Essays on Digital Markets

Shuoqi Sun

Submitted in partial fulfillment of the requirements for

the degree of Doctor of Philosophy in Economics

Tepper School of Business

Carnegie Mellon University

November, 2025

Acknowledgement

Contents

Chapter 1

Privacy Choice in Personalized Search markets

Abstract

We study consumer's incentives to give up their privacy in exchange for better personalization in a competitive search marketplace. The platform acts as an information intermediary, who personalizes consumer's search environment at the consent of consumers. Consumers, heterogeneous in search costs, are facing the following trade-off: enabling cookies help them find better match, but it also shapes the degree of price competition among sellers which could be harmful to consumers. I show that no privacy for all consumers improves match efficiency, expands market demands, and despite its ambiguous effect on equilibrium price, a platform should enable cookies for everyone, while pooling everyone's cost type to maximize consumer welfare. However, when leaving the choice to consumers, a separating equilibrium may arise, in which consumers with low search cost remain anonymous in order to separate themselves from high search cost consumers. Compared to the consumer optimal outcome, almost all consumers are worse off, suggesting that the prevalent consent-based approach to privacy regulation may lead to unintended welfare loss.

1.1 Introduction

The advances of information technologies have fundamentally reshaped the collection, storing, and sharing of personal data, which is essential capital to business activities. Nowadays these data are extensively used by online marketplaces for personalized recommendations, targeted advertising, and improvement of user experience. Concurrent with these developments, consumers have been increasingly concerned about their personal information being intruded or misused.

In response to consumers' privacy concerns, governments and firms have already taken steps to protect privacy, such as the General Data Protection Regulation (GDPR) of EU and the California Consumer Privacy Act (CCPA) of the state of California. Apple, among other big tech companies, also raised lots of attention for their pro-privacy moves. In 2021, Apple launched App Tracking Transparency Initiative which requires every app to seek users' explicit consent before tracking them across other apps and websites. Alternatively, users can reject all tracking in advance by disabling "Allowing Apps to request to track" in system preferences. A report from Flurry Analytics(Laziuk (2021)) suggests that only a fraction of 4% of iOS users have disabled data tracking on a system level, and around 40% of users allow cookies app-by-app.

At the front and center of all these moves are the take-it-or-leave-it consent requirements that give consumers some control over whether and how their personal information are to be utilized. These policies aim to protect privacy on the ground of property rights. However, the real consequences of these consent requirements are not well understood. Who benefits and who suffers when individuals can decide whether to give up their privacy? Is it indeed the case that every consumer benefits from it? How does consumers' control over their privacy affect market competition and thereby consumer welfare? These questions are also relevant for the design of regulation policy. For example, if consumer control over privacy does not guarantee a Pareto-improving outcome all by itself, could there be any policy instruments to complement it and restore the welfare loss?

In this paper I aim to investigate these questions in a search framework where consumers' search environment is personalized by an online shopping platform, whose goal is to maximize consumer surplus. Personalization is only possible if the consumer allows the platform to access her personal preference information from some outset. Sellers offer horizontally differentiated products on the platform and engage in price competition. They are not ex ante informed about consumers' match values, but they may partially infer that from the platform's personalization policy and the consumer's equilibrium privacy choice. Consumers are facing the following trade-off: enabling cookies help them find better match, while this may also signal their 'type', which would be fully internalized by the competing sellers when setting prices.

To illustrate heterogeneous privacy choices in equilibrium, I further introduce consumer heterogeneity in terms of search costs (hereafter referred to as a consumer's "type") into the model. Search costs could be interpreted as a cognitive or attention cost to inspect the product information, or the opportunity cost of time for such inspection. It's well observed that attention is scarce for online shoppers, though the degree of scarcity may be different across individuals. Our model demonstrates that such heterogeneity alone could be enough for revealed preference for privacy in equilibrium, intrinsic preference for privacy is not needed.

With a single consumer of a known type, the result is straightforward: better personalization encourages the consumer to participate in search, improve match efficiency, and thereby increases social surplus. Under some well-behaved specifications of the match distribution, this also leads to a lower equilibrium price. As a result, consumer surplus increases unambiguously as the level of personalization increases. I also study the social planner's problem, i.e. what if the platform knows every consumer's search cost and can segment the market by search cost in any way to maximize consumer surplus. It turns out that there is no way the social planner could improve total consumer surplus by segmenting the market.

Next I study the case with endogenous privacy choice. I show that, when the level of personalization is sufficiently high (but not too high to forego active search), consumers with lower search costs have incentive to separate themselves by disabling cookies. Such separating equilibrium results in loss of efficiency and lower consumer surplus compared to the optimal"centralized" outcome. Moreover, despite driven by the incentive to get higher surplus, among the consumers who voluntarily disable cookies, only a fraction of them actually get higher surplus compared to the centralized outcome. Banning price discrimination based on privacy choice does not automatically improve efficiency, since firms may coordinate on the equilibrium they prefer, which is often the consumer-worst equilibrium.

Our model suggests that it's possible to have consumers that value privacy purely for instrumental reasons and still derive economically intuitive and interesting results. This complements the popular debate that often views privacy as one's intrinsic rights.

The paper unfolds as follows. Section 1.2 summarizes the related literature. Section 1.3 presents the main model. In section 1.4.1, I study the case of exogenous privacy choice. I conduct comparative statics on social, consumer and industry welfare with respect to the level of personalization. Next in section 1.4.3, I take consumers' choice of privacy to be endogenous, and identify the conditions under which a separation equilibrium arises. The outcome is to be compared with the centralized outcome in section 1.4.1 in terms of consumer surplus. Section 1.6 contains discussion on equilibrium selection and policy implications. Finally, section 1.7 concludes the paper by identifying a couple of interesting ventures for future research.

1.2 Related Literature

This paper belongs to the literature on the economics of privacy, personalization, and their welfare implications. Within the privacy literature, two broad views prevail: privacy as an intrinsic good, in which consumers derive direct utility from withholding information, and privacy as an instrumental choice, in which the costs and benefits of disclosure are determined endogenously in equilibrium (see Acquisti et al. (2016) for a survey). This paper adopts the latter view. Consumers partially—or fully—internalize how their disclosure decision shapes sellers' behavior, which in turn feeds back into their own payoffs.

Most existing models of instrumental privacy treat consumers as ex ante identical,thereby eliminating equilibrium heterogeneity in privacy decisions (e.g., De Corniere and De Nijs (2016); Ichihashi (2020))—or allow heterogeneity only in willingness to pay. By contrast, we show that heterogeneity in search costs, a central determinant of consumer decision-making in many markets, can itself endogenously generate diverse privacy choices in equilibrium.

Notably, Ichihashi (2020) also studies a setting under which consumers choose a disclosure rule before observing her valuations. He considers a setup where a multi-product monopolist seller recommends one and only one product to a representative consumer, based on information collected from the consumer. The result would not hold if the consumer could sample more than one product on their own, which is the setup of this paper.

Hidir and Vellodi (2021) study consumers' incentives to voluntarily disclose preference information to a multiproduct monopolist in a cheap-talk environment. In contrast, Ali et al. (2020) examine consumers' incentives to disclose hard evidence about their vertical match values, anticipating that firms will infer willingness to pay in equilibrium. In both settings, consumers are privately informed about their valuations. By contrast, in my model consumers do not know their valuations ex ante. This distinction allows the analysis to cleanly isolate the role of search costs—rather than private information about valuations—in shaping consumers' privacy decisions.

The idea that platforms leverage consumer data to shape consumers' search environments also situates this paper in the literature on information design for online platforms (see, e.g., Armstrong and Zhou (2022); Au and Whitmeyer (2018); Board and Lu (2018); Bergemann et al. (2021)). This literature studies the optimal design of general information structures across diverse market environments and typically abstracts away from price competition, consumer heterogeneity, or search frictions. By contrast, we allow for price competition and heterogeneous consumers in a model with search frictions—-at the cost of restricting the feasible set of information policies. We adopt Zhong (2022)'s approach to platform personalization, which yields clean closed-form characterizations.

More broadly, this paper joins a growing literature on how the design of two-sided marketplaces' intermediary policies shapes participants' beliefs about the opposite side and, in turn, equilibrium behavior. Notably, Condorelli and Szentes (2022) study buyer-optimal matching policies when sellers are uninformed about consumer valuations and the matching algorithm can reveal information to them. They abstract away from price competition to focus on surplus redistribution within one-to-one matches, whereas our focus is on how policy-induced belief changes operate at the market level under price competition.

This paper builds on the consumer-search framework with product differentiation in Wolinsky (1986); Perloff and Salop (1985); Anderson and Renault (1999). In particular, it relates to work on consumer heterogeneity and inter-consumer externalities in search markets (see Anderson and Renault (2000); Gamp et al. (2018); Moraga-Gonz'alez et al. (2017), among others). A core insight in this literature is that better-informed consumers impose a negative externality on uninformed consumers: greater information increases horizontal differentiation, softens price competition, and raises prices. By contrast, in my setting the externality runs in the opposite direction: when some high search-cost consumers enter without privacy (i.e., share data), they impose a negative externality on incumbent low search-cost consumers.

Finally, the possibility that platforms condition prices on observed privacy choices links this paper to the extensive literature on the welfare consequences of price discrimination and market segmentation. Rather than characterizing outcomes under arbitrary information structures, we analyze a setting in which price discrimination is endogenously limited by consumers' strategic behavior. Two related contributions also study segmentation via privacy choice: Rhodes and Zhou (2021) and Mauring (2021). In Rhodes and Zhou (2021), privacy preferences are modeled as purely intrinsic; by contrast, our framework emphasizes how privacy choices interact with search costs and competition to shape equilibrium outcomes. In Mauring (2021), they explored the similar idea that firms could partially infer consumers' search cost from data tracking cookies and discriminate against search cost. However, their focus is on the effects of enabling data tracking on equilibrium price dispersion and abstract away from product differentiation.

1.3 The model

The model is built on a limiting version of the Wolinsky-Anderson-Renault model of random consumer search with horizontal product differentiation. A monopolist platform hosts an infinite number of sellers[1] indexed by j each supplying a product of zero marginal production cost. We assume the platform's objective is to maximize consumer welfare. On the other side of the platform, there is a unit mass of consumers indexed by i with unit demand. They must search on the platform in order to make a purchase.

Consumer i' utility from buying firm j's product is given by: $u_{i,j} = \theta_{i,j} - p_j$, where $\theta_{i,j}$ is a consumer-firm specific match value. They are independently and identically distributed across consumers and firms, drawn from a distribution with c.d.f F supported on $[0, 1]$. Sellers are not informed about the realizations of θ_{ij}.

[1]I work with infinite number of sellers to avoid the extra complication in demand function caused by heterogeneous consumers.

Assumption 1. *f is smooth, strictly positive and -1-concave on* $[0, 1]$.[2]

This assumption guarantees that the first order condition to the monopoly pricing problem $p(1 - F(p))$ has a unique solution. -1-concavity also implies the standard monotone hazard rate property: $\frac{f(\theta)}{1-F(\theta)}$ is non-decreasing.

Consumers are heterogeneous in terms of their search cost. Consumers are aware of their own type, whereas ex ante the platform and the firms only know the cumulative distribution. Assume that consumer i's search cost s_i is drawn identically and independently from a differentiable c.d.f distribution $G(\cdot)$ with support $(0, \bar{s}]$. Unless otherwise specified, we assume that the search cost distribution is uniform. This captures a natural benchmark in which sellers believe that all consumer types are equally likely. Note that search cost is the only ex-ante source of heterogeneity among consumers.

Consumers have to search for both product prices and match values. Before search, consumers make privacy choices: whether to enable or disable cookies. By enabling cookies, consumers allow the platform to access to their personal information and use that to personalize their search environment. We parameterize the personalization level by $\underline{v}$, so that the match distribution of the new pool is a lower truncation of the original one at $\underline{v}$.[3] The interpretation is that consumers will only be able to reach the product of which the match value is higher than $\underline{v}$. It's straight-forward to verify that the personalization match distribution $F([\underline{v}, 1])$ F.O.S.D $F([0, 1])$. As such, it's a specific way to improve consumers' search environment.

1.3.1 Timing of events

1. The platform announces its personalization level $\underline{v}$.

[2]-1-concavity is a weaker concept than the more-widely used log-concavity: log-convex distributions Nevertheless, it preserves the increasing hazrd rate property which is essential to guarantee the existence of a pure strategy equilibrium of a . See for detailed discussion.

[3]By Bagnoli and Bergstrom (2006), truncation of log-concave density function will also have a log-concave density function.

2. Consumers decide whether or not to participate in search. if they don't participate, they exit the market immediately and get an outside option of zero; if they choose to participate, they further decide whether or not to enable cookies. They then conduct sequential and random search for both price and product match in the tailored search environment determined by their privacy choice.

3. Firms simultaneously and independently set prices, based on beliefs on the type of consumers that reached through the privacy/no privacy channels.

The equilibrium concept that I use is perfect Bayesian equilibrium with passive beliefs. This means that if a consumer observes a price deviation by one firm, she does not change her belief about the pricing behavior of other firms. Similarly, if a firm observes a deviation of cookie choice by one consumer, it also does not change its belief about the type of other consumers. As is standard within the literature, I focus on symmetric equilibria in which all firms charge the same price.

1.3.2 Discussion of modeling assumptions

I discuss several important assumptions here.

First, the assumption that consumers are not informed of their own preference is suitable for many online situations where consumers are asked for their permission to enable cookies before searching on a website. Consumers base their privacy choice on their rational expectation of equilibrium price and match value followed by their privacy choice.

Second, I assume that by enabling cookies, the consumers are essentially facing a new utility match distribution that is a lower truncation of the original one. An interpretation is that enabling cookies give the platform access to a signal that reveals whether the match value of product is above certain threshold or not, and the platform only display the products to consumers whose valuation is above a certain threshold. This captures the idea that disclosure helps the platform with more personalized recommendations and thus make the search pool more relevant, while at the same time consumers still have the incentive to search. It's not derived from the platform objectives. Rather, it is a theoretically elegant and concise way to capture the consequence of the platform's use of consumers' personal information. Admittedly, there are many ways that online platforms could leverage the data collected from consumers. A more intrinsic way of modeling is beyond the scope of this study and is left for further research.

Lastly, I assume that search cost, despite being heterogeneous across consumers, are constant across each individual consumer's search attempts. Note that this is not without generality. For instance, if the first search is free and the second beyond is extremely costly, then in equilibrium every consumer will search once and only once. However the focus of this paper is on the externality across consumers, so I assume away such complications.

1.4 Equilibrium Analysis

1.4.1 Exogenous privacy choice

In this section, I present the benchmark model where the privacy choices are not made by consumers. The platform either 1) use the same personalization technology for all entering consumers because each individual consumer's search cost is not observable to her; or 2) as a social planner, observes all consumers' search cost and segments the market accordingly to maximize welfare. As is standard in the literature, we focus on pure-strategy symmetric Perfect Bayesian equilibrium, which is characterized by a price p^* for all firms and a stationary stopping rule characterized by a reserve match utility θ^* for consumers such that:

1. the consumers' stopping rule is optimal given their rational expectations of firms' pricing strategy that all firms charge p^*;

2. firms optimally charge price p^* and no firm has an incentive to deviate given the consumers' stopping rule and the belief that all firms charge p^*.

