Countering Price Discrimination with Buyer Information*

Philipp Strack† Kai Hao Yang‡

February 25, 2024

Abstract

We consider the welfare implications of buyer's information in a monopoly pricing setting, where the seller privately observes a signal for the buyer's value. The seller can commit to any mechanism that depends on the realizations of the seller's signal and the buyer's reported message. The buyer privately observes a signal about their own value, and then makes a participation decision and reports a message in the mechanism if participating. We characterize the buyer signals that maximize the buyer's surplus. This buyer optimal signal is (i) *privacy-preserving*. Namely, it is independent of the seller's signal, so that the buyer is immune to any form of price discrimination; and (ii) unit-elastic, so that any posted price mechanism with prices in the support of the distribution of the buyer's posterior expected value is optimal for the seller. We further use these signals to characterize the welfare outcomes that can be induced by some buyer signal, and show that the set of feasible welfare outcomes becomes smaller as the seller's signal becomes Blackwell more informative.

Keywords: Monopoly pricing, buyer information, privacy-preserving signal, price discrimination, surplus extraction.

JEL classification: C11, D42, D83,

*We thank Dirk Bergemann and Jidong Zhou for their valuable comments and suggestions. All errors are our own.

†Department of Economics, Yale University, Email: philipp.strack@yale.edu

‡School of Management, Yale University, Email: kaihao.yang@yale.edu

1 Introduction

Advancements of information technology facilitates information revelation to buyers and sellers prior to trading. Buyers have many sources of information, such as product ratings and platforms' product recommendation, which allow them to obtain product information and to be better-informed about their values. Meanwhile, sellers can receive information about buyers' values by collecting or purchasing consumer data, which contains buyers' characteristics and purchase histories that are informative about their tastes. The information structure between a buyer and a seller shapes their incentives, and thus the welfare outcomes, in many ways. On the one hand, buyers' information allow them to be more informed when making purchasing decisions, but at the same time changes the demand elasticity and thus the seller's pricing strategy. On the other hand, sellers' information about the buyer allows them to price discriminate and extract more surplus, which may or may not benefit the buyers. Moreover, the *correlation* between buyers' and sellers' information could further affect the degree in which buyers are price discriminated and how much information rent they can retain.

Specifically, consider a monopoly pricing setting where a buyer and a seller trade a single product. The buyer is risk-neutral with a quasi-linear preference, and has a unit demand while the seller has zero marginal cost. Suppose that a seller has access to certain consumer data (e.g., each consumer's website browsing history)—and hence some information about the buyer's value—and suppose that the seller can commit to any selling mechanism to use this information. From the buyer's perspective, the information they have about the product—and hence their own value—could affect the seller's mechanism in various ways.

For example, the seller could potentially use their information to engage in price discrimination. The degree of price discrimination depends on the type of information the buyer has about their value. If the buyer knows exactly what the seller knows, the seller can perfectly price discriminate them and extract all the surplus. If the buyer knows more than the seller does, the seller can still price discriminate but would sometimes leave the buyer some information rent. If the buyer's posterior expected value is a sufficient statistic for their belief about the seller's signal, and the buyer's and the seller's signals are correlated in a certain way, then the seller can extract (almost) all the surplus using a Crémer-McLean mechanism (Crémer and McLean 1988; McAfee and Reny 1992). In the meantime, if the buyer is completely uninformed, then there is no need for the seller to price discriminate, but can still fully extract surplus by simply charging a posted price that equals the buyer's expected value. In addition, the buyer's information determines the distribution of their posterior expected value, which in turn determines the demand elasticity and hence the prices charged by the seller. Given the various ways in which the buyer's information could affect the seller's mechanism, and hence the welfare outcome, it is thus natural to ask what kind of information is the best for the buyer, and more generally, what are the welfare implications of the buyer's information when the seller is informed about the buyer's value.

The main result (Theorem 1) of this paper shows that a certain class of buyer information performs particularly well. We say that a signal for the buyer is *privacy-preserving* if the buyer's signal, as a random variable, is independent of the seller's signal. Under any privacy-preserving signal, (almost) full surplus extraction in the sense of Crémer and McLean (1988) and McAfee and Reny (1992) is impossible, as there is no correlation between the buyer's posterior expected value and the buyer's private signal. Moreover, it is optimal for the seller to use a uniform posted price mechanism without any form of price discrimination. In other words, privacy-preserving signals allow the buyer to be "immune" to any form of price discrimination, as well as the Crémer-McLean-McAfee-Reny type surplus extraction. This is because realizations of seller's signal, despite being informative about the buyer's true value, is not informative at all about the buyer's *posteriors expected values* when the buyer receives a privacy-preserving signal. Furthermore, under a privacy-preserving signal, the buyer's incentives under any given mechanism does not depend on the seller's private signal either, as the seller's private signal is independent of the buyer's signal. Risk neutrality of the buyer then implies that it is without loss for the seller to use a mechanism that only depends on the buyer's reported expected value, which in turn implies that a uniform posted price mechanism is always optimal.

Theorem 1, together with a characterization of distributions of posterior expected values induced by privacy-preserving signal, which is given by Theorem 3 of Strack and Yang (2024), then allows us to completely characterize the buyer's optimal signal (Proposition 1), as well as the set of feasible welfare outcomes (Proposition 2). The buyer's optimal signal is a privacy-preserving signal that induces a distribution of posterior expected values corresponding to a unit-elastic demand curve, so that the seller is indifferent in charging any posted prices in the support. By selecting the smallest posted price, the seller's profit is minimized, and trade occurs with probability one, which implies that the buyer's surplus is maximized.

Proposition 2 further exploits the unit-elastic nature to generate a privacy-preserving signal that induces any welfare outcome where the seller's profit is above the minimized profit identified by Proposition 1, and the sum of the buyer's surplus and the seller's profit is at most the total surplus. Finally, Theorem 3 shows that the seller's minimized profit increases—and hence the set of feasible welfare outcomes shrinks—as the seller's signal becomes more informative.

