

Data Provision to an Informed Seller*

Shota Ichihashi

Alex Smolin

August 25, 2023

Abstract

A monopoly seller is privately and imperfectly informed about the buyer's value of the product. The seller uses information to price discriminate the buyer. A designer offers a mechanism that provides the seller with additional information based on the seller's report about her type. We establish the impossibility of screening for welfare purposes—i.e., the designer can attain any implementable combination of buyer surplus and seller profit by providing the same signal to all seller types. We use this result to characterize the set of implementable welfare outcomes, study the seller's incentive to acquire third-party data, and demonstrate the trade-off between buyer surplus and efficiency.

*Ichihashi: Queen's University, shotaichihashi@gmail.com. Smolin: Toulouse School of Economics, University of Toulouse Capitole and CEPR, alexey.v.smolin@gmail.com. For valuable suggestions and comments, we would like to thank Nageeb Ali, Ricardo Alonso, James Best, Teck Yong Tan, Jidong Zhou, and especially our discussant Daniele Condorelli, as well as seminar participants at Carnegie Mellon University (Tepper), University of California Irvine, Concordia University, Monash-NTU-RUC joint seminar, University of Illinois Urbana-Champaign, Penn State University, Queen's University, CETC 2022, EARIE 2022, Conference on Mechanism and Institution Design, 2nd Workshop on Contracts, Incentives and Information in Collegio Carlo Alberto, and TSE Digital Conference 2023. Smolin acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program (grant ANR-17-EURE-0010).

1 Introduction

The use of consumer data has become an important and ubiquitous aspect of the interactions between companies and consumers in the digital economy. Companies collect personal information and track consumers as they browse entertainment portals, post content on social media, and buy products on e-commerce websites. This data is used to adapt online services directly at the point of collection or later if transferred elsewhere. One prominent case is the use of consumer data for price discrimination: A seller may use data on its customers to learn about their preferences and tailor the prices of services and products via outright personalized pricing, providing discount coupons, or steering customers to a more expensive version of similar products.

As the importance of consumer data grows, various parties—such as policymakers, platforms, and consumers themselves—are attempting to control the flow of consumer data to sellers. However, the ability of these parties to control the allocation of data across sellers is likely to be limited because of information asymmetry—i.e., sellers may be privately informed about their market demands and customers. For example, a regulator might prefer that sellers within a certain group not use certain data, but the regulator may not know which sellers belong to that group. This raises questions about the scope of data provision to privately informed sellers and its potential impact on consumer surplus, seller profits, and overall efficiency.

Motivated by the above discussion, we study a model of data provision to a privately informed seller. The seller is a monopolist and has a product for sale. The buyer has a binary uncertain value for the product, which is either high or low. The seller’s interim belief that the value is high, which we call the seller’s type, is her private information. The seller may improve her pricing by obtaining additional information, which we model as a statistical signal that is informative about the value. We take a mechanism-design approach and consider a designer who offers a menu of signals to the seller. In practice, a menu corresponds to a restriction imposed by a regulator or a platform regarding what data sellers can use. Given a menu, the seller selects a signal and learns about the buyer’s value. Finally, the seller sets a price and the buyer decides whether to purchase the product.

Our focus is the set of all possible welfare outcomes the designer can implement by offering an arbitrary menu of signals. If the designer observes the seller’s type, the set of possible outcomes becomes the “surplus triangle,” as described by [Bergemann, Brooks, and Morris \(2015\)](#)—i.e., any division of the total surplus between the buyer and the seller can be achieved provided that the total surplus is not higher than the efficient surplus, the buyer surplus is nonnegative, and the seller’s profit is no lower than the profit she can achieve without additional information.

When the seller’s type is private, the designer can no longer attain the surplus triangle. For example, the first-best buyer surplus, which attains efficiency and gives no extra rents to the seller, requires that different types of sellers obtain different signals. However, if the designer were to implement such an outcome by offering a menu of those signals, some types would choose signals intended for other types in order to extract surplus from the buyer.

Our first main result shows that the designer cannot effectively screen the seller at all: An outcome is implementable through some menu of signals if and only if it is implementable by providing the same public signal to all seller types. This result illustrates the difficulty of eliciting the seller’s private information and providing personalized data. To prove the result, we develop novel structural properties of how personalized data provision can impact prices across different seller types and buyer values. We then use these properties to explicitly construct a public signal that replicates the equilibrium pricing behavior of any menu of signals.

We build on this result and geometrically characterize the set of implementable outcomes as the convex hull of an aggregate surplus function that averages across seller types. We show that not only can the designer focus on public signals, but these signals do not need to be complex: Any implementable outcome can be achieved using a signal with at most three signal realizations, and any extreme implementable outcome requires at most two signal realizations.

We use the above results to derive the surplus set in a closed form when the seller’s type is uniformly distributed. [Figure 1](#) depicts the set of implementable outcomes for observable types (light blue triangle) and unobservable types (dark blue area). The seller’s private information shrinks the implementable set. In this example, providing any signal reduces




Figure 1: The sets of implementable outcomes for observable and unobservable seller types when low value equals 1, high value equals 2, and types are uniformly distributed on $[0, 1]$. The horizontal axis and vertical axis represent the ex ante expected payoffs of the buyer and the seller, respectively.

the buyer's ex ante expected payoff because a large mass of seller types will use it to extract surplus from the buyer. Moreover, the only efficient outcome leads to zero buyer surplus. The right boundary of the surplus set is spanned by signals that reveal that the buyer has the high value with some probability, and symmetrically, the left boundary is spanned by signals that reveal the low value.

Our result on the impossibility of screening has three implications. First, the result highlights a tension between consumer protection and efficiency regarding the provision of consumer data to sellers. The provision of data may enable sellers to tailor pricing, benefit consumers, and enhance efficiency. At the same time, how much and what kind of data sellers should be allowed to use for such beneficial pricing will depend on the specific market conditions and the prior information each seller faces. However, in reality, a regulator would not be able to apply different rules to different firms, partly because of the informational friction we highlight. In such a situation, allowing some sellers to use a certain piece of information could increase total surplus, but other sellers may use the same information to merely extract consumer surplus. In some settings, this trade-off can be so stark that it may be impossible to attain efficiency without giving the entire total surplus to the seller.