Consumer behavior

It is well known that when the match distribution F is i.i.d across products and consumers, under a symmetric-price equilibrium, the optimal strategy of a consumer with search cost s is simply characterized by an optimal stopping threshold x, which is the solution to the following equation:

$$\int_0^1 \max\{0, \theta - x\} f(\theta) d\theta = \int_{[x,1]} (\theta - x) f(\theta) d\theta = s \tag{1.4.1}$$

The left-hand side denotes consumers' additional gain from one more search with a match utility of x in hand, and the right-hand side is her search cost. x, as a function of s, defines the reservation utility $\theta^*(s)$, at which the consumer is indifferent between one more search and acquiring the current product. It can be readily seen that there exists a unique solution $\theta^*(s)$ to equation 1.4.1.

Lemma 1. *θ^* is continuously differentiable, strictly decreasing and strictly convex in s.*

Importantly, it pins down the consumer's equilibrium search behavior if the consumer indeed has active search motive. It's possible that $\theta^*(s)$ falls below the lower bound of the match distribution. Equivalently, this is the case when $s \geq \mathrm{E}_F(\theta) - \underline{\mathrm{v}}$. Intuitively, the consumer never searches when the lowest match value yields a sufficiently high payoff; even if she knew she would sample the highest value in the next period, she would prefer to accept $\underline{\mathrm{v}}$ today. There is no equilibrium with trade under this scenario unless the first search is free, in which case each firm effectively becomes a local monopolist conditional on being visited by the consumer. As will be specified later, we eliminate this possibility by putting an upper bound on s.

Consumers will participate in search only if she will get positive surplus from participation. If this is indeed the case, then she will continue searching until she finds a product of reservation utility $\theta^*(s)$. If not, she will not enter the market in the first place. The surplus $v(s)$ of a consumer with search cost s is $\theta^*(s) - p^*$[4], where p^* is the equilibrium price which will be derived later. Setting $v(s)$ to zero and using equation 1.4.1, we have the threshold search cost $s(p^*)$ above which consumers will **not** participate in search:

$$s(p^*) := \int_{p^*}^{\bar{s}} (\theta - p^*) f(\theta) d\theta \tag{1.4.2}$$

Note that $s(p^*)$ depends on p^*, which in turn depends on $s(p^*)$. Standard search literature assumed that the market is fully covered, in the sense that $s(p^*) > \bar{s}$. If on the other hand in equilibrium $s(p^*) < \bar{s}$, then not everyone search actively and some will drop out the market immediately.

To summarize, consumers' search strategy is a stationary optimal stopping rule characterized by a reservation value, subject to participation constraints.

Firm behavior

Having characterized consumers' search behavior, we proceed to compute equilibrium price p. We do this by deriving firm j's payoff of charging price p' while all other firms charge p^*, then compute the first-order condition and apply the symmetry condition $p' = p^*$. Firm j's payoff per consumer type s from deviating to price p' is the product of: the density of consumer type s that visits firm j; conditional on visiting firm j, the probability that she will buy from firm j; and the deviating price p'.

Conditional on reaching firm j, consumer i'of type s will stop searching at firm j if and only if $\epsilon_j - p' > \theta^*(s) - p^*$. Thus the conditional probability that consumer with search

[4]To see this, note that a consumer's surplus must satisfy the following recursive condition: $CS(s) = -s + (1 - F(\theta^*(s)) \frac{\int_{\theta^*(s)}^{1} (\theta - p^*) f(\theta) d\theta}{1 - F(\theta^*(s))} + F(\theta^*(c)) CS(s)$. Plugging s from equation 1 yields the formula.

cost s stops searching and buys at firm j is:

$$\Pr(\epsilon_j - p' > \theta^*(s) - p^*) = 1 - F(\theta^*(s) + p' - p^*)$$

Since there are infinite number of firms, the unconditional probability that she stops at firm j and acquire the product of firm j is:

$$\frac{1 - F(\theta^*(s) + p' - p^*)}{1 - F(\theta^*(s))} \tag{1.4.3}$$

Let $\tilde{s}(p^*) := \min\{s(p^*), \bar{s}\}$. $[\underline{s}, \tilde{s}(p^*)]$ is the set of consumers that enter the search market. The payoff to a deviating firm is:

$$\pi(p'; p^*) = p'D(p^{'}; p^*) = p' \int_{\underline{s}}^{\tilde{s}(p^*)} \frac{1 - F(\theta^*(s) + p' - p^*)}{1 - F(\theta^*(s))} dG(s) \tag{1.4.4}$$

Taking F.O.C with respect to p', we get:

$$\int_{\underline{s}}^{\tilde{s}(p^*)} \frac{1 - F(\theta^*(s) + p' - p^*)}{1 - F(\theta^*(s))} dG(s) - p' \int_{\underline{s}}^{\tilde{s}(p^*)} \frac{f(\theta^*(s) + p' - p*)}{1 - F(\theta^*(s))} dG(s) = 0 \tag{1.4.5}$$

Applying the symmetry condition $p^* = p'$, we get the formula for equilibrium price:

$$p^* = \frac{G(\tilde{s}(p^*))}{\int_{\underline{s}}^{\tilde{s}(p^*)} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} dG(s)} \tag{1.4.6}$$

If in equilibrium $s(p^*) > \bar{s}$, then $G(\tilde{s}(p^*)) = 1$. It's straightforward to verify that there exists an unique candidate equilibrium price

$$p^* = \frac{1}{\int_{\underline{s}}^{\bar{s}} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} dG(s)} \tag{1.4.7}$$

where $\theta^*(s)$ is given by equation 1.4.1.

If the search cost distribution $G(s)$ is degenerate at a single point s, we can show that the equilibrium price is:

$$p^* = \frac{1 - F(\theta^*(s))}{f(\theta^*(s))} \tag{1.4.8}$$

By Assumption 1, p^* is decreasing in $\theta^*(s)$.

The next statement characterizes the condition under which in equilibrium it is indeed the case that $s(p^*) > \bar{s}$. Notably, this cutoff value determined by the match distribution $F(\cdot)$ and the search cost distribution $G(\cdot)$ only, and does not depend on the equilibrium price.

Assumption 2. *In absence of personalization, the market is not fully covered. The upper bound of the search cost distribution satisfies*

$$\bar{s} > \int_{\int_{\underline{s}}^{\bar{s}} \frac{1-f(\theta^*(s))}{1-F(\theta^*(s))} g(s)ds}^{1} (\theta - \int_{\underline{s}}^{\bar{s}} \frac{1 - f(\theta^*(s))}{1 - F(\theta^*(s))} g(s)ds) f(\theta) d\theta$$

where F is supported on $[0, 1]$.

The following lemma characterize the equilibrium price in a market equilibrium with partial market coverage.

Lemma 2. *For any level of personalized match distribution $F[\underline{v}, 1]$, only a fraction of consumers with search cost in $[0, (\theta^*)^{-1}(\underline{v})]$ have active search incentive. Among them, only those whose search cost fall below $\tilde{s}(p^*)$ get non-negative expected surplus from searching(thus participate), where the equilibrium price p^* is implicitly given by*

$$p^* = \frac{G(\tilde{s}(p^*))}{\int_{\underline{s}}^{\tilde{s}(p^*)} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} dG(s)} \tag{1.4.9}$$

in which

$$\tilde{s}(p^*) = \min\{s(p^*), \bar{s}\}.$$

The result is reminiscent of Theorem 1 in Moraga-Gonz'alez et al. (2017). Here we provide a proof sketch. First, it can be shown that there exists a unique solution to equation

$L(p) := p\int_{\underline{s}}^{\tilde{s}(p)} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} g(s)ds - G(\tilde{s}(p)) = 0$ by showing that $L(0) > 0$, $L(p^m) < 0$ (p^m is the monopoly price) and $L(p)$ is strictly decreasing in p. To establish sufficiency, since a direct verification of second order conditions are hard to carry out, they instead proceed by showing that the demand function of an individual firm is log-concave in its own price under the assumption that $g(\theta^{*,-1})$ is log-concave. In order to prove that the demand function is log-concave in its own price, they show that the demand from a single consumer is log-concave both in price and reservations values. By Theorem 6 in Prekopa(1973), integration over consumer reservation value preserves log-concavity. Once it is proved, the profit function of an individual seller is quasi-concave in its own price so the candidate equilibrium price is indeed an equilibrium. Importantly, if the search cost distribution is uniform(which is the distributional assumption we will keep throughout the paper), the log-concavity of $g(\theta^{*,-1})$ is guaranteed by any match distribution.

Mathematically, the equilibrium price is inversely proportional to the weighted average value of function $\frac{f(\theta^*(s))}{1-F(\theta^*(s))}$ over interval $[\underline{s}, \tilde{s}(p^*)]$(the set of participating consumers in this market segment). Equivalently, it is the weighted average of equilibrium price with a single consumer, given by 1.4.7, with higher weight being put on equilibrium price with lower search cost consumers.

1.4.2 The effect of personalization

Before introducing consumers' privacy choice, we study the effect of changing the level of personalization on equilibrium price and welfare. For ease of exposition, let H be the match distribution of the improved search pool. By assumption, H F.O.S.D. F. We have

$$\int_0^1 \max\{0, \theta - x\} h(\theta) d\theta \geq \int_0^1 \max\{0, \theta - x\} f(\theta) d\theta, \forall x \in [0, 1] \tag{1.4.10}$$

That is, for any current match utility in hand, the incremental benefit of one more search is larger in the case of H than in the case of F. It can readily be seen that $\theta^*_H(s) \geq \theta^*_F(s)$[5], which implies that consumer with a given search cost s have a higher reservation value in the improved search pool.

Recall that equilibrium price is given by the inverse hazard rate of F evaluated at reservation match value $\theta^*(s)$. In the special case where H is a truncation of F at $\underline{\text{v}}$, we have

$$H(x) = \frac{F(x) - F(\underline{\text{v}})}{1 - F(\underline{\text{v}})} \Rightarrow \frac{1 - H(x)}{h(x)} = \frac{1 - F(x)}{f(x)}$$

which implies that the two distributions have the same inverse hazard rate. Given that $\theta^*_H(s) \geq \theta^*_F(s)$ and the non-decreasing hazard rate assumption, we can conclude that in equilibrium $p^*_H < p^*_F$. That is, equilibrium price is lower under higher level of personalization, and consumer surplus must also increase accordingly. Intuitively, consumers who actively search in the refined search pool would view products as closer substitutes, which intensify the price competition among firms admitted into the search pool.

To fix ideas, we now parameterize F as $Uniform[\underline{\text{v}}, 1]$. By straight calculation, the reservation value $\theta^*(s)$ of a consumer of type s is $1 - \sqrt{2s(1 - \underline{\text{v}})}$. The equilibrium price when only consumer with search cost s is present in the market is:

$$p^*(s) = \frac{1 - F(\theta^*(s))}{f(\theta^*(s))} = 1 - \theta^*(s) = \sqrt{2s(1 - \underline{v})} \tag{1.4.11}$$

Let's take a step back and drop Assumption 2 for now to see how consumers' two participation constraints evolve as search cost changes. To ensure non-negative surplus from search and active search incentive, we must have $1 - \sqrt{2s(1 - \underline{\text{v}})} \geq \underline{\text{v}}$ and $\theta^*(s) - p^*(s) \geq 0$, which imply $\underline{\text{v}} < 1 - 2s$ and $1 - \frac{1}{8s} < \underline{\text{v}}$ respectively. Consumers with search cost higher than $\frac{1-\underline{\text{V}}}{2}$ will automatically leave the market. Furthermore, when $s \geq \frac{1}{4}$, no $\underline{v}$ will satisfy both constraints. For these consumers with very high search cost, personalization plays

[5]By implicit function theorem, let $f = \int_{\theta^*}^{1} (\theta - \theta^*) h(\theta) d\theta - s$. We have $\frac{\theta^*(s)}{\partial \underline{\text{V}}} = -\frac{\partial f / \partial \underline{\text{V}}}{\partial f / \theta^*(s)} = \frac{f(\underline{\text{V}}) s}{1 - F(\theta^*)} > 0$.

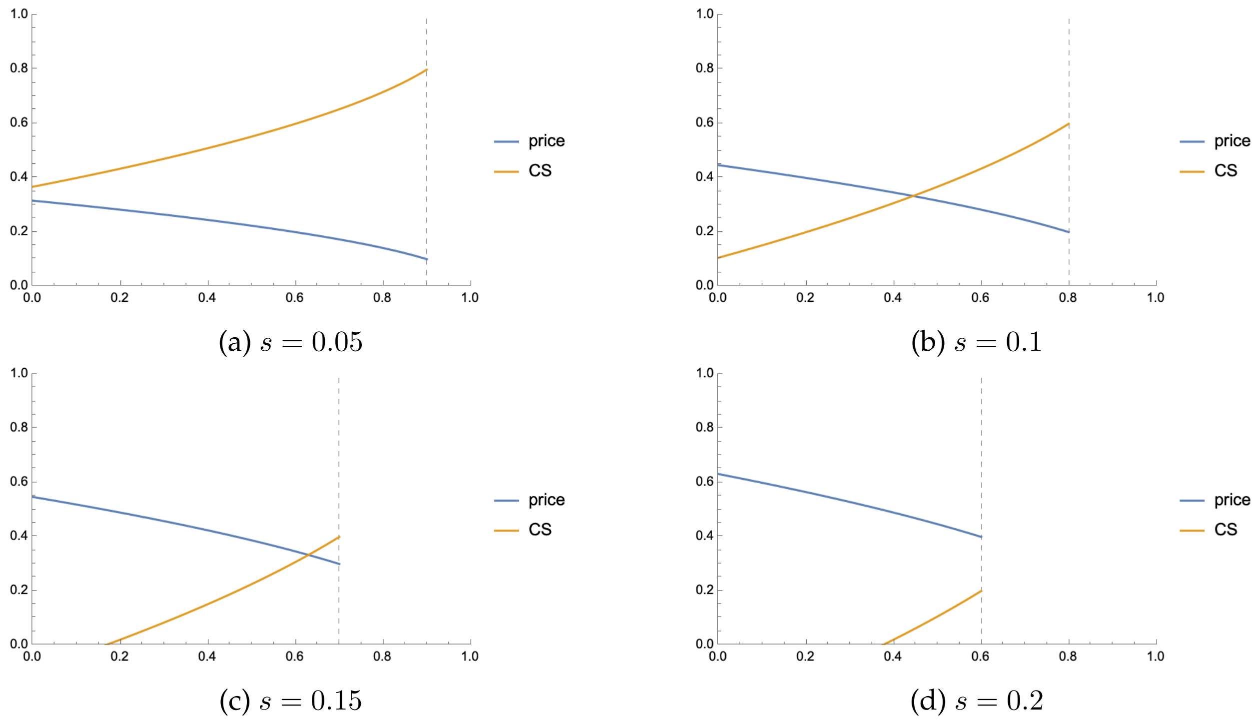

(a) $s = 0.05$ (b) $s = 0.1$ (c) $s = 0.15$ (d) $s = 0.2$

Figure 1.1: Comparative statics w.r.t $\underline{v}$: single consumer

no role since it won't give positive surplus to her while keeping active search motive. For $s < \frac{1}{8}$, the latter constraints never binds for any choice of $\underline{v}$, meaning that consumers with very low search cost always have active search motive regardless of the personalization level. For $\frac{1}{8} < s < \frac{1}{4}$, these conditions summarize to $1 - \frac{1}{8s} \leq \underline{v} \leq 1 - 2s$, meaning that these consumer will only participate given a moderate range of personalization level.

