help consumers discover items that they might not have found otherwise. However, these sellers could also use such information to obtain estimates of consumers' willingness to pay and, in turn, set prices on this basis.

The model consists of a monopolistic seller of K products and a consumer who has a unit demand and independently and identically distributed (IID) values for the products. At the beginning of the game, the consumer chooses a disclosure rule, which determines the information that the seller learns about the consumer's values. After learning about the values, the seller recommends one of K products. Finally, the consumer observes the value of the recommended product and decides whether to buy it.

In line with Bayesian persuasion ([Kamenica and Gentzkow, 2011](#)), I assume that the consumer discloses information without observing his values. In online marketplaces, it is often difficult for consumers to determine which item in the set of available products is most appropriate for them; however, sellers can often do this using personal data. For example, sellers might analyze browsing histories by using their knowledge of the products' characteristics, the prior experiences of other consumers, and their computing power. Sellers can then map a given consumer's data into estimates of his values for products. In such a case, even though the consumer himself cannot evaluate all products, his privacy choice affects what sellers can learn about the values.

A novel aspect of the paper is that I consider two settings that differ in the timing at which the seller sets prices. Under the *no-commitment* regime, the seller sets prices after learning about the consumer's values. Under the *commitment* regime, the seller sets prices up front without observing the consumer's information disclosure. The commitment regime captures the seller's *commitment to not use consumer information for pricing*.

Note that, in the classical theory of (third-degree) price discrimination, it is profitable for a seller to set prices on the basis of information about consumers' willingness to pay. Also, consumers are worse off if the seller can tailor prices to fine-grained information about consumers' values.

My model generates different predictions. The first main finding is that *the seller is better off by committing to not use the consumer's information for pricing*. The key is the consumer's endogenous disclosure: By making such a commitment, the seller can induce the consumer to disclose more information, which enables the seller to make more accurate recommendations. This shifts up the consumer's demand for recommended products and increases revenue. The commitment regime has an obvious downside that the seller cannot base prices on information. I provide conditions on the set of available disclosure rules under which the seller's benefit from accurate recommendations dominates the potential loss. For example, the result holds if the consumer can disclose *any* information about the value of each product. The result gives a potential explanation of an observed 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 tailor prices.¹

The second main finding is that *the consumer is worse off if the seller commits to not use information for pricing*. A key observation is that, under the no-commitment regime, the consumer can induce the seller to set lower prices by withholding information about which product is most valuable to the consumer. Although this leads to less accurate product recommendations, the consumer's gain from low prices can exceed the loss from potential product mismatch. In contrast, if the seller commits to prices in advance, the consumer misses the opportunity to influence prices by strategically concealing information. Thus, the commitment makes the consumer worse off.

¹There have been several attempts by researchers to detect price discrimination by e-commerce websites. For instance, Iordanou et al. (2017) examine around 2,000 e-commerce websites and "conclude that the specific e-retailers do not perform PDI-PD" (personal-data-induced price discrimination).

Finally, I show that equilibrium is inefficient under a mild distributional assumption, regardless of the pricing regime. The result contrasts with the single-product case of [Bergemann et al. \(2015\)](#), in which equilibrium is fully efficient under the no-commitment regime. In my model, the seller’s lack of commitment discourages information disclosure, and this leads to inefficiency due to inaccurate recommendations.

The main insights are also applicable to offline transactions. For example, consider a consumer looking for a car. The consumer may talk to a salesperson and reveal some information—such as lifestyle and preferences for fuel efficiency versus horsepower—which is indicative of his tastes; even his clothes may reveal his preferences. Based on this information and her knowledge about available cars, the salesperson gives recommendations. On the one hand, the consumer benefits from the recommendations because he can avoid extra research and test-driving. On the other hand, disclosing too much information may put him in a disadvantageous position in price negotiation, because knowing that he loves a particular car, the salesperson would be unwilling to compromise on prices. My result says that car dealers should eliminate this trade-off by committing to prices up front. This encourages consumers to disclose information; at the same time, dealers can set relatively high prices to extract surplus created by better product matches. In practice, the commitment appears as no-haggle policies adopted by some firms such as Fiat, Tesla, and Toyota ([Zeng et al., 2013, 2016](#)).

The remainder of the paper is organized as follows. After discussing related work in the next subsection, in [Section 2](#), I present the baseline model. In [Section 3](#), I restrict the consumer to choosing from a simple class of disclosure rules. This section shows that the seller is better off and the consumer is worse off under the commitment regime. [Section 4](#) allows the consumer to choose any disclosure rule. I show that equilibrium is typically inefficient and then use the inefficiency result to establish the welfare comparisons. This section also shows that the commitment can increase total surplus. [Section 5](#) discusses extensions and [Section 6](#) concludes.

1.1 Related Work

This paper relates to three strands of literature: The literature on monopoly price discrimination, the economics of privacy, and information design. Starting with the work of [Pigou \(1920\)](#), the lit-

erature on third-degree price discrimination examines what happens to consumer surplus, producer surplus, and total surplus as a market is segmented.² Relative to this literature, my model differs in two ways. First, market segmentation is endogenous. This is especially relevant when a seller uses consumer information to segment a market but consumers can influence how much information is available to the seller. Second, the seller can use consumer information to recommend products. In this case, finer segmentation enables the seller to recommend a more relevant product to consumers in each segment, and this can shift up the demands for all products. These two points are crucial to why a seller can be better off if she cannot tailor prices to segments.

The economic forces of this paper are reminiscent of those in the durable goods monopoly problem, such as [Stokey \(1979, 1981\)](#). In that setting, a seller prefers to commit to not price discriminate based on the timing of purchase. This is because if the seller engages in such discrimination, consumers would delay the timing of their purchase. In my model, the seller prefers to commit to not lower prices even if the consumer discloses information that makes it (ex post) optimal for the seller to do so. That way, the consumer discloses as much information as possible about which product is most valuable. Despite this high-level connection, I argue that this paper provides novel insights. First, the effect of sellers' commitment power in how consumers reveal information is fundamental and economically important. This economic force relates not only to the recent discussion on consumer privacy but also to broader settings such as markets for cars and houses as well as bargaining between workers and employers. Second, the result that the seller's commitment can increase total welfare is unique to my model: In the durable goods monopoly problem, the seller's inability to commit to prices leads to efficient outcomes if the length of the time period is short.

In terms of modeling, one related work is [Bergemann et al. \(2015\)](#), who consider a single-product monopoly pricing in which a seller has additional information about a consumer's value. I consider a multi-product seller with product recommendations, which renders information useful not only for pricing but also for improving product match quality. These components lead to different welfare consequences of information disclosure and price discrimination.

In the economics of privacy literature, recent papers have devoted considerable attention to the relationship between personal data and intertemporal price discrimination ([Acquisti and Varian](#),

²For more recent papers, see, for example, [Aguirre et al. \(2010\)](#), [Chen and Schwartz \(2015\)](#), and [Cowan \(2016\)](#).

2005; Conitzer et al., 2012; Fudenberg and Tirole, 2000; Fudenberg and Villas-Boas, 2006; Taylor, 2004; Villas-Boas, 1999, 2004). In these models, sellers learn about consumers' preferences from their purchase record, which arises endogenously as a history of a game. In my model, the consumer is endowed with his personal data at the outset.

Hidir and Vellodi (2018) consider a model in which a buyer communicates his preferences to the seller, who uses the information to tailor product offerings and prices. Their focus and formulation differ from mine in at least two ways. First, their main focus is to understand the consumer's trade-off between better product match and better prices. In contrast, on top of this trade-off, I ask what could lead sellers to commit to not use consumer data for price discrimination. Thus, Hidir and Vellodi (2018) focus on the no-commitment regime, whereas I compare two pricing regimes. Second, the intended applications are different. The main application of Hidir and Vellodi (2018) is a situation where a consumer looks for a product knowing which product he wants to buy. This is suitable, for example, when a consumer performs a search query on an e-commerce website. In contrast, I consider a situation where a consumer makes privacy choices without knowing exactly the desired product. For example, when a consumer decides whether or not to reveal his browsing activities (by accepting cookies), he may not have a particular product in mind; however, the consumer expects that by sharing more data, he will likely see more relevant products through product recommendations or targeted ads. Consequently, Hidir and Vellodi (2018) adopt a cheap talk setting, whereas I employ a formulation à la Bayesian persuasion, where a consumer decides what information to disclose without knowing values.

Several papers, such as Conitzer et al. (2012) and Montes et al. (2017), examine consumers' endogenous privacy choices. Braghieri (2017) studies a consumer search model in which a consumer can choose to be targeted, which enables him to find appropriate products at low cost but can lead to higher prices. A similar trade-off arises in De Corniere and De Nijs (2016), who study a platform's choice of disclosing consumers' preferences to advertisers. In contrast to these papers, I consider the seller's commitment of not using information for pricing. Also, they assume that disclosure decision is binary, whereas I consider a broader set of disclosure rules.

Beyond the context of online disclosure, this paper relates to voluntary information disclosure in bilateral transactions (Glode et al., 2018) and information gathering by buyers before trade (Roesler, 2015; Roesler and Szentes, 2017). Finally, several papers, such as Calzolari and Pavan

(2006a,b), and Dworczak (2017), study the privacy of agents in mechanism design problems.

2 Baseline Model

There is a monopolistic seller of $K \in \mathbb{N}$ products with the set of products denoted by $\mathcal{K} = \{1, \dots, K\}$. There is a single consumer with 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 some non-degenerate probability distribution supported on a compact set $V \subset \mathbb{R}_+$.³ Let $u := (u_1, \dots, u_K)$ denote the vector of values.

The consumer's preferences are quasi-linear: If the consumer buys product k at price p , his ex post payoff is $u_k - p$. Otherwise, the payoff is zero. The seller's payoff is her revenue. The consumer and the seller are risk-neutral.

At the beginning of the game, before observing u , the consumer chooses a *disclosure rule* (M, ϕ) from an exogenously given set \mathcal{D} .⁴ 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 all probability distributions over M . After the consumer chooses a disclosure rule, nature draws $u \in V^K$ and a message $m \in M$ according to $\phi(\cdot|u) \in \Delta(M)$. In the application of online disclosure, \mathcal{D} corresponds to the set of consumers' privacy choices, such as whether or not to share one's browsing history. If \mathcal{D} consists of all disclosure rules, information disclosure takes the form of Bayesian persuasion (Kamenica and Gentzkow, 2011). Section 4 considers such a case.

Next, I describe the seller's pricing. I consider two games that differ in the timing at which the seller sets prices. Under the *no-commitment regime*, the seller sets prices *after* observing a disclosure rule (M, ϕ) and a realized message m . Under the *commitment regime*, the seller sets the price of each product *simultaneously* with the consumer's choice of a disclosure rule.⁵ In this case, the seller does not base prices on a realized message or on a disclosure rule.⁶ The commitment

³See Remark 2 for how the results extend to correlated values.

⁴I will impose more structure on \mathcal{D} in Sections 3 and 4.

⁵Alternatively, I can assume that under the commitment regime, the seller sets prices first, and after observing them, the consumer chooses a disclosure rule. This assumption does not change the equilibrium if the consumer can only reveal information about which product is most valuable as in Section 3. In contrast, it could change the equilibrium if the consumer can disclose information in an arbitrary way as in Section 4. However, the main result continues to hold: The seller is better off and the consumer is worse off under the commitment regime. This is because the seller setting prices strictly before the consumer only increases the seller's revenue under the commitment regime.

⁶For example, an e-commerce firm that adopts this regime sets prices based on neither browsing histories nor

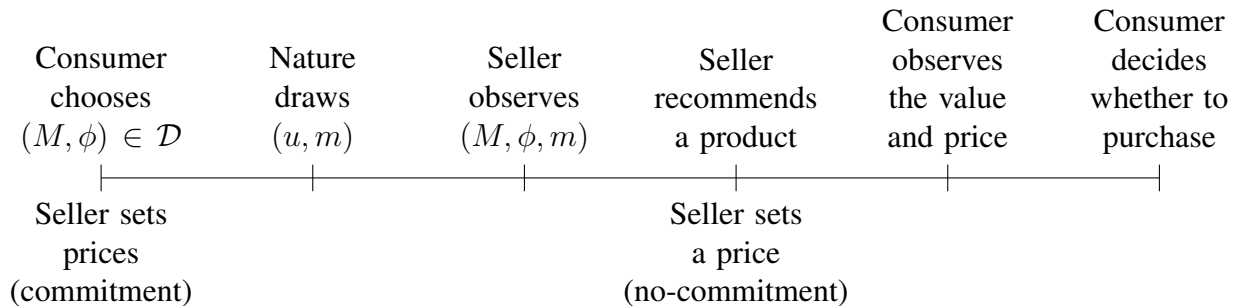


Figure 1: Timing of moves under each pricing regime.

regime captures the seller’s commitment of not using information to set prices.