Second, we study the seller's incentive to acquire third-party data. Here, third-party

data—such as the data a seller may purchase from a data broker—refers to the source of information the seller can acquire but is out of the designer’s control. We can model this as the seller’s choice to be privately informed about the buyer. From the designer’s perspective, the seller’s private information shrinks the set of implementable outcomes, and in particular, shifts the Pareto frontier downward. Thus, third-party data typically hurts the designer who cares about social welfare. In contrast, the seller’s incentives to acquire third-party data depend on which implementable outcome is selected. In particular, if the seller could propose to acquire additional “first-party” data directly from the buyer but the buyer has the authority to veto the proposal, the seller may prefer not to acquire third-party data. Indeed, because of adverse selection, the buyer may optimally veto acquisition of any additional data by a privately informed seller. As a result, the third-party data could crowd out the first-party data and reduce the overall information available to the seller.

Third, we show that the surplus set we characterize subsumes the set of equilibrium outcomes of various games in which the buyer and the seller exchange information with each other, including cheap talk communication, voluntary disclosure of the buyer’s value, and a request-consent protocol in which a privately informed seller chooses a signal subject to the buyer’s consent. This observation underlies the fact that our mechanism-design approach is useful for evaluating various communication protocols and data collection policies.

Our baseline model assumes that the buyer’s value is binary. On the one hand, this assumption is crucial for our result whereby the designer cannot effectively screen the seller. On the other hand, many of our key insights extend to the case of general multiple values. In particular, we show that the seller’s private information generally shrinks the set of implementable outcomes and creates a tension between consumer protection and efficiency.

Related Literature First and foremost, our paper relates to the recent theoretical literature on the impact of information under third-degree price discrimination or, equivalently, of market segmentation on market outcomes. In their seminal paper, [Bergemann, Brooks, and Morris \(2015\)](#) show that all individually rational outcomes can arise in a single-product monopoly setting. Their analysis was later extended to multiproduct markets ([Haghpanah and Siegel \(2022, forthcoming\)](#)); competitive markets ([Shi and Zhang \(2020\)](#)); [Rhodes and](#)

Zhou (2022); Elliott, Galeotti, Koh, and Li (2022)); and two-sided markets (Condorelli and Szentes (2022)).¹ We contribute to this literature by highlighting the importance of the seller’s private information, which is absent in those papers. The presence of such private information is relevant in practice for regulators and platforms, which aim to control the flow of consumer data but are likely to face information asymmetry vis-a-vis sellers. We show that the seller’s private information limits possible welfare outcomes and introduces the trade-off between consumer welfare and efficiency.

Second, our paper contributes to the literature on consumer privacy and privacy regulation (Acquisti, Taylor, and Wagman, 2016; Choi, Jeon, and Kim, 2019; Fainmesser, Galeotti, and Momot, 2022; Argenziano and Bonatti, 2021; Bergemann, Bonatti, and Gan, 2022). We highlight the difficulty of tailoring a privacy regulation to unknown market conditions, assess the impact of access to third-party data on first-party data collection, and compare the performance of several communication protocols. To focus on the role of the seller’s private information, we abstract away from other economic forces studied in the above papers, such as the use of data for product selection and service improvement as well as information externalities between consumers.

Our mechanism-design framework to study the regulation of a privately informed monopolist follows that of Baron and Myerson (1982). We employ this framework in the context of restricting the use of consumer data by a monopolist. Similar mechanism-design machinery in the context of information provision has recently been used by Kolotilin, Mylovanov, Zapechelnuyk, and Li (2017), Bergemann, Bonatti, and Smolin (2018); Smolin (forthcoming); and Yang (2022). All of these allow for a fully flexible way of designing information, following the Bayesian persuasion literature (Rayo and Segal (2010); Kamenica and Gentzkow (2011)).

¹Relatedly, Roesler and Szentes (2017) and Deb and Roesler (2022) analyze the informational impacts in single-product monopoly pricing and second-degree price discrimination, respectively.

2 Model

There is a seller and a buyer. The seller has a unit good for sale. The buyer's value v for the good is uncertain and either high or low: $v \in V \triangleq \{L, H\}$ with $H > L > 0$. The seller is privately informed about the value. The seller's private information is captured by $\theta \in \Theta \subseteq [0, 1]$ and represents the seller's belief that $v = H$. The type is distributed according to measure $F \in \Delta([0, 1])$.²

To improve her pricing, the seller seeks to acquire additional data. We model data as a statistical signal that can be arbitrarily informative about the value. Formally, a *signal* $\mathcal{I} = (S, \pi)$ consists of a set S of signal realizations s and a family of likelihood functions $\{\pi(\cdot|v)\}_{v \in V}$ over S . Where it does not cause confusion, we write conditional distribution $\pi(\cdot|v)$ as $\pi(v)$.

We adopt a mechanism-design approach and assume that signals are provided by a designer. At the outset, the designer posts a menu \mathcal{M} of signals to the seller. Then the game between the seller and the buyer proceeds as follows. First, the nature draws the seller's type θ according to prior distribution F and the buyer's value v according to θ .³ Second, the seller privately observes her type θ and chooses a signal $\mathcal{I} = (S, \pi) \in \mathcal{M}$. Third, the seller observes signal realization s drawn according to $\pi(v)$ and posts a price $p \in \mathbb{R}$ for the product. Finally, the buyer observes value v and price p and decides whether to buy the product. If the trade occurs, the buyer obtains payoff $v - p$ and the seller obtains payoff p . Otherwise, both players obtain zero payoffs.

For any given menu \mathcal{M} of signals, the solution concept is a perfect Bayesian equilibrium. Any equilibrium induces an *allocation rule* $a : V \times \Theta \rightarrow [0, 1] \times \mathbb{R}$, which specifies, for each pair of value and type, the probability of a trade and the expected payment from the buyer to the seller. We call the corresponding ex ante expected payoffs of the buyer and the seller as the *buyer surplus* and *seller profit*, respectively. A welfare outcome, or simply *outcome*, refers to a pair of buyer surplus and seller profit. Each allocation rule leads to a unique outcome, but a given outcome may come from multiple allocation rules. An allocation rule

²Given set X , we write $\Delta(X)$ for the set of all probability distributions on X .

³By Bayes' rule, this timing is equivalent to one in which the value is drawn before the type.

and an outcome are *implementable* if they can arise in an equilibrium of some menu.