Related Literature This paper is related to multiple streams of literature. In a monopoly setting,[1] several recent papers have explored the welfare implications of the information structure between the seller and the buyer. Bergemann, Brooks and Morris (2015) study the seller's signal about the buyer's value, assuming the buyer is fully informed, and characterizes welfare outcomes that can arise under all seller signals. Roesler and Szentes (2017) study the buyer's signals about their own values, assuming the seller is completely uninformed, and characterize welfare outcomes that can arise under all buyer signals. In contrast, we consider the welfare implications of the buyer's signal with an arbitrarily fixed seller signal. In a setting that features second-degree price discrimination (Mussa and Rosen 1978), Yang (2021) studies the welfare frontier among all buyer signal, assuming that the seller is completely uninformed and that the value is binary. Bergemann, Heumann and Morris (2023) characterize the buyer signal that maximizes the seller's profit. Bergemann, Heumann and Wang (2024) study the welfare implication of the seller's signal, assuming that the buyer is fully informed. In a general multi-product monopoly setting, Haghpanah and Siegel (2022) and Haghpanah and Siegel (2023) study the welfare implications of the seller's signal, assuming the buyer is fully informed. Deb and Roesler (forthcoming) study the max-min mechanism for the seller playing against the Nature who chooses a buyer signal to minimize the seller's profit, assuming the seller is completely uninfomed.

An important feature of our setting is that some buyer signal might be correlated with the seller's signal. If these signals are correlated in a way that satisfies the conditions identified by Crémer and McLean (1988) and McAfee and Reny (1992), the the buyer's surplus can be fully extracted. In this regard, the result that privacy-preserving signals are optimal for the buyer is thus of the same flavor as the min-max information structure in a informally

[1]See Elliott, Galeotti, Koh and Li (2023) and Armstrong and Zhou (2022) for the same types of questions in an oligopoly setting.

robust auction setting characterized by Brooks and Du (2021), in which the bidders' signals are independent of each other in order to avoid full surplus extraction.

Our results are built upon the characterization of privacy-preserving signals, which is due to Strack and Yang (2024). Relatedly, He, Sandomirskiy and Tamuz (2023) also study information structures where agents' signals are independent, which they call private private information structure. Our comparative statics result is based on a connection between the Blackwell order and the majorization order among the expected quantile functions, which in turn is based on Strassen's theorem (Strassen 1965) and is related to the majorization theory of Lorenz (1949), and, in particular Gutmann, Kemperman, Reeds and Shepp (1991).

Outline The rest of this paper is organized as follows. Section 2 sets up the model. Section 3 uses an example to illustrate the main trade-off of the problem. Section 4 characterizes the buyer optimal signal while Section 5 characterizes the feasible welfare outcomes and presents a comparative statics. Section 6 concludes.

2 Model

Preferences A monopolist sells a single object to a buyer with unit demand. The buyer's payoff is quasi-linear, so that their payoff is $pv - t$ when the probability of obtaining the product is $p \in [0,1]$, the payment to the seller is t, and the value is $v \in [0,1] =: V$. The seller has zero marginal cost, and seeks to maximize expected transfer.

Signals Both the buyer and the seller are informed about v through some private signals. A signal is a random variable defined on the probability space $(V \times [0,1]^2, \mathcal{B}, \mathbb{P})$, where $\mathcal{B}$ is the Borel σ-algebra on $V \times [0,1]^2$ and $\mathbb{P}$ is a product of some probability measure with CDF F and two Lebesgue measures on $[0,1]$. The first component v of the state $(v, \rho, r) \in V \times [0,1]^2$ denotes the buyer's value, whereas the second and the third components ρ and r are randomization devices. The seller's signal is a random variable $\tilde{\theta}$ that depends only on the buyer's value v and the randomization device ρ. The buyer's signal is a random variable $\tilde{s}$ that can depend on v and *both* randomization devices ρ and r.[2] Since V is a Polish

[2]In other words, we allow the buyer to always have access to an extra randomization device that is independent of the seller's. Note that although the buyer can always use a randomization device that is

space, so is the set of posteriors $\Delta(V)$ over V (under the weak-* topology). It is thus without loss to assume that $\tilde{s}$ and $\tilde{\theta}$ take values in $[0,1]$. A class of buyer signals will be of particular importance for our results, which we define below.

Definition 1. A buyer signal $\tilde{s}$ is said to be *privacy-preserving* if $\tilde{s}$ is independent of $\tilde{\theta}$.

A privacy-preserving signal $\tilde{s}$ allows the buyer to obtain information about v in a way that the seller cannot make any inferences about. For example, the fully informative signal is in general not privacy-preserving, as that signal must be correlated with the seller's signal $\tilde{\theta}$, unless $\tilde{\theta}$ is completely uninformative. In contrast, a completely uninformative signal is privacy-preserving. Furthermore, as shown in Strack and Yang (2024), the *quantile signal*

$$\tilde{q} := rF(v \mid \tilde{\theta}) + (1-r)F^{-}(v \mid \tilde{\theta}) \tag{1}$$

is privacy-preserving and reveals the buyer some information about v, where $F(\cdot \mid \theta)$ is the conditional distribution of v given the seller's signal realization θ, and $F^{-}(\cdot \mid \theta)$ is its left-limit. Let $\overline{F}$ be the distribution of the buyer's posterior expected value induced by the quantile signal $\tilde{q}$:

$$\overline{F}(x) := \inf\left\{q \in [0,1] : \mathbb{E}\left[F^{-1}(q \mid \tilde{\theta})\right] \geq x\right\},$$

for all $x \in [0,1]$.

Mechanisms Prior to observing the realization of $\tilde{\theta}$, the seller commits to a mechanism. By the revelation principle, it is without loss to restrict attention to incentive compatible and individually rational direct mechanisms, which asks the buyer to report the signal realization they observes and, conditional on the seller's signal realization, decides the probability of allocating the object and the transfer. That is, a mechanism is a pair (p,t), where $p : [0,1]^2 \to [0,1]$ and $t : [0,1]^2 \to \mathbb{R}$, so that the probability of selling the object equals $p(\theta, s)$ if the buyer's report is s and the seller's signal realization is θ, whereas the amount of transfer is $t(\theta, s)$ when the buyer's report is s and the seller's signal realization is θ.