When the market is consists of a continuum of consumers with search cost distribution G, the maximum search cost to induce participation is given by:

$$s(p^*) := \int_{p^*}^{\bar{s}} (\theta - p^*) f(\theta) d\theta = \frac{(1-p^*)^2}{2(1-\underline{v})} \tag{1.4.12}$$

Equilibrium price is given by:

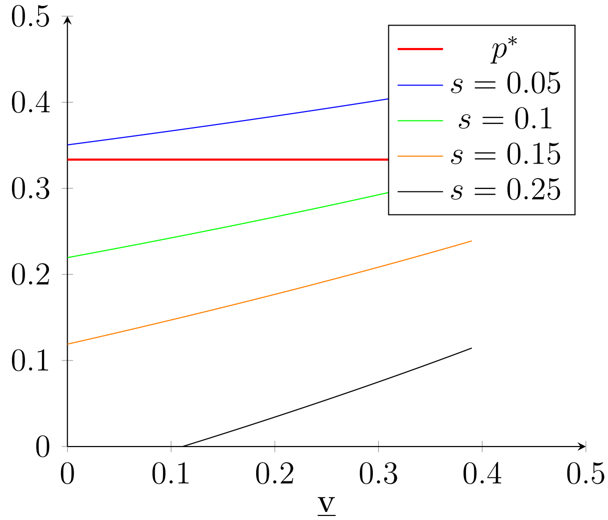

Figure 1.2: Equilibrium price and consumer surplus with $F = Uniform[\underline{v}, 1]$

$$p^* = \frac{G(\frac{(1-p^*)^2}{2(1-\underline{v})})}{\int_0^{\frac{(1-p^*)^2}{2(1-\underline{v})}} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} g(s)ds} = \frac{\frac{(1-p^*)^2}{2(1-\underline{v})}}{\int_0^{\frac{(1-p^*)^2}{2(1-\underline{v})}} \frac{1}{\sqrt{2s(1-\underline{v})}} ds} = \frac{\frac{(1-p^*)^2}{2}}{\sqrt{2s}|_0^{\frac{(1-p^*)^2}{2(1-\underline{v})}}} = \frac{(1-p^*)}{2} \Rightarrow \quad p^* = \frac{1}{3}$$

(1.4.13)

In this example, p^* is constant at $\frac{!}{3}$, regardless of the value of $\underline{v}$. The type of consumers that enter the search market is $(0, \min\{\frac{2}{9(1-\underline{v})}, \bar{s}\}]$. As personalization level increases, there will be more consumers entering the market. Consumer surplus is given by $\max\{1 - \sqrt{2s(1-\underline{v})} - p^*, 0\}$. Figure 1.2 depicts the equilibrium price and consumer surplus in this example.

The effects of moving v on equilibrium price is summarized in the following proposition:

Proposition 1 (Effect of $\underline{v}$ on the equilibrium price). *Let F be log-concave and let g be any nonnegative density on $[0, 1]$. In any symmetric equilibrium $p^*(\underline{v})$ we have $p^*(\underline{v}) \geq \underline{v}$ (price-binds regime), and the comparative statics are:*

(i) If g is uniform on $[0, 1]$, then $p^*(\underline{v})$ is constant in $\underline{v}$ throughout the admissible region.

(ii) With non-uniform g, $\frac{dp^*(\underline{v})}{d\underline{v}}$ is, in general, ambiguous: changes in $\underline{v}$ also affect p^* via induced

changes in the valuation tail F that enter $\bar{s}(p^*)$.

Proof. See Appendix. □

Improving the personalization technology affects equilibrium price in both the intensive and the extensive margins. On the one hand, consumers who were present in the market prior to the launch of personalization would view the products as closer substitutes, which intensifies price competition among the firms. On the other hand, some consumers who previously refrain from searching due to high search cost now enters the market, since the incremental benefit of search now becomes higher. This drives up market demand and pushes the sellers to raise prices. In this specific example, however, competition effect is exactly offset by market expansion effect, thus the price remains unchanged before and after the under introduction of personalization technology.

We next study the effect on total surplus. Note that social surplus is just the sum of reservation value of all participating consumers(since price is a surplus that transfers from consumers to firms). The total derivative of the marginal type $s^*(\underline{v})$ along the equilibrium path satisfies

$$\frac{ds^*(\underline{v})}{d\underline{v}} = -\frac{U_{\underline{v}}\Big(s^*;\underline{v},p^*(\underline{v})\Big) + U_p\Big(s^*;\underline{v},p^*(\underline{v})\Big)\,\frac{dp^*(\underline{v})}{d\underline{v}}}{U_s\Big(s^*;\underline{v},p^*(\underline{v})\Big)}.$$

We already know that $U_s < 0$, $U_{\underline{v}} > 0$, and $U_p < 0$. Hence the sign of $ds^*/d\underline{v}$ depends on the sign of the term

$$U_{\underline{v}} + U_p\,\frac{dp^*(\underline{v})}{d\underline{v}}.$$

The first term captures the direct increase in expected match surplus from a higher truncation level, whereas the second term reflects how the equilibrium price reacts to changes in $\underline{v}$. As shown in Proposition 1, the sign of $dp^*(\underline{v})/d\underline{v}$ is ambiguous, To ensure that the direct gain in match surplus is not fully offset by the price response, we impose the following mild condition.

Assumption 3 (Bounded price response)**.** *Along the equilibrium path,*

$$U_{\underline{v}}\Big(s^*(\underline{v}); \underline{v}, p^*(\underline{v})\Big) + U_p\Big(s^*(\underline{v}); \underline{v}, p^*(\underline{v})\Big) \frac{dp^*(\underline{v})}{d\underline{v}} > 0.$$

Equivalently, sellers cannot fully absorb the additional match surplus generated by a higher truncation level through higher prices.

This assumption allows us to get a clean result on the effect of $\underline{v}$ on total surplus:

Proposition 2. *Under Assumption 3, total surplus $TS(\underline{v})$ is strictly increasing in the truncation level $\underline{v}$.*

Proof. See Appendix. □

As personalization level becomes higher, total surplus becomes unambiguously higher, because the consumers who were present in the market now have a higher reservation value, and the set of active searchers expands. Nevertheless, as stated in the following proposition, the effect on total consumer surplus is ambiguous.

Proposition 3 (Effect of $\underline{v}$ on Consumer Surplus)**.** *Fix a truncation level $\underline{v}$ and let $p^*(\underline{v})$ be the corresponding symmetric equilibrium price. Consumer surplus can be written as*

$$CS(\underline{v}) = \int_0^{s^*(\underline{v})} U\Big(s; \underline{v}, p^*(\underline{v})\Big)\, dG(s),$$

where $s^*(\underline{v})$ is the marginal search-cost type defined by $U(s^*(\underline{v}); \underline{v}, p^*(\underline{v})) = 0$.

1. *Consumer surplus $CS(\underline{v})$ need not be monotone in $\underline{v}$. A higher truncation level improves the match distribution ($U_{\underline{v}} > 0$), but the net price response $dp^*(\underline{v})/d\underline{v}$ may have either sign, the effect on $CS(\underline{v})$ is in general ambiguous.*

2. *If the price response satisfies*

$$U_{\underline{v}}(s; \underline{v}, p^*(\underline{v})) + U_p(s; \underline{v}, p^*(\underline{v})) \frac{dp^*(\underline{v})}{d\underline{v}} \geq 0, \qquad \forall\, s \in [0, s^*(\underline{v})], \tag{M}$$

then $CS(\underline{v})$ is strictly increasing in $\underline{v}$.

Proof. See Appendix. □

Condition (M) requires that, for every active search-cost type, the direct gain from a better match distribution dominates any loss coming from the equilibrium price adjustment. When the search-cost distribution G is uniform and the match-value distribution is uniform on $[\underline{v}, 1]$, condition (M) holds automatically. In this benchmark the equilibrium price is constant ($dp^*(\underline{v})/d\underline{v} = 0$), and hence the direct improvement in the match distribution implies that $CS(\underline{v})$ is strictly increasing in $\underline{v}$.

Under this scenario,introducing a second personalization level cannot raise total consumer surplus beyond the best pooled truncation. Intuitively, if every consumer strictly prefers the higher truncation level given the corresponding regime price, then putting anyone in the low-truncation regime only makes that consumer worse off. Since nobody benefits from being placed in the low-truncation regime, segmenting the market cannot outperform simply placing everyone in the high-truncation regime.

Likewise, once all consumers face the same truncation level, further segmenting the market by search cost—while keeping the truncation fixed—cannot improve consumer surplus. The effect is purely redistributive: it leaves total surplus unchanged but weakly reduces consumer surplus relative to pooling. The next proposition proves this intuition:

Proposition 4. *Given $\underline{v}$, Consumer surplus under pooling is weakly higher than under any partition of the search cost distribution.*

Proof. See Appendix. □

We conclude this section with a natural question: Is the consumer-optimal outcome always achievable when privacy choices are left to consumers? Section 1.4.3 demonstrates that this need not be the case.

1.4.3 Endogenous privacy choice

In this section, we study consumers' endogenous privacy choice. We restrict attention to a simple type of separating equilibrium where consumers with low search cost participate in search but chooses to disable cookies, consumers with intermediate search cost participate in search and chooses to enable cookies, and consumers with high search cost does not participate. Formally, this is defined as

Definition 1. $(p_1^{eq}, p_0^{eq}, s^*, s^{**})$ *is a monotone separating equilibrium if there exists an equilibrium in which:*

- *Low-search-cost consumers* $(s \in (0, s^*])$ *find it optimal to **remain anonymous** and search. Firms, observing anonymity, infer a low search cost and charge a competitive price* p_0^{eq}*.*

- *Medium-search-cost consumers* $(s \in (s^*, s^{**}])$ *find it optimal to **enable cookies** to avoid search costs. Firms, observing that cookies are enabled, infer a higher search cost and charge a higher price* p_1^{eq} *(where* $p_1^{eq} > p_0^{eq}$*).*

- *High-search-cost consumers* $(s \in (s^{**}, \bar{s}])$ *find neither searching nor enabling cookies worthwhile and leave the market immediately.*

- *Firms' pricing strategies are a best response to these consumer segments, and their beliefs about consumer types are consistent with the equilibrium strategies.*

We now proceed to establish the existence of such a segmented equilibrium. Crucially, the existence of a separating equilibrium hinges on both the shape of the search cost distribution and the incremental benefit of personalization technology uncovered by enabling cookies, because separation requires a well-defined marginal type of consumer—one for whom the gains from joining a more efficient segment exactly offset the higher price it entails. We are able to construct an example of such equilibrium under the assumption that G and F are uniform. [6].

[6]While the uniform assumption is primarily adopted for tractability and is not required for the existence

Proposition 5. *There exists a unique separating equilibrium as described in Definition 1. In particular, with uniform G and F, there exists v_{min} that sustains a separating equilibrium whenever $\underline{v} >= v_{min}$.*

Proof. See appendix. □

Intuitively, when personalization level is relatively low, market expansion effect is very small, plus the marginal benefit of personalization for high cost types are so low that they would better pooled with low cost types to get better price. When personalization level rises above a certain threshold, consumers in the lower cost end and in the higher cost end would have misaligned interests: consumers with high search cost prefer better match anyway despite they know they will face a higher price, whereas low cost consumers stay anonymous because the extra benefit from getting a better match does not offset that loss from being recognized as high search cost consumers.

The following two examples illustrates the (non-)existence of separating equilibrium under various levels of personalization. In both cases $F \sim U[0, 1]$.

Example 1. *Under $\underline{v} = 0.45$,in the unique separating equilibrium, $p_1^{eq} = 0.47$,$p_0^{eq} = 0.29$, $s^* = 0.21$,$s^{**} = 0.23$. In contrast, Consumer-optimal outcome is to enable cookies for everyone.*

Example 2. *Suppose $\underline{v}$=0.1. There does not exist a separating equilibrium: in equilibrium everyone disable cookies. Consumer-optimal outcome is to enable cookies for everyone.*

1.5 Welfare Analysis

Our main result in this section compares aggregate consumer welfare under the two regimes.

---

of a separating equilibrium, it conveniently encodes a regularity condition: the search-cost distribution must be sufficiently diffuse to sustain such an equilibrium. To see this, suppose the distribution is left-skewed (i.e., low-cost types are more prevalent), low-cost types have weaker incentives to separate because they expect the relatively scarce high-cost consumers to exert little influence on the equilibrium price.