Our goal is twofold. First, we study how the seller’s private information limits the designer’s ability to implement certain outcomes. To do so, in Section 3, we characterize the set of implementable outcomes and compare it with the case in which the seller’s type is observed. Second, we aim to derive implications on data policies for regulators and sellers when the sellers may have private information or face heterogeneous demand conditions. To this end, in Section 4, we use the characterization result to highlight a general tension between consumer protection and efficiency, and we examine the seller’s incentives to acquire external data that is outside the control of the designer. Finally, we show that the set of implementable outcomes we characterize subsumes the equilibrium outcomes of various communication protocols motivated by applications.

2.1 Discussion of Modeling Assumptions

Before proceeding with the analysis, we briefly discuss several modeling choices that are characteristic of our model.

Data as a Signal. As is common in the literature on information design, we model data as a statistical signal that is informative about the buyer’s value. This approach bypasses the technical aspects of data analysis and algorithmic implementation and focuses directly on the seller’s economic assessments of values. Different kinds of data, such as web cookies, geolocation data, and consumer behavior on the seller’s website, are indicative of consumer preferences and thus can be considered to be signals. In turn, a signal realization refers to the specific instance of data, such as the contents of a consumer’s cookies, his geolocation, or his web-browsing history. If we view the model as consisting of a continuum of buyers with heterogeneous values and sellers with heterogeneous types, any given signal induces a market segmentation with buyers belonging to the same segment if they have the same realized signal. However, since sellers have different types, the same signal can lead to different market segmentations depending on their types.

Role of the Designer. We abstract away from the designer’s objective and focus on all implementable outcomes. This enables us to make predictions that do not depend on the designer’s preferences. This approach is also useful when considering a player who controls consumer data with a particular objective. For example, the designer may be a regulator trying to maximize consumer welfare by restricting the data that sellers can use about consumers. Alternatively, the designer could be a platform or an information intermediary that provides a seller with information about buyers. Specifying the designer’s objective or a specific data collection protocol will select certain implementable outcomes.

Informed Seller. The seller’s private information about the value of the product to the buyer can come from two sources. First, the seller may have some information about the buyer that is beyond the designer’s control. Information could come from technological constraints like knowledge of the consumer’s IP address or from uncontrolled consequences of online interactions such as purchasing decisions. Second, the seller may be more informed than the designer about the general quality of their product. If the seller has a high-quality product, they believe that the consumer’s value is more likely to be high.

Price Discrimination. We assume that the seller uses data for third-degree price discrimination. In practice, while sellers may be reluctant to display different base prices to different consumers, there are at least two common and indirect ways to price discriminate in the digital economy. First, sellers may offer personalized discounts to consumers by setting the base price high and changing the size or frequency of discounts to implement discriminatory pricing. Second, sellers may offer personalized recommendations that point consumers toward similar products that vary in price. By maintaining a large inventory of such products, sellers can effectively engage in price discrimination.

3 Surplus Set Characterization

To study how the seller’s private information affects implementable outcomes, we begin our analysis with the benchmark case in which the seller’s type is observable to the designer. We then proceed to the primary scenario of unobservable types and present our main results.

3.1 Observable Seller Type

First, suppose that the seller’s type is deterministic, i.e., $\Theta = \{\theta_0\}$. Denote the maximum feasible total surplus by $\overline{W}(\theta_0) \triangleq \theta_0 H + (1 - \theta_0)L$ and the seller’s profit from optimal uniform pricing by $\underline{\Pi}(\theta_0) \triangleq \max\{\theta_0 H, L\}$. Clearly, any feasible welfare outcome (U, Π) must belong to the “surplus triangle” characterized by constraints $U \geq 0$, $\Pi \geq \underline{\Pi}(\theta_0)$, and $U + \Pi \leq \overline{W}(\theta_0)$. [Bergemann et al. \(2015\)](#) demonstrate that *any* such welfare outcome can be achieved by some signal.

Claim 1. (Bergemann, Brooks, Morris (2015)) *If the seller type is commonly known to be θ_0 , then outcome (U, Π) is implementable if and only if $U \geq 0$, $\Pi \geq \underline{\Pi}(\theta_0)$, and $U + \Pi \leq \overline{W}(\theta_0)$.*

This result immediately generalizes to the case in which the seller’s type is drawn according to distribution F but the designer observes the realized type. The aggregate set of implementable outcomes is a Minkowski average of the surplus sets for each realized type. For each type θ , by [Claim 1](#), the implementable outcomes satisfy three linear constraints, two of which feature type-dependent terms $\underline{\Pi}(\theta)$ and $\overline{W}(\theta)$. Denote their aggregate values averaged across types by

$$\underline{\Pi} \triangleq \int_0^1 \underline{\Pi}(\theta) dF(\theta), \tag{1}$$

$$\overline{W} \triangleq \int_0^1 \overline{W}(\theta) dF(\theta). \tag{2}$$

From the ex ante perspective, $\underline{\Pi}$ is the profit the seller can guarantee, \overline{W} is the maximum feasible total surplus, and 0 is the welfare level the buyer can guarantee. Hence, the aggregate welfare outcome must belong to a triangle outlined by these constraints. At the same time, as in the case of a known type, the converse is also true (the omitted proofs are in Appendix).

Claim 2. (Observable type) *If the seller type is commonly known and distributed according to F , then outcome (U, Π) is implementable if and only if $U \geq 0$, $\Pi \geq \underline{\Pi}$, and $U + \Pi \leq \overline{W}$.*

The result implies that any efficient outcome is implementable as long as the seller’s profit exceeds the profit under no additional data. One notable implementable outcome is the buyer-optimal outcome $(\overline{W} - \underline{\Pi}, \underline{\Pi})$, under which the allocation rule is efficient but the

seller's profit stays at the level of no additional information. At this outcome the buyer obtains all surplus created by the data. As a result, if the seller's type is observable, there is no inherent trade-off between consumer protection and efficiency.

However, the buyer-optimal outcome generally requires that different seller types obtain different signals. When the seller type is her private information this outcome may not be implementable, and the trade-off between consumer surplus and efficiency reappears. To see this, suppose that the seller's type is either $\theta_1 \in (0, L/H)$ or $\theta_2 \in (L/H, 1)$ —i.e., the optimal uniform price is L for type θ_1 and H for type θ_2 . At the buyer-optimal outcome, type θ_1 must not benefit from the data, because it is already willing to set price L . In contrast, for the outcome to be efficient, type θ_2 must be provided with an informative signal that reveals that the buyer's value is H with a positive probability.⁴ However, in that case type θ_1 could strictly benefit from the signal of θ_2 , which leads to a contradiction.

This example shows that the seller's private information restricts the set of implementable outcomes. In the next section, we analyze the structure of the seller's incentive compatibility constraints and characterize the implementable allocation rules and welfare outcomes.