A mechanism is incentive compatible if reporting the signal realization truthfully is opti-

independent of the seller's, the buyer's signal might still be correlated with the seller's signal, as both depend on v. Namely, this formulation only requires a *conditionally* independent signal to be feasible for the buyer.

mal for the buyer. That is,

$$\mathbb{E}[p(\tilde{\theta}, s) \mid s] \cdot \mathbb{E}[v \mid s] - \mathbb{E}[t(\tilde{\theta}, s) \mid s] \geq \mathbb{E}[p(\tilde{\theta}, s') \mid s] \cdot \mathbb{E}[v \mid s] - \mathbb{E}[t(\tilde{\theta}, s') \mid s]$$

for all $s, s' \in [0, 1]$, and is individually rational if

$$\mathbb{E}[p(\tilde{\theta}, s) \mid s] \cdot \mathbb{E}[v \mid s] - \mathbb{E}[t(\tilde{\theta}, s) \mid s] \geq 0\,,$$

for all $s \geq 0$. Given signals $\tilde{\theta}$ and $\tilde{s}$, the seller chooses an incentive compatible and individually rational mechanism to maximize $\mathbb{E}[t(\tilde{\theta}, \tilde{s})]$.

3 An Example

In what follows, we use a simple example to illustrate the nature of various types of the buyer's signals. Suppose that the buyer's value v is uniformly distributed on $[0, 1]$, and suppose that the seller receives a binary signal $\tilde{\theta} := \mathbf{1}\{v > 1/2\}$. Suppose first that the buyer's signal equals the seller's signal, so that the buyer also observes (and only) observes $\tilde{\theta}$. Then the seller can fully extract the buyer's surplus via perfect price discrimination. That is, the seller can charge a price of $1/4$ when $\tilde{\theta} = 0$, and charge a price of $3/4$ when $\tilde{\theta} = 1$.

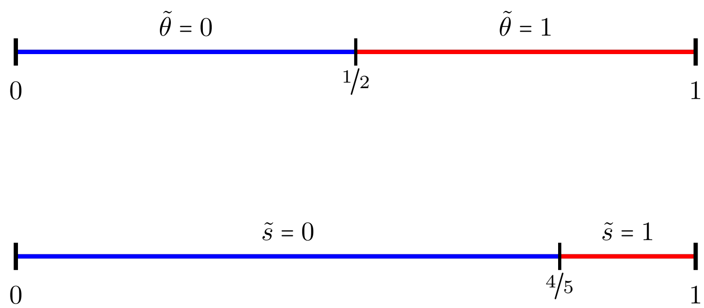

Figure 1: Signals $\tilde{\theta}$ and $\tilde{s}_1$

In the meantime, consider another signal $\tilde{s}$ for the buyer, where $\tilde{s} := \mathbf{1}\{v > 4/5\}$. Note that under this buyer signal, $\mathbb{E}[v \mid \tilde{s} = 0] = 2/5$ and $\mathbb{E}[v \mid \tilde{s} = 1] = 9/10$. Moreover, according to the buyer's posterior beliefs, $\mathbb{P}[\tilde{\theta} = 1 \mid \tilde{s} = 0] = 5/8$, and $\mathbb{P}[\tilde{\theta} = 1 \mid \tilde{s} = 1] = 1$. This signal is illustrated by Figure 1. Note that the full-surplus extraction condition in Crémer and McLean (1988) is satisfied under this signal, and hence the buyer's surplus is zero under any seller-optimal mechanism.[3]

To avoid surplus extraction, one could consider a buyer signal $(\tilde{s}, \tilde{\theta})$, where the buyer combines to two aforementioned signals together. In this case, the seller's private signal $\tilde{\theta}$ becomes public, and thus the full surplus extraction mechanism is not incentive compatible. Moreover, the buyer knows more than the seller does and thus the seller cannot perfectly price discriminate. Given this buyer signal, the seller's optimal mechanism is simply choosing an optimal posted price for each realized $\tilde{\theta}$. Note that $\mathbb{P}[\tilde{s} = 1 \mid \tilde{\theta} = 0] = 0$, and $\mathbb{P}[\tilde{s} = 1 \mid \tilde{\theta} = 1] = {}^1\!/_5$. Therefore, when $\tilde{\theta} = 0$, the seller optimally charges the buyer's expected value conditional on $\tilde{s} = 0$, which leaves the buyer zero surplus; whereas when $\tilde{\theta} = 1$, the seller optimally charges a price of ${}^2\!/_5$, so that trade occurs with probability 1, which leaves the buyer a surplus of ${}^1\!/_5 \cdot ({}^9\!/_{10} - {}^2\!/_5) = {}^1\!/_{10}$. Therefore, the buyer's total surplus equals $\mathbb{P}[\tilde{\theta} = 0] \cdot 0 + \mathbb{P}[\tilde{\theta} = 1] \cdot {}^1\!/_{10} = {}^1\!/_{20} > 0$ when their signal is $(\tilde{\theta}, \tilde{s})$.

The buyer can do better. Consider the signal that fully informs the buyer, so that the buyer learns v. Under this signal, the buyer is also fully informed about the seller's signal $\tilde{\theta}$. This becomes a standard price discrimination problem, where the seller chooses an optimal posted price conditional on each realization of $\tilde{\theta}$. When $\tilde{\theta} = 0$, the buyer's value is uniformly distributed on $[0, {}^1\!/_2]$, and the seller's optimal price is ${}^1\!/_4$. When $\tilde{\theta} = 1$, the buyer's value is uniformly distributed on $[{}^1\!/_2, 1]$, and the seller's optimal price is ${}^1\!/_2$. Together, the buyer's surplus is ${}^1\!/_2 \cdot {}^1\!/_{16} + {}^1\!/_2 \cdot {}^1\!/_4 = {}^5\!/_{32} > {}^1\!/_{20}$.

Since the seller essentially engages in third-degree price discrimination, and charges optimal posted prices separately when the buyer is fully informed, the buyer's surplus can be further improved by garbling v conditional on $\tilde{\theta}$. Conditional on $\tilde{\theta} = 0$, one can garble v using the method developed by Roesler and Szentes (2017), which leads to a signal under which trade occurs with probability 1, and the seller's optimal price is approximately 0.102. The buyer's surplus is thus approximately ${}^1\!/_4 - 0.102 = 0.148$. Together, the buyer's surplus is

[3] More precisely, let $p(\theta, s) = 1$, for all $\theta, s \in \{0, 1\}^2$ and let

$$t(\theta, s = 0) = \begin{cases} -\frac{5}{8}, & \text{if } \theta = 0 \\ \frac{3}{8}, & \text{if } \theta = 1 \end{cases} + \frac{2}{5}, \quad \text{and} \quad t(\theta, s = 1) = \begin{cases} 0, & \text{if } \theta = 0 \\ -1, & \text{if } \theta = 1 \end{cases} + \frac{9}{10}.$$

It then follows that (p, t) is incentive compatible and individually rational when the buyer's signal is $\tilde{s}$, and $\mathbb{E}[t(\tilde{\theta}, \tilde{s})] = {}^1\!/_2$.