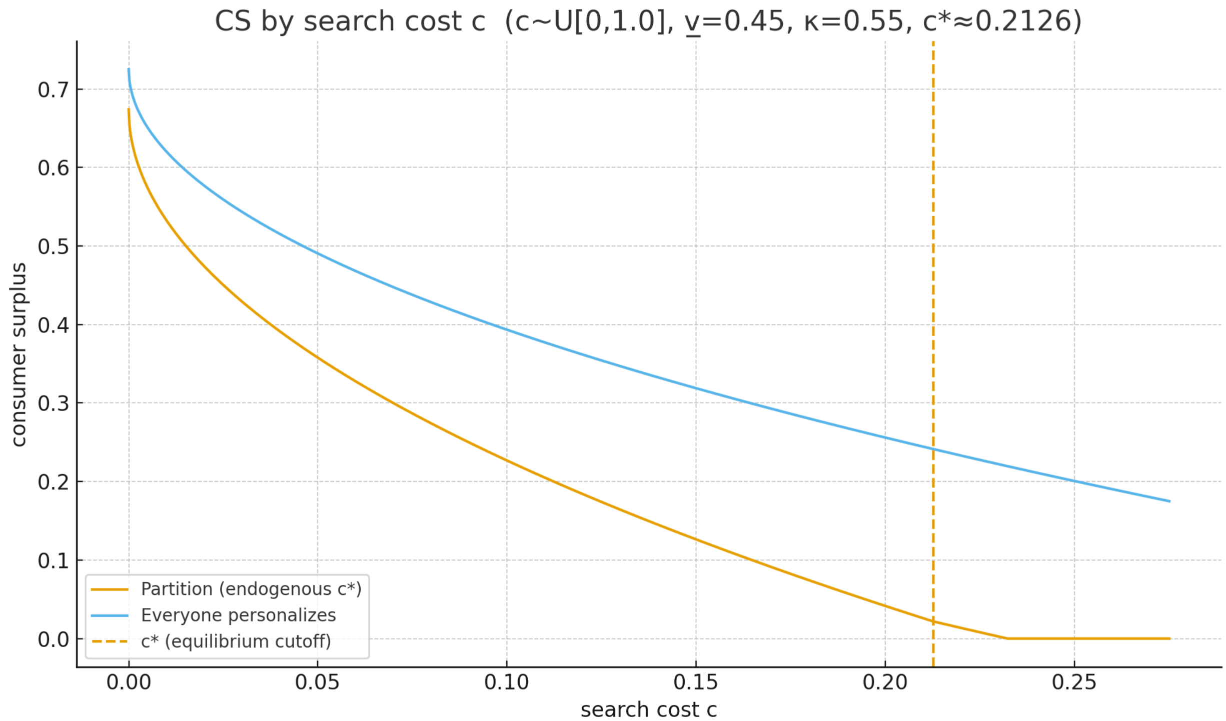

Figure 1.3: Comparison of equilibrium welfare by cost type

Proposition 6 (Welfare comparison)**.** *Let CS_{part} and CS_{all} be the total consumer surplus under the decentralized and centralized regime, respectively. Whenever a separating equilibrium exists,*

$$CS_{part} < CS_{all}$$

Proof. See the appendix. □

Our next result characterizes the distributional consequences of the decentralized privacy regime. We show that, relative to the pooling benchmark in which all consumers permit personalization, any separating equilibrium—whenever it exists—makes almost all consumers worse off. In particular, all consumers except possibly a vanishingly small set of low–search-cost types in the interval $(0, \hat{s}]$ suffer a welfare loss when privacy choice is decentralized. Moreover, $\hat{s} \leq s^*$.

Corollary 1. *Suppose a separating equilibrium exists. Then:*

1. *For any* $s \in (s^*, s^{**}]$

$$\theta^*(\kappa s) - p_{\text{all}} \;\geq\; \theta^*(\kappa s) - p_1^{eq}$$

2. *For any* $s \in [0, s^*]$*, define*

$$\Delta\text{CS}(s) := \Big(\theta^*(\kappa s) - p_{\text{all}}^{\text{eq}}\Big) - \Big(\theta^*(s) - p_0^{\text{eq}}\Big) = \underbrace{\Big(\theta^*(\kappa s) - \theta^*(s)\Big)}_{=:G(s)\ \nearrow} + \underbrace{(p_0^{\text{eq}} - p_{\text{all}}^{\text{eq}})}_{\textit{constant in } s}.$$

Then $\Delta\text{CS}(s)$ *is strictly increasing in* s*. In particular:*

$$\begin{cases} p_{\text{all}}^{\text{eq}} \leq p_0^{\text{eq}} \;\Rightarrow\; \Delta\text{CS}(s) \geq 0 \;\; \forall s \in [0, s^*]; \\ p_{\text{all}}^{\text{eq}} > p_0^{\text{eq}} \;\Rightarrow\; \exists\, \hat{s} \in [0, s^*] \textit{ s.t. } \Delta\text{CS}(\hat{s}) = 0, \;\; s < \hat{s} \textit{ gain}, \;\; s > \hat{s} \textit{ lose}. \end{cases}$$

Proof. See Appendix. □

See Figure Figure 1.3 for graph illustrations. Intuitively, the marginal benefit of improved personalization is minimal for those consumers with lowest search costs, since they will find almost perfect match by themselves anyway. So they may lose the least, or even benefit from foregoing the benefit of better personalization. For the consumers with very high search costs, they have to rely on the platform's personalization to find a good match despite being recognized as high type consumers. For consumers with intermediate level of search costs, however, the trade-off between lower price and better match turns out to be more subtle. In this particular case, they choose to disable cookies because any unilateral deviation to enable cookies will reduce their payoff.

1.6 Discussion

Should price discrimination be banned? Suppose firms are not allowed to charge consumers different prices based on privacy choice because of a policy ban, then privacy choice will lose its value for signaling. Depending on their beliefs of equilibrium price, either everyone enable cookies or no one does. It is “as if” consumers are given no choice, except that firms may coordinate on the equilibrium they prefer(which is often harmful to consumers). As such, a straight ban of price discrimination does not automatically restore welfare-maximal outcome.

Equilibrium Multiplicity This game naturally has multiple equilibria. In particular, there exists pooling equilibria where for all participating consumers, either 1)all participating consumers enable cookies or 2) all participating consumers disable cookies. Firms collectively charge the profit-maximizing price, anticipating the search cost distribution of active consumers and set a price high enough that discourages all types of consumers from deviation. Such equilibra could be sustained because consumers make privacy choice before they observe the actual price(including any deviations by the firms in the unreached market), hence such beliefs could never be disconfirmed in equilibrium. The equilibra in which all consumers enable cookies are particularly worth contemplating, since it (often) coincides with the consumer optimal outcome. However, these equilibra does not pass the intuitive criterion of Cho and Kreps (1987): to sustain such equilibra, sellers must hold the off-path belief that the deviating consumers have high search costs. Such beliefs are unreasonable since only consumers with low search costs have incentive to deviate. On the other hand, the equilibra in which all disable cookies are never consumer-optimal: equilibrium price is highest, participation is lowest, and reservation value is also lowest.

As such, among all credible equilibra of this privacy game, separating equilibrium, if exists, dominates every other in terms of consumer surplus. In other words, int serves as an upper bound of consumer surplus under the decentralized regime, which strengthens our key finding that decentralized outcome is often inferior to the centralized outcome in terms of consumer welfare.

1.7 Concluding remarks

In this paper, I study consumers' incentive to share their preference information in a setting where each consumer has to search costly to make a purchase, and the platform can personalize her search environment at the permission by consumers. Enabling cookies help consumers find better match, while the sellers may infer information relevant to pricing decision from their privacy choice thus opens the door for price discrimination.

With a single search cost, personalization improves social welfare by facilitating matching between consumers and products, and by intensifying the price competition among sellers under mild search costs. However, the presence of consumers with heterogeneous search costs create inter-consumer externality. As such, not every consumer benefit from better personalization. In equilibrium, some consumers may have incentive to stay anonymous if they have the choice. The underlying mechanism is strikingly different from the case with the monopolist seller: here the option to disable cookies serves as a signaling device of low search cost, thereby pushing up price competition among sellers at the cost of lower expected match value. Compared to the consumer-optimal outcome, in the unique separating equilibrium of the privacy choice game, almost all consumers are worse off.

There are several interesting avenues for future research. In the main body of this paper, we assume that every inspection of price and valuation is costly, and in order to make a purchase the buyer has to search in the first place. This is not always the case in reality. There are two reasons why consumers may forego searching in the first place: first, she anticipates the equilibrium price exceeds the reservation value; second, she anticipates the match distribution is too concentrated such that the incremental benefit per search is always lower than the search cost. In a market with single cost type the consumer will never get postive surplus without searching first because of the Diamond Paradox logic. However, when the market consists of sufficiently many consumers with low search cost, these buyers may free-ride on the low cost types, thus may make a purchase even without searching in the first place. To accommodate such scenario, it could be interesting to modify the assumption by allowing immediate purchase: the buyer has the option to buy a random product without inspection. A complete characterization of equilibrium outcome is left for future study. Nevertheless, the basic insights that high-search cost consumers leave an negative externality on low search cost consumers carry over.

A natural direction for future work is to introduce heterogeneity in consumers' intrinsic privacy concerns. Some consumers may incur a fixed utility loss from enabling cookies, while others are essentially privacy-neutral. Incorporating this fixed privacy cost into the model would generate richer sorting patterns—for example, partial and non-monotone adoption of personalization—and would allow the analysis to speak more directly to the observed diversity in privacy choices. It would also sharpen the welfare trade-offs, as mandated personalization may benefit low-privacy types while harming those with strong privacy concerns. Exploring this extension would make the model more realistic and open new policy-relevant questions about defaults, opt-in mechanisms, and privacy-preserving personalization.

1.8 Appendix

Proof of Lemma 1

Proof.

$$\theta^{*\prime}(x) \;=\; -\,\frac{1}{1 - F(\theta^*(x))} \;<0, \qquad \theta^{*\prime\prime}(x) \;=\; \frac{f(\theta^*(x))}{\Big(1 - F(\theta^*(x))\Big)^3} \;>\; 0.$$

Convexity follows since $\theta^{*\prime\prime} > 0$. □

Proof of Proposition 1

Proof. To see the comparative statics of v on equilibrium price, we start first by showing that p^* is a harmonic mean of the single type equilibrium price $p_{\underline{v}}(s)$ over the active set of consumers, with weights $g(s)$:

$$p^*(\underline{v}) \;=\; \frac{\displaystyle\int_0^{s_0(\underline{v},p^*)} g(s)\,ds}{\displaystyle\int_0^{s_0(\underline{v},p^*)} \frac{g(s)}{p_{\underline{v}}(s)}\,ds} \;=\; H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[0,\,s_0(\underline{v},p^*)]} \tag{1.8.1}$$

The following two properties of harmonic mean are helpful with the proof:

Lemma 3 (Pointwise order lowers the harmonic mean)**.** *Let $w \geq 0$. If $X_2(c) \leq X_1(c)$ for all c in $[0, b]$, then*

$$H_w(X_2;\,[0,b]) \;\leq\; H_w(X_1;\,[0,b]).$$

Lemma 4 (Upper-limit monotonicity for increasing integrands)**.** *Let $w \geq 0$ and X be (weakly) increasing in c. Then*

$$b_2 \leq b_1 \quad\Longrightarrow\quad H_w(X;\,[0,b_2]) \;\leq\; H_w(X;\,[0,b_1]).$$

Proof. Write $H_w(X) = A/B$ with $A(b) = \int_0^b w$, $B(b) = \int_0^b w/X$. Lemma 3 is immediate since $X_2 \leq X_1$ implies $B_2 \geq B_1$ while $A_2 = A_1$. For Lemma 4, differentiate: $\frac{d}{db}H_w = \frac{w(b)}{B(b)^2}\Big(B(b) - A(b)/X(b)\Big)$. Because X is increasing, $X(b) \geq H_w(X)$, hence $B(b) \geq A(b)/X(b)$ and the derivative is nonnegative. $\square$

Next, we show that for a given type s, truncating the match distribution is equivalent to letting her search in the untruncated distribution with a new search cost scaled by $\kappa := 1 - F(\underline{v})$. A left truncation at $\underline{v}$ satisfies $1 - F_{\underline{v}}(\theta) = \frac{1-F(\theta)}{1-F(\underline{v})}$. The reservation value M with truncation solves:

$$s = \int_M^1 \Big(1 - F_{\underline{v}}(t)\Big)\, dt = \frac{1}{1 - F(\underline{v})} \int_M^1 \Big(1 - F(t)\Big)\, dt.$$

Rearranging yields:

$$\underbrace{(1 - F(\underline{v}))\, s}_{\kappa s} = \int_M^1 \Big(1 - F(t)\Big)\, dt.$$

Clearly, $\kappa \leq 1$. This allows us to rewrite the reservation value of type s under personalization as $\theta(\kappa s)$.

We are now ready to prove the comparative statics result of $\underline{v}$ on equilibrium price. By direct consequence of 3and 4, for each fixed s we have $p_{\underline{v}_2}(s) \leq p_{\underline{v}_1}(s)$ whenever $\underline{v}_2 > \underline{v}_1$, and, for each fixed $\underline{v}$, the map $s \mapsto p_{\underline{v}}(s)$ is increasing.

Here the active endpoint satisfies $\theta^*(\kappa(\underline{v}\bar{s}) = p^*$, i.e. $\kappa(\underline{v})\, s_0(\underline{v}, p^*)$ is constant when p^* is held fixed. As $\underline{v}$ increases, $\kappa(\underline{v})$ falls and $\bar{s}$ rises. Two forces affect the harmonic mean: (i) for existing s, $p_{\underline{v}}(s)$ falls, which tends to lower the harmonic mean; (ii) the interval extends to higher s where $p_{\underline{v}}(s)$ is larger, which tends to raise the harmonic mean. With uniform search cost distribution g these two forces exactly offset at the symmetric solution, yielding a locally constant $p^*(\underline{v})$; with a general (non-uniform) g, the reweighting of cost types may tilt the balance either way, rendering the sign ambiguous. $\square$

Proof of Proposition 2

Proof. Total surplus at truncation level $\underline{v}$ is

$$TS(\underline{v}) = \int_0^{s^*(\underline{v})} \theta^*(s;\underline{v})\, dG(s),$$

where $s^*(\underline{v})$ is defined by $U(s^*(\underline{v});\underline{v},p^*(\underline{v})) = 0$.

Totally differentiating the marginal-indifference condition with respect to $\underline{v}$,

$$U_s \frac{ds^*(\underline{v})}{d\underline{v}} + U_{\underline{v}} + U_p \frac{dp^*(\underline{v})}{d\underline{v}} = 0,$$

so that

$$\frac{ds^*(\underline{v})}{d\underline{v}} = -\frac{U_{\underline{v}} + U_p\, dp^*(\underline{v})/d\underline{v}}{U_s}.$$

From the search problem we have $U_s < 0$, $U_{\underline{v}} > 0$, and $U_p < 0$ for all active types. By Assumption 3, $U_{\underline{v}} + U_p\, dp^*/d\underline{v} > 0$, and hence $ds^*(\underline{v})/d\underline{v} > 0$.

Applying Leibniz's rule,

$$\frac{dTS(\underline{v})}{d\underline{v}} = \theta^*\Big(s^*(\underline{v});\underline{v}\Big)\, g\Big(s^*(\underline{v})\Big) \frac{ds^*(\underline{v})}{d\underline{v}} + \int_0^{s^*(\underline{v})} \frac{\partial\theta^*(s;\underline{v})}{\partial\underline{v}}\, dG(s).$$

At the cutoff $s^*(\underline{v})$ the consumer is indifferent, so her expected utility is zero, but the firm earns strictly positive profit whenever $p^*(\underline{v}) > 0$, implying $\theta^*(s^*(\underline{v});\underline{v}) > 0$. Together with $g(s^*(\underline{v})) > 0$ and $ds^*/d\underline{v} > 0$, the boundary term is strictly positive.

Finally,per-type surplus increases with the truncation level,

$$\frac{\partial\theta^*(s;\underline{v})}{\partial\underline{v}} > 0 \qquad \text{for all active } s,$$

so the integral term is also strictly positive. Hence $dTS(\underline{v})/d\underline{v} > 0$, and total surplus is strictly increasing in the truncation level $\underline{v}$. □

Proof of Proposition 3

Proof. By Leibniz' rule,

$$\frac{dCS(\underline{v})}{d\underline{v}} = \int_0^{s^*(\underline{v})} \left[U_{\underline{v}} + U_p \frac{dp^*(\underline{v})}{d\underline{v}} \right] dG(s).$$

The direct term $U_{\underline{v}} > 0$ because a higher truncation level improves the match distribution in the sense of first-order stochastic dominance. The second term captures the price response.