Under both pricing regimes, after observing a disclosure rule (M, ϕ) and a realized message m , the seller recommends one of K products. The consumer observes the value and price of the recommended product and decides whether to buy it.

The timing of the game under each pricing regime, summarized in Figure 1, is as follows. First, the consumer chooses a disclosure rule $(M, \phi) \in \mathcal{D}$. Under the commitment regime, the seller simultaneously sets the price of each product. Then nature draws the consumer’s values u and a message $m \sim \phi(\cdot|u)$. After observing (M, ϕ) and m , the seller recommends a product. Under the no-commitment regime, the seller sets the price of the recommended product at this point. Finally, the consumer decides whether to buy the recommended product.

My solution concept is pure-strategy perfect Bayesian equilibrium (PBE) with four restrictions. First, at each information set where the seller recommends a product (and sets a price under the no-commitment regime), the seller forms her belief about values u according to the prior distribution, disclosure rule (M, ϕ) , realized message m , and Bayes’ rule.⁷ Second, the seller breaks a tie in favor of the consumer whenever she is indifferent among multiple prices or recommendations. Third, under the commitment regime, I focus on equilibrium in which each product has the same price. Fourth, if there are still multiple equilibria given previous restrictions, I focus on those that maximize the seller’s payoff. This condition eliminates the multiplicity of equilibria due to the consumer’s indifference among disclosure rules. Hereafter, “equilibrium” refers to PBE that satisfies these restrictions.

whether consumers share their browsing histories.

⁷In particular, this uniquely pins down the seller’s beliefs after the consumer deviates from the equilibrium disclosure rule.

As discussed in the introduction, the model has many applications beyond online privacy choices. Consider markets for cars, houses, and financial products, in which the variety of available products is large. In these markets, consumers often reveal information to sellers and obtain product recommendations, which enable consumers to focus on a small subset of products; however, sellers may also base prices on the information. The model captures the interaction between consumers’ incentives to reveal information and sellers’ pricing strategies in those markets.

Indeed, the application is not even restricted to buyer-seller interactions. Consider the following situation: An employer assigns her worker one of K tasks, the completion of which delivers a fixed value to the employer. The worker can disclose information about the cost of completing each task. For instance, he might communicate what kind of tasks he is good at. The employer wants to maximize the value from a task minus wage payment, whereas the worker wants to maximize wage minus cost. This is mathematically equivalent to the baseline model, where two pricing regimes correspond to whether the employer commits to wages up front.

2.1 Discussion of Modeling Assumptions

Before proceeding to the analysis, I provide a detailed discussion of modeling assumptions.

Pricing Regimes

The motivation behind each pricing regime is as follows. First, the no-commitment regime is a tractable benchmark where the consumer faces the trade-off between accurate recommendations and low prices. Second, the commitment regime captures the seller’s commitment to not use the consumer’s information to set prices. Note that this is not the only form of commitment that yields a greater revenue than the no-commitment regime. For example, the seller could do even better if she could commit to any contingent schedule of prices as a function of disclosure rules and realized messages.⁸ I exclude this kind of commitment power and focus on the commitment of not using information for pricing. This is because the latter seems more practical and relevant to sellers or policymakers that consider whether to engage in or regulate personalized pricing.

Information Disclosure

⁸If the seller could commit to any mechanism that maps a disclosure rule and a realized message into prices, then the seller would extract full surplus with a mechanism that sets price $+\infty$ for all products if and only if the consumer does not disclose full information.

I formulate the consumer’s privacy choice as a Blackwell experiment about his values. This formulation calls for several implicit assumptions; for example, the consumer understands how his privacy choice affects the seller’s posterior belief. Although such an assumption might be restrictive, it enables us to draw general insights on consumers’ informational incentives in online and offline transactions without referring to specific disclosure technologies.

Relatedly, it is crucial to my results that the consumer chooses a disclosure rule before observing his values. This would be suitable, for instance, if the consumer is not informed of the existence or characteristics of products, but understands that his personal data enable the seller to learn about his value of each product. In [Section 5](#), I provide a microfoundation for this idea in a model of two-sided private information, where the consumer is informed of his subjective taste and the seller is informed of the products’ characteristics. It would also be natural to assume that the consumer cannot manipulate message realizations ex post, 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.

Product Recommendation and Purchase Decision

There are also two substantial assumptions on product recommendation and purchasing decision. First, the seller recommends a single product and the consumer decides whether to buy it. In other words, the consumer cannot purchase products not recommended by the seller. This formulation captures situations where there are constraints on how many products can be marketed to a given consumer. Such constraints are natural if the variety of available products is large but the consumer has a limited ability to assess products because of his time or cognitive constraints.⁹ It is important to note that this assumption is *not* a result of the seller’s revenue maximization: The seller is assumed to recommend a single product (as opposed to recommending multiple products) not because it is optimal for the seller to do so, but because the assumption concisely captures the consumer’s limited attention. Indeed, in a hypothetical scenario where the consumer can evaluate all products and choose which one to buy, the seller prefers to offer and price all products.¹⁰ The

⁹Several papers, such as [Salant and Rubinstein \(2008\)](#) and [Eliaz and Spiegel \(2011\)](#), formulate consumers’ limited attention in a similar way.

¹⁰If the consumer can evaluate all products and the seller can offer all products, then I obtain the following results: Under the commitment regime, the seller sets a price p_k of each product k to maximize $\sum_{k=1}^K \mathbf{P}(u_k - p_k \geq \max\{u_j - p_j, 0\}, \forall j \in \mathcal{K})$. This is because the consumer can always find and purchase the product that maximizes his net payoff $u_k - p$. In this case, information disclosure and product recommendations play no role. In particular,

model excludes such a mechanism to incorporate the consumer's limited attention.

Second, the consumer observes the value of the recommended product when he decides whether to buy it. One way to interpret this assumption is that the consumer does not know what products exist and has not thought about how much he would be willing to pay for each possible bundle of characteristics; however, once he is shown a particular product and sees its characteristics, he is able to compute a value for it. In practice, the assumption is reasonable if a consumer can learn the value after the purchase and return it for a refund whenever the price exceeds the value.

Production Costs

It is *not* without loss of generality to assume that production costs are equal across products. (Assuming that they are equal, it is without loss to normalize the costs to zero.) For example, if the seller can produce product 1 more cheaply, then she has a greater incentive to recommend product 1 even when 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.

3 Restricted Model

To illustrate the main idea in a simple way, this section imposes the following structure on the baseline model. First, assume that the seller sells two products ($K = 2$). Then, identify \mathcal{D} with $[1/2, 1]$. Each $\delta \in [1/2, 1]$ is called a *disclosure level*, which represents the amount of information that the consumer discloses about *which product is more valuable*. As in [Figure 2](#), each δ represents a disclosure rule that draws either message 1 or 2: It draws message $m \in \{1, 2\}$ with probability δ if product m is strictly more valuable than the other product, and it draws messages 1 and 2 with equal probability if products 1 and 2 have the same value. If the seller observes message m , she infers that the consumer has a higher value for product m with probability δ . Thus, a greater δ implies that the seller can more accurately learn about which product has a higher value.

even if the seller could restrict the number of products offered to the consumer, the seller prefers to offer all products so that the consumer is more likely to find high-value products. Under the no-commitment regime, product recommendations continue to play no role. In contrast, depending on the set of disclosure rules \mathcal{D} , the consumer might reveal some information to induce the seller to set lower prices.

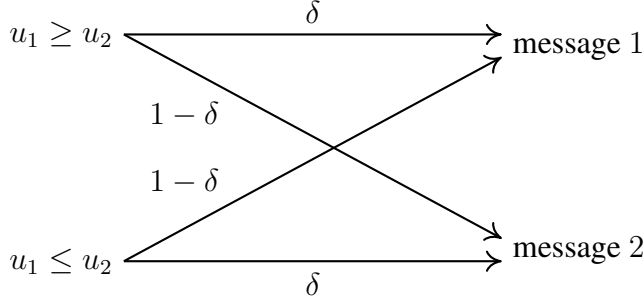


Figure 2: Disclosure rule for $\delta \in [1/2, 1]$.

3.1 Equilibrium Analysis

First, consider the seller's recommendation strategy. For a given price, the seller prefers to recommend the product that is more likely to have a higher value, because it maximizes the probability of purchase. This leads to the following lemma (see [Appendix A](#) for the proof).

Lemma 1. *Fix a pricing regime and take any equilibrium. Suppose that the consumer chooses a disclosure level $\delta > 1/2$. After observing message $k \in \{1, 2\}$, the seller recommends product k .*

[Lemma 1](#) implies that in equilibrium, a disclosure level δ is equal to the probability of the seller recommending the most valuable product (see [Figure 2](#)). Thus, greater disclosure makes it more likely that the consumer sees his preferred product.

Next, consider the equilibrium pricing. Suppose that the consumer chooses a disclosure level δ , and message $k \in \{1, 2\}$ is realized. Then, the consumer's value for the recommended product, which is product k by [Lemma 1](#), is drawn from $\delta F^{MAX} + (1 - \delta)F^{MIN}$, where F^{MAX} and F^{MIN} denote the cumulative distribution functions (CDFs) of $\max(u_1, u_2)$ and $\min(u_1, u_2)$, respectively.¹¹ This leads to the following lemma.

Lemma 2. *Consider the no-commitment regime and take any equilibrium. Suppose that the consumer has chosen a disclosure level of δ . Then, the seller sets a price of*

$$p(\delta) := \min \left(\arg \max_{p \in \mathbb{R}} p [1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)] \right) \quad (1)$$

¹¹In this paper, I define a CDF as a left-continuous function. For example, $F^{MAX}(p)$ is the probability of $\max(u_1, u_2)$ being *strictly* lower than p .

for the recommended product.¹²

It is crucial to note that $p(\delta)$ in Lemma 2 is also the optimal price under the *commitment regime* when the seller anticipates the equilibrium choice of δ . The reason is as follows. The price of product k affects revenue only when the seller recommends product k , or equivalently, only when message k is realized (Lemma 1). Thus, under the commitment regime, the seller chooses the price of product k to maximize the expected revenue *conditional on* message k . Given a disclosure level δ , this maximization problem is identical to the one in the right hand side of (1).

These observations lead to the first main result (see Appendix B for the proof).

Theorem 1. *In any equilibrium, the seller obtains a higher payoff and the consumer obtains a lower payoff under the commitment regime than under the no-commitment regime.*¹³

The intuition is as follows. Under the commitment regime, more disclosure leads to better recommendations without affecting prices. Thus, the consumer prefers the highest disclosure level $\delta = 1$. Anticipating the choice of $\delta = 1$, the seller sets a price of $p(1)$ for each product up front. In contrast, under the no-commitment regime, the consumer takes into account the effect of disclosure on prices. In particular, he can always choose $\delta = 1$ to induce the equilibrium outcome of the commitment regime, because $p(\cdot)$ in (1) describes the optimal pricing under both pricing regimes. Thus, the consumer is (weakly) better off under the no-commitment regime.

In contrast, the seller prefers committing to prices up front. First, the optimal price $p(\delta)$ depends on δ but not on a realized message. Then, the seller is indifferent between the two pricing regimes if the consumer chooses the same disclosure level δ . Moreover, the revenue is increasing in δ , because greater disclosure leads to more accurate recommendations and shifts up the consumer's demand. Thus, the seller prefers the commitment regime, where the consumer chooses $\delta = 1$.

Theorem 1 gives an economic explanation of the observed puzzle: Online sellers seem to not use individual data to price discriminate,¹⁴ and consumers seem to casually share their information

¹²Appendix H, which shows the existence of an equilibrium, proves that $p(\delta)$ is well-defined.

¹³Given my equilibrium restriction, the equilibrium is unique up to which product the seller recommends when the consumer chooses $\delta = 1/2$.

¹⁴For papers that provide empirical evidence in support of this, see the discussion in the introduction. In another instance of this, in 2000 Amazon CEO Jeff Bezos said, "We never have and we never will test prices based on customer demographics." (<http://www.e-commercetimes.com/story/4411.html>). Of course, there can be other

despite the growing concerns for personalized pricing. In light of the theorem, one may view this puzzle as sellers' strategic commitment and consumers' best response to it. The theorem further shows that this outcome might not be desirable for consumers. Indeed, the result that consumers are better off under the no-commitment regime suggests that they could be better off were sellers less than fully informed about which products are most relevant.