$1/2 \cdot 0.148 + 1/2 \cdot 1/4 \approx 0.198 > 5/32$.

In the meantime, the buyer's can completely avoid price discrimination and induces the seller to charge a uniform posted price. Consider the quantile signal $\tilde{q}$ defined by (1). That is, $\tilde{q} = v$ if $\tilde{\theta} = 0$ and $\tilde{q} = v + 1/2$ if $\tilde{\theta} = 1$. In this case, the buyer's signal is privacy-preserving, and the seller's private signal $\tilde{\theta}$ becomes useless. The seller's optimal mechanism is therefore a uniform posted price, and thus the seller does not price discriminate at all under the signal $\tilde{q}$. Nonetheless, the buyer's surplus would be lower under signal $\tilde{q}$ than under the previous signal obtained by garbling v conditional on $\tilde{\theta}$. The optimal posted price under signal $\tilde{q}$ is determined by the distribution of the buyer's posterior expected value $\overline{F}$, which is a uniform distribution on $[1/4, 3/4]$. Therefore, under signal $\tilde{q}$, the seller optimally charges a price of $3/8$, and the buyer's surplus is $9/64 < 0.198$.

Nonetheless, the seller can do better by garbling the signal $\tilde{q}$. Since the distribution of posterior expected values under signal $\tilde{q}$ is uniform on $[1/4, 3/4]$, we can again apply the method of Roesler and Szentes (2017) and garble this signal so that the demand becomes unit-elastic. The garbled signal, call it s^*, induces trade with probability one, and gives a seller profit of approximately 0.274, and gives the buyer approximately a surplus of $1/2 - 0.274 = 0.225 > 0.198$. Thus, the buyer's surplus is higher under s^*. In fact, our main results implies that this is a signal that maximizes the buyer's surplus.

4 Buyer-Optimal Signal

Since the seller's signal $\tilde{\theta}$ might be informative about v, buyer's signal $\tilde{s}$ could be correlated with $\tilde{\theta}$. Existing results on monopoly pricing with a private value buyer may not apply. However, when the buyer's signal is privacy-preserving, standard results can be restored. In particular, as Lemma 1 below shows, whenever the buyer's signal $\tilde{s}$ is privacy-preserving, there exists a posted price mechanism that is optimal for the seller.

Lemma 1. *For any buyer signal $\tilde{s}$ that is privacy-preserving, there exists a posted price mechanism that maximizes the seller's profit.*

Proof. Consider any privacy-preserving signal $\tilde{s}$ for the buyer and consider any incentive compatible and individually rational mechanism (p,t). Let $\bar{p}$ and $\bar{t}$ be defined as

$$\bar{p}(s) \coloneqq \mathbb{E}[p(\tilde{\theta}, s)] \quad \text{and} \quad \bar{t}(s) \coloneqq \mathbb{E}[t(\tilde{\theta}, s)],$$

for all $s \in [0,1]$. Note that $(\bar{p}, \bar{t})$ is also a mechanism. Since $\tilde{s}$ is privacy-preserving, for any $s, s' \in [0,1]$,

$$\mathbb{E}[p(\tilde{\theta}, s') \mid s] \cdot \mathbb{E}[v \mid s] - \mathbb{E}[t(\tilde{\theta}, s')] = \mathbb{E}[p(\tilde{\theta}, s')] \cdot \mathbb{E}[v \mid s] - \mathbb{E}[t(\tilde{\theta}, s')] = \bar{p}(s') \cdot \mathbb{E}[v \mid s] - \bar{t}(s').$$

Therefore, the mechanism $(\bar{p}, \bar{t})$ is incentive compatible and individually rational if and only if (p,t) is incentive compatible and individually rational. Moreover, since $(\bar{p}, \bar{t})$ depends only on the buyer's reported signal realization, and since the buyer's payoff only depends on the true signal realization s through $\mathbb{E}[v \mid s]$ given the reports, any incentive compatibility implies that $(\bar{p}, \bar{t})$ must only depend on the buyer's expected posterior value $\mathbb{E}[v \mid s]$. The seller's problem is then equivalent to finding an incentive compatible and individually rational mechanism $(\tilde{p}, \tilde{t})$ that only asks the buyer to report their posterior expected value $\mathbb{E}[v \mid s]$ to maximize revenue. Standard arguments (see, for instance Myerson 1981; Riley and Zeckhauser 1983; Börgers 2015) then imply that there exists a posted price mechanism that is optimal. □

From Lemma 1, it follows that the seller would optimally choose not to price discriminate the buyer using their private signal, if the buyer's signal is privacy-preserving. In particular, while it is possible in general that the buyer's signal $\tilde{s}$ and the seller's signal $\tilde{\theta}$ might be correlated in a way that satisfies the condition identified by Crémer and McLean (1988) and McAfee and Reny (1992), and hence allow the seller to (almost) fully extract the buyer's surplus, this is not possible if the buyer's signal is privacy-preserving. More generally speaking, Lemma 1 highlights a particular benefit of privacy-preserving signals: they allow the buyer to be immune to price discrimination. In fact, as shown by Theorem 1 below, this implies that privacy-preserving signals are optimal for the buyer.

Theorem 1. *For any buyer signal $\tilde{s}$ and any seller-optimal mechanism (p,t), there exists another buyer signal s^* that is privacy-preserving, as well as a seller-optimal mechanism (p^*, t^*), such that the buyer's surplus is weakly higher and the seller's profit is weakly lower.*

Proof. Consider any signal $\tilde{s}$ and any optimal mechanism (p,t) for the seller, let $\pi(\tilde{s})$ and $\sigma(\tilde{s})$ be the seller's profit and the buyer's expected surplus under this signal and this mechanism, respectively. We show that for any signal $\tilde{s}$, there exists a privacy-preserving signal s^* and a seller-optimal mechanism (p^*, t^*) such that $\sigma(s^*) \geq \sigma(\tilde{s})$ and $\pi(s^*) \leq \pi(\tilde{s})$

Let $G(\cdot \mid \theta)$ be the distribution of posterior means $\mathbb{E}[v \mid \tilde{s}]$ conditional on the realization θ of $\tilde{\theta}$, it then follows that $G(\cdot \mid \theta)$ must be a mean-preserving contraction of $F(\cdot \mid \theta)$,[4] and hence, $G^{-1}(\cdot \mid \theta)$ is majorized by $F^{-1}(\cdot \mid \theta)$, for all $\theta \in [0,1]$ (see, e.g., Shaked and Shanthikumar 2007, Section 3.A). Let $\overline{G}^{-1}$ be defined as $\overline{G}^{-1}(q) := \mathbb{E}[G^{-1}(q \mid \tilde{\theta})]$, for all $q \in [0,1]$. Together with Fubini's theorem, it follows that $\overline{G}^{-1}$ is majorized by $\overline{F}^{-1}$, and thus $\overline{G}$ is a mean-preserving contraction of $\overline{F}$.