(1) In general, $dp^*(\underline{v})/d\underline{v}$ is ambiguous due to competing competition and market-expansion effects. Hence the sign of $(*)$ is not determined.

(2) Under condition (M), each integrand is non-negative, so $CS(\underline{v})$ is strictly increasing.

□

Proof of Proposition 4

Proof. Following immediately from Proposition 1, total surplus is maximized at setting $v = v_{max}$ for every consumer. From the previous proof we know that p^* is a harmonic mean of single type equilibrium price $p^*(s)$, where the expectation is taken with respect to G truncated on $[\underline{s}, \tilde{s}(p^*)]$. Also note that p^* is a convex function of $\mathrm{E}[1/p^*(s)]$ on $(0, +\infty)$. By Jensen Inequality, we have

$$p^* \leq \mathrm{E}[p^*(s)] = \frac{\int_{\underline{s}}^{\tilde{s}(p^*)} p(s) dG(s)}{G(\tilde{s}(p^*))}$$

Note that the RHS is the weighted arithmetic mean of equilibrium price paid by consumers under perfect market segmentation, i.e. the case when search cost of every single consumer is perfectly observed by all firms and firms optimally set prices accordingly.

Therefore, total consumer surplus is higher under complete pooling of search cost than under perfect segmentation.

Furthermore, we can show that this relation can be generalized to any measurable partitions of the consumer type space, i.e., expected price paid by consumers in equilibrium would not be lower under arbitrary way of partition than under complete pooling. To see this, let $\mathcal{G}$ be a sub-σ-algebra generated by the partition. The equilibrium price within each partition is given by

$$Z := \Big(\mathbb{E}[\,1/p^*(s) \mid \mathcal{G}\,]\Big)^{-1}.$$

Let $U := \mathbb{E}[1/p^*(s) \mid \mathcal{G}] > 0$. Since $\phi(u) = 1/u$ is convex on $(0, \infty)$, Jensen's inequality gives

$$\frac{1}{\mathbb{E}[1/p^*(s)]} = \phi\Big(\mathbb{E}[U]\Big) \leq \mathbb{E}[\phi(U)] = \mathbb{E}\Big[\frac{1}{U}\Big] = \mathbb{E}\Big[\Big(\mathbb{E}[1/p^*(s) \mid \mathcal{G}]\Big)^{-1}\Big].$$

That is, the expected sum of price paid by consumers under complete pooling is weakly lower than the expected sum of price paid by consumers under any arbitrary segmentation. Together with the fact that total surplus is maximized under complete pooling, we can conclude that consumer surplus is also maximized under complete pooling. □

Proof of Proposition

Proof. The existence of equilibrium is proved by construction.

We first prove a useful property about reservation value $\theta(s)$:

Lemma 5 (strict single-crossing of the reservation uplift)**.** *For $0 < \kappa \leq 1$, the function*

$$G(s) \ := \ \theta^*(\kappa s) - \theta^*(s)$$

is <u>strictly increasing</u> in s on $[0, \bar{s}]$.

Proof. Let $H(s) := -\theta^{*\prime}(s) = \dfrac{1}{1 - F(\theta^*(s))} > 0$. By Lemma 1,

$$H'(s) \;=\; \frac{f(\theta^*(s))}{\Big(1 - F(\theta^*(s))\Big)^3} \;>\; 0,$$

so H is strictly increasing. Using $\theta^{*\prime} = -H$,

$$G'(s) \;=\; \kappa\, \theta^{*\prime}(\kappa s) - \theta^{*\prime}(s) \;=\; H(s) - \kappa\, H(\kappa s).$$

Since H is increasing and $\kappa \leq s$, we have $H(s) \geq H(\kappa s)$, hence

$$G'(s) \;\geq\; H(\kappa s) - \kappa\, H(\kappa s) \;=\; (1-\kappa)\, H(\kappa s) \;>\; 0 \quad \text{for } \kappa < 1.$$

□

Intuitively, the lemma says for any marginal improvement in personalization, the higher the search cost is, the more incremental gains from search.

We want to show that there exists a unique solution of s^* to the indifference condition on $(0, s^{**})$:

$$\theta^{NT}(s^*) - p_0^{eq} = \theta^T(s^*) - p_1^{eq} > 0 \tag{1.8.2}$$

where s^{**} is endogenously determined by s^* and $\underline{\text{v}}$ by solving $\theta^T(s^{**}) - p_1^{eq} = 0$.

For ease of exposition, let's now define the surplus difference function on $[0, s^* *]$:

$$\Delta(s) \;:=\; \Big(\theta^*(\kappa s) - p_1\Big) \;-\; \Big(\theta^*(s) - p_0\Big) \;=\; G(s) \;-\; \Big(p_1^{eq} - p_0^{eq}\Big). \tag{5}$$

Note that for any given $\underline{v}$, the above equation can be written exclusively in terms of s^* since by the harmonic mean argument,

$$p_1^{eq} = H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s \in [s^*,\, s^{**}(\underline{v}, p_1^{eq})]}, \quad p_0^{eq} = H_g\Big(p_(s)\Big)\Big|_{s \in [0,\, s^*]}$$

By Lemma 2, G is strictly increasing, hence so is $\Delta(s) = G(s) - (p_1^{eq} - p_0^{eq})$. Since a strictly increasing continuous function crosses a given level at most once, it follows immediately that there is at most one $s^* \in [0, \bar{s}]$ solving $\Delta(s^*) = 0$ (the indifference condition (1.8.2)). Applying Intermediate Value Theorem on $[0, \bar{s}]$, it follows that if in addition $\Delta(0) \leq 0$ and $\Delta(\bar{s}) \geq 0$ (e.g.$p_1^{eq} \geq p_0^{eq}$ and $G(\bar{s}) \geq p_1^{eq} - p_0^{eq}$), then there exists a unique $s^* \in [0, \bar{s}]$ with $\Delta(s^*) = 0$.

When $s = 0$,we have $\theta^*(s) = \theta^*(\kappa s) = 1$, $p_0 = 0$(because search is free, the consumer will not stop until he finds perfect match. Plus competition drives price to marginal cost, which we assume to be zero here), $p_1 > 0$, thus $\Delta(0) < 0$.

Under Assumption 2, the market is not fully covered which implies there exists $\hat{s} \in (0, \bar{s}]$ such that $\theta(\hat{s}) - p_0 < 0$. Then a sufficient condition for $G(\hat{s}) > 0$ to hold is $\theta^*(\kappa\hat{s}) - p_1 > 0$, where $p_1 = \frac{1-F(\theta^*(\kappa s))}{f(\theta^*(\kappa s))}$. Note that $\theta^*(\kappa\hat{s}) - p_1$ is decreasing/increasing in κ/$\underline{\text{v}}$ with opposite signs at the endpoints 0 and 1. Then there exists $\underline{\text{v}}$, such that $\theta^*(\kappa\hat{s}) - p_1 > 0$ for any $\underline{v} \geq v_{min}$.

It then remains to establish the conditions under which s^* is well defined, that is $s^* \in (0, s^{**})$tbd.

Next, fixing $\underline{\text{v}}$, it's is straightforward to verify that the cutoffs specified above indeed constitute an equilibrium. It suffices to show that the following IC constraints hold, i.e., no consumer has incentive to deviate:

$$\theta^{NT}(s) - p_0^{eq} \geq \theta^T(s) - p_1^{eq}, \quad \forall s \in (0, s^*];$$
$$\theta^T(s) - p_1^{eq} \geq \theta^{NT}(s) - p_0^{eq}, \quad \forall s \in (s^*, s^{**}]$$

From equation 1.8.2 we have $p_1^{eq} - p_0^{eq} = \theta_1^{eq}(s^*) - \theta_0^{eq}(s^*)$. Substituting into the IC constraints and rearranging, the constraints boil down to

$$\theta^T(s^*) - \theta^{NT}(s^*) \geq \theta^T(s) - \theta^{NT}(s), \quad \forall s \in (0, s^*];$$
$$\theta^T(s) - \theta^{NT}(s) \geq \theta^T(s^*) - \theta^{NT}(s^*), \quad \forall s \in (s^*, s^{**}]$$

Define $G(s) := \theta^T(s) - \theta^{NT}(s)$. According to Lemma 2, $G(s)$ is monotonically increasing in s. Therefore the above inequalities hold.

Finally, note that in this equilibrium all actions in the consumers' action space are played with positive probability, thus no need to specify off-path beliefs for the firms. By convention, the consumers would hold a passive off-path belief when seeing deviating firms. □

Proof of Proposition 6

Proof. Let $[0, s^{all}]$ be the set of active cost types of consumers under the pooling regime. By Proposition 4, we know that for any $a < b \in$[0,s^{all}],

$$\text{CS}_{\text{all}} = \int_0^{s^{\text{all}}} \Big(\theta^*(\kappa s) - p^{\text{eq}}_{\text{all}}\Big)\, ds \geq CS_{[0,a]} + CS_{(a,b]} + CS_{(b,s^{all})}$$

Note that $p_1^{eq} = H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[s^*,\, s^{**}]}$, $p_{all}^{eq} = H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[0,\, s^{all}]}$. By Lemma 2, $p_{\underline{v}}(s)$ increases in s; hence, by Lemma 3(b), removing low-s types raises the harmonic mean: $H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[0,\,\cdot]} \leq H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[s^*,\,\cdot]}$ Adjusting upper limits preserves the inequality, producing $p^{\text{eq}}_{\text{all}} \leq p^{\text{eq}}_1$. Since

$$\theta(\kappa s^{**}) - p_1^{eq} = 0, \qquad \theta(\kappa s^{all}) - p_{all}^{eq} = 0$$

, together with the fact that $\theta(\cdot)$ is decreasing in s, it must be the case that $s^{**} < s^{all}$. Therefore we can properly define the parameters as $a = s^*$, $b = s^{**}$, the above inequality becomes:

$$\text{CS}_{\text{all}} = \int_0^{s^{\text{all}}} \Big(\theta^*(\kappa s) - p^{\text{eq}}_{\text{all}}\Big)\, ds \geq CS_{[0,s^*]} + CS_{(s^*,s^{**}]} + CS_{(s^{**},s^{all})} \geq CS_{[0,s^*]} + CS_{(s^*,s^{**}]}$$

The latter inequality follows from the fact that consumer surplus must be non-negative for active consumers. Furthermore,

$$CS_{(s^*,s^{**}]} = \int_s^{s^{**}} \Big(\theta^*(\kappa s) - p^{\text{eq}}_1\Big)\, ds$$

$$CS_{[0,s^*]} = \int_0^{s^*} \Big(\theta^*(\kappa s) - p\Big)\, ds$$

,where

$$p = H_g\Big(p_{\underline{v}}(s)\Big)\Big|_{s\in[0,\,s^*]} < H_g\Big(p(s)\Big)\Big|_{s\in[0,\,s^*]} = p^{eq}_0$$

Therefore

$$CS_{[0,s^*]} > \int_0^{s^*} \Big(\theta^*(\kappa s) - p^{eq}_0\Big)\, ds > \int_s^{s^{**}} \Big(\theta^*(s) - p^{\text{eq}}_0\Big)\, ds$$

$$CS_{all} > \int_0^{s^*} \Big(\theta^*(s) - p^{eq}_0\Big)\, ds + \int_{s^*}^{s^{**}} \Big(\theta^*(\kappa s) - p^{\text{eq}}_1\Big)\, ds = CS_{part}$$

(For simplicity of notation, the proof assumes a uniform cost distribution G, though the result generalizes immediately to non-uniform distributions.) □

Proof of Corollary 1

Proof. (A) It follows straightly from the proof of 6 that $p^{\text{eq}}_{\text{all}} \leq p^{\text{eq}}_1$.

(B) By Lemma 1, θ^* is convex and decreasing; hence $G(s) = \theta^*(\kappa s) - \theta^*(s)$ is strictly increasing in s (single-crossing). Therefore $\Delta\text{CS}(s) = G(s) - (p^{\text{eq}}_{\text{all}} - p^{\text{eq}}_0)$ is strictly increasing. □

An numerical example with uniform G and F

I explicitly solve for the separate equilibrium under the additional assumption that F is uniform on $[\underline{v}, 1]$ and G is uniform on $[0, 1]$. Let s^* be the consumer that is indifferent between enabling cookies and disabling cookies.

$$\theta^{NT}(s^*) = 1 - \sqrt{2s^*}; \quad \theta^{T}(s^*) = 1 - \sqrt{2s^*(1-\underline{v})};$$

$$\theta^{T}(s^*) - \theta^{NT}(s^*) = \sqrt{2s^*} - \sqrt{2s^*(1-\underline{v})};$$

$$p_0^{eq} = \frac{s^*}{\int_0^{s^*} \frac{f(\theta(s))}{1-F(\theta(s))} dG(s)} = \frac{\sqrt{2s^*}}{2}$$

$$p_1^{eq} = \frac{G(\frac{(1-p_1^{eq})^2}{2(1-\underline{v})}) - G(s^*)}{\int_{s^*}^{\frac{(1-p_1^{eq})^2}{2(1-\underline{v})}} \frac{f(\theta^*(s))}{1-F(\theta^*(s))} g(s) ds} = \frac{\frac{(1-p_1^{eq})^2}{2(1-\underline{v})} - s^*}{\int_{s^*}^{\frac{(1-p_1^{eq})^2}{2(1-\underline{v})}} \frac{1}{\sqrt{2s(1-\underline{v})}} ds} = \frac{\frac{(1-p_1^{eq})^2}{2} - s^*(1-\underline{v})}{1 - p_1^{eq} - \sqrt{2s^*(1-\underline{v})}}$$

$$\Rightarrow p_1^{eq} = \frac{1 + \sqrt{2s^*(1-\underline{v})}}{3} \tag{1.8.3}$$

Condition 1.8.2 can then be written as

$$1 - \frac{3\sqrt{2s^*}}{2} = 1 - \sqrt{2s^*(1-\underline{v})} - \frac{1 + \sqrt{2s^*(1-\underline{v})}}{3} = \frac{2}{3} - \frac{4\sqrt{2s^*(1-\underline{v})}}{3}$$

$$\Rightarrow \frac{1}{3} = \frac{3\sqrt{2s^*}}{2} - \frac{4\sqrt{2s^*(1-\underline{v})}}{3}$$

The solution of s^* to the above equation is decreasing in $\underline{v}$. At $\underline{v} = 0.40487$, we have $s^* = 0.25$. It can be readily verified that for $\underline{v} \in (0.40487, 0.5]$, there exists a unique separating equilibrium.