3.2 Unobservable Seller Type

From now on we focus on the case in which the seller's type is unobservable and characterize the set of implementable outcomes. We show that the seller's private information leads to an extreme adverse selection: Any outcome with heterogeneous data allocation, in which different types self-select into different signals, can be implemented by providing all types with the same public signal.

An important class of menus is a class of *direct mechanisms*. A direct mechanism is a menu of *direct signals* indexed by θ so that $\mathcal{I}(\theta)$ sends two signal realizations, $S = \{s_L, s_H\}$, and the likelihood functions $\pi(\theta)$ are such that each type θ is willing to choose signal $\mathcal{I}(\theta)$ and to set prices $p = L$ and $p = H$ after observing signal realizations s_L and s_H of $\mathcal{I}(\theta)$, respectively. Interpreting signal realization s_v as a recommendation to set price $v \in \{H, L\}$, we can view a direct mechanism as sending a value-dependent price recommendation based

⁴By Bayes plausibility, some signal realization s will cause type θ_2 to believe that value H is even more likely than θ_2 . To cause type θ_2 to price efficiently, signal realization s must reveal that the value is H .

on the seller’s reported type. The following result shows that the designer can without loss of generality focus on direct mechanisms (the proof is standard and omitted; it follows the revelation principle argument of [Myerson \(1982\)](#) and [Bergemann et al. \(2018\)](#)).

Claim 3. (Direct Mechanisms) *An allocation rule is implementable if and only if it is implementable by a direct mechanism.*

Given a direct mechanism, we can parameterize the direct signal $\mathcal{I}(\theta)$ for each type θ by probabilities $\alpha(\theta)$ and $\beta(\theta)$ with which the signal sends realization s_H conditional on values H and L , respectively. Without loss of generality, we assume $\beta(\theta) \geq \alpha(\theta)$.⁵ We can then express signal $\mathcal{I}(\theta)$ in matrix form as

$$\begin{array}{c|cc} \mathcal{I}(\theta) & s_L & s_H \\ \hline v = L & 1 - \alpha(\theta) & \alpha(\theta) \\ v = H & 1 - \beta(\theta) & \beta(\theta) \end{array} . \quad (3)$$

For truth-telling to be optimal, each type should prefer her own signal $\mathcal{I}(\theta)$ to all other alternatives. The value of a signal depends on the type’s response to recommendations, which in turn depends on the posterior beliefs the recommendations induce. Because $\beta(\theta) \geq \alpha(\theta)$, the posterior belief rank is the same for all types: The posterior probability of $v = H$ is higher after observing s_H than after observing s_L . As such, no type would be willing to swap the pricing decisions—i.e., set $p = H$ after s_L and set $p = L$ after s_H . Hence, the relevant incentive constraints are those under which the seller misreports the type and follows the recommendation. Type θ ’s profit after such a deviation to type θ' is

$$\begin{aligned} \Pi(\theta, \theta') &\triangleq (1 - \theta)(1 - \alpha(\theta'))L + \theta((1 - \beta(\theta'))L + \beta(\theta')H) \\ &= (1 - \alpha(\theta'))L + \theta(\alpha(\theta')L + \beta(\theta')(H - L)). \end{aligned} \quad (4)$$

Incentive compatibility requires that $\Pi(\theta, \theta) \geq \Pi(\theta, \theta')$ for all $\theta, \theta' \in \Theta$. This property must hold in any direct mechanism and allows us to pin down the structural properties of any implementable allocation rule.

⁵Also, without loss of generality, in what follows we set $\alpha(1) = \beta(1) = 1$.

Proposition 1. (Allocation Properties) *In any direct mechanism, for any $\theta_1, \theta_2, \theta_3 \in \Theta$ such that $\theta_1 \leq \theta_2 \leq \theta_3$, the following hold:*

1. (Monotonicity) $\alpha(\theta_1) \leq \alpha(\theta_2)$ and $\beta(\theta_1) \leq \beta(\theta_2)$; and
2. (Relative Impact) $(\beta(\theta_3) - \beta(\theta_2))(\alpha(\theta_2) - \alpha(\theta_1)) \leq (\beta(\theta_2) - \beta(\theta_1))(\alpha(\theta_3) - \alpha(\theta_2))$.

The equilibrium properties in [Proposition 1](#) are direct consequences of the seller’s local incentive constraints and do not necessarily hold if the seller’s type is observable. Both properties can be interpreted in terms of comparative statics of the seller’s behavior with respect to her type. The first property means that the higher the type—i.e., the more likely high-value buyers are in the seller’s market—the greater the probability of a higher price for all buyers, irrespective of the seller’s additional information. In other words, regardless of the data-provision mechanism, high-value buyers impose negative externalities on low-value buyers because the presence of high-value buyers increases the likelihood that low-value buyers will face higher prices.

The second property evaluates the relative price impact of an increase in the seller’s type on low-value and high-value buyers, and is more easily interpreted when viewed as ratio monotonicity—i.e., $\frac{\alpha(\theta_3) - \alpha(\theta_2)}{\beta(\theta_3) - \beta(\theta_2)} \geq \frac{\alpha(\theta_2) - \alpha(\theta_1)}{\beta(\theta_2) - \beta(\theta_1)}$. Essentially, it states that the seller type’s increasing disproportionately affects low-value buyers at higher types compared with lower types. This property follows from the optimal change in the seller’s data strategy. Higher seller types value and primarily choose signals in the menu that help identify low-value buyers, because such signals are more likely to change the types’ prior behavior. These signals pool low-value buyers with high-value buyers at high-price recommendations, thus exacerbating the price impact of the type’s increase on those buyers. In contrast, lower seller types value and choose signals that help identify high-value buyers. These signals pool low-value buyers with high-value buyers at low-price recommendations, thus mitigating the price impact of the type’s increase on those buyers.