Meanwhile, note that the seller's profit under any mechanism must be bounded from below by the profit induced by choosing an optimal posted price conditional on each realization of θ. That is,

$$\pi(\tilde{s}) \geq \mathbb{E}\left[\max_{p \geq 0} p(1 - G^{-}(p \mid \theta))\right],$$

Moreover, for each realization θ of $\tilde{\theta}$, choosing an optimal posted price is equivalent to choosing an optimal quantity, and therefore,

$$\begin{aligned}
\pi(\tilde{s}) \geq \mathbb{E}\left[\max_{p \geq 0} p(1 - G^{-}(p \mid \theta))\right] = \mathbb{E}\left[\max_{q \in [0,1]} q G^{-1}(1 - q \mid \theta)\right] &\geq \max_{q \in [0,1]} q \mathbb{E}[G^{-1}(1 - q \mid \theta)] \\
&= \max_{q \in [0,1]} q \overline{G}^{-1}(1 - q) \\
&= \max_{p \geq 0} p(1 - \overline{G}^{-}(p)) \\
&=: \pi
\end{aligned}$$

Now consider the family distributions $\{G_{\pi}^{b}\}$, defined as

$$G_{\pi}^{b}(z) := \begin{cases} 0, & \text{if } z \leq \pi \\ 1 - \frac{\pi}{z}, & \text{if } z \in [\pi, b) \\ 1, & \text{if } z \in [b, 1] \end{cases},$$

for all $z \in [0,1]$, for all $\pi \in [0,1]$, and for all $b \in [\pi, 1]$. Since $\pi = \max_{p \geq 0} p(1 - \overline{G}^{-}(p))$, it

[4] To see this, note that for any realization θ of $\tilde{\theta}$, the distribution $G(\cdot \mid \theta)$ of $\mathbb{E}[v \mid \tilde{s}]$ is a mean-preserving contraction of the distribution of $\mathbb{E}[v \mid \theta, \tilde{s}]$, which in turn is a mean-preserving contraction of the distribution $F(\cdot \mid \theta)$ of v, conditional on θ.

follows that $G^1_\pi(z) \le \overline{G}(z)$ for all $z \in [0,1]$. Therefore, there exists a unique $b \in [\pi, 1]$, such that G^b_π is a mean-preserving contraction of $\overline{G}$, and thus is a mean-preserving contraction of $\overline{F}$. As a result, by Theorem 3 of Strack and Yang (2024), there exists a privacy-preserving signal s^* such that G^b_π is the distribution of $\mathbb{E}[v \mid s^*]$. Furthermore, by Lemma 1, a posted price mechanism is optimal, and, in particular, charging a posted price $p = \pi$ is an optimal mechanism for the seller, under which the outcome is efficient and hence $\sigma(s^*) = \mathbb{E}[v] - \pi$ and $\pi(s^*) = \pi \le \pi(\tilde{s})$.

Together,

$$\sigma(\tilde{s}) \le \mathbb{E}[v] - \pi(\tilde{s}) \le \mathbb{E}[v] - \pi = \sigma(s^*)\,,$$

as desired. $\square$

From Theorem 1, the problem of finding the buyer-optimal signal can then be reduced to finding the buyer-optimal privacy-preserving signals. Using the characterization of distribution of posterior expected values given by Strack and Yang (2024), which we state below in Theorem 2 for completion, this problem becomes very tractable.

Theorem 2. *Fix any seller signal $\tilde{\theta}$, a CDF G is the distribution of the buyer's posterior expected value $\mathbb{E}[v \mid \tilde{s}]$ under some privacy-preserving signal $\tilde{s}$ if and only if G is a mean-preserving contraction of $\overline{F}$.*

According to Theorem 1 and Theorem 2, finding the buyer-optimal signal is thus equivalent to finding a mean-preserving contraction G of $\overline{F}$ to maximize the buyer's surplus

$$\int_{p(G)}^{1} (1 - G(x))\,\mathrm{d}x$$

under the posted price mechanism with price $p(G)$, where $p(G)$ is the smallest optimal posted price for the seller when the buyer's distribution of posterior expected value equals G. Note that this problem has the same structure as the uninformed-seller problem analyzed by Roesler and Szentes (2017), except that the mean-preserving spread upper bound is $\overline{F}$ instead of F. As a result, we obtain the following:

Proposition 1 (Buyer-Optimal Signal). *Consider any privacy-preserving signal s^* under* *which the distribution of the buyer's posterior expected value* $\mathbb{E}[v \mid s^*]$ *is given by*

$$G_{\pi^*}^{b^*}(z) = \begin{cases} 0, & \text{if } z < \pi^* \\ 1 - \frac{\pi^*}{z}, & \text{if } z \in [\pi^*, b^*) \\ 1, & \text{if } z > b^* \end{cases},$$

where π^* *is the smallest* π *such that* $G_\pi^b \preceq_{\mathrm{MPS}} \overline{F}$ *for some* $b \geq \pi$*, and* b^* *is the unique* b *for which* $G_{\pi^*}^b \preceq_{\mathrm{MPS}} \overline{F}$*. A posted price mechanism with a price* π^* *is optimal for the seller. Under this mechanism, trade occurs with probability 1 and the buyer's surplus is higher than under any buyer signal and any seller-optimal mechanism.*

Proof. Consider any buyer signal $\tilde{s}$ and any seller-optimal mechanism. Let $\sigma(\tilde{s})$ be the buyer's surplus under this signal and this mechanism. Theorem 1 implies that there exists a privacy-preserving signal $\hat{s}$ such that the buyer's surplus is higher under $\hat{s}$ than under $\tilde{s}$. Let G be the distribution of the buyer's posterior expected value under signal $\hat{s}$. Since $\hat{s}$ is privacy-preserving, Theorem 2 implies that $G \preceq_{\mathrm{MPS}} \overline{F}$. Since $\hat{s}$ is privacy-preserving, Lemma 1 implies that a posted price mechanism is optimal for the seller. Let $\pi := \max_{p \geq 0} p(1 - G^-(p))$ be the seller's optimal profit given that the buyer's signal is $\hat{s}$. Then note that $G_\pi^1(p) \leq G(p)$ for all $p \in [0,1]$, and thus there exists a unique $b \in [0,1]$ such that $G_\pi^b \preceq_{\mathrm{MPS}} G \preceq_{\mathrm{MPS}} \overline{F}$. By the definition of π^*, it then follows that $\pi^* \leq \pi$.