Chapter 2

Selling a New Product to the Crowd

Abstract

We consider a monopolist's problem of selling a new product with common value to unit-demand buyers who privately observe conditionally independent signals. The seller commits to a mechanism that elicits information from buyers, chooses what to disclose, and sets prices. We first show that any use of private, buyer-specific messages can be replicated—and weakly improved—by a single public disclosure via a convexity/Blackwell argument. Second, under BIC, the mechanism reduces to a threshold allocation with flat transfers in the buy region, resulting in take-it-or-leave-it posted pricing; a DSIC implementation follows through leave-one-out prices that ignore a buyer's own report, maintaining the posted-price structure while making truth-telling dominant. These conclusions are robust in allowing seller-private information and in adding idiosyncratic taste components to buyers' valuations. Finally, we identify two economically relevant departures that break the convexity logic—and can make coarser disclosure strictly optimal even without capacity constraints—namely (i) resale/stockpiling (an intertemporal arbitrage "min" bottleneck) and (ii) message-dependent participation margins that render revenue non-affine in beliefs.

2.1 Introduction

In many digital markets, such as API access to large language models, real-time risk now-casting—the seller offers an effectively unlimited-supply product whose payoff to every buyer depends on the same uncertain state (model quality or outage risk). Each buyer privately observes noisy evidence about that state, while the seller can elicit information from buyers, choose what to disclose, and how prices are set. Should the seller aggregate and reveal the information or keep it coarse? Should posted prices depend on reported signals or be insulated from them? And how does the value of disclosure scale with the size of the customer base?

This paper studies these questions in a simplified environment with quasi-linear utility and price-only transfers, asking: which information policy and pricing rule maximize revenue when buyers hold private, conditionally independent signals about a common value. We explicitly solve the monopolist's revenue maximizing information and pricing mechanism.

Our analysis rests on two pillars: (1)any private message scheme can be replicated by a public message (indeed, full revelation) with weakly higher expected revenue, backed by the observation that revenue after a message equals a supremum of linear functionals of the induced posterior distribution and is therefore convex in beliefs. Thus, expected revenue is Blackwell-increasing in the informativeness of the public disclosure. (2) under BIC and quasilinear preferences, transfers are constant on any sure-sale set, so the optimal mechanism is a posted price conditional on the public message.

Our analysis draws on standard IC/envelope tools (Myerson, 1981; Rochet, 1987) but focuses on an unlimited-supply environment rather than auctions. The observation that more informative public signals weakly increase expected revenue connects directly to Johnson–Myatt's insight that information and positioning rotate demand by changing dispersion (Johnson and Myatt, 2006, 2015). Our convexity result further supplies the pricing geometry for why finer public information profitably "picks the right supporting price" state by state.

Relatedly, Ottaviani and Prat (2001) show in a nonlinear (second-degree) pricing model that, under affiliation between the state, buyers' private information, and a public signal, making the public experiment more informative weakly raises profit. Our posted-price setting delivers a clean, sufficient condition—convexity of R—that yields the same monotonicity without invoking affiliation. Lewis and Sappington (1994) analyze the supply of information to facilitate price discrimination: disclosure improves matching and segmentation but can erode extractable rents; our focus on public experiments and posted pricing provides a complementary sufficient condition under which public refinement is weakly beneficial. Finally, on the dual side, Roesler and Szentes (2017) study buyer-optimal information that disciplines seller pricing; we analyze the seller-optimal side under price-only transfers and show that full public revelation is weakly optimal in this posted-price environment.

2.2 Model and General Mechanism

There is a common value $v \in$ with prior π_0. Buyers $i = 1, \ldots, n$ have unit demand and observe private signals $s_i \in \mathcal{S}_i$ drawn i.i.d. from a density $f(s_i \mid v)$; conditional on v, signals are independent across i. The seller may elicit reports $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n)$ and optionally her own signal $\hat{t}$. The seller can send private, possibly randomized messages $m_i \in \mathcal{M}_i$ to buyers and choose allocations x_i and transfers t_i, where $(x,t) \sim \kappa : \prod_{i=1}^{n} \mathcal{S}_i \to [0,1] \times +$. We restrict attention to price-only transfers[1] Buyers' payoffs are quasi-linear. There is no capacity constraint.

[1]That is, a buyer pays a price if served and zero otherwise. This rules out report-contingent side bets or cross-report lotteries that can exploit correlation to extract rents in the sense of the Cremer–McLean framework; see Crèmer and McLean (1985, 1988). Formally, we assume $t_i = 0$ whenever $x_i = 0$. It's worthwhile to note that we do not assume away the possibility of randomization, though later we prove that randomization is not needed.

Mechanism. By direct revelation, buyers report $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n) \in \mathcal{S}^n$. A (possibly randomized) information policy

$$\Phi : \mathcal{S}^n \;\rightrightarrows\; \Delta(\mathcal{M}), \qquad \mathcal{M} := \mathcal{M}^0 \times \mathcal{M}_1 \times \cdots \times \mathcal{M}_n,$$

maps reports $\hat{s}$ into a distribution over messages $M = (M^0, M_1, \ldots, M_n)$, where M^0 is public and M_i is privately sent to i. An outcome rule

$$\kappa : \mathcal{M} \times \mathcal{S}^n \;\rightrightarrows\; \Delta\Big([0,1]^n \times \mathbb{R}^n\Big)$$

maps $(M, \hat{s})$ into allocations $x = (x_1, \ldots, x_n) \in [0,1]^n$ and transfers $t = (t_1, \ldots, t_n) \in \mathbb{R}^n_+$. Randomization inside Φ and κ is allowed.

Timing. Nature draws (v, s); buyers observe s_i and report $\hat{s}_i$; the mechanism draws $M \sim \Phi(\cdot \mid \hat{s})$; finally $(x, t) \sim \kappa(\cdot \mid M, \hat{s})$ are realized.

Seller's program. By Revelation-Principle, any implementable mechanism is outcome-equivalent to a direct, truthful-report, BIC mechanism that (possibly) uses messages derived from $\hat{s}$. Therefore, seller's problem can be written as:

$$\begin{aligned}
\text{maximize}_{\Phi,\kappa} \quad & \mathbb{E}\Big[\sum_{i=1}^{n} t_i\Big] \\
\text{subject to} \quad & U_i(s_i; \Phi, \kappa) \;\geq\; U_i(\tilde{s}_i \mid s_i; \Phi, \kappa), \;\; \forall i, \; \forall s_i, \tilde{s}_i, \\
& U_i(s_i; \Phi, \kappa) \;\geq\; 0, \;\; \forall i, \; \forall s_i, \\
& x_i \in [0,1] \;\; t_i \in \mathbb{R}_+
\end{aligned}$$

where the ex–interim utility under truthful reporting is

$$U_i(s_i; \Phi, \kappa) = \mathbb{E}\Big[v\, x_i - t_i \;\Big|\; s_i \text{ true}, \; \hat{s}_i = s_i, \; \hat{s}_{-i} = s_{-i}\Big],$$

and the utility from deviating to $\tilde{s}_i$ (holding others truthful) is

$$U_i(\tilde{s}_i \mid s_i; \Phi, \kappa) = \mathbb{E}\Big[\, v\, x_i - t_i \;\Big|\; s_i \text{ true}, \hat{s}_i = \tilde{s}_i,\ \hat{s}_{-i} = s_{-i} \Big].$$

The expectations integrate over v, $s_{-i} \sim f(\cdot \mid v)$, the induced message draw $M \sim \Phi(\cdot \mid \hat{s})$, and the outcome $(x, t) \sim \kappa(\cdot \mid M, \hat{s})$.

Remark on Commitment Assumption We assume the seller can commit ex ante to an information policy and to a pricing rule that maps the realized public message into a price. This commitment renders the problem static and underpins our geometry. Without commitment, the seller can re-optimize prices after observing messages, buyers anticipate this and shade or suppress information, and coarse disclosure may be optimal. Our results thus provide commitment benchmarks.

2.3 Analysis and Results

Our first result shows that it is W.L.O.G to restrict the monopolist's message space to public message. Here, a public message M is a single realization observed by all buyers. We first make explicit the convexity of the revenue functional that underlies the publicization.

Lemma 6 (Convexity of revenue in posterior beliefs)**.** *For a fixed buyer i and message, let $\theta = \mathrm{E}[v \mid s_i, m_i]$ and F be the cdf of θ. The maximal per-buyer revenue after the message is*

$$R(F) = \sup_{p\in} \; p\left(1 - F(p)\right), \tag{2.3.1}$$

which is a convex functional of F (a supremum of linear functionals).

Proof. Let $L_p(\mu)$ be the revenue at a fixed posted price p when the buyer's post-message type $\theta := \mathrm{E}(v|m_i, s_i)$ has posterior distribution μ. Then

$$L_p(\mu) \;=\; p\,\mu([p,\infty)) \;=\; p\int \mathbf{1}\{\theta \geq p\}\,d\mu(\theta).$$

Integration against a fixed measurable function is linear in the measure, hence for $\mu_\lambda := \lambda\mu_1 + (1-\lambda)\mu_2$,

$$L_p(\mu_\lambda) \;=\; \lambda L_p(\mu_1) + (1-\lambda)L_p(\mu_2).$$

(Equivalently in cdf form: with $F_\lambda = \lambda F_1 + (1-\lambda)F_2$, $p[1-F_\lambda(p)] = \lambda p[1-F_1(p)] + (1-\lambda)p[1-F_2(p)]$.) That is, for fixed p, $L_p(\mu)$ is linear in posterior distribution μ.

Using Step 1, and the fact that point-wise supremum of linear functionals is convex, we have

$$\begin{aligned} R(\mu_\lambda) &= \sup_{p\geq 0} L_p(\mu_\lambda) = \sup_{p\geq 0}\Big(\lambda L_p(\mu_1) + (1-\lambda)L_p(\mu_2)\Big) \\ &\leq \lambda \sup_{p\geq 0} L_p(\mu_1) + (1-\lambda)\sup_{p\geq 0} L_p(\mu_2) = \lambda R(\mu_1) + (1-\lambda)R(\mu_2), \end{aligned}$$

so R is convex. □

We are now ready to show it's W.L.O.G. to reduce the monopolist's information disclosure policy to public message. Moreover, fully revealing everyone's report is weakly optimal.

Lemma 7 (Publicization)**.** *For any BIC mechanism possibly with private messages m_i, there exists a BIC mechanism that uses only a public message $M = \Phi(\hat{s})$ and yields weakly higher expected revenue. In particular, full revelation is weakly optimal.*

Proof. As the first step, we show that public message can replicate the outcome of any private message policy. Consider the original (direct) mechanism (Φ_0, κ_0) that, given the profile of buyers' reports $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n)$, draws

$$m \;\sim\; \Phi_0(\cdot \mid \hat{s}) \;\in\; \Delta\Big(\textstyle\prod_{i=1}^n \mathcal{M}_i^{\mathrm{priv}}\Big),$$

where $m = (m_1, \ldots, m_n)$ are private messages. Then it implements $(x, t) \sim \kappa_0(m)$.

Define a new mechanism $(\Phi^{\text{pub}}, \kappa^{\text{pub}})$ with only a public message by publishing everything that was previously private:

$$M \;:=\; (m_1, \ldots, m_n) \;\in\; \mathcal{W} \;:=\; \prod_{i=1}^{n} \mathcal{M}_i^{\text{priv}}.$$

Formally, let $\Phi^{\text{pub}}(\cdot \mid \hat{s})$ be the pushforward of $\Phi_0(\cdot \mid \hat{s})$ under $m \mapsto M$, and set

$$\kappa^{\text{pub}}(\cdot \mid M) \;:=\; \kappa_0(\cdot \mid m).$$

Let $\Gamma_0(\cdot \mid \hat{s})$ and $\Gamma_{\text{pub}}(\cdot \mid \hat{s})$ denote the induced distributions of the outcome (x, t) under (Φ_0, κ_0) and $(\Phi^{\text{pub}}, \kappa^{\text{pub}})$, respectively. Then for every $\hat{s}$,

$$\Gamma_{\text{pub}}(\cdot \mid \hat{s}) = \int \kappa^{\text{pub}}(\cdot \mid M)\, \Phi^{\text{pub}}(dM \mid \hat{s}) = \int \kappa_0(\cdot \mid m)\, \Phi_0(dm \mid \hat{s}) = \;\Gamma_0(\cdot \mid \hat{s}).$$

Hence, conditional on any report profile, the distribution of outcomes (x, t) does not change after making every buyer's private message public. Therefore, all BIC/IR inequalities and expected revenue are preserved.

Fix a buyer i. After a realized public message $M = m$, buyer i's post-message type is $\theta \;\equiv\; \mathrm{E}[v \mid M = m, s_i]$ with conditional cdf F_m. Then the maximal per-buyer revenue achievable at message m equals

$$R(F_m) \;:=\; \sup_{p \in_+} \; p\left(1 - F_m(p)\right), \tag{2.3.2}$$

As the next step, consider any other public information policy $\tilde{\Phi}$ that is more informative than Φ^{pub} in the Blackwell sense: there exists a garbling device K such that $M \sim \Phi^{\text{pub}}(\cdot \mid \hat{s}, \hat{t})$ and $\tilde{M} \sim K(\cdot \mid M)$. For the buyer's post-message type we have that the conditional cdf given $\tilde{M} = \tilde{m}$ can be written as a mixture of the finer posteriors given M:

$$F_{\tilde{m}}(\cdot) \;=\; \int F_m(\cdot)\,\Pi(dm \mid \tilde{m}),$$

for the Bayes kernel $\Pi(dm \mid \tilde{m})$. Using (2.3.2) and the convexity of $R(\cdot)$ in F (Lemma 6), Jensen's inequality yields, conditional on $\tilde{M}$,

$$R(F_{\tilde{m}}) \;\geq\; \int R(F_m)\,\Pi(dm \mid \tilde{m}).$$

Taking expectations over $\tilde{M}$ gives

$$\mathrm{E}\Big[R(F_{\tilde{M}})\Big] \;\geq\; \mathrm{E}\Big[R(F_M)\Big]. \tag{2.3.3}$$

Therefore, any Blackwell refinement of a public message weakly increases the monopolist's optimal expected revenue when she re-optimizes prices after the message. Let $M^{\max} = (\hat{s}, \hat{t})$ be the finest public statistic (full revelation of reports). Since $M^{\max}$ Blackwell-refines every other public message, repeated application of (2.3.3) shows

$$\mathcal{R}(M^{\max}) \;\geq\; \mathcal{R}(M) \quad \text{for any public message } M.$$

Combining with the fact that any private message scheme can be replicated by a public one, expected revenue under the best public-message mechanism is weakly greater than under the original mechanism with private messages. This proves the lemma. □

Remark on smooth pricing rule In many markets, prices cannot (or should not) track signals finely even if richer signals exist. For example, anti-discrimination and transparency rules often restrict individual-level or highly personalized pricing; firms commit to group- or catalog-level prices. If the pricing rule must be a coarse function of M , i.e., there exists $m \neq m'$ such that $p(m) = p(m')$, then full revelation still weakly dominates because any coarse rule available under a coarse M can be replicated after full revelation by composing the fine signal with the original coarse map.