[Theorem 1](#) also has policy implications: Consumers may benefit from regulations that restrict the amount of information sellers can expect to acquire. To see this, consider the equilibrium under the commitment regime, where the consumer chooses $\delta = 1$. Now, suppose that a regulator restricts the set of available disclosure levels to $[1/2, \delta^*]$, where δ^* is the equilibrium choice under the no-commitment regime. Then, the consumer chooses disclosure level δ^* regardless of the pricing regime, and obtains a greater payoff than without the regulation.

One may think that [Theorem 1](#) is driven by the particular restriction on the set \mathcal{D} of available disclosure rules. This is partly true and actually an important insight. The result shows that if the consumer's endogenous disclosure is crucial for the seller to give accurate recommendations, then the seller can be better off by committing to not use information for pricing. Focusing on disclosure rules parametrized by δ is a simple way to capture such a situation. In contrast, if endogenous disclosure is not important—for example, if the seller knows the consumer's willingness to pay at the outset—then the seller may prefer the no-commitment regime.¹⁵

Nonetheless, the current restriction on \mathcal{D} is one of various conditions under which [Theorem 1](#) holds. For example, \mathcal{D} can be any set of disclosure levels such as $\mathcal{D} = \{0.5, 0.8\}$, which consists of, say, enabling cookies ($\delta = 0.8$) and disabling cookies ($\delta = 0.5$). In [Section 4](#), I establish the *strict* welfare comparisons, assuming that the consumer can choose *any* disclosure rule. [Subsection 5.1](#) discusses more general conditions on \mathcal{D} under which the (weak) welfare comparisons hold.

explanations for sellers not price discriminating. For instance, sellers may think that price discrimination would infuriate consumers who have fairness concerns. The explanation in [Theorem 1](#) suggests that, even if sellers can frame personalized pricing in a way that consumer backlash is less likely to occur, sellers may still find it profitable to refrain from price discrimination.

¹⁵We can capture such a situation by assuming that \mathcal{D} consists only of the disclosure rule that reveals the exact value vector. That is, the consumer does not have a choice of concealing information. In other contexts, the preferences of the seller and the consumer over pricing regimes could be aligned. For instance, if there is a single product and the value is drawn from $U[0, 1]$, then both the seller and the consumer prefer the no-commitment regime if the consumer can only disclose whether $v \geq 1/2$ or not.

3.2 Impact of Information Disclosure on Prices

To complement the above analysis, I provide comparative statics on how information disclosure affects prices. To do so, I first introduce a stochastic order that compares the hazard rates of two distributions. For the sake of generality, I employ the definition that does not require distributions to have densities.¹⁶

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 $z \in (-\infty, \max(s_1, s_0))$.¹⁷ 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, [Definition 1](#) is equivalent to saying that the hazard rate of G_1 is point-wise smaller than that of G_0 .¹⁸ The next lemma follows from Theorem 1.B.26 of [Shaked and Shanthikumar \(2007\)](#).

Lemma 3. F^{MAX} is greater than F^{MIN} in the hazard rate order.

The intuition is as follows. Suppose that the consumer's value for some product exceeds p . Conditional on this event, how likely is the consumer to stop buying the product if the seller marginally increases the price by ε ? If the product is the consumer's preferred one, whose value is $\max(u_1, u_2)$, then he stops buying only when both u_1 and u_2 are below $p + \varepsilon$; if the product is his less preferred one so that the value is $\min(u_1, u_2)$, then he stops buying whenever one of u_1 and u_2 is below $p + \varepsilon$. Thus, the consumer is more likely to stop buying the less preferred product than the more preferred product. This implies that the value distribution F^{MAX} has a lower hazard rate than F^{MIN} .

¹⁶First-order stochastic dominance, which is weaker than the hazard rate order, is not sufficient for the analysis of this subsection, because it has no implications on the behavior of the monopoly price. For example, suppose that distribution F_0 puts equal probability on values 1 and 3 and that distribution 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.

¹⁷ $a/0$ is taken to be equal to $+\infty$ whenever $a > 0$.

¹⁸Suppose that G_0 and G_1 have densities g_0 and g_1 . By differentiating $\log \frac{1-G_1(z)}{1-G_0(z)}$ in z , we can show that the above definition is equivalent to

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

As this intuition suggests, the hazard rate order relates to demand elasticity.¹⁹ As in [Bulow and Roberts \(1989\)](#), the demand curve for a CDF F is given by $D(p) = 1 - F(p)$, and thus the demand elasticity is $-\frac{d \log D(p)}{d \log p} = \frac{f(p)}{1-F(p)}p$. Then, a CDF F_1 is greater than a CDF F_0 in the hazard rate order if and only if the demand curve for F_1 has a lower price elasticity of demand than F_0 . Thus, [Lemma 3](#) states that the consumer's demand for the more preferred product is less elastic.

Recall that, by [Lemma 1](#), a disclosure level δ is equal to the probability that the seller recommends the consumer's preferred product. Thus, a greater δ implies that the consumer is more likely to have less elastic demand for the recommended product. As a result, the seller prefers to set a higher price for it. The following result formalizes this intuition (see [Appendix C](#) for the proof).

Proposition 1. *If the consumer chooses a greater disclosure level, then the seller sets higher prices: $p(\delta)$ in (1) is increasing in δ .*

[Lemma 1](#) and [Proposition 1](#) clarify the consumer's trade-off: If the consumer chooses a greater disclosure level, then the seller can recommend the best product with a higher probability. However, this makes the consumer's demand for the recommended product less elastic. As a result, the seller sets a higher price for the recommended product. While this trade-off exists under both pricing regimes, the consumer takes it into account only under the no-commitment regime. Thus, the consumer discloses more information and obtains a lower payoff under the commitment regime.

3.3 [Theorem 1](#) as a Tragedy of the Commons

The economic mechanism of [Theorem 1](#) is similar to a tragedy of the commons due to a “negative externality” associated with information disclosure. To illustrate this idea, I consider a slight variant of the restricted model described below.

Suppose that there is now a unit mass of consumers, each of whom chooses a disclosure level. The value of each product is independent across consumers.²⁰ Under both pricing regimes, the seller sets prices *after* observing the disclosure level and realized message of each consumer. However, under the commitment regime, the seller has to set a single price for each product that applies

¹⁹Demand refers to the ex ante demand (i.e., probability of purchase) unconditional on the value instead of the ex post demand given the value, which is a discontinuous function of price.

²⁰[Sun \(2006\)](#) studies a continuum of IID random variables for which the law of large numbers holds.

to all consumers. Under the no-commitment regime, the seller can charge different prices to different consumers. As before, the seller can always recommend different products to different consumers.

The equilibrium prediction in [Theorem 1](#) persists. To see this, consider the commitment regime. In equilibrium, each consumer $i \in [0, 1]$ chooses a disclosure level δ_i *taking prices as given*, because the choice of a single consumer in a large population does not affect the seller's subsequent pricing decision. Thus, every consumer chooses $\delta_i = 1$, following which the seller sets a price of $p(1)$ for each product.

According to this interpretation, we can view [Theorem 1](#) as a tragedy of the commons: Under the commitment regime, if some (positive mass of) consumers disclose more information, the seller prefers to increase prices, as she can offer accurate recommendations to a greater fraction of consumers. But then, *all* consumers face higher prices. That is, under the commitment regime, greater disclosure by some consumers lowers the welfare of other consumers through higher prices. Since consumers do not internalize this negative impact, they choose the highest disclosure level, although they could be better off by collectively withholding information. The problem does not arise under the no-commitment regime, because each consumer i internalizes the impact of disclosure on prices, $p(\delta_i)$. [Appendix D](#) formalizes this observation.

Remark 1. One might think that the consumer's inability to evaluate all products is inconsistent with his ability to make an optimal privacy choice. I argue that these assumptions are not necessarily contradictory. First, e-commerce firms, such as Amazon and eBay, sell more products than one can exhaustively examine. Then, it has to be an institutional feature of these platforms to display only a subset of the products. In such cases, we cannot conclude that consumers who know how to use privacy tools should also be able to find relevant products without the help of search engines and recommendations. Second, there are situations where it is not difficult to figure out how to withhold information. For example, on the Internet it is increasingly common that pop-up windows ask users whether to enable cookies due to recent legislation in the European Union.²¹

Remark 2. [Theorem 1](#) is robust to a variety of extensions.

Correlated Values: [Theorem 1](#) holds if (u_1, u_2) has an exchangeable distribution.

²¹For example, Article 5(3) of the ePrivacy directive states that websites must ask users if they agree to most cookies and similar technologies (e.g. web beacons, Flash cookies, etc.) before the site starts to use them.

Costly Disclosure: The result regarding the consumer's welfare in [Theorem 1](#) holds even if the consumer incurs a cost $c(\delta)$ to choose a disclosure level δ .²² For instance, the consumer may incur privacy cost, in which case $c(\delta)$ is increasing. For another instance, the consumer may need some effort to figure out how to conceal information, in which case $c(\delta)$ is decreasing. For any shape of $c(\cdot)$, the consumer is worse off under the commitment regime.

Informational Externality: In practice, online sellers may infer the preference of a consumer from the preferences of other consumers if they share similar characteristics. To incorporate this, consider the model with a continuum of consumers, and assume that a true disclosure level for consumer i is $\Delta(\delta_i, \bar{\delta})$, which is increasing in i 's own choice δ_i and the average disclosure level of the population $\bar{\delta} = \int_{i \in [0,1]} \delta_i di$. This captures the idea that the seller can learn about i 's preferences from information disclosed by others. In this case, a similar result to [Theorem 1](#) holds.

4 Unrestricted Model

This section assumes that the consumer can disclose *any* information about the value vector $u = (u_1, \dots, u_K)$. Formally, suppose that the seller sells $K \geq 2$ products, and \mathcal{D} consists of all disclosure rules with finite message spaces.²³ In this unrestricted model, the consumer can disclose not only information about which product is most valuable (as in [Section 3](#)) but also information about the value of a particular product, such as whether u_k exceeds some threshold.

Two differences between the unrestricted model and the restricted model in [Section 3](#) are worth mentioning. First, compared to the restricted model, the unrestricted model a priori favors the no-commitment regime in terms of revenue. This is because, given general disclosure rules, the no-commitment regime often yields a higher revenue, as opposed to the restricted model where the two pricing regimes yield equal revenue for any disclosure level δ .²⁴ Second, the unrestricted model relates to [Bergemann et al. \(2015\)](#). Their results imply that a single-product monopolist is

²²If the consumer chooses a disclosure level δ and purchases product k at price p , his payoff is $u_k - p - c(\delta)$. If he buys nothing, the payoff is $-c(\delta)$.

²³Since I will assume that the value distribution has a finite support V , this restriction is without loss of generality. Indeed, the consumer can maximize his payoff if he can choose any disclosure rule with a message space M such that $|M| \leq |V| \times K$. This is because the consumer can always pool two message realizations that lead to the same recommended product and the same price.

²⁴For example, if the consumer chooses a disclosure rule that reveals exact value vector u , then the seller can extract full surplus only under the no-commitment regime.

indifferent between the two pricing regimes in equilibrium, and the equilibrium is efficient under the no-commitment regime. In contrast, I will show that the equilibrium is typically inefficient and the seller strictly prefers the commitment regime.

The consumer's values u_1, \dots, u_K are IID draws from the prior distribution $x_0 \in \Delta(V)$. I assume that x_0 has a finite support $V = \{v_1, \dots, v_N\}$ with $0 < v_1 < \dots < v_N$ and $N \geq 2$. For any $x \in \Delta(V)$, $x(v)$ denotes the probability that x puts on $v \in V$. Abusing notation slightly, let $p(x)$ denote the lowest optimal price given $x \in \Delta(V)$:

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

Note that $p(x)$ does not depend on the number of products (K). To focus on the most interesting case, I impose the following assumption. Loosely speaking, it requires that the prior x_0 does not put too much weight on the lowest value of its support.

Assumption 1. The lowest optimal price at the prior value distribution strictly exceeds the lowest value of its support: $p(x_0) > \min V$.

As the consumer can access a rich set of disclosure rules, the analysis is more involved than before. However, there turns out to be a clear relationship between pricing regimes and the kinds of information disclosed by the consumer. The next subsection illustrates this by showing that different pricing regimes lead to different kinds of inefficiency. I use these results to show that the seller is strictly better off and the consumer is strictly worse off under the commitment regime.

4.1 Inefficiency of Equilibrium

In this model, an equilibrium can be inefficient in two ways: One is when the consumer decides not to buy any products; the other is when the consumer buys some products that are not the most valuable ones. In principle, these kinds of inefficiencies can coexist. However, I show that each pricing regime is associated with only one type of inefficiency. The following result states that the commitment regime leads to the first type of inefficiency (see [Appendix E](#) for the proof).