The allocation properties of [Proposition 1](#) outline the constraints imposed by the seller’s private information in direct mechanisms. By [Claim 3](#), this result applies to any mechanism, with $\alpha(\theta)$ and $\beta(\theta)$ being interpreted as the equilibrium pricing probabilities of different types. In particular, it implies that if the designer provides a single public signal to all

types, they optimally respond in such a way that the induced pricing behavior conforms with [Proposition 1](#). It turns out that the opposite is also true: The allocation rule of *any* mechanism can be replicated by providing a single public signal.⁶

Proposition 2. (Public Signals) *Any implementable allocation rule can be implemented by a menu with a single signal in it.*

Proof. The proof is constructive. Take any implementable allocation rule. [Claim 3](#) ensures that the allocation rule can be implemented by direct mechanism $(\alpha(\theta), \beta(\theta))_{\theta \in \Theta}$. We construct a public signal $\hat{\mathcal{I}}$ that implements the same outcome as $(\alpha(\theta), \beta(\theta))_{\theta \in \Theta}$. The signal realization space of $\hat{\mathcal{I}}$ is $S = [0, 1]$. The likelihood function π is such that for all $x \in \Theta$, $\Pr(s \leq x \mid v = L) = \alpha(x)$ and $\Pr(s \leq x \mid v = H) = \beta(x)$. Doing so is possible, because [Proposition 1](#) ensures that functions α and β are increasing, take values in $[0, 1]$, and equal 1 at $\theta = 1$. If $\alpha(\theta)$ and $\beta(\theta)$ are not right continuous, π employs their right-continuous modifications.

The defining feature of signal $\hat{\mathcal{I}}$ is that pooling signal realizations s below and above θ results in signal $\mathcal{I}(\theta) = (\alpha(\theta), \beta(\theta))$. That is, $\hat{\mathcal{I}}$ is Blackwell more informative than any of the signals in the original mechanism. However, no type can benefit from the extra informativeness of $\hat{\mathcal{I}}$. Indeed, the second property of [Proposition 1](#) implies that signal $\hat{\mathcal{I}}$ satisfies a monotone likelihood ratio property.⁷ As such, lower signal realizations induce higher posterior beliefs across all types, and for each type, a best response to $\hat{\mathcal{I}}$ is characterized by a threshold \tilde{s} such that the type sets price $p = H$ for all $s \leq \tilde{s}$ and price $p = L$ for all $s > \tilde{s}$. By construction, the choice between different thresholds is equivalent to the choice between different signals in the original mechanism. Hence, by the incentive compatibility of the original mechanism, type θ optimally chooses the threshold $\tilde{s} = \theta$. The resulting allocation rule mimics the allocation rule in the original direct mechanism type by type. \square

⁶[Proposition 2](#) is reminiscent of the equivalence between the experiments and persuasion mechanisms of [Kolotilin et al. \(2017\)](#). However, there are important differences between the settings and the results. First, in our setting, the seller’s private information is correlated with the value, which affects the seller’s assessment and response to data. Second, we obtain a stronger equivalence result, which applies to allocation rules and not just to the seller’s interim utility profile. In our setting, this stronger result is necessary for welfare analysis, because the seller’s interim utility profile alone does not determine the buyer’s surplus.

⁷The condition $(\alpha(\theta_3) - \alpha(\theta_2))(\beta(\theta_2) - \beta(\theta_1)) \geq (\alpha(\theta_2) - \alpha(\theta_1))(\beta(\theta_3) - \beta(\theta_2))$ is the general definition for CDF α dominating CDF β in the monotone likelihood ratio order (see Theorem 1.C.5 of [Shaked and Shanthikumar \(2007\)](#)).

3.3 Implementable Outcome Characterization

Proposition 2 implies that the set of all implementable outcomes coincides with the set of all outcomes that are implementable by public signals. We derive the implementable outcomes under public signals in two steps: First, for each signal, we derive the payoffs of the buyer and the seller conditional on each signal realization but unconditional on the seller’s type. Second, we then aggregate these payoffs across all signal realizations.

Specifically, for any given signal \mathcal{I} and realization s , define by $\mu \in [0, 1]$ the posterior belief that s induces in a hypothetical uninformed seller with a prior belief $\mu_0 \triangleq \mathbb{E}[\theta]$, so that $\mu \triangleq \Pr(v = H \mid s, \mu_0)$. We call this belief μ a *basic* posterior belief. By Bayes’ rule, the same signal realization observed by the seller of type θ results in the posterior belief $t \triangleq \Pr(v = H \mid s, \theta)$ equal to⁸

$$t(\mu, \theta) = \frac{\theta\mu(1 - \mu_0)}{\theta\mu(1 - \mu_0) + (1 - \theta)(1 - \mu)\mu_0}. \quad (5)$$

The posterior belief increases both in μ and θ , equals zero if either of the arguments equals zero, and equals one if either of the arguments equals one. Formula (5) applies to all types. Therefore, the basic posterior belief μ fully determines the distribution of the seller’s posterior beliefs. The higher the μ , the higher the corresponding distribution of the seller’s posterior beliefs in the sense of first-order stochastic dominance—i.e., the beliefs of different seller types move in concordance upon observing the same signal realization.

In turn, the seller’s posterior belief determines her pricing decision: The seller sets price $p = H$ if $t > L/H$ and price $p = L$ if $t < L/H$. Equivalently, when observing a signal that induces a basic posterior belief μ , the informed seller with type θ sets a high price if $\theta > \tilde{\theta}(\mu)$ and a low price if $\theta < \tilde{\theta}(\mu)$, where the threshold type $\tilde{\theta}$ is uniquely defined by the condition $t(\mu, \tilde{\theta}(\mu)) = L/H$.

By Bayes’ consistency, if the seller’s posterior belief is t , then the probability that the consumer’s value is H is indeed t . Consequently, the players’ expected payoffs after signal

⁸Alonso and Camara (2016) use this property to analyze Bayesian persuasion with heterogeneous priors.

realization s that induces a basic posterior belief μ can be written as

$$U(\mu) = \int_0^{\tilde{\theta}(\mu)} t(\mu, \theta)(H - L)dF(\theta), \quad (6)$$

$$\Pi(\mu) = \int_0^{\tilde{\theta}(\mu)} LdF(\theta) + \int_{\tilde{\theta}(\mu)}^1 t(\mu, \theta)HdF(\theta). \quad (7)$$

The welfare functions $U(\mu)$ and $\Pi(\mu)$ aggregate consumer surplus and seller profit respectively across seller types for any basic posterior belief. We can use these functions to characterize the set of implementable outcomes. Define by $\text{graph}(U, \Pi)$ a graph of the vector function that keeps track of players' payoffs $(U(\mu), \Pi(\mu))$ at different basic beliefs μ . Any public signal is characterized by a distribution of the basic beliefs, which by Bayes' rule must average to the prior belief μ_0 . Therefore, any public signal implements an outcome in a convex hull of $\text{graph}(U, \Pi)$ with the first component equal to μ_0 . Vice versa, [Aumann and Maschler \(1995\)](#) and [Kamenica and Gentzkow \(2011\)](#) show that any belief distribution that averages to the prior can be induced by some signal. Therefore, any point in the convex hull of $\text{graph}(U, \Pi)$ such that the first component is μ_0 can be implemented by some signal (cf. [Doval and Smolin \(2023\)](#)).