Fix a public message M. For ease of exposition, we can parametrize the buyer's (interim) type as a one-dimensional,scalar variable θ as

$$\theta \equiv \mu_i(M, s_i) := \mathrm{E}[v \mid M, s_i], \tag{2.3.4}$$

A direct mechanism then specifies $x(\cdot \mid M) : \Theta \to [0,1]$, a mapping from each buyer's interim type to the probability she gets the item, and $t(\cdot \mid M)$, the transfer.

The next lemma shows that with price-only transfers(so that $t_i = 0$ whenever $x_i = 0$, the optimal mechanism is post-price, so no randomization is needed to maximize revenue. The idea is to invoke Bauer's Maximum Principle which states that any continuous convex function defined on a compact, convex set attains its maximum at some extreme point of that set.

Lemma 8. *With unlimited supply and quasi-linear payoffs, the seller's objective is linear in $x(\cdot \mid M)$. Under monotonicity, an optimal x can be chosen as a $\{0,1\}$ step function: there exists a cutoff $\theta^*(M)$ with $x(\theta \mid M) = \mathbf{1}\{\theta \geq \theta^*(M)\}$.*

Proof. Fix a realized public message M and write $x(\theta) = x(\theta \mid M)$, $t(\theta) = t(\theta \mid M)$, and $U(\theta) = \theta x(\theta) - t(\theta)$. Let F denote the cdf of $\theta \mid M$.

Step 1: the objective is linear(thereby convex) in x. Under BIC in a one-dimensional type space Θ, U is absolutely continuous and satisfies the envelope $U'(\theta) = x(\theta)$ almost everywhere. Fix a baseline $\underline{\theta}$ and impose IR so $U(\underline{\theta}) = 0$. Then

$$U(\theta) = \int_{\underline{\theta}}^{\theta} x(z)\, dz, \qquad t(\theta) = \theta x(\theta) - U(\theta).$$

Take two feasible allocation rules x_1, x_2 and $\lambda \in [0,1]$, and define $x_\lambda = \lambda x_1 + (1-\lambda)x_2$. By the envelope, $U_\lambda = \lambda U_1 + (1-\lambda)U_2$, hence

$$t_\lambda(\theta) = \theta x_\lambda(\theta) - U_\lambda(\theta) = \lambda t_1(\theta) + (1-\lambda)t_2(\theta).$$

Therefore, the expected state-wise revenue of the seller $\mathcal{R}[x] := \mathbb{E}_F[t(\theta)]$ satisfies $\mathcal{R}[x_\lambda] = \lambda\mathcal{R}[x_1] + (1-\lambda)\mathcal{R}[x_2]$: the objective is linear in x.

Step 2: the feasible set of $x(\cdot)$ is convex. Let $\mathcal{X} := \{x : \Theta \to [0,1] \text{ non-decreasing}\}$ be the set of all feasible allocation rules. Since any point-wise convex combinations of x preserves $[0,1]$ and monotonicity, by definition, $\mathcal{X}$ is a convex set. By Claim 1(see Appendix), the extreme points of $\mathcal{X}$ is a step function valued on $\{0,1\}$.

Step 3: optimality at an extreme point. By Bauer's Maximum Principle, a convex objective over a convex feasible set attains its maximum at an extreme point. Therefore an optimal allocation rule is a step function valued on $\{0,1\}$. □

Remark 1 (Flatness on constant-x regions). By the envelope $U'(\theta) = x(\theta \mid M)$ and $t(\theta) = \theta x(\theta \mid M) - U(\theta \mid M)$. Hence on any interval where $x \equiv 1$ (buy region) we have $t'(\theta) = 0$, so t is constant there. On any interval where $x \equiv 0$ (no-trade region), $U'(\theta) = 0$ so U is constant; with price-only transfers and IR we set $t \equiv 0$ there.

A direct implication of Lemma 8 is that the optimal mechanism takes the form of a take-it-or-leave-it pricing conditional on M.

Corollary 2 (Take-it-or-leave-it pricing conditional on M)**.** *For a fixed public message M the mechanism reduces to a single price vector $p(M)$ on the buy-region $\{\theta \geq \theta^*(M)\}$; the payment cannot vary with a buyer's own report inside that region.*

Having pinned down the information disclosure and the allocation rule, we now proceed to solve for the optimal pricing rule. We show that the revenue-maximizing mechanism has a DSIC implementation.

Proposition 7 (DSIC implementation via leave-one-out posted prices)**.** *Let M_{-i} be a public statistic built from reports/signals of $-i$ only. For each buyer i, post the self-insensitive take-it-or-leave-it price*

$$p_i^*(M_{-i}) \in \arg\max_{p\geq 0} \underbrace{p\,[1 - F_{i,M_{-i}}(p)]}_{\text{per-buyer revenue given } M_{-i}},$$

where $F_{i,M_{-i}}$ is the cdf of $\theta_i = \mathbb{E}[V \mid M_{-i}, s_i]$ conditional on M_{-i}. Allocate iff $\theta_i \geq p_i^*(M_{-i})$ and charge $p_i^*(M_{-i})$ if allocated. Then truthful reporting is a dominant strategy (DSIC).

Proof. Apply cor:posted with M replaced by M_{-i}. Since the price shown to i ignores s_i, i cannot improve their terms by misreporting. Any information loss from excluding i can be mitigated by re-optimizing $p(\cdot)$ given the chosen statistic. □

Combining all the results we have so far, the following theorem provides a complete characterization of the optimal mechanism:

Theorem 3 (Optimal disclosure and pricing)**.** *One can w.l.o.g. restrict attention to mechanisms that (i) use a public disclosure rule M (indeed, full revelation of $(\hat{s}, \hat{t})$ is weakly optimal) and (ii) given M, employ take-it-or-leave-it pricing with a cutoff rule in $\theta = \mathrm{E}[v \mid M, s_i]$. The revenue after the message M is*

$$R(F_M) = \sup_p \; p\,[1 - F_M(p)], \qquad F_M \text{ the cdf of } \theta \mid M,$$

which is convex in F_M; hence, the expected revenue is (Blackwell) increasing in the informativeness of M. A DSIC implementation is obtained via prop:dsic.

In appendix B, we explicitly solve the optimal price schedule in uniform posterior example.

2.4 Robustness

This section shows that our baseline conclusions survive two natural generalizations. First, when the seller observes a private signal about the common value; Second, when the valuation of each buyer includes an idiosyncratic component beyond the common part (e.g., additive noise, multiplicative "taste" shifters, or more generally $v_i = f(\theta_i, \eta_i)$ with ID η independent of θ_i).

2.4.1 Seller-private information

Suppose, in addition, the seller observes a private signal Z (with realization z) about the common value V; buyers observe (M, s_i) where M is the public experiment of the designer (chosen ex ante) and s_i are conditionally independent signals given (V, Z). Mechanisms remain price-only, and the seller can commit ex ante to an information policy and pricing rule that may depend on z.

For any realized public state, per-buyer posted-price revenue is $R(F) = \sup_{p \geq 0} p\,[1 - F(p)]$, which is convex in the posterior law F (it is the supremum of affine maps $F \mapsto p[1 - F(p)]$). Write F_Z for the cdf of $\theta = \mathbb{E}[V \mid Z, s_i]$ conditional on Z.

If the seller reveals and conditions on Z, expected revenue is $\mathbb{E}\Big[R(F_Z)\Big]$. If the seller ignores Z, revenue equals $R(\bar{F})$ where $\bar{F} := \mathbb{E}[F_Z]$. By Jensen's inequality and convexity of R,

$$\mathbb{E}\Big[R(F_Z)\Big] \;\geq\; R\Big(\mathbb{E}[F_Z]\Big) \;=\; R(\bar{F}),$$

with strict inequality whenever the z-specific optimal prices differ with positive probability.

2.4.2 General valuation functions

We next consider the case when buyers' valuation also depends on a private component beyond just common value θ. Let each buyer's value be given by

$$v_i \;=\; u(\theta,\eta_i), \qquad \theta_i := \mathbb{E}[V \mid M, s_i], \quad \eta_i \overset{\text{i.i.d.}}{\sim} G \text{ independent of } (V, s_1, \ldots, s_n).$$

This nests additive tastes $v = \theta + \eta_i$, multiplicative tastes $v = a\,\theta$, and mixed forms $v = a_i\,\theta + \eta_i$. Fix a realized public message M and denote by F_M the cdf of $\theta \mid M$. The buy probability at a posted price p is

$$1 - H_{F_M}(p), \qquad H_{F_M}(p) \;:=\; \Pr\Big(u(\theta,\eta) \le p \,\Big|\, M\Big) = \int \underbrace{\Pr\Big(u(\theta,\eta) \le p\Big)}_{=:\,h_p(\theta)} \, dF_M(\theta).$$

Hence, the seller posts

$$p^*(M) \in \arg\max_{p \ge 0} \; p\,[1 - H_{F_M}(p)],$$

and implements the threshold rule $x_i = \mathbf{1}\{u(\theta_i,\eta_i) \ge p^*(M)\}$ with transfer $t_i = p^*(M)\,x_i$.

Convex revenue in beliefs $\Rightarrow$ Blackwell/Jensen monotonicity. For each fixed p, the map $F \mapsto p[1 - H_F(p)] \;=\; p\int(1 - h_p(\theta))\,dF(\theta)$ is affine in F. Therefore, the state-wise revenue functional

$$R_f(F) \;:=\; \sup_{p \ge 0} \; p\,[1 - H_F(p)]$$

is convex in F (supremum of affine maps). If a public experiment M' is a (Blackwell) garbling of M, then $F_{M'} = \mathbb{E}[\,F_M \mid M'\,]$ and Jensen's inequality yields

$$\mathbb{E}\Big[\,R_f(F_M)\,\Big] \;\ge\; \mathbb{E}\Big[\,R_f(F_{M'})\,\Big],$$

with strict inequality whenever optimal supporting prices differ on a set of positive probability. Thus, more informative public signals about the common component weakly increase expected revenue, for any f of the above form with i.i.d. η independent of θ.

DSIC implementation (leave-one-out). Define a self-insensitive price for buyer i by $p_i(M_{-i})$, where M_{-i} excludes i's report. Then buyer i's report cannot affect her own price and allocation, so truth-telling is a dominant strategy (DSIC). Acceptance remains the threshold event $\{f(\theta_i, \eta_i) \geq p_i\}$, and the posted-price/threshold structure is preserved. The DSIC revenue is weakly below the all-reports benchmark, but the gap vanishes as n grows (each single report has vanishing price influence).

2.5 When Full Disclosure Is Suboptimal

In the baseline model, expected revenue is a convex functional of the posterior law of the common component, so Blackwell refinement is weakly optimal. Here we identify two scenarios under which, even with unlimited supply, full revelation can be suboptimal when messages directly shift the extensive margin(participation), or when there is a secondary resale market.

2.5.1 Message-dependent participation

Consider a case when the message itself enters participation constraints: certain disclosed states (e.g., a "red" status) trigger compliance or operational lockouts that set demand to zero regardless of price. Because finer disclosure raises the frequency of such lockout states, coarser message can increase expected revenue by sustaining positive participation in marginal states. Let $\theta = \mathrm{E}[V \mid M]$ and suppose the message itself induces a participation multiplier $\sigma(M) \in [0, 1]$ (e.g., status tiers where "red" enforces a freeze: $\sigma = 0$, "amber" dampens entry: $\sigma = \alpha \in (0, 1)$, "green": $\sigma = 1$). For a posted price p, the statewise revenue becomes

$$\mathrm{Rev}(M) \;=\; \sigma(M)\cdot p\,[1-F_M(p)].$$

Even with unlimited supply, the mapping $F_M \mapsto \sigma(M)\,p[1-F_M(p)]$ is not affine in F_M because $\sigma(M)$ depends on the message. Full revelation tends to label more "bad" messages (low θ) as red (i.e., $\sigma = 0$), creating zero-revenue states that dominate the Jensen gain from finer pricing. By coarsening M (merging low and medium signals into an "amber" tier with $\sigma = \alpha > 0$), the seller trades some pricing precision for a first-order increase in participation on previously locked-out states:

$$\mathrm{E}[\mathrm{Rev}(\tilde{M})] \;=\; \mathrm{E}\Big[\sigma(\tilde{M})\,R(F_{\tilde{M}})\Big] \;\;>\;\; \mathrm{E}\Big[\sigma(M)\,R(F_M)\Big]$$

whenever the reduction in "red" mass is large enough. Intuitively, aggressive triage (full revelation) maximizes informativeness but destroys demand on bad news. Therefore a coarser map sustains market thickness and can raise expected revenue.

2.5.2 Intertemporal Resale

Consider a case when units are transferable across two dates, the seller sets prices sequentially after public messages, and a competitive secondary market allows frictionless resale. To preserve incentives, we implement leave-one-out (LOO) posted prices: buyer i's price in period t depends only on others' reports, never on i's own report, so truthful reporting is dominant (DSIC). We show that, using a simple example, that full-revelation is strictly dominated by partial revelation.

Claim 1. *Two periods $t \in \{1,2\}$; common-state qualities $\theta_t \in \{L,H\}$ i.i.d. with $\Pr(\theta_t = H) = \pi \in (0,1)$ and $0 < L < H$; unit mass of buyers each period; quasilinear utility; unlimited supply. In period t, a public message M_t is realized and the seller offers LOO posted prices: for each buyer* *i,*

$$p_{it} \;=\; p^*\Big(F_t^{(-i)}\Big), \qquad p^*(F) \in \arg\max_{p\geq 0} p\,[1-F(p)],$$

where $F_t^{(-i)}$ is the cdf of $\theta_{it} = \mathrm{E}[v_t \mid M_t^{(-i)}, s_{it}]$ conditional on $M_t^{(-i)}$, the statistic that excludes i's own report; allocation is $x_{it} = \mathbf{1}\{\theta_{it} \geq p_{it}\}$. A competitive secondary market implies the effective consumption price is

$$q \;=\; \min\{\,\bar{p}_1, \bar{p}_2\,\},$$

where $\bar{p}_t$ is the common LOO price in period t (with sufficiently many or a continuum of buyers, p_{it} is almost surely equal across i). Then:

1. *If M_t fully reveals θ_t, then $M_t^{(-i)} = M_t = \theta_t$ for all i, so $p_{it} = p^*(F_{\theta_t}) = \theta_t$ and*

$$\mathrm{Rev}^{\mathrm{FR}}(\theta_1,\theta_2) = 2\,\min\{\theta_1,\theta_2\}, \qquad \mathbb{E}[\mathrm{Rev}^{\mathrm{FR}}] = 2\Big(L + \pi^2(H-L)\Big).$$

2. *If M_t reveals nothing (same prior both periods), period-2 LOO best replies imply $\bar{p}_2 \geq \bar{p}_1$ (undercutting lowers q and total revenue); by symmetry $\bar{p}_1 \geq \bar{p}_2$, hence $\bar{p}_1 = \bar{p}_2 = p^\dagger \in \{L, H\}$ with*

$$\mathbb{E}[\mathrm{Rev}^{\mathrm{Pool}}] = 2\,\max\{L,\ \pi H\}.$$

For some parameters (e.g. $L=1$, $H=10$, $\pi=0.2$), $\mathbb{E}[\mathrm{Rev}^{\mathrm{Pool}}] > \mathbb{E}[\mathrm{Rev}^{\mathrm{FR}}]$, so garbling strictly dominates full revelation.