Proposition 2. *Consider the commitment regime. In any equilibrium, the seller recommends the most valuable product with probability 1. However, trade fails to occur with a positive probability.*

The intuition is as follows. Once the seller commits to prices, the consumer prefers to fully disclose the highest-value product so that the seller can recommend it. Given the efficient recommendations, the consumer is less likely to have low values for the recommended product than when the value is drawn from the prior value distribution. Anticipating this, the seller commits to prices strictly greater than $\min V$ under [Assumption 1](#). With a positive probability, the consumer's value for the best product falls below the price, and trade does not occur.

Next, I show that the no-commitment regime exhibits a different kind of inefficiency. The proof needs some work, which is contained in [Appendix F](#). As the existence of an equilibrium is non-trivial, I separately prove it in [Appendix H](#).

Proposition 3. *Consider the no-commitment regime. In any equilibrium, trade occurs with probability 1. However, for generic priors x_0 satisfying [Assumption 1](#), in any equilibrium, the consumer purchases some products other than the most valuable ones with a positive probability.*²⁵

An intuition for the first part is as follows. Suppose that the seller charges a price that exceeds the consumer's value with a positive probability. Suppose that, on such an event, the consumer discloses whether his value exceeds the price. If the seller learns that the value falls below the original price, she prefers to revise the price downward or recommend another product, which benefits both the consumer and the seller relative to no trade. Importantly, on the complementary event where the seller learns that the value exceeds the price, she prefers to recommend the same product at the same price as before.²⁶ Overall, the consumer is always willing to disclose whether his value exceeds the price, which ensures that trade occurs with probability 1.

The second part of [Proposition 3](#) implies that any equilibrium is generically inefficient, because the consumer fails to purchase the highest-value product with a positive probability. The interpretation is that the seller's lack of commitment incentivizes the consumer to withhold information about which product is most valuable. The result states that, under the current distributional assumption, the consumer *always* finds it profitable to conceal some information.

²⁵“Generic priors x_0 satisfying [Assumption 1](#)” means that there is a Lebesgue measure zero set $X_0 \subset \Delta(V) \subset \mathbb{R}^N$ such that for any $x_0 \in \Delta(V) \setminus X_0$ satisfying [Assumption 1](#), any equilibrium has a positive probability of product mismatch.

²⁶This hinges on the following observation in the (single product) monopoly pricing: If the seller optimally sets a price given some value distribution, then she does not revise the price even after learning that the value exceeds the price. Indeed, if the value $w \in \mathbb{R}$ is drawn from $\mathbf{P}(\cdot)$, then the optimal price p^* maximizes $p \cdot \mathbf{P}(w \geq p)$. If the seller additionally learns that $w \geq p^*$, then her new problem is $\max_{p \geq p^*} p \cdot \frac{\mathbf{P}(w \geq p)}{\mathbf{P}(w \geq p^*)}$, which is maximized at $p = p^*$.

It would need more work to derive concrete policy implications, but [Proposition 3](#) has the following takeaway. Consider a regulator or an Internet intermediary, who cares about consumers and wants to release their information to sellers in order to improve welfare. The analysis suggests that releasing information about consumers who have low values (i.e., values below the monopoly price) for all products is good for the welfare of sellers and consumers. In contrast, a regulator or an intermediary should be careful about releasing information about consumers who have low values for some products and high values for other products. Whereas releasing such information may increase total welfare, it can hurt consumers to the extent that the loss from higher prices dominates the benefit from the improved match quality.

I sketch the proof of [Proposition 3](#). For ease of exposition, I use the following terminologies.

Definition 2. An equilibrium is *vertically efficient* if trade occurs with probability 1. An equilibrium is *horizontally efficient* if the seller recommends the most valuable products with probability 1.

[Proposition 3](#) is now rephrased as follows: Under the no-commitment regime, any equilibrium is vertically efficient but generically horizontally inefficient. The proof of vertical efficiency follows the previous intuition: If an equilibrium is vertically inefficient, we can construct another equilibrium in which the consumer discloses more information and both the seller and the consumer are better off. This contradicts my equilibrium restriction.

Horizontal inefficiency is challenging to prove for at least two reasons. First, the consumer’s disclosure problem is a Bayesian persuasion ([Kamenica and Gentzkow, 2011](#)) with multidimensional state and action spaces, which is known to be difficult.²⁷ Second, there may be multiple equilibria and I want to prove horizontal inefficiency for all of them.

To prove horizontal inefficiency without characterizing equilibrium, I take the following two-step approach. First, solve a “constrained Bayesian persuasion” in which the consumer chooses a disclosure rule subject to the constraint that the outcome is horizontally efficient (given the seller’s optimal behavior). Characterizing such a disclosure rule, denoted by ϕ^* , turns out to be simpler than the unconstrained maximization problem that the consumer faces in equilibrium. Second,

²⁷Precisely, the model is slightly different from a Bayesian persuasion because the sender (consumer) also takes an action. However, we can regard the consumer’s payoffs from his optimal purchase behavior as the sender’s payoffs, which depend only on the state (product values) and the receiver’s action (seller’s recommendation and pricing).

modify ϕ^* to create disclosure rule ϕ^I that leads to horizontal inefficiency but gives the consumer a strictly greater payoff than ϕ^* . These two steps imply that any equilibrium is associated with horizontal inefficiency. The following example illustrates these two steps.

Example 1. Suppose that $K = 2$, $V = \{1, 2\}$, and $(x_0(1), x_0(2)) = (1/3, 2/3)$.

Step 1: Consider disclosure rule ϕ in Table 1. The first column shows possible value vectors, and each row shows the distribution over messages 1 and 2 given each value vector. Note that ϕ

Table 1: Disclosure rule ϕ revealing product ranking.

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

only discloses which product is more valuable, the information necessary and sufficient to achieve horizontally efficient outcomes. I characterize ϕ^* by maximizing the consumer's payoff among all disclosure rules weakly more informative than ϕ . To do so, for each k , I first calculate the posterior distribution of u_k conditional on message $k \sim \phi(\cdot|u)$. Then, I apply Bergemann et al.'s (2015) *consumer surplus maximizing segmentation (CSMS)* to each posterior distribution. In the single-product case, a CSMS discloses information about the consumer's value to maximize consumer surplus. In the current context, applying a CSMS to each posterior enables the consumer to disclose information about the value of the best product to maximize his payoff. ϕ^* constructed in this way is a solution of the constrained problem.

Table 2 presents disclosure rule ϕ^* obtained in this way. The CSMS decomposes each message k (of ϕ) into messages $k1$ and $k2$.²⁸ The seller's best responses are as follows: After observing message $k1$ ($k = 1, 2$), the seller recommends product k at price 1, being indifferent between prices 1 and 2. After observing message $k2$ ($k = 1, 2$), the seller recommends product k at price 2.

Step 2: I modify ϕ^* twice to create ϕ^I in Table 3: First, at $(u_1, u_2) = (2, 1)$, ϕ^I sends message 21 instead of 12 with a small probability $\varepsilon > 0$. This creates horizontal inefficiency, because once the “new” message 21 is realized, the seller recommends product 2 even though $(u_1, u_2) = (2, 1)$ with

²⁸In general, a CSMS is not unique. In the proof and this example, I use a CSMS obtained by the *greedy algorithm* of Bergemann et al. (2015).

Table 2: Efficient disclosure rule ϕ^* .

	$\phi^*(11 u_1, u_2)$	$\phi^*(12 u_1, u_2)$	$\phi^*(21 u_1, u_2)$	$\phi^*(22 u_1, u_2)$
(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 3: Horizontally inefficient disclosure rule ϕ^I .

	$\phi^I(11 u_1, u_2)$	$\phi^I(12 u_1, u_2)$	$\phi^I(21 u_1, u_2)$	$\phi^I(22 u_1, u_2)$
(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

a positive probability. However, this modification does not affect the consumer's payoff, because at message 12, the consumer continues to obtain a payoff of zero. Importantly, this modification relaxes the seller's incentive, as she now *strictly* prefers to set price 1 at message 21. Second, I further modify ϕ^* so that, at $(u_1, u_2) = (2, 2)$, ϕ^I sends message 21 instead of 22 with a small probability $\varepsilon' > 0$. This strictly increases the consumer's payoff: At $(u_1, u_2) = (2, 2)$, where the consumer obtains a payoff of zero at the original ϕ^* , he now obtains a strictly positive payoff when message 21 is realized. To sum up, ϕ^I leads to horizontal inefficiency but gives the consumer a strictly greater payoff than ϕ^* .

Finally, I discuss how to generalize the proof strategy for arbitrary parameters (K and x_0). Generalizing Step 1 is straightforward. For Step 2, I first prove that disclosure rule ϕ^* obtained in Step 1 (generically) sends messages m_0 and m_1 with the following properties: The seller recommends different products at m_0 and m_1 ; conditional on message m_0 , the consumer obtains a payoff of zero and has the lowest value v_1 for all the products that are not recommended; conditional on message m_1 , the seller prefers to set the lowest possible price v_1 , being indifferent to setting any prices in V . I modify ϕ^* so that it sends m_1 instead of m_0 with a small positive probability, in order to give the seller a strict incentive to set price v_1 at the new m_1 . This does not lower the consumer's payoff. Finally, I use the seller's strict incentive to show that I can modify the disclosure rule to increase the consumer's payoff.

4.2 Welfare Comparisons in the Unrestricted Model

In the restricted model, the seller is better off and the consumer is worse off under the commitment regime ([Theorem 1](#)). We might think that such a result no longer holds in the current setting, because the no-commitment regime has a greater probability of trade ([Proposition 3](#)).

The following result, however, shows that the seller still prefers to commit to not use the consumer's information for pricing, and the commitment hurts the consumer. To state the result, let R_C and U_C denote the equilibrium payoffs of the seller and the consumer, respectively, under the commitment regime. Similarly, let R_N and U_N denote the payoffs of the seller and the consumer, respectively, in any equilibrium under the no-commitment regime.

Theorem 2. *Suppose that the consumer can choose any disclosure rule and [Assumption 1](#) holds. Generically, the seller is strictly better off and the consumer is strictly worse off under the commitment regime: $R_C > R_N$ and $U_C < U_N$.*

I prove the result without characterizing equilibria. The key is to use ϕ^* , the disclosure rule that maximizes consumer surplus subject to the efficiency constraint, which is characterized in the proof of [Proposition 3](#). I compare ϕ^* with the equilibrium disclosure rule of each pricing regime in terms of the welfare of the seller and the consumer.

Proof. Let ϕ^C denote any equilibrium disclosure rule under the commitment regime, where the seller recommends the most valuable products ([Proposition 2](#)). Recall that ϕ^* is the disclosure rule constructed in the proof of [Proposition 3](#): ϕ^* maximizes the consumer's payoff among all the disclosure rules achieving efficient recommendations. Given either disclosure rule, conditional on the event that the seller recommends product k , the value distribution of product k is equal to the distribution of $\max_k u_k$, denoted by x^{MAX} .

Let p^* denote the equilibrium price of each product under ϕ^C . One observation is that p^* also maximizes revenue under any posteriors drawn by ϕ^* . (This is because I construct ϕ^* from ϕ^C using the [Bergemann et al.'s \(2015\)](#) consumer surplus maximizing segmentation.) In other words, under ϕ^* , the seller can achieve the highest revenue by posting price p^* up front for all products. Denoting the optimal revenue under ϕ^C and ϕ^* by R_C and R^* respectively, I obtain $R_C = R^*$.

As ϕ^* is efficient, it can never consist of an equilibrium ([Proposition 3](#)). That is, the consumer's equilibrium payoff under the no-commitment regime (U_N) is strictly greater than the one from ϕ^* .

Since the consumer's payoff is greater under ϕ^* than ϕ^C , $U_N > U_C$ holds. Also, $R_N < R^*$ holds because the total surplus is strictly smaller under ϕ^N than ϕ^* but the consumer strictly prefers ϕ^N . As $R^* = R_C$, it holds $R_N < R_C$. \square

As [Proposition 3](#) shows, the consumer under the no-commitment regime obfuscates which product has the highest value. However, the seller might still prefer the no-commitment regime, because it makes trade more likely to occur. The reason why this argument fails is that, if the consumer can choose any disclosure rule, he can disclose partial information to increase the probability of trade without increasing the seller's payoff. (This logic is from [Bergemann et al. \(2015\)](#).) In other words, the seller does not lose from not being able to base prices on information. As a result, the seller prefers the commitment regime, which leads to more accurate recommendations.