Moreover, by Theorem 2 again, since $G_{\pi^*}^{b^*} \preceq_{\mathrm{MPS}} \overline{F}$, there exists a privacy-preserving signal s^* under which the buyer's posterior expected value $\mathbb{E}[v \mid s^*]$ is distributed according to $G_{\pi^*}^{b^*}$. Lemma 1 then implies that a posted price mechanism is optimal for the seller when the buyer's signal is π^*. Furthermore, since

$$\pi^* = \pi^*(1 - G_{\pi^*}^{b^*-}(\pi^*)) \geq p(1 - G_{\pi^*}^{b^*-}(p)),$$

for all $p \in [0,1]$, a posted price mechanism with price π^* is optimal for the seller. Under this mechanism, trade occurs with probability 1, and thus the buyer's surplus equals

$$\mathbb{E}[v] - \pi^* \geq \mathbb{E}[v] - \pi \geq \sigma(\tilde{s}).$$

Therefore, the signal s^* is optimal. □

Proposition 1 characterizes the buyer-optimal signals. According to Proposition 1, this buyer-optimal signal induces a distribution of posterior expected values under which the seller is indifferent in charging any posted prices $p \in [\pi^*, b^*]$. If the seller chooses the smallest optimal posted price π^* under this signal, then the buyer's payoff is maximized. If the seller chooses the highest posted price b^*, the the buyer's surplus is zero.

In the example discussed in Section 3, $\overline{F}(x) = 2(x - 1/4)$ for all $x \in [1/4, 3/4]$. In this case $\pi^* \approx 0.274$, $b^* \approx 0.625$, it is optimal for the seller's to charge a posted price π^*, in which case the buyer's surplus is approximately $\mathbb{E}[v] - \pi^* \approx 0.225$.

5 Feasible Welfare Outcomes

With Proposition 1, we could further characterize the set of feasible welfare outcomes that can be induced by a buyer signal $\tilde{s}$ for any fixed seller signal $\tilde{\theta}$. A welfare outcome is a pair (σ, π), where σ denotes the buyer's surplus and π denotes the seller's profit. For any fixed seller signal $\tilde{\theta}$, a welfare outcome (σ, π) is said to be *feasible* if there exists a buyer signal $\tilde{s}$ and an optimal mechanism (p, t) for the seller, such that

$$\sigma = \mathbb{E}[p(\tilde{\theta}, \tilde{s})v - t(\tilde{\theta}, \tilde{s})] \quad \text{and} \quad \pi = \mathbb{E}[t(\tilde{\theta}, \tilde{s})].$$

Proposition 2 below characterizes the feasible welfare outcomes.

Proposition 2 (Feasible Welfare Outcomes)**.** *A pair $(\sigma, \pi) \in [0,1]^2$ is a feasible welfare outcome if and only if $\pi \geq \pi^*$ and $\sigma + \pi \leq \mathbb{E}[v]$.*

Proof. For any buyer signal $\tilde{s}$ and for any seller-optimal mechanism, let $\sigma(\tilde{s})$ be the buyer's surplus and $\pi(\tilde{s})$ be the seller's profit. By Proposition 2, it must be that $\pi(\tilde{s}) \geq \pi^*$. Moreover, since total surplus is $\mathbb{E}[v]$, it must be that $\pi(\tilde{s}) + \sigma(\tilde{s}) \leq \mathbb{E}[v]$.

Conversely, for any (σ, π) such that $\pi \geq \pi^*$ and $\sigma + \pi \leq \mathbb{E}[v]$. By the definition of π^*, there exists $b \geq \pi$ such that $G_\pi^b \preceq_{\text{MPS}} \overline{F}$, and hence, by Theorem 2, there exists a privacy-preserving signal s^* such that the buyer's posterior expected value $\mathbb{E}[v \mid s^*]$ equals G_π^b under signal s^*. Lemma 1 then implies that a posted price mechanism is optimal for the seller. Then, by the definition of G_π^b, any $p \in [\pi, b]$ is an optimal posted price for the seller and the seller's profit is $\underline{\pi}$. Thus, $\pi(s^*) = \underline{\pi}$. Moreover, since any posted price $p \in [\underline{\pi}, b]$ is optimal for the seller, any $\sigma \in [0, \mathbb{E}[v] - \underline{\pi}]$ can be induced by some seller-optimal mechanism under the signal s^*. $\square$

From Proposition 2, any pair $(\sigma, \pi) \in [0,1]^2$ such that $\pi \geq \pi^*$ and $\sigma + \pi \leq \mathbb{E}[v]$ is a feasible welfare outcome. Compared to Corollary 1 of Roesler and Szentes (2017), when the seller is completely uninformed, the feasible welfare outcomes becomes smaller when the seller is informed by some signal $\tilde{\theta}$. In particular, the smallest possible seller profit π^* when the seller is informed by a private signal $\tilde{\theta}$ is higher than the smallest profit $\underline{\pi}$ when the seller is uninformed. In the example in Section 3, the lowest possible seller profit is $\pi^* = 1/4$ when the seller receives signal $\tilde{\theta}$, whereas the lowest possible seller profit $\underline{\pi}$ when the seller is completely uninformed, according to Roesler and Szentes (2017), is approximately $1/5$.

In fact, it is generally the case that when the seller's signal $\tilde{\theta}$ becomes more informative in Blackwell's sense, the profit lower bound π^* increases, as shown by Theorem 3 below.