Intuitively,leave-one-out (LOO) pricing makes each buyer's offer depend only on others' reports, so a buyer cannot move her own price by misreporting (DSIC); conditional on whatever public statistic is used, the seller still picks the statewise monopoly price. With frictionless resale/stockpiling, consumption clears at the cheapest date, so the effective price across periods is

$$q = \min\{p_1, p_2\}.$$

Under full revelation the message equals the state each date, and the best response is

$p_t = \theta_t$; the minimum then becomes a bottleneck—one low realization forces a low q and depresses revenue in both periods. Under pooling, messages do not separate high from low, and best replies equalize prices across dates: any undercut lowers q (and total revenue), any overcut is irrelevant, so equilibrium has $p_1 = p_2 = p^\dagger$. With a constant price there is no intertemporal arbitrage (still $q = p^\dagger$), and expected revenue can exceed full revelation when state dispersion is sufficiently large. In short, resale imposes a law of one price across dates; full revelation makes buyers pay the minimum, whereas pooling sustains a stable common price that removes the arbitrage channel.

2.6 Concluding Remarks

In this paper, we study a monopolist's problem of selling an unlimited-supply good with common value v to unit-demand buyers who privately observe conditionally independent signals. We show that the mechanism design problem can be collapsed to choosing a public experiment and then posting a state-contingent price.

A key insight is that the state-wise posted-price revenue $R(F) = \sup_{p \geq 0} p[1 - F(p)]$ is convex in the posterior distribution F. Therefore, information improvement in Blackwell sense weakly raises expected revenue, with strict gains whenever different posterior states call for different supporting prices. This clarifies the role of information: IR caps per-buyer payments at $\mathbb{E}[V]$, but public refinement systematically removes selection losses created by privately held signals and pushes revenue up to (but never beyond) that cap. These conclusions are robust to adding seller-private information: publicly revealing the seller's signal and conditioning the posted price on it weakly increases expected revenue by the same convexity logic.

Two economically salient departures overturn this geometry. First, intertemporal resale couples dates through the no-arbitrage condition, making revenue non-affine in beliefs across messages;garbling can then strictly dominate full disclosure. Second, messagetriggered participation margins render demand discontinuous in disclosure, again breaking convexity and favoring coarser signals. These examples highlight that it is the cross-message nonlinearity—not capacity constraints per se—that restores a role for garbling and (potentially) lotteries. Future work should relax price-only transfers, examine limited commitment and dynamic reputation, introduce supply coupling or capacity, and endogenize buyers' information acquisition; each of these may similarly break convexity/publicization and reshape the optimal disclosure–pricing trade-off.

2.7 Appendix

2.7.1 Additional Proof Details

Lemma 9 (Extreme points). *If $x \in \mathcal{X}$ is an extreme point of $\mathcal{X}$, then x is $\{0,1\}$-valued almost everywhere.*

Proof. Suppose, to the contrary, that x takes values in $(0,1)$ on a set of positive measures. Because x is nondecreasing, there exist numbers $a < b$ with $0 < x(\theta) \in (0,1)$ for a.e. $\theta \in [a,b]$. Pick two disjoint subintervals $I_1 < I_2 \subset [a,b]$ and choose a small $\varepsilon > 0$ so that $x \pm \varepsilon$ stays within $[0,1]$ on $I_1 \cup I_2$ and does not violate monotonicity at the endpoints. Define

$$x^{\pm}(\theta) = \begin{cases} x(\theta) \mp \varepsilon, & \theta \in I_1, \\ x(\theta) \pm \varepsilon, & \theta \in I_2, \\ x(\theta), & \text{otherwise.} \end{cases}$$

Then $x^{\pm} \in \mathcal{X}$, $x^+ \neq x^-$, and $x = \frac{1}{2}(x^+ + x^-)$, contradicting extremality. Hence, any extreme x is $\{0,1\}$-valued a.e. $\square$

2.7.2 Analytical solution with uniform posteriors

We can solve analytically the optimal Leave-One-Out price schedule with uniform posterior distribution.

First, suppose there are two buyers, $i \in \{1,2\}$. The seller posts to buyer i a price that depends only on the other buyer's report, $p_i(s_{-i})$. Buyer i accepts iff her posterior expected value after combining her own signal and the other buyer's report is at least the posted price. For tractability, take each buyer's individual posterior $s_i \in (0,1)$ to be i.i.d. $U[0,1]$. With conditionally i.i.d. information, posterior odds multiply, so the combined posterior for buyer i has odds

$$\frac{\Pr(V=1 \mid s_i, s_{-i})}{\Pr(V=0 \mid s_i, s_{-i})} = \frac{s_i}{1-s_i} \cdot \frac{s_{-i}}{1-s_{-i}}.$$

Thus buyer i accepts p given $s_{-i} = s$ iff $s_i \geq T(p,s)$ where the acceptance threshold is

$$T(p,s) \;:=\; \frac{p(1-s)}{p(1-s)+(1-p)s} = \frac{1}{1+\frac{1-p}{p}\cdot\frac{s}{1-s}}. \tag{2.7.1}$$

Since $s_i \sim U[0,1]$, the conditional buy probability is $1 - G(T) = 1 - T$. Hence the seller's statewise revenue (conditioning on $s := s_{-i}$) is

$$r(p;s) \;=\; p\,[1-T(p,s)] \;=\; p\,\frac{(1-p)s}{p(1-s)+(1-p)s} \;=\; \frac{s\,p(1-p)}{s+p(1-2s)}. \tag{2.7.2}$$

Optimal statewise price. Fix $s \in (0,1)$. Maximizing $r(p;s)$ over $p \in (0,1)$ gives a simple closed form. Differentiate (2.7.2) and set the first-order condition to zero:

$$\frac{\partial r}{\partial p}(p;s) = 0 \quad\Longleftrightarrow\quad (1-2p)\Big[s+p(1-2s)\Big] \;=\; (1-2s)\Big(p-p^2\Big).$$

Solving this quadratic for p yields two roots; the unique root in $(0,1)$ is

$$p^*(s) \;=\; \frac{s-\sqrt{s(1-s)}}{2s-1}\;; \qquad p^*\Big(\tfrac{1}{2}\Big) := \tfrac{1}{2} \;\text{(continuous extension)} \tag{2.7.3}$$

It's straightforward to check that $p^*(s)$ is strictly increasing in s, with limits $p^*(0) = 0$ and $p^*(1) = 1$.

Ex-ante revenue. Ex-ante (integrating over $s_{-i} \sim U[0,1]$) the per-buyer revenue is

$$\mathbb{E}\Big[r\Big(p^*(s_{-i});\; s_{-i}\Big)\Big] \;=\; \int_0^1 r(p^*(s);\; s)\; ds,$$

and total revenue is twice this value by symmetry.

In general, if $s_i \sim G$ with density g on $[0,1]$, the statewise FOC becomes

$$1 - G\Big(T(p,s)\Big) \;=\; p\, g\Big(T(p,s)\Big) \cdot \frac{\partial T}{\partial p}(p,s), \qquad \frac{\partial T}{\partial p}(p,s) = \frac{s(1-s)}{\Big[s + p(1-2s)\Big]^2},$$

which typically requires numerical solution for $p^*(s)$.

We can easily extend the same case to N Buyers. Let the product of others' odds be

$$C_{-i} \;:=\; \prod_{j \neq i} \frac{s_j}{1-s_j}.$$

Under conditionally i.i.d. signals, Bayes' rule in odds form implies that independent pieces of evidence multiply their posterior odds. Hence C_{-i} is the odds contribution from all buyers other than i: up to the (constant) prior-odds factor, buyer i's combined posterior odds are $\frac{s_i}{1-s_i} \times C_{-i}$. If buyer i faces a posted price $p \in (0,1)$, she buys iff

$$s_i \;\geq\; T_N(p, C_{-i}) \;:=\; \frac{p}{p + (1-p)C_{-i}}.$$

With $s_i \sim U[0,1]$, the (statewise) per-buyer revenue is

$$r_i(p; C_{-i}) \;=\; p\,[1 - T_N(p, C_{-i})] \;=\; \frac{p(1-p)\,C_{-i}}{p + (1-p)C_{-i}}.$$

Maximizing $r_i(\cdot; C_{-i})$ over $p \in (0,1)$ yields

$$p_i^*(s_{-i}) \;=\; \frac{\sqrt{C_{-i}}}{1 + \sqrt{C_{-i}}} \qquad \text{and} \qquad r_i^*(s_{-i}) \;=\; \left(\frac{\sqrt{C_{-i}}}{1 + \sqrt{C_{-i}}}\right)^2 .$$

Symmetric specialization and limits. If $s_j = s$ for all $j \neq i$, then

$$C_{-i} = \left(\tfrac{s}{1-s}\right)^{N-1} \quad \Rightarrow \quad p_i^*(s;N) = \frac{\left(\tfrac{s}{1-s}\right)^{\frac{N-1}{2}}}{1+\left(\tfrac{s}{1-s}\right)^{\frac{N-1}{2}}} = \Lambda\!\left(\tfrac{N-1}{2}\log\tfrac{s}{1-s}\right), \;\; \Lambda(z) = \frac{1}{1+e^{-z}}.$$

Comparative statics in N. As N increases, the log-odds aggregate linearly: $p_i^*(s_{-i})$ becomes more extreme (closer to 0 or 1) whenever most others tilt the odds in the same direction; under informative signals, $p_i^* \to 1$ in the high state and $p_i^* \to 0$ in the low state as $N \to \infty$.

Chapter 3

The Quality-Variety Tradeoff in Algorithmic Content Curation

Abstract

We study a content platform's strategy for algorithmically screening user-generated content when the platform faces a fundamental trade-off between elevating high-quality content and preserving sufficient variety in the recommendation pool. In settings such as TikTok or Instagram—where prices are absent and content curation is achieved through algorithmic ranking rather than explicit market exchange—the platform's intervention operates entirely through admission and exposure decisions. Creators may invest in manipulative effort that raises their algorithmic score without improving underlying quality, and the platform's ability to deter such behavior depends critically on whether it can credibly commit to a screening policy. We compare three benchmark regimes—static selection, optimal screening under full commitment, and screening without commitment—and characterize the corresponding quality–variety frontiers. Embedding this framework into a competitive environment, we show that the value of commitment is shaped systematically by market structure: intensified on-platform competition reduces creators' post-admission rents and thereby mitigates manipulation incentives even without commitment, while stronger cross-platform competition raises creators' outside options and amplifies the platform's ex-post temptation to loosen standards. Together, these results provide a unified theoretical foundation for understanding how algorithmic governance, strategic creator behavior, and competitive pressures jointly shape content quality on modern digital platforms.

Bibliography

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman, "The economics of privacy," *Journal of economic Literature*, 2016, *54* (2), 442–92.

Ali, S Nageeb, Greg Lewis, and Shoshana Vasserman, "Voluntary disclosure and personalized pricing," in "Proceedings of the 21st ACM Conference on Economics and Computation" 2020, pp. 537–538.

Anderson, Simon P and Regis Renault, "Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model," *The RAND Journal of Economics*, 1999, pp. 719–735.

_ **and Régis Renault**, "Consumer information and firm pricing: negative externalities from improved information," *International economic review*, 2000, *41* (3), 721–742.

Armstrong, Mark and Jidong Zhou, "Consumer information and the limits to competition," *American Economic Review*, 2022, *112* (2), 534–77.

Au, Pak Hung and Mark Whitmeyer, "Attraction Versus Persuasion," *arXiv preprint arXiv:1802.09396*, 2018.

Bagnoli, Mark and Ted Bergstrom, "Log-concave probability and its applications," in "Rationality and Equilibrium," Springer, 2006, pp. 217–241.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, "Search, information, and prices," *Journal of Political Economy*, 2021, *129* (8), 2275–2319.

Board, Simon and Jay Lu, "Competitive information disclosure in search markets," *Journal of Political Economy*, 2018, *126* (5), 1965–2010.

Cho, In-Koo and David M Kreps, "Signaling games and stable equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

Condorelli, Daniele and Balazs Szentes, "Buyer-Optimal Platform Design," 2022.

Corniere, Alexandre De and Romain De Nijs, "Online advertising and privacy," *The RAND Journal of Economics*, 2016, *47* (1), 48–72.

Crèmer, Jacques and Richard P. McLean, "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 1985, *53* (2), 345–361.

_ and _ , "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1988, *56* (6), 1247–1277.

Gamp, Tobias, Daniel Kraehmer et al., "Deception and competition in search markets," Technical Report, University of Bonn and University of Mannheim, Germany 2018.

Hidir, Sinem and Nikhil Vellodi, "Privacy, Personalization, and Price Discrimination," *Journal of the European Economic Association*, 2021, *19* (2), 1342–1363.

Ichihashi, Shota, "Online privacy and information disclosure by consumers," *American Economic Review*, 2020, *110* (2), 569–95.

Johnson, Justin P. and David P. Myatt, "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 2006, *96* (3), 756–784.

_ and _ , "The Properties of Product-Line Prices," *International Journal of Industrial Organization*, 2015, *40*, 1–19.

Laziuk, Estelle, "iOS 14 Opt-in Rate-Weekly Updates Since Launch," https://www.flurry.com/blog/ios-14-5-opt-in-rate-idfa-app-tracking-transparency-weekly/ 2021.

Lewis, Tracy R. and David E. M. Sappington, "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 1994, *35* (2), 309–327.

Mauring, Eeva, "Search and price discrimination online," *CEPR Discussion Paper No. DP15729*, 2021.

Moraga-Gonz'alez, Jos'e Luis, S'andor Zsolt, and Wildenbeest Matthijs R., "Prices and heterogeneous search costs," *The RAND Journal of Economics*, 2017, *48* (1), 125–146.

Myerson, Roger B., "Optimal Auction Design," *Econometrica*, 1981, *49* (5), 1089–1122.

Ottaviani, Marco and Andrea Prat, "The Value of Public Information in Monopoly," *Econometrica*, 2001, *69* (6), 1673–1683.

Perloff, Jeffrey M and Steven C Salop, "Equilibrium with product differentiation," *The Review of Economic Studies*, 1985, *52* (1), 107–120.

Rhodes, Andrew and Jidong Zhou, "Personalized pricing and privacy choice," Technical Report, Working paper 2021.

Rochet, Jean-Charles, "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 1987, *16* (2), 191–200.

Roesler, Anne-Katrin and Balázs Szentes, "Buyer-Optimal Information Structures in Auctions," *Econometrica*, 2017, *85* (1), 67–92.

Wolinsky, Asher, "True monopolistic competition as a result of imperfect information," *The Quarterly Journal of Economics*, 1986, *101* (3), 493–511.

Zhong, Zemin, "Platform Search Design: The Roles of Precision and Price," *Marketing Science*, 2022.