Proposition 3. (Implementable Outcomes) *The set of all implementable outcomes is*

$$\mathcal{E} = \{(x, y) : (\mu_0, x, y) \in \text{co}(\text{graph}(U, \Pi))\}, \quad (8)$$

where functions U and Π are given by (6) and (7).

For any given type distribution, [Proposition 3](#) enables a geometric characterization of the set of equilibrium outcomes. The result highlights the fact that the type distribution F affects the set of implementable outcomes via its impact on the indirect payoffs (6) and (7).

To simplify exposition, for the rest of this section we assume that θ is continuously distributed over $[0, 1]$. The indirect payoffs are then continuous functions of basic belief μ , because $\tilde{\theta}(\mu)$ is continuous and the types are continuously distributed over $[0, 1]$. Hence, by [Proposition 3](#), the set \mathcal{E} is compact and convex, and we can characterize its extreme points via supporting Bayesian persuasion problems. Indeed, any extreme point of \mathcal{E} maximizes

some linear combination of buyer surplus and seller profit. This maximization problem corresponds to a Bayesian persuasion problem over a public split of basic posterior beliefs to maximize

$$W^\lambda(\mu) \triangleq \lambda_U U(\mu) + \lambda_\Pi \Pi(\mu) \tag{9}$$

for some $(\lambda_U, \lambda_\Pi) \in \mathbb{R}^2$. We then obtain the following result:

Proposition 4. (Extreme Outcomes) *If the seller type is continuously distributed over $[0, 1]$, then the set of extreme implementable outcomes is spanned by solutions to Bayesian persuasion problems parameterized by $\lambda = (\lambda_U, \lambda_\Pi) \in \mathbb{R}^2$:*

$$\max_{\tau \in \Delta(\Delta(V))} \mathbb{E}_{\mu \sim \tau} [W^\lambda(\mu)] \quad \text{subject to} \quad \int_{\Delta(V)} \mu \tau(d\mu) = \mu_0. \tag{10}$$

Propositions 3 and 4 imply bounds on the necessary complexity of the signals used. When the type is continuously distributed, the indirect payoffs are continuous functions of basic belief μ and $\text{graph}(U, \Pi)$ features a single connected component. Thus, by Fenchel-Bunt's theorem, every point in the convex hull of $\text{graph}(U, \Pi)$ can be generated by a randomization over at most $|V| - 1 + 2 = 3$ of its points. Such randomization corresponds to a signal with at most 3 signal realizations. Moreover, any extreme point of \mathcal{E} is implemented by a solution of a Bayesian persuasion problem that can be set to contain at most $|V| = 2$ signal realizations.

Corollary 1. (Signal Complexity) *If the seller type is continuously distributed over $[0, 1]$, then any implementable outcome is implementable by a public signal with at most 3 signal realizations. Any extreme point of the set of implementable outcomes is implementable by a public signal with at most 2 signal realizations.*

Corollary 1 shows that even when there are many seller types, the ex ante consumer and seller payoffs can be obtained if the seller obtains coarse data that features at most 3 labels. Furthermore, extreme outcomes, such as a buyer-optimal outcome, can be obtained with even coarser data that features at most 2 labels.

3.4 Uniform Type Distribution

So far we have established the general properties of implementable allocations, characterized the set of implementable outcomes, and placed upper bounds on the complexity of necessary signals. However, we have been silent with respect to a few important questions, such as when providing some data can benefit the buyer or how to achieve a buyer-optimal outcome. To tackle these questions and to further illustrate our results, we now consider the special case of the seller's type being uniformly distributed over the unit interval. This case captures a natural benchmark in which all seller types are possible and equally likely.

To state the next result, we introduce two classes of signals. Consider a signal whose likelihood function in tabular form is

$$\begin{array}{c|cc}
 \mathcal{I} & s_L & s_H \\
 \hline
 v = L & 1 - \alpha & \alpha \\
 v = H & 1 - \beta & \beta
 \end{array} \tag{11}$$

for some $\alpha, \beta \in [0, 1]$. A signal is *high-value flagging* if $\alpha = 0$, so that realization s_H can arise only if $v = H$. A signal is *low-value flagging* if $\beta = 1$, so that realization s_L can arise only if $v = L$. Note that the fully informative signal ($\alpha = 0$ and $\beta = 1$) and the uninformative signal ($\alpha = \beta = 0$) belong to these classes.

Proposition 5. (Uniform Types) *Let the seller type be uniformly distributed on $[0, 1]$. The full-information outcome and the no-information outcome are extreme points of \mathcal{E} . The left and right boundaries of \mathcal{E} connect these two outcomes and are spanned by low-value and high-value flagging signals, respectively. Moreover, the following hold:*

1. *If $\frac{L}{H} \geq \frac{1}{2}$, then the buyer-optimal outcome in \mathcal{E} is generated by the uninformative signal.*
2. *If $\frac{L}{H} < \frac{1}{2}$, then the buyer-optimal outcome in \mathcal{E} is generated by the high-value flagging signal with flagging rate $\beta = \frac{H-2L}{H-L}$.*

The proof of [Proposition 5](#) builds on [Proposition 4](#). Any extreme point of \mathcal{E} solves a Bayesian persuasion problem [\(10\)](#) and is generated by a public signal with two signal




Figure 2: Implementable outcomes when seller types are uniformly distributed on $[0, 1]$. Light blue denotes the case of observable type. Dark blue denotes the case of unobservable type. Black boundaries are spanned by flagging signals. Points C indicate buyer-optimal outcomes in \mathcal{E} .

realizations. We can thus present the persuasion problem as a maximization problem with respect to α and β and solve it in closed form.

Intuitively, data provision affects the buyer surplus in two ways. On the one hand, it can benefit the buyer by persuading some seller types $\theta > L/H$ to set price $p = L$ even when the buyer has value H . On the other hand, data provision may harm the buyer if $\theta < L/H$, because such seller types would set price $p = L$ in the absence of data. Because all implementable outcomes are spanned by public signals, any data provision increases the buyer's surplus for some seller types and decreases it for other types. If L/H is high—i.e., a large fraction of seller types set price L in the absence of data—then the negative effect dominates and any data provision is detrimental for the buyer (Part 1). If L/H is low, then some data provision is Pareto improving (Part 2).