Theorem 3. *Consider any pair of seller signals $\tilde{\theta}_1$ and $\tilde{\theta}_2$, where $\tilde{\theta}_2$ Blackwell dominates $\tilde{\theta}_1$. Let*

$$\overline{F}_i(x) := \inf\left\{q \in [0,1] : \mathbb{E}\left[F^{-1}(q \mid \tilde{\theta}_i)\right] \geq x\right\},$$

for all $x \in [0,1]$ and for all $i \in \{1,2\}$ be the distribution of the buyer's posterior belief induced by the quantile signal

$$\tilde{q}_i := rF(v \mid \tilde{\theta}_i) + (1-r)F^{-}(v \mid \tilde{\theta}_i),$$

for $i \in \{1,2\}$, respectively. Then

$$\overline{F}_2 \preceq_{\text{MPS}} \overline{F}_1 .$$

Proof. Since $\tilde{\theta}_2$ Blackwell dominates $\tilde{\theta}_1$, by Strassen's theorem (Strassen 1965), there exists a pair of seller signals $\hat{\theta}_1$ and $\hat{\theta}_2$ such that $\tilde{\theta}_i$ is Blackwell equivalent to $\hat{\theta}_i$ for $i \in \{1,2\}$, and that

$$F(x \mid \theta_1) = \mathbb{E}[F(x \mid \hat{\theta}_2) \mid \theta_1],$$

for all $x \in [0,1]$ and for all $\theta_1 \in [0,1]$, where $F(\cdot \mid \theta_i)$ is the conditional distribution of v conditional on the realization θ_i of $\hat{\theta}_i$, for $i \in \{1,2\}$. Now let

$$\hat{q} := rF(v \mid \hat{\theta}_2) + (1-r)F^{-}(v \mid \hat{\theta}_2)$$

be the quantile signal for $\hat{\theta}_2$. By Lemma 3 of Strack and Yang (2024), $\hat{q}$ is independent of $\hat{\theta}_2$. Moreover, according to Proposition 1 of Strack and Yang (2024), since $\hat{\theta}_2$ Blackwell dominates $\hat{\theta}_1$, $\hat{q}$ is independent of $\hat{\theta}_1$ as well. Then, conditional on any realization of θ_1 of $\hat{\theta}_1$ and any realization q of $\hat{q}$, the posterior expected value upon observing (q, θ_1) is given by

$$\mathbb{E}[F^{-1}(q \mid \hat{\theta}_2) \mid \theta_1] .$$

Therefore, the distribution of $\mathbb{E}[F^{-1}(\hat{q} \mid \hat{\theta}_2) \mid \theta_1]$ must be a mean-preserving contraction of $F(\cdot \mid \theta_1)$, for all $\theta_1 \in [0,1]$, which implies that (see, e.g., Shaked and Shanthikumar 2007, Section 3.A),

$$\int_0^q \mathbb{E}[F^{-1}(z \mid \hat{\theta}_2) \mid \theta_1] \, \mathrm{d}z \geq \int_0^q F^{-1}(z \mid \theta_1) \, \mathrm{d}z ,$$

for all $q \in [0,1]$ and for all $\theta_1 \in [0,1]$, which in turn implies that

$$\int_0^q \mathbb{E}[\mathbb{E}[F^{-1}(z \mid \hat{\theta}_2) \mid \hat{\theta}_1]] \, \mathrm{d}z \geq \int_0^q \mathbb{E}[F^{-1}(z \mid \hat{\theta}_1)] \, \mathrm{d}z, .$$

Moreover, by the law of iterated expectation, $\mathbb{E}[\mathbb{E}[F^{-1}(q \mid \hat{\theta}_2) \mid \hat{\theta}_1]] = \mathbb{E}[F^{-1}(q \mid \hat{\theta}_2)]$ for all $q \in [0,1]$. Together, $\mathbb{E}[F^{-1}(\cdot \mid \hat{\theta}_1)]$ majorizes $\mathbb{E}[F^{-1}(\cdot \mid \hat{\theta}_2)]$. Since $\hat{\theta}_i$ is Blackwell equivalent to $\tilde{\theta}_i$, $\mathbb{E}[F^{-1}(q \mid \hat{\theta}_i)] = \mathbb{E}[F^{-1}(q \mid \tilde{\theta}_i)]$ for all $q \in [0,1]$ and for all $i \in \{1,2\}$. As a result, $\overline{F}_1^{-1}$ majorizes $\overline{F}_2^{-1}$, which implies that $\overline{F}_2$ is a mean-preserving contraction of $\overline{F}_1$. □

From Theorem 3 and Theorem 1, it immediately follows that the seller's profit lower bound is higher when the seller's signal is Blackwell-more informative, as summarized below.

Corollary 1. *Consider any pair of seller signals $\tilde{\theta}_1$ and $\tilde{\theta}_2$, where $\tilde{\theta}_2$ Blackwell dominates $\tilde{\theta}_1$. Let π_1^*, π_2^* be the seller's lowest feasible profit given by Proposition 2 under signals $\tilde{\theta}_1$ and $\tilde{\theta}_2$, respectively. Then $\pi_1^* \leq \pi_2^*$.*

To illustrate, consider a numerical example where v is uniformly distributed on $[0,1]$. The seller's signal $\tilde{\theta}_p$ equals v with probability $p \in [0,1]$, and equals ε with probability $1-p$, where ε is independently and uniformly drawn from $[0,1]$.[5]

[5]Formally, let ξ, ζ be two independently and uniformly distributed random variables on $([0,1], \mathcal{B}, \lambda)$, with λ being the Lebesgue measure, and let

$$\tilde{\theta}_p(v, \rho) := \begin{cases} v, & \text{if } \xi(\rho) \leq p \\ F^{-1}(\zeta(\rho)), & \text{if } \xi(\rho) > p \end{cases}$$

In this setting, the seller's signal $\tilde{\theta}_p$ is increasing in p under the Blackwell order, where the seller fully learns the buyer's value when $p = 1$, and learns nothing when $p = 0$. The distribution $\overline{F}$ is given by

$$\overline{F}(z) = \begin{cases} (1-p)(1-\sqrt{1-2z}), & \text{if } z \in [0, 1/2) \\ p + (1-p)\sqrt{2z-1}, & \text{if } z \in [1/2, 1] \end{cases},$$

if $p > 1/2$, and

$$\overline{F}(z) = \begin{cases} (1-p)(1-\sqrt{1-2z}), & \text{if } z \in \left[0, \frac{1}{2}\left(1 - \frac{(1-2p)^2}{(1-p)^2}\right)\right) \\ \frac{(1-p)^2 z - \frac{1}{2}p^2}{1-2p}, & \text{if } z \in \left[\frac{1}{2}\left(1 - \frac{(1-2p)^2}{(1-p)^2}\right), \frac{1}{2}\left(1 + \frac{(1-2p)^2}{(1-p)^2}\right)\right) \\ p + (1-p)\sqrt{2z-1}, & \text{if } z \in \left[\frac{1}{2}\left(1 + \frac{(1-2p)^2}{(1-p)^2}\right), 1\right] \end{cases},$$

if $p < 1/2$.