4.3 Commitment and Efficiency

Which pricing regime achieves greater total surplus? If there is a single product, the answer is simple because the no-commitment regime achieves full efficiency. If there are multiple products, the answer is not obvious, as neither of the pricing regimes achieves full efficiency under [Assumption 1](#). Indeed, the answer generally depends on the prior value distribution x_0 of each product and the number K of products.

The next result shows that if there is a large number of products, the commitment regime is more efficient. To focus on the interesting case where it leads to a *strictly* greater total surplus, I assume that x_0 does not put too much weight on the highest value of its support V . In the following analysis, I no longer impose [Assumption 1](#).

Assumption 2. The optimal price at the prior distribution is strictly lower than the highest value of its support: $p(x_0) < \max V < +\infty$.

In the following result, I fix x_0 and take a large K (see [Appendix G](#) for the proof).

Proposition 4. *Under the commitment regime, as $K \rightarrow +\infty$, the equilibrium total surplus converges to $\max V$. Under the no-commitment regime, if [Assumption 2](#) holds, then there is $\varepsilon > 0$ such that for any K , the equilibrium total surplus is at most $\max V - \varepsilon$.*

The intuition is as follows. Under the commitment regime, the seller can recommend the best product with probability 1 (Proposition 2). Then, as K grows large, the consumer's value for the recommended product (i.e., $\max(u_1, \dots, u_K)$) becomes nearly degenerate at $\max V$. The seller can then set prices near $\max V$ to obtain the revenue close to $\max V$. As a result, consumer surplus converges to zero and total surplus converges to $\max V$. In contrast, under the no-commitment regime, the consumer can disclose no information and secure a positive payoff that is independent of the number of products (given Assumption 2). This implies that the consumer's value for the recommended product, which is weakly greater than total surplus, never approaches $\max V$.

5 Extensions

5.1 General Conditions for Welfare Ranking Results

This extension provides more general conditions on the set \mathcal{D} of available disclosure rules under which the seller is better off and the consumer is worse off under the commitment regime. Define ϕ^E as a disclosure rule that only discloses the highest-value product: For any realized $u \in V^K$, ϕ^E draws message $k \in \arg \max_{\ell} u_{\ell}$ with probability $\frac{1}{|\arg \max_{\ell} u_{\ell}|}$. (I can alternatively define ϕ^E as a disclosure rule that discloses the ranking of the products in terms of their values.) The following result provides a condition under which the seller prefers to commit to not use information for pricing. The condition subsumes the one in Section 3.

Proposition 5. *If ϕ^E belongs to \mathcal{D} and is more informative than any other disclosure rule in \mathcal{D} , then the seller obtains a higher payoff under the commitment regime.*

Proof. As in Theorem 1, under the commitment regime, the consumer chooses ϕ^E and the seller recommends the best product. Thus, the consumer's value for the recommended product is $\max_k u_k$. Consider the no-commitment regime. Let $(M, \phi) \in \mathcal{D}$ denote an equilibrium disclosure rule. Take any realized message m drawn by ϕ . Conditional on m , the consumer's value distribution for the recommended product (say k^*) is lower than that for $\max_k u_k$ in the sense of first-order stochastic dominance. Indeed, since ϕ is less informative than ϕ^E , the value distribution of product k^* conditional on m is a convex combination of the distributions of $u_{k^*} \mid \{u_{k^*} = \max_k u_k\}$ and $u_{k^*} \mid \{u_{k^*} \neq \max_k u_k\}$. Thus, the seller is better off under the commitment regime. \square

The next result provides a condition under which the consumer is worse off under the seller's commitment. The condition subsumes the ones in Sections 3 and 4. However, in contrast to Theorem 2, it only states that the consumer is *weakly* better off under the no-commitment regime.

Proposition 6. *If $\phi^E \in \mathcal{D}$, the consumer is better off under the no-commitment regime.*

Proof. As in Theorem 1, under the commitment regime, the consumer chooses ϕ^E . Under the no-commitment regime, the consumer can achieve the same outcome by choosing ϕ^E . Thus, the consumer is weakly better off under the no-commitment regime. \square

5.2 Market for Personal Data

Institutions within which consumers can sell their information have been discussed as market-based solutions to privacy problems. In my model, such a market for data could benefit both the seller and the consumer.

To see this, consider the following extension: At the beginning of the game, the seller can offer to buy information. Formally, the seller chooses a disclosure rule $\phi \in \mathcal{D}$ and a transfer $t \in \mathbb{R}$. Then, the consumer decides whether to accept it. If the consumer accepts, he then reveals values according to ϕ and receives t ; otherwise, he can choose any disclosure rule in \mathcal{D} but receives no transfer. In either case, this is followed by a product recommendation and a purchasing decision. Again, I consider the two pricing regimes, where the commitment regime means that the seller sets a price for each product before she chooses (ϕ, t) .

How does this market for data affect equilibrium outcomes? First, it has no impact under the commitment regime, because the consumer is willing to disclose full information without compensation. In contrast, under the no-commitment regime, the market for data can benefit the seller without affecting the consumer. For example, suppose that \mathcal{D} contains a disclosure rule ϕ^* that fully reveals values u . Then, the seller offers (ϕ^*, t) , where t makes the consumer indifferent between accepting and rejecting the offer. In equilibrium, the consumer accepts the offer and the seller engages in perfect price discrimination with efficient recommendations. Recall that without compensation, the consumer typically hides some information, which leads to product mismatch (Proposition 3). Therefore, with the market for data, not only the consumer but the seller may prefer the no-commitment regime.

5.3 A Model of Two-Sided Private Information

It is crucial to my results that the consumer chooses a disclosure rule without observing his values of products. As discussed, this is suitable if the consumer is initially uninformed of product characteristics necessary to calculate his willingness to pay. I provide a microfoundation for this idea, focusing on the restricted model in [Section 3](#).

For ease of exposition, label two products as 1 and -1 . At the beginning of the game, the consumer privately observes his *taste* $\theta \in \{1, -1\}$. Also, the seller privately observes *product characteristics* $\pi \in \{1, -1\}$. Each pair of (θ, π) is equally likely. Given a realized (θ, π) , the consumer draws values of products $\theta \cdot \pi$ and $-\theta \cdot \pi$ from (the distributions of) $\max\{u_1, u_2\}$ and $\min\{u_1, u_2\}$, respectively. Note that θ or π alone is not informative of product values, but (θ, π) is. This formulation captures situations in which sellers have to combine information about the tastes of consumers and product characteristics in order to learn about preferences.

The game proceeds as follows. *After privately observing θ* , the consumer (publicly) chooses a disclosure level δ : With probabilities δ and $1 - \delta$, messages θ and $-\theta$ are realized, respectively. The seller observes δ and a realized message, and then recommends a product. As before, I consider two pricing regimes. Note that, once the consumer observes θ , the game looks identical with the original restricted model. Thus, this setting leads to the same welfare comparisons as in [Theorem 1](#).

5.4 Alternative Interpretation: Online Advertising Platform

We can view the restricted model in [Section 3](#) as a reduced form of a game among a consumer, an online advertising platform, and two advertisers. Advertisers 1 and 2 sell products 1 and 2, respectively. The platform displays an ad of a product to the consumer based on the outcome of an ad auction.

In this interpretation, first, the consumer chooses a disclosure level δ (e.g., whether to accept a cookie) and visits the platform. Each advertiser $k \in \{1, 2\}$ chooses a price of product k and a bidding rule $b_k : \{1, 2\} \rightarrow \mathbb{R}$. Here, $b_k(j)$ is the bid by advertiser k for the impression of the consumer with a realized message $j \in \{1, 2\}$. I assume that advertisers choose bidding rules after observing δ and a realized message. If advertiser k wins the auction, the consumer sees the ad for

product k . After seeing an ad, the consumer learns the value and price of the advertised product, and then decides whether to buy it.

I obtain the same result as [Theorem 1](#). Suppose that the consumer chooses a disclosure level δ . First, if advertisers can base product prices on disclosure levels, each advertiser chooses price $p(\delta)$ and bidding rule b_k , where $b_k(k) = p(\delta)[1 - \delta F^{MAX}(p(\delta)) - (1 - \delta)F^{MIN}(p(\delta))]$ and $b_k(j) < b_k(k)$ for $j \neq k$. The platform runs a first or second price auction with a reserve price $p(\delta)[1 - \delta F^{MAX}(p(\delta)) - (1 - \delta)F^{MIN}(p(\delta))]$ to extract full surplus from advertisers. Given these strategies, the consumer sees the ad for his preferred product with probability δ . Second, if advertisers have to set prices without observing δ , the consumer chooses disclosure level 1 and each advertiser sets price $p(1)$. Thus, I obtain the same welfare comparisons as in [Theorem 1](#): The platform benefits from committing to not reveal δ to advertisers, and the commitment hurts the consumer. For example, platforms might classify consumers into two segments, each of which consists of consumers who are more likely to prefer one product than the other. Then, the platform can encourage consumers to disclose more information by committing to not reveal the exact probability that each consumer prefers one product to the other.

6 Concluding Discussion

This paper studies consumers' privacy choices, their price implications, and welfare consequences. The key to the analysis is the following trade-off: A consumer may benefit from revealing information, because a seller can then offer more accurate product recommendations. However, the seller may also use the information to tailor prices.

The main finding, however, is that the consumer may not encounter this trade-off in equilibrium. I identify a condition under which the seller prefers to commit to not use the consumer's information for pricing. The commitment encourages the consumer to disclose more information that is useful for accurate recommendations. Moreover, this commitment can hurt the consumer: The consumer could be better off if the seller had less information and could only make noisy recommendations. In contrast to the standard single-product model, the seller's commitment can enhance total welfare at the expense of consumer welfare. Finally, the model of unrestricted disclosure reveals that even with fine-grained control of information, we cannot simultaneously achieve

efficient price discrimination and efficient matching of products without sacrificing consumer welfare.

References

- Acquisti, Alessandro and Hal R Varian (2005), “Conditioning prices on purchase history.” *Marketing Science*, 24, 367–381.
- Aguirre, Inaki, Simon Cowan, and John Vickers (2010), “Monopoly price discrimination and demand curvature.” *American Economic Review*, 100, 1601–15.
- Aliprantis, Charalambos D and Kim Border (2006), *Infinite dimensional analysis: a hitchhiker’s guide*. Springer Science & Business Media.
- 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.
- Calzolari, Giacomo and Alessandro Pavan (2006a), “Monopoly with resale.” *The RAND Journal of Economics*, 37, 362–375.
- Calzolari, Giacomo and Alessandro Pavan (2006b), “On the optimality of privacy in sequential contracting.” *Journal of Economic theory*, 130, 168–204.
- Chen, Yongmin and Marius Schwartz (2015), “Differential pricing when costs differ: a welfare analysis.” *The RAND Journal of Economics*, 46, 442–460.
- 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 (2016), “Welfare-increasing third-degree price discrimination.” *The RAND Journal of Economics*, 47, 326–340.

- De Corniere, Alexandre and Romain De Nijs (2016), “Online advertising and privacy.” *The RAND Journal of Economics*, 47, 48–72.
- Dworczak, Piotr (2017), “Mechanism design with aftermarkets: Cutoff mechanisms.”
- Eliasz, Kfir and Ran Spiegler (2011), “Consideration sets and competitive marketing.” *The Review of Economic Studies*, 78, 235–262.
- 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 Xingtian Zhang (2018), “Voluntary disclosure in bilateral transactions.” *Journal of Economic Theory*, 175, 652–688.
- Hidir, Sinem and Nikhil Vellodi (2018), “Personalization, discrimination and information revelation.” Mimeo, University of Warwick.
- Iordanou, Costas, Claudio Soriente, Michael Sirivianos, and Nikolaos Laoutaris (2017), “Who is fiddling with prices?: Building and deploying a watchdog service for e-commerce.” In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 376–389, ACM.
- 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.” Mimeo, Toulouse School of Economics.
- Narayanan, Arvind (2013), “Online price discrimination: conspicuous by its absence.” URL <https://33bits.wordpress.com/2013/01/08/online-price-discrimination-conspicuous-by-its-absence/>.
- Pigou, Arthur (1920), *The Economics of Welfare*. London: Macmillan.

- Ponomarev, Stanislav P (1987), “Submersions and preimages of sets of measure zero.” *Siberian Mathematical Journal*, 28, 153–163.
- 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.
- Salant, Yuval and Ariel Rubinstein (2008), “(a, f): choice with frames.” *The Review of Economic Studies*, 75, 1287–1296.
- Shaked, Moshe and George Shanthikumar (2007), *Stochastic orders*. Springer Science & Business Media.
- Stokey, Nancy L (1979), “Intertemporal price discrimination.” *The Quarterly Journal of Economics*, 355–371.
- Stokey, Nancy L (1981), “Rational expectations and durable goods pricing.” *The Bell Journal of Economics*, 112–128.
- 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.