Figure 2 depicts the sets of implementable outcomes for two concrete cases of value distribution. When the seller type is observable, any individually rational outcome—i.e., the whole surplus triangle—can be implemented by data provision. In contrast, when the seller's type is unobserved, the sets of implementable outcomes shrink. The uninformative signal implements the “south” end of \mathcal{E} as an extreme point, which is also the seller-worst

outcome. The right boundary is spanned by high-value flagging and the left boundary by low-value flagging. As we move along each boundary from south to north, the corresponding signals have higher flagging rates and become more informative.⁹ The two boundaries meet at the “north” end of \mathcal{E} , which is the seller-optimal outcome and implemented by providing full information. In fact, this is the only efficient outcome and gives zero surplus to the buyer. The buyer-optimal outcome differs in the two cases. Figure 2(a) depicts part 2 of the proposition: As $L/H = 1/3 < 1/2$, we have $\beta = 1/2$, so the buyer surplus is maximized by flagging one-half of high-value buyers. Figure 2(b) depicts part 1: As $L/H = 2/3 > 1/2$, the buyer surplus is maximized by providing no data. In this case, adverse selection is so severe that any additional information would on average hurt the buyer.

4 Implications

We derive several implications of the above results. First, we show that the seller’s private information creates a trade-off between buyer protection and efficiency. Second, we study the impact of the seller’s use of third-party data, which we capture as the seller’s choice to acquire private information. Third, we demonstrate that the implementable set subsumes the equilibrium outcomes of various games in which the buyer and the seller exchange information with each other. Furthermore, we show that our key insights extend beyond binary consumer values.

4.1 The Trade-off Between Buyer Protection and Efficiency

If the seller’s type is observable, the designer can distribute the efficient total surplus flexibly between the buyer and the seller ([Claim 2](#)). However, implementing an efficient outcome typically requires providing different signals to different types, which is infeasible when the seller’s type is her private information ([Proposition 2](#)). The impossibility of heterogeneous data provision introduces a tension between enhancing consumer surplus and total surplus—and in many cases, achieving efficiency implies that the buyer receives no surplus at all.

⁹This observation is general—see the proof in Appendix, which characterizes those flagging signals in closed form.

Below, the seller’s “rent” means the seller’s expected profit under a given outcome minus her profit in the absence of additional data.

Proposition 6. (Efficiency and Buyer Surplus) *Assume that the type distribution is non-degenerate and places probability 1 on $(0, 1)$. Then the following hold:*

1. *If there is a positive measure of types strictly above L/H , then any implementable efficient outcome gives a strictly positive rent to the seller.*
2. *If there is a strictly positive measure of types in any neighborhood of 1, then the only implementable efficient outcome is the outcome in which all seller types obtain full information and perfectly price discriminate the buyer.*

The intuition is as follows. To incentivize type $\theta > L/H$ to price efficiently, a signal has to reveal the buyer’s value with some probability $\beta > 0$ when $v = H$, so that in the remaining event, the seller sets price L , believing that the value is likely to be L . The higher the θ , the higher the probability β must be and the more informative the signal becomes. While it is possible for the designer to choose β so that the seller with a known type achieves the same profit as with no additional information, this is not the case for a privately informed seller. With the privately informed seller, a signal provided to one type for efficient pricing will also be used by lower types to earn positive rents. Moreover, a type that is arbitrarily close to 1 must receive the (almost) fully informative signal to price efficiently. The impossibility of screening implies that all seller types receive full information and the buyer receives no surplus.

The result has an implication for policies regarding the use of consumer data by firms. Consumer data may enable sellers to tailor pricing, which can benefit consumers and enhance efficiency. At the same time, how much and what kind of data sellers should be allowed to use for such beneficial pricing will depend on the specific market conditions and the initial endowment of data and technology for each seller. However, in practice, a regulator would not be able to tailor a regulation to each firm, partly because of the informational friction highlighted in this paper. In such a case, the same regulation is applied to heterogeneous sellers. Our result yields two relevant distortions in such a situation: First, some sellers obtain too little information, which leads to inefficient pricing. Second, some sellers obtain

too much information, which not only allows them to create surplus but also to extract surplus from consumers. Our result clarifies this economic force and further shows that the distortion can lead to an extreme outcome in which ensuring efficient pricing erodes the entire buyer surplus.

4.2 Third-party Data

In some applications, the seller can obtain third-party data at her will from outside sources. For example, a seller may learn about buyers from tracking tools or data brokers. We can analyze the acquisition of such data by interpreting it as the seller’s choice to be privately informed.

Specifically, we can study the welfare impact of third-party data by comparing the following two kinds of sellers: One is the informed seller, who has private type $\theta \sim F$ that can take values in $(0, L/H)$ and $(L/H, 1)$ with positive probabilities. The other is the uninformed seller, whose type is deterministic and known to be $\theta_0 \triangleq \mathbb{E}_{\tilde{\theta} \sim F}[\tilde{\theta}]$. In general, the seller’s private information shrinks the set of implementable outcomes: This set shrinks as we move from the uninformed seller to the one with observable types, and it shrinks even further when the seller is privately informed (an example of the latter change is captured by [Figure 2](#)).¹⁰ For example, if the type distribution has full support over $[0, 1]$, then no efficient outcome, except for the outcome in which the seller extracts full surplus, is implementable when the seller has private information ([Proposition 6](#)).

A natural question is how the seller’s acquisition of third-party data affects the seller and the buyer. The answer depends on which implementable outcome is selected. As an illustration, suppose that the selected outcome is the *constrained seller-optimal outcome*—i.e., the implementable outcome that maximizes seller surplus subject to the constraint whereby the buyer must weakly benefit from the data provision in terms of his ex ante expected surplus. This outcome reflects a commonly observed feature of the digital economy, in which firms specify the terms of data collection and usage, but harmful data collection

¹⁰The set of implementable outcomes for a seller with observable types is indeed a subset of the implementable outcomes for an uninformed seller. This is because the designer can mimic any outcome from the observable type scenario by providing some information to the uninformed seller.

may be prevented by a user’s opting out or by a regulator’s intervention.