Figure 2 plots the feasible welfare outcomes for $p \in \{0, 0.6, 0.8\}$. When $p = 0$, the feasible welfare outcomes is the set characterized by Roesler and Szentes (2017). As demonstrated by Figure 2, as p increases, the feasible welfare set becomes smaller. As $p \to 1$, the feasible welfare outcome converges to the full-surplus extraction outcome.

6 Conclusion

We study the welfare implications of information in a monopoly pricing setting. The seller privately observes a signal for the buyer's value, which the mechanism can condition upon. The buyer privately observes a signal, makes participation decisions, and reports messages in the mechanism if participating. For a fixed seller signal, a buyer signal may be correlated with the seller's signal in a way that allows for (almost) full surplus extraction. In addition, even without being fully extracted, the seller's signal potentially allows them to price discriminate the buyer. We show that the buyer's optimal signal must thus be privacy-preserving signal, under which the seller's signal becomes useless, and the buyer can be immune to price discrimination. We further characterize the buyer-optimal signal, as well as the feasible welfare outcomes for an arbitrarily fixed seller signal.

Several questions naturally follow after establishing the aforementioned results. First,

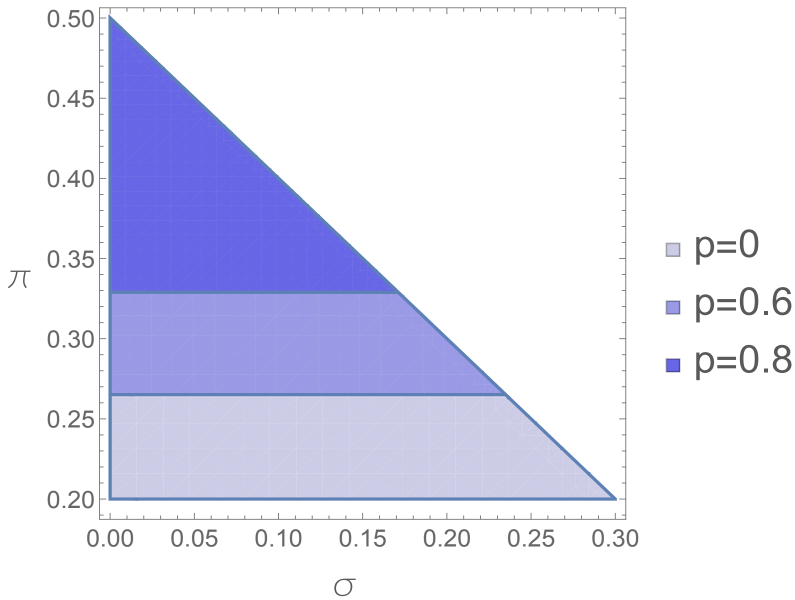

Figure 2: Feasible Welfare Outcomes

the buyer is allowed to be completely uninformed here. A natural question would be about the buyer-optimal signal when the buyer cannot choose to be completely uninformed and must receive some signal. Secondly, it is also relevant to explore a setting that feature competition. Finally, exploring the implications of information received by some agents while fixing information of other agents in other settings could also be economically relevant.

References

Armstrong, M. and J. Zhou (2022) "Consumer Information and the Limits to Competition," *American Economic Review*, 112 (2), 534–77.

Bergemann, D., B. Brooks, and S. Morris (2015) "The Limits of Price Discrimination," *American Economic Review*, 105 (3), 921–957.

Bergemann, D., T. Heumann, and S. Morris (2023) "Screening with Persuasion," Working Paper.

Bergemann, D., T. Heumann, and M. Wang (2024) "A Unified Approach to Second and Third Degree Price Discrimination," Working Paper.

Börgers, T. (2015) *An Introduction to the Theory of Mechanism Design*: Oxford University Press.

Brooks, B. and S. Du (2021) "Optimal Auction Design with Common Values: An Informationally-Robust Approach," *Econometrica*, 89 (3), 1313–1360.

Crémer, J. and R. P. McLean (1988) "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56 (6), 1247–1257.

Deb, R. and A.-K. Roesler (forthcoming) "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Review of Economic Studies*.

Elliott, M., A. Galeotti, A. Koh, and W. Li (2023) "Market Segmentation through Information," Working Paper.

Gutmann, S., J. H. B. Kemperman, J. A. Reeds, and L. A. Shepp (1991) "Existence of Probability Measures with Given Marginals," *The Annals of Probability*, 19 (4), 1781–1797.

Haghpanah, N. and R. Siegel (2022) "The Limits of Multi-Product Price Discrimination," *American Economic Review: Insight*, 4 (4), 443–458.

——— (2023) "Pareto Improving Segmentation of Multi-product Markets," *Journal of Political Economy*, 131 (6), 1546–1575.

He, K., F. Sandomirskiy, and O. Tamuz (2023) "Private Private Information," arXiv preprint arXiv:2112.14356.

Lorenz, G. G. (1949) "A Problem of Plane Measure," *American Journal of Mathematics*, 71 (2), 417–426.

McAfee, R. P. and P. J. Reny (1992) "Correlated Information and Mechanism Design," *Econometrica*, 60 (2), 395–421.

Mussa, M. and S. Rosen (1978) "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–307.

Myerson, R. (1981) "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1), 58–73.

Riley, J. and R. Zeckhauser (1983) "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, 98 (2), 267–289.

Roesler, A.-K. and B. Szentes (2017) "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107 (7), 2072–2080.

Shaked, M. and J. G. Shanthikumar (2007) *Stochastic Orders*: Springer.

Strack, P. and K. H. Yang (2024) “Privacy Preserving Signals,” Working Paper.

Strassen, V. (1965) “The Existence of Probability Measures with Given Marginals,” *Annuals of Mathematical Statistics*, 36, 423–439.

Yang, K. H. (2021) “Efficient Market Demands in a Multi-Product Monopoly,” *Journal of Economic Theory*, 197.