Zeng, Xiaohua, Srabana Dasgupta, and Charles B Weinberg (2013), “The competitive implications of a no-haggle pricing policy: The access Toyota case.” Technical report, Working paper, College of Business, City University of Hong Kong, Hong Kong .

Zeng, Xiaohua, Srabana Dasgupta, and Charles B Weinberg (2016), “The competitive implications of a no-haggle pricing strategy when others negotiate: Findings from a natural experiment.” *International Journal of Research in Marketing*, 33, 907–923.

Appendix For Online Publication

A Proof of Lemma 1

Let F^{MAX} and F^{MIN} denote the CDFs of $\max(u_1, u_2)$ and $\min(u_1, u_2)$, respectively. Without loss of generality, suppose that message 1 is realized. If the seller recommends products 1 and 2, then the consumer draws values from $\delta F^{MAX} + (1 - \delta)F^{MIN}$ and $\delta F^{MIN} + (1 - \delta)F^{MAX}$, respectively. The former first order stochastically dominates the latter, because $\delta > 1/2$ and F^{MAX} first order stochastically dominates F^{MIN} . Then, under both pricing regimes, recommending product 1 maximizes revenue given any prices. (Here, I use the equilibrium restriction that all products have the same price under the commitment regime.) Note that the seller is indifferent between recommending two products if and only if $p(\delta) = \min V$. In this case, the equilibrium restriction requires that the seller recommends product 1, because it uniquely maximizes the consumer’s payoff given that the value distribution is non-degenerate.

B Proof of Theorem 1

If the seller recommends the consumer's preferred and less preferred products at price p , then the expected payoffs are respectively $u^{MAX}(p) := \int_p^{+\infty} (v - p) dF^{MAX}(v)$ and $u^{MIN}(p) := \int_p^{+\infty} (v - p) dF^{MIN}(v)$.

Consider the commitment regime. Let p^* denote the equilibrium price of each product. If the consumer chooses δ , then his expected payoff is $\delta u^{MAX}(p^*) + (1 - \delta)u^{MIN}(p^*)$. This is maximized

at $\delta = 1$, which is also a unique disclosure level consistent with my equilibrium restriction. Anticipating $\delta = 1$, the seller sets $p^* = p(1)$ up front and thus the consumer's payoff is $u^{MAX}(p(1))$.

Under the no-commitment regime, the consumer's payoff from δ is $\delta u^{MAX}(p(\delta)) + (1 - \delta)u^{MIN}(p(\delta))$. Thus, his equilibrium payoff is

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

where the existence of a maximizer is shown in [Appendix H](#). That is, the consumer is worse off under the commitment regime.

Next, consider the seller's payoff. Since F^{MAX} first order stochastically dominates F^{MIN} , it holds $F^{MAX}(p) \leq F^{MIN}(p)$, and thus $\delta F^{MAX}(p) + (1 - \delta)F^{MIN}(p)$ is decreasing in δ for any p . Thus, $p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$ is increasing in δ for any p . Then, $\max_p p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$ is maximized at $\delta = 1$. Therefore, the seller is better off under the commitment regime.

C Proof of [Proposition 1](#)

Given δ , the set of the optimal prices is $P(\delta) := \arg \max_p p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)]$. By the tie-breaking rule, the seller sets a price of $p(\delta) := \min P(\delta)$, which is well-defined as shown in [Appendix H](#). I show that $p(\delta)$ is increasing in δ . Note that

$$\begin{aligned} & \log p[1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)] - \log p[1 - \delta' F^{MAX}(p) - (1 - \delta')F^{MIN}(p)] \\ &= \log \frac{1 - \delta F^{MAX}(p) - (1 - \delta)F^{MIN}(p)}{1 - \delta' F^{MAX}(p) - (1 - \delta')F^{MIN}(p)}. \end{aligned} \quad (2)$$

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, (2) 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. Therefore, $p(\delta)$ is increasing in δ .

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

In the alternative interpretation of the model, information disclosure by some consumers lowers the welfare of other consumers. To see this, note that if each consumer $i \in [0, 1]$ chooses δ_i and the seller 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 the commitment regime is $p(\bar{\delta})$, where $p(\cdot)$ is defined by (1). Now, if a positive mass of consumers chooses strictly greater disclosure levels, then $\bar{\delta}$ increases. This increases $p(\bar{\delta})$ and decreases the payoffs of other consumers who have not changed disclosure levels.

In contrast, under the no-commitment regime, consumer i is charged a price of $p(\delta_i)$ for recommended products. Thus, each consumer’s problem is identical with the one in the original formulation. Thus, consumers disclose less information and are better off under the no-commitment regime.

E Proof of Proposition 2

Take any equilibrium under the commitment regime. Since prices are fixed and the same across products, it is optimal for the consumer to disclose information so that the seller recommends the most valuable products with probability 1.²⁹ Now, the seller sets a price of product k' to maximize the expected revenue conditional on recommending product k' . Conditional on this event, the seller’s posterior belief for $u_{k'}$ is equal to the distribution of $\max_{k \in \mathcal{K}} u_k$, denoted by x^{MAX} . Note that the argument does not depend on k' because values are IID across products. Then, it holds that

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

²⁹For the sake of completeness, I present an example. Consider disclosure rule (ϕ^*, M^*) such that $M^* = \mathcal{K}$ 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^* is a symmetric disclosure rule that reveals the name of the most valuable product. An equilibrium disclosure rule is not unique; we can consider any disclosure rules weakly more informative than (ϕ^*, M^*) .

where the last inequality follows from [Assumption 1](#). Thus, the price for each product is strictly greater than v_1 , and the consumer buys no products with a probability of at least $x(v_1)^K > 0$.

F Proof of [Proposition 3](#)

[Proposition 3](#) follows from a series of lemmas. [Lemma 4](#) proves vertical efficiency. [Lemma 5](#) proves that any equilibrium is horizontally inefficient whenever x^{MAX} has a unique monopoly price. [Lemma 6](#) proves that this condition is true for generic x_0 .

Lemma 4. *Under the no-commitment regime, any equilibrium is vertically efficient.*

Proof. Take any disclosure rule (M^*, ϕ^*) that leads to a vertically inefficient allocation given the seller's best response and the consumer's optimal purchase decision with the tie-breaking. (Hereafter, I omit the caveat "given the tie-breaking rule.") Then, ϕ^* draws a posterior $x \in \Delta(V^K)$ at which trade fails to occur with positive probability.³⁰ Without loss of generality, suppose that given x , the seller recommends product 1 at price v_n . Consider the following disclosure rule ϕ^{**} : On top of the information that ϕ^* discloses, ϕ^{**} also discloses $u_1 \geq v_n$ or $u_1 < v_n$ whenever posterior x is realized.

I show that ϕ^{**} yields a weakly greater consumer surplus and a strictly greater total surplus than ϕ^* does. Let x^+ and $x^- \in \Delta(V^K)$ denote the posterior beliefs of the seller when the consumer discloses $u_1 \geq v_n$ and $u_1 < v_n$ (following x), respectively. Note that for some $\alpha \in (0, 1)$, $x = \alpha x^+ + (1 - \alpha)x^-$. First, consider the consumer's payoff and total surplus conditional on x^- . The consumer obtains a greater payoff under ϕ^{**} than under ϕ^* because consumer surplus is zero under ϕ^* . Total surplus is strictly greater under ϕ^{**} because trade occurs with a positive probability under ϕ^{**} but occurs with zero probability under ϕ^* . Second, I show that the seller continues to recommend product 1 at price v_n given x^+ . Suppose to the contrary that the seller strictly prefers to recommend product m at price v_ℓ where $(m, \ell) \neq (1, n)$. Let $x_1^+ \in \Delta(V)$ and $x_m^+ \in \Delta(V)$ denote the marginal distributions of u_1 and u_m given x^+ , respectively. Because the seller strictly

³⁰Because $|V^K| < +\infty$, without loss of generality, I can assume $|M^*| < +\infty$. Then, each message is realized with a positive probability from the ex-ante perspective. This implies that there is an ex-ante positive probability event such that some posterior $x \in \Delta(V^K)$ is realized and trade fails to occur.

prefers recommending product m at price v_ℓ to recommending product 1 at price v_n , we get

$$v_\ell \sum_{j=\ell}^N x_m^+(v_j) > v_n \sum_{j=n}^N x_1^+(v_j),$$

which implies

$$v_\ell \sum_{j=\ell}^N [\alpha x_m^+(v_j) + (1 - \alpha)x_m^-(v_j)] \geq v_\ell \sum_{j=\ell}^N \alpha x_m^+(v_j) > v_n \sum_{j=n}^N \alpha x_1^+(v_j) = v_n \sum_{j=n}^N [\alpha x_1^+(v_j) + (1 - \alpha)x_1^-(v_j)]. \quad (3)$$

The last equality follows from $x_1^-(v) = 0$ for any $v \geq v_n$. [Inequality \(3\)](#) contradicts the fact that the seller prefers to recommend product 1 at price v_n at x . Thus, the seller continues to recommend the same product at the same price between x and x^+ . Overall, disclosing $u_1 \geq v_n$ or $u_1 < v_n$ at x leads to a weakly greater consumer surplus and a strictly greater total surplus.

To show that any equilibrium is vertically efficient, take any equilibrium disclosure rule ϕ^* . [Appendix H](#) proves the existence of an equilibrium. Suppose to the contrary that ϕ^* is vertically inefficient. Then, I can apply the modification described above to create ϕ^{**} . ϕ^{**} gives the consumer a weakly greater payoff than ϕ^* . Because ϕ^* is optimal for the consumer, he is indifferent between ϕ^* and ϕ^{**} . However, ϕ^{**} yields a strictly greater total surplus, which implies that the seller strictly prefers ϕ^{**} . This contradicts the tie-breaking rule, which requires that ϕ^* maximizes the seller's payoffs among all consumer-optimal disclosure rules. Therefore, ϕ^* is vertically efficient. \square

Lemma 5. *Suppose that the prior distribution x^0 satisfies [Assumption 1](#) and there is a unique monopoly price given value distribution F^{MAX} , which is the CDF of $\max(u_1, \dots, u_K)$ where each u_k is an IID draw from x^0 . Then, any equilibrium is horizontally inefficient.*

Proof. I construct a disclosure rule that maximizes the consumer's ex ante expected payoff among $\mathcal{E} \subset \mathcal{D}$, where \mathcal{E} is the set of all disclosure rules that lead to horizontally efficient outcomes given the optimal behavior of each player. Take any disclosure rule $\phi \in \mathcal{E}$. Since the recommended product belongs to $\arg \max_{\ell \in \mathcal{K}} u_\ell$ with (ex ante) probability 1, the consumer's value of the recommended product (unconditional on which product is recommended) is drawn according to F^{MAX} . This implies that under ϕ , the seller can obtain a revenue of at least $\underline{R} := \max_{p \in V} p[1 - F^{MAX}(p)]$

by setting a price of $\arg \max_{p \in V} p[1 - F^{MAX}(p)]$ for all realized posteriors. This implies that if ϕ^* achieves an efficient allocation and gives the seller a payoff of \underline{R} , then ϕ^* maximizes the consumer's payoff among \mathcal{E} .

Consider disclosure rule $\phi^E \in \mathcal{E}$ such that for any realized $u \in V^K$, ϕ^E draws message $k \in \arg \max_{\ell} u_{\ell}$ with probability $\frac{1}{|\arg \max_{\ell} u_{\ell}|}$. Two remarks are in order. First, u_k is distributed according to F^{MAX} conditional on message k . Second, the seller prefers to recommend product k after observing message k no matter what additional information she learns, because she can maximize the probability of trade by recommending product k .

Next, I create $\phi^* \in \mathcal{E}$ by modifying ϕ^E as follows: For each $k \in \mathcal{K}$, conditional on that message k is realized under ϕ^E , ϕ^* discloses additional information about u_k according to a *consumer surplus maximizing segmentation (CSMS)* characterized by [Bergemann et al. \(2015\)](#).³¹ In our context, the information disclosed according to (any) CSMS ensures that the trade occurs with probability 1 whereas the seller's resulting revenue is \underline{R} . Thus, under ϕ^* , the seller recommends the highest-value product and the trade occurs with probability 1, whereas the seller's revenue is \underline{R} . Thus, ϕ^* maximizes the consumer's payoff among \mathcal{E} .

Hereafter, I focus on a particular ϕ^* where the additional information about the highest-value product is disclosed according to a CSMS constructed by the *greedy algorithm* in [Bergemann et al. \(2015\)](#). This has the following implication. Let $\{x_{S_1}^k, \dots, x_{S_L}^k\}$ denote the set of posteriors induced by ϕ^* conditional on ϕ^E drawing message k . Without loss of generality, regard $x_{S_1}^k, \dots, x_{S_L}^k$ as messages drawn by ϕ^* . Let us also regard each $x_{S_{\ell}}^k$ as a marginal distribution of u_k instead of a joint distribution of (u_1, \dots, u_K) . The greedy algorithm guarantees that each $x_{S_{\ell}}^k$ has support $S_{\ell} \subset V$, $S_1 \subset S_2 \subset \dots \subset S_L = V$, and the set of all optimal prices against $x_{S_{\ell}}^k$ is S_{ℓ} . Moreover, it holds that $S_1 = \{v^*\}$ with $v^* > v_1$. To see this, note that $|S_1| \geq 2$ implies that two prices in S_1 are optimal against all posteriors in $\{x_{S_1}^k, \dots, x_{S_L}^k\}$, which in turn implies that these prices are optimal under F^{MAX} because the expected revenue is linear in the value distribution. This contradicts the assumption that there is a unique optimal price under F^{MAX} . Thus, $|S_1| = 1$, which implies that $S_1 = \{v^*\}$. Following the proof of [Proposition 2](#), we can show that [Assumption 1](#) implies $v^* > v_1$.

³¹In single product monopoly pricing, a consumer surplus maximizing segmentation is equivalent to a disclosure rule that has the following property. First, at each realized posterior, the seller is willing to set the price equal to the minimum of its support, which implies that the trade occurs with probability 1. Second, at each posterior, the seller is indifferent between charging the minimum of each posterior and charging the monopoly price for the prior.

I modify ϕ^* to create a horizontally inefficient ϕ^I that yields a strictly greater consumer surplus than ϕ^* . From now on, I treat each $x_{S_L}^k$ as a joint distribution of (u_1, \dots, u_K) . To simplify exposition, I use the following terminologies. First, I regard a distribution $x \in \Delta(V^K)$ as consisting of a unit mass of consumers, where mass $x(u)$ of consumers have value vector u . Second, I call any set of (a continuum of) consumers a “segment.”

To construct ϕ^I , I make three observations. First, a positive mass of consumers in $x_{S_1}^1$ has value v^* for product 1 and the lowest possible value $v_1 < v^*$ for product 2. Call this mass of consumers “segment (v^*, v_1) .” Second, a positive mass of consumers in $x_{S_1}^1$ has value v^* for both products 1 and 2. Call these consumers “segment (v^*, v^*) .”

First, I take a small but positive (say ε_1) mass of segment (v^*, v_1) from $x_{S_1}^1$ and pool this segment with $x_{S_L}^2$.³² Let $\hat{x}_{S_L}^2$ denote the posterior created by this pooling. For a sufficiently small $\varepsilon_1 > 0$, at $\hat{x}_{S_L}^2$, the seller recommends product 2, and she *strictly* prefers to set price v_1 for product 2. The reason is as follows. Under the original posterior $x_{S_L}^2$, it is optimal for the seller to recommend product 2 at any price in V because $S_L = V$. After the modification, $\hat{x}_{S_L}^2$ contains a strictly greater mass of consumers who have value v_1 for product 2 (i.e., segment (v^*, v_1)). Thus, the seller strictly prefers to set price v_1 for product 2. Moreover, for a small $\varepsilon_1 > 0$, the seller does not strictly prefer to recommend other products. Indeed, if the seller recommended product $k \neq 2$ at $x_{S_L}^2$, then she would strictly prefer to set price v_1 . Thus, for a small $\varepsilon_1 > 0$, the seller’s pricing incentive does not change under $\hat{x}_{S_L}^2$. This implies that the optimal revenue from recommending other products (at the new posterior $\hat{x}_{S_L}^2$) is v_1 , which is no greater than the revenue from recommending product 2. Importantly, this modification does not change the consumer’s payoff, because consumers in segment (v^*, v_1) obtain zero payoffs under $x_{S_1}^1$. Let ϕ^H denote the resulting disclosure rule.

Finally, I modify ϕ^H by pooling a small but positive (say ε_2) mass of segment (v^*, v^*) in $x_{S_1}^1$ with $\hat{x}_{S_L}^2$. Let $\tilde{x}_{S_L}^2$ denote the posterior following this pooling. If ε_2 is small, the seller continues to recommend product 2 at price v_1 under $\tilde{x}_{S_L}^2$, because she strictly prefers to set price v_1 for product 2 at $\hat{x}_{S_L}^2$. This modification strictly increases the consumer’s payoff relative to ϕ^* , because consumers in segment (v^*, v^*) obtain a positive payoff $v^* - v_1$ while they obtain zero payoff under ϕ^* . Let ϕ^I denote the resulting disclosure rule.

³²In terms of a disclosure rule, this means that I modify ϕ^* so that it draws message $x_{S_L}^2$ with probability $\varepsilon_1 > 0$ not only following message 2 but also when ϕ^* draws segment (v^*, v_1) in $x_{S_1}^1$.

ϕ^I gives the consumer a strictly greater expected payoff than ϕ^* but leads to an inefficient recommendation at $\tilde{x}_{S_L}^2$. This implies that any equilibrium is horizontally inefficient, because for any disclosure rule leading to horizontal efficiency, the consumer can find strictly more profitable disclosure rules that lead to horizontal inefficiency. \square

Lemma 6. *Fix a finite support $V \subset \mathbb{R}_+$ with $|V| \geq 2$. There is a Lebesgue measure zero set $X_0 \subset \Delta(V)$ such that, for any $x_0 \in \Delta(V) \setminus X_0$, the induced F^{MAX} has a unique monopoly price (i.e., F^{MAX} satisfies the condition of [Lemma 5](#)).*

Proof. First, let D_2 denote the set of all distributions $x \in \Delta(V)$ such that there are two or more optimal prices. I show that D_2 has zero measure. Let $D_{v,v'}$ denote the set of all distributions in $\Delta(V)$ such that both prices v and v' are optimal. I can write $D_{v,v'} = \{x \in \Delta(V) : v \sum_{v_n \geq v, v_n \in V} x(v_n) = v' \sum_{v_n \geq v', v_n \in V} x(v_n)\}$. $D_{v,v'}$ is a subset of $N-1$ -dimensional hyperplane, which has zero measure in \mathbb{R}^N . Thus, $D_2 = \cup_{(v,v') \in V^2} D_{v,v'}$ has zero measure.

Consider a function φ that maps any distribution $x = (x_1, \dots, x_N) \in \Delta(V)$ to the distribution of $\max(u_1, \dots, u_K)$, where each u_k is an IID draw from x . φ is written as follows:

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

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

$$\{(x_1, \dots, x_N) \in \Delta(V) : \exists n, x_n = 0\}. \quad (4)$$

By Theorem 1 of [Ponomarev \(1987\)](#), $\varphi : \mathbb{R}^N \rightarrow \mathbb{R}^N$ has the “0-property”: The inverse image of

any measure zero set by φ has zero measure. In particular, $X_0 := \varphi^{-1}(D_2)$ has zero measure. Clearly, X_0 has the desired property. \square

G Proof of Proposition 4

Proposition 4 relies on the following lemma.

Lemma 7. *Under the commitment regime, as $K \rightarrow +\infty$, the seller's equilibrium payoff converges to $\max V$ and the consumer's equilibrium payoff converges to 0. Under the no-commitment regime, if Assumption 2 holds, then there is $\underline{u} > 0$ such that the consumer's equilibrium payoff is at least \underline{u} for any K .*

Proof. By the same argument as Proposition 2, the seller under the commitment regime recommends the most valuable product with probability 1. Let F denote the CDF of the value for each product (induced by x_0). Take any $\varepsilon > 0$. Suppose that the seller sets $p = \max V - \varepsilon/2$ for each product up front. As $K \rightarrow +\infty$, the probability $1 - F(p)^K$ that the consumer buys the recommended product goes to 1. Thus, there is \underline{K} such that the seller's revenue is at least $\max V - \varepsilon$ 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 the no-commitment regime, observe that the consumer can choose to disclose no information and obtain a payoff of $\int_{p(x_0)}^{\max V} v - p(x_0) dF(v) > 0$, which is positive and independent of K . \square

Proof of Proposition 4. The result under the commitment regime follows from the previous result, as total surplus is weakly greater than the seller's revenue.

I show that total surplus under the no-commitment regime is uniformly bounded away from $\max V$. Suppose to the contrary that for any $n \in \mathbb{N}$, there exists K_n such that when the seller sells K_n products, some equilibrium under the no-commitment regime achieves total surplus of at least $\max V - \frac{1}{n}$. Then, I can take a subsequence $(K_{n_\ell})_\ell$ such that $K_{n_\ell} < K_{n_{\ell+1}}$ for any $\ell \in \mathbb{N}$. Next, I show that for any $p < \max V$ and $\varepsilon < 1$, there exists $\ell^* \in \mathbb{N}$ such that for any $\ell \geq \ell^*$,

$$\mathbb{P}_\ell(\text{the consumer's value for the recommended product} \geq p) \geq \varepsilon. \quad (5)$$

where $P_\ell(\cdot)$ is the probability measure on the consumer's value for the recommended product in equilibrium of K_{n_ℓ} -product model. To show [inequality \(5\)](#), suppose to the contrary that there is some (p, ε) and a subsequence $(K'_m)_m$ of $(K_{n_\ell})_\ell$ such that the inequality is violated. Then, given any K'_m in this subsequence, the total surplus is at most $\varepsilon p + (1 - \varepsilon) \max V < \max V$. This contradicts the assumption that the equilibrium total surplus converges to $\max V$ as $K'_m \rightarrow +\infty$.

Now, I use [inequality \(5\)](#) to show that the seller's equilibrium revenue converges to $\max V$ along $(K_{n_\ell})_\ell$. Take any $r < \max V$. If the seller sets price $\frac{r + \max V}{2}$, then for a sufficiently large ℓ , the consumer accepts the price with probability greater than $\frac{2r}{r + \max V} < 1$. That is, for a large ℓ , the seller's expected revenue exceeds r . Since this holds for any $r < \max V$, the seller's revenue converges to $\max V$ as $\ell \rightarrow +\infty$. This contradicts the observation that the consumer's payoff is bounded from below by a positive number independent of K , as in [Lemma 7](#). \square

H Existence of Equilibrium under the No-commitment Regime

I prove the existence of an equilibrium under the no-commitment regime. Recall that for the commitment regime, I have proved the existence by explicitly constructing an equilibrium.

Restricted Model

Claim 1. In the restricted model, there exists an equilibrium under the no-commitment regime.

The result follows from two lemmas.

Lemma 8. *Given a disclosure level δ , the lowest optimal price $p(\delta)$ in (1) exists and is lower semicontinuous in δ .³³*

Proof. Define $G(p) := \delta F^{MAX}(p) + (1 - \delta)F^{MIN}(p)$. Recall that I define a CDF as a left-continuous function. Take any p^* and c with $G(p^*) > c$. For some $\delta > 0$ and for all $p \in (p^* - \delta, p^*]$, $G(p) > c$. Since $G(p)$ is increasing in p , for all $p \in (p^* - \delta, p^* + \delta)$, $G(p) > c$. Thus, G is lower semicontinuous. Then, $1 - G$ is upper semicontinuous, and thus $p[1 - G(p)]$, which is a product

³³Depending on the context, I use one of the following two equivalent conditions as a definition of lower semicontinuity. Given a (first countable) topological space X and $f : X \rightarrow \mathbb{R}$, f is lower semicontinuous if for each $c \in \mathbb{R}$ the set $\{x \in X : f(x) \leq c\}$ is closed (or equivalently, the set $\{x \in X : f(x) > c\}$ is open). Equivalently, f is lower semicontinuous if $x_n \rightarrow x$ implies $\liminf_n f(x_n) \geq f(x)$. For the equivalence of the two conditions, see Lemma 2.42 of [Aliprantis and Border \(2006\)](#). f is upper semicontinuous if $-f$ is lower semicontinuous.

of two nonnegative upper semicontinuous functions, is upper semicontinuous in p .³⁴ This implies that $P(\delta) := \arg \max_{p \in V} p[1 - G(p)]$ is nonempty and compact (Theorem 2.43 of Aliprantis and Border (2006)). Thus, $p(\delta) := \min P(\delta)$ exists.

Next, suppose to the contrary that $p(\delta)$ is not lower semicontinuous at some δ^* . Then, there is $\varepsilon > 0$ such that we can construct a sequence $\delta_n \rightarrow \delta^*$ so that $p(\delta_n) < p(\delta^*) - \varepsilon$ for all n .³⁵ Then, we can find a convergent subsequence of $(p(\delta_n))_n$ because $p(\delta_n) \in V$ and V is compact. Without loss of generality, assume that $(p(\delta_n))_n$ itself converges, so that there exists $p^* := \lim_n p(\delta_n) < p(\delta^*)$. Define $Y(p, \delta)$ as

$$Y(p, \delta) := p [1 - \delta F^{MAX}(p) - (1 - \delta) F^{MIN}(p)] - p(\delta^*) [1 - \delta F^{MAX}(p(\delta^*)) - (1 - \delta) F^{MIN}(p(\delta^*))].$$

Because $p(\delta_n)$ is optimal given δ_n , it holds $Y(p(\delta_n), \delta_n) \geq 0$. Also, $Y(p, \delta)$ is upper semicontinuous in (p, δ) .³⁶ This implies that $Y^* := \{(p, \delta) : Y(p, \delta) \geq 0\}$ is closed. Thus, $(p^*, \delta^*) = \lim_n (p(\delta_n), \delta_n) \in Y^*$, or equivalently,

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

which implies $p^* \in P(\delta^*)$. This contradicts $p(\delta^*) = \min P(\delta^*)$ because $p^* < p(\delta^*)$. Therefore, $p(\delta)$ is lower semicontinuous. \square

Lemma 9. $\delta u^{MAX}(p(\delta)) + (1 - \delta) u^{MIN}(p(\delta))$ is upper semicontinuous in δ .

Proof. $u^{MAX}(p) = \int_p^{+\infty} (x - p) dF^{MAX}(x) = \int_p^{+\infty} 1 - F^{MAX}(x) dx$ is continuous and decreasing in p . Since $p(\delta)$ is lower semicontinuous, $u^{MAX}(p(\delta))$ is upper semicontinuous in δ . (To see this, if g is continuous and decreasing, and f is lower semicontinuous, then for $x_n \rightarrow x$,

³⁴To see this, if $f : X \rightarrow \mathbb{R}$ and $g : X \rightarrow \mathbb{R}$ are nonnegative and upper semicontinuous, for any $x_n \rightarrow x$, we obtain $\limsup_n f(x_n)g(x_n) \leq \limsup_n f(x_n) \limsup_n g(x_n) \leq f(x)g(x)$. Thus, fg is upper semicontinuous.

³⁵Formally, if $p(\cdot)$ is not lower semicontinuous, then for some c the set $S := \{\delta : p(\delta) > c\}$ is not open. This implies that we can take some $\delta^* \in S$ such that there is $\delta_n \rightarrow \delta^*$ with $p(\delta_n) \leq c < p(\delta^*)$. Define $\varepsilon := \frac{p(\delta^*) - c}{2}$, then $p(\delta_n) < p(\delta^*) - \varepsilon$.

³⁶This follows from the fact that the product of two nonnegative upper (lower) semicontinuous functions is upper (lower) semicontinuous, and the same thing holds for the sum. Since $F^{MAX}(p)$ and $F^{MIN}(p)$ are lower semicontinuous in p , $\delta F^{MAX}(p) + (1 - \delta) F^{MIN}(p)$ is lower semicontinuous. Thus, $p [1 - \delta F^{MAX}(p) - (1 - \delta) F^{MIN}(p)]$ is upper semicontinuous. Because the second term of $Y(p, \delta)$, $p(\delta^*) [1 - \delta F^{MAX}(p(\delta^*)) - (1 - \delta) F^{MIN}(p(\delta^*))]$, is continuous in (p, δ) , overall, $Y(p, \delta)$ is upper semicontinuous.

we have $\lim_{k \rightarrow +\infty} \sup_{n > k} g(f(x_n)) = \lim_{k \rightarrow +\infty} g(\inf_{n > k} f(x_n)) = g(\lim_{k \rightarrow +\infty} \inf_{n > k} f(x_n)) \leq g(f(x))$. The last inequality is from the lower semi-continuity of f .) Similarly, we can show that $u^{MIN}(p(\delta))$ is upper semicontinuous. Therefore, $\delta u^{MAX}(p(\delta)) + (1 - \delta)u^{MIN}(p(\delta))$ is upper semicontinuous in δ . \square

Proof of Claim 1. $\delta u^{MAX}(p(\delta)) + (1 - \delta)u^{MIN}(p(\delta))$ is upper semicontinuous, and the set of disclosure levels, $[1/2, 1]$, is compact. Thus, $D^* := \arg \max_{\delta \in [1/2, 1]} \delta u^{MAX}(p(\delta)) + (1 - \delta)u^{MIN}(p(\delta))$ is nonempty and compact (Theorem 2.43 of [Aliprantis and Border \(2006\)](#)). Thus, $\delta^* := \max D^*$ combined with the optimal on and off path actions consists of an equilibrium. \square

Unrestricted Model

Claim 2. In the unrestricted model, there exists an equilibrium under the no-commitment regime.

Proof. I prepare some notations. Let $A := \mathcal{K} \times V$ denote the seller's (finite) action space, i.e., the set of all pairs of recommended products and prices. When the seller recommends product $k \in \mathcal{K}$ at price $p \in V$, I say that the seller chooses $a = (k, p) \in A$. Given $(a, b) \in A \times \Delta(V^K)$, let $U(a, b)$ and $R(a, b)$ denote the expected payoffs of the consumer and the seller, respectively, when the seller chooses a , the consumer's values are drawn from b , and the consumer takes an optimal purchase decision (breaking ties in favor of the seller). Given the seller's belief $b \in \Delta(V^K)$, let $a(b) \in A$ denote the seller's optimal recommendation and price that break ties in favor of the consumer. $a(b)$ exists because A is finite.

The proof consists of two steps. First, I show that the consumer's payoff is upper semicontinuous in disclosure rules. Second, I show that the set of all disclosure rules is compact. I use weak* topology in $\Delta(\Delta(V^K))$.

Consider an information set where the seller sets a price and recommends a product. Let $b \in \Delta(V^K)$ denote the seller's belief about the value vector. If the seller and the consumer take optimal actions following the information set, the consumer's expected payoff is given by $U(a(b), b)$. I show that $U(a(b), b)$ is upper semicontinuous in $b \in \Delta(V^K)$. Suppose to the contrary that there exists $\varepsilon > 0$ and $(b_n)_{n=1}^{+\infty} \subset \Delta(V^K)$ such that $\lim_n b_n = b$ but $U(a(b_n), b_n) \geq U(a(b), b) + \varepsilon$ for all n . Because A is finite, we can choose a subsequence $(b_{n(m)})_{m=1}^{+\infty}$ so that for some $a' \in A$, $a(b_{n(m)}) = a'$ for all m . Without loss of generality, assume that $a(b_n) = a'$ for all n . Note that

$R(a', b_n) \geq R(a(b), b_n)$ because $a' = a(b_n)$ is optimal for the seller given its belief b_n . Note also that $R(a, b)$ is continuous in b with a fixed a . Indeed, suppose that a is such that the seller recommends product k at price p , where the consumer's value for product k is distributed according to $b^k = (b_1^k, \dots, b_N^k) \in \Delta(V)$ under $b \in \Delta(V^K)$. Then, $R(a, b) = p \cdot \sum_{\ell=1}^N \mathbf{1}_{\{v_\ell \geq p\}} b_\ell^k$, which is continuous in b . ($\mathbf{1}_{\{v_\ell \geq p\}}$ is the indicator function that takes value 1 or 0 if $v_\ell \geq p$ or $v_\ell < p$, respectively.) Given the continuity of $R(a, b)$ in b , $R(a', b_n) \geq R(a(b), b_n)$ for all n implies $R(a', b) \geq R(a(b), b)$. Thus, a' is optimal for the seller given b . This implies that $U(a(b), b) \geq U(a', b)$ by the seller's tie-breaking rule. Also, $U(a(b_n), b_n) \geq U(a(b), b) + \varepsilon$ for all n implies $U(a', b) \geq U(a(b), b) + \varepsilon$, because $a(b_n) = a'$, and $U(a, b) = \sum_{\ell=1}^N \mathbf{1}_{\{v_\ell \geq p\}} (v_\ell - p) b_\ell^k$ is continuous in b . However, these two inequalities lead to $U(a(b), b) \geq U(a', b) \geq U(a(b), b) + \varepsilon$, which is a contradiction. Thus, $U(a(b), b)$ is upper semicontinuous in $b \in \Delta(V^K)$. By Theorem 15.5 of [Aliprantis and Border \(2006\)](#), $\int_{\Delta(V^K)} U(a(b), b) d\tau(b)$ is upper semicontinuous in $\tau \in \Delta(\Delta(V^K))$ when $\Delta(\Delta(V^K))$ is endowed with weak* topology. This completes the first part.

Next, I show that the set \mathcal{D} of all disclosure rules is weak* compact. Let $b_0 := x_0 \times \dots \times x_0$ denote the prior distribution of value vector. Also, for any disclosure rule $\phi \in \mathcal{D}$, define $\hat{\phi} \in \Delta(\Delta(V^K))$ as the distribution over posterior beliefs about value vector u induced by ϕ and b_0 . Moreover, let $\hat{\mathcal{D}} := (\hat{\phi})_{\phi \in \mathcal{D}} \subset \Delta(\Delta(V^K))$. Proposition 1 (iii) \Rightarrow (ii) of [Kamenica and Gentzkow \(2011\)](#) implies that if the consumer's payoff $\int_{\Delta(V^K)} U(a(b), b) d\tau(b)$ is maximized at some $\tau^* \in \hat{\mathcal{D}} = \left\{ \tau \in \Delta(\Delta(V^K)) : \int_{\Delta(V^K)} b d\tau(b) = b_0 \right\}$, then there is a disclosure rule (with a finite message space) that maximizes his payoff among all available disclosure rules. Now, $\Delta(\Delta(V^K))$ is weak* compact because $\Delta(V^K)$ is compact (e.g., Theorem 15.11 of [Aliprantis and Border \(2006\)](#)). Also, $\hat{\mathcal{D}}$ is closed. This is because if $\tau_\ell \in \hat{\mathcal{D}}$ for each $\ell \in \mathbb{N}$ and $\tau_\ell \rightarrow \tau$ in weak* topology, then $b_0 = \int_{\Delta(V^K)} b d\tau_\ell(b) \rightarrow \int_{\Delta(V^K)} b d\tau(b)$, which implies $\tau \in \hat{\mathcal{D}}$. Thus, $\hat{\mathcal{D}}$, which is a closed subset of a compact set $\Delta(\Delta(V^K))$, is weak* compact.

Finally, in equilibrium, the consumer solves $\max_{\tau \in \hat{\mathcal{D}}} \int_{\Delta(V^K)} U(a(b), b) d\tau(b)$. Since the objective function is upper semicontinuous in τ and $\hat{\mathcal{D}} \subset \Delta(\Delta(V^K))$ is compact with respect to weak* topology, the set $\mathcal{D}^* := \arg \max_{\tau \in \hat{\mathcal{D}}} \int_{\Delta(V^K)} U(a(b), b) d\tau(b)$ is nonempty and weak* compact. Moreover, as I prove below, the seller's expected payoff $\int_{\Delta(V^K)} R(a(b), b) d\tau(b)$ is upper semicontinuous in τ . Therefore, $\max_{\tau \in \mathcal{D}^*} \int_{\Delta(\Delta(V^K))} R(a(b), b) d\tau(b)$ has a maximizer. Any maximizer τ^* , combined with the optimal on and off path behavior of the seller and the consumer, consists of an

equilibrium.

To see that $\int_{\Delta(\Delta(V^K))} R(a(b), b) d\tau(b)$ is upper semicontinuous, I show that $R(a(b), b)$ is upper semicontinuous in b . Suppose to the contrary that there exists $\varepsilon > 0$ and $(b_n)_{n=1}^{+\infty} \subset \Delta(V^K)$ such that $\lim_n b_n = b$ but $R(a(b_n), b_n) \geq R(a(b), b) + \varepsilon$ for all n . Because A is finite, we can choose a subsequence $(b_{n(m)})_{m=1}^{+\infty}$ so that for some $a' \in A$, $a(b_{n(m)}) = a'$ for all m . As $R(a, b)$ is continuous in b , we obtain $R(a', b) \geq R(a(b), b) + \varepsilon$. However, this contradicts $R(a(b), b) \geq R(a', b)$. Thus, $R(a(b), b)$ is upper semicontinuous in b , and thus $\int_{\Delta(\Delta(V^K))} R(a(b), b) d\tau(b)$ is upper semicontinuous by Theorem 15.5 of [Aliprantis and Border \(2006\)](#). \square