Proposition 7. (Third-party Data) *Let (U, Π) and (U', Π') be the constrained seller-optimal outcomes given the uninformed seller and the informed seller. Then the following hold:*

1. *If $\theta_0 \in (\frac{L}{H}, 1)$, then $\Pi > \Pi'$ and $U < U'$ —i.e., the seller is worse off and the buyer is better off when the seller has private information.*
2. *If $\theta_0 \in (0, \frac{L}{H})$, then $\Pi < \Pi'$ and $U > U'$ —i.e., the seller is better off and the buyer is worse off when the seller has private information.*

The intuition for Part 1 is as follows. The condition $\theta_0 > L/H$ means that the uninformed seller believes that the buyer’s value is likely to be high, so she optimally sets price H , which leads to zero buyer surplus. In such a case, the constrained seller-optimal outcome is that the seller obtains full information and extracts the efficient total surplus. In contrast, the constrained seller-optimal outcome for the privately informed seller dictates that the seller never obtains full information, because the buyer surplus at no additional data provision is still positive. Consequently, the seller’s private information decreases the seller’s profit and increases buyer surplus. Part 1 of the result suggests that the seller may choose not to be privately informed about the buyer in order to eliminate adverse selection and obtain superior first-party data. This choice, while seemingly privacy friendly, could in fact harm the buyer.¹¹

The argument for Part 2 is symmetric: Given that the uninformed seller sets price L , the constrained seller-optimal outcome is to provide no data. Compared with this outcome, the seller’s private information increases her profit at the expense of buyer surplus and efficiency.

4.3 Communication Protocols

So far, we have adopted a mechanism-design approach and studied the set of implementable outcomes. We now show that this set contains the equilibrium outcomes in a large class of communication protocols with an informed buyer and seller, such as cheap-talk communication, voluntary disclosure, and the collection of verifiable data.

¹¹In other words, third-party data may crowd out first-party data. A similar interplay between several sources of information is the main focus of [Alonso and Zachariadis \(2021\)](#) in an investment setting.

Formally, assume that the buyer and seller are privately informed about value v , having types $t \in T$ and $\theta \in \Theta$, respectively, which are independent conditional on v . Before trade, the players can communicate via a predefined two-stage protocol.¹² A *protocol* \mathcal{P} is a triple $(A_S, \{A_B(t)\}_{t \in T}, \Psi)$. The sets A_S and $\{A_B(t)\}_{t \in T}$ denote the seller’s action space and the buyer’s action space, respectively. The buyer’s action space $A_B(t)$ can depend on his type t , which enables us to capture, e.g., voluntary disclosure settings. The signal scheme is denoted by $\Psi = (S, \pi)$. Here, S is the set of signal realizations the seller can possibly observe in the protocol and $\pi : V \times A_S \times A_B \rightarrow \Delta(S)$ is the likelihood function of different signal realizations, which can depend on the value and, importantly, on the players’ actions. In the protocol, first the seller takes an action. Second, the buyer observes the seller’s action and chooses his own action. A seller strategy is $\sigma_S : \Theta \rightarrow \Delta(A_S)$ and a buyer strategy is $\sigma_B : T \times A_S \rightarrow \Delta(A_B(t))$. At the end of the communication, the seller observes signal realization s according to Ψ .

After communicating within the protocol, the seller sets the price p . The buyer learns the value v , observes the price, and decides whether to purchase the good. If trade occurs, the buyer’s ex post payoff is $v - p$ and the seller’s ex post payoff is p . If trade does not occur, both players obtain a zero payoff. The solution concept is a perfect Bayesian equilibrium.

Any equilibrium in a protocol results in an allocation rule $a : V \rightarrow [0, 1] \times \mathbb{R}$, which as before specifies the probability of a trade and the expected payment from the buyer to the seller for each value. We continue to use (welfare) outcome, buyer surplus, and seller profit to mean relevant ex ante expected payoffs.

Proposition 8. (Menu Mechanisms) *An allocation rule can arise in an equilibrium with some communication protocol if and only if it can arise in an equilibrium of a menu mechanism.*

Proof. The proof follows the revelation principle argument of Myerson (1982, 1983), but is adapted to fit the specific details of our environment. For the “if” direction, note that any menu mechanism is itself an example of a communication protocol with $A_B(t) \equiv \{a_0\}$. For the “only if” direction, consider any protocol and an equilibrium in it. In this equilibrium,

¹²This is a large class of protocols that are natural in applications. However, like a revelation principle, our results can be stated with respect to a broader class of protocols at the expense of additional notation.

each action of the seller induces a signal $\mathcal{I}(a) = (S, \pi(a))$ with $\pi(a) : V \rightarrow \Delta(S)$, where the likelihood function averages over the buyer’s equilibrium strategy and the signal scheme. We can replace this protocol with a menu mechanism $\mathcal{M} = \{\mathcal{I}(a)\}_{a \in A_S}$. This change does not alter the seller’s equilibrium strategy and results in the same allocation rule as the original protocol. Indeed, the only thing that matters at the trading stage is the seller’s estimate of the buyer’s value. It does not matter whether this estimate is obtained through direct data provision or equilibrium inference. Moreover, the buyer type is not useful for screening the seller type because they are conditionally independent. As a result, menu mechanisms with direct data provision implement all equilibrium allocation rules. \square

Below, we describe protocols that select different subsets of the implementable outcomes, as depicted in [Figure 3](#). For simplicity, assume that the seller’s type distribution F has a full support on $[0, 1]$.

Cheap Talk. As a benchmark, suppose that the buyer knows the value and can inform the seller about it via an unverifiable, cheap-talk message—i.e., $t \equiv v$, $A_S = \{a_0\}$, $A_B(H) = A_B(L) = A_B$, and $s \equiv a_B$. This protocol leads to a no-information outcome. Indeed, if an equilibrium entailed some nontrivial communication—i.e., if $v = L$ were strictly more likely after message $m \in A_B$ than after message $m' \in A_B$ —then sending m would lead to lower prices for all seller types and strictly so, given the full-support assumption. An H -buyer would then never send message m' , which would contradict the presumption that message m indicates a higher likelihood of an L -buyer. The result implies that for informative communication to occur, information provision must be verifiable to some extent.

Voluntary Disclosure. We turn to the buyer’s voluntary disclosure of verifiable information.¹³ In this protocol, given the full-support assumption, the buyer is never willing to verify that he has value H , because the buyer would then face a high price and lose a positive surplus. As a result, any equilibrium strategy of the buyer generates a low-value flagging signal.

¹³If the seller is uninformed, this setting is a special case of [Ali, Lewis, and Vasserman \(forthcoming\)](#).




Figure 3: Equilibrium outcomes selected by different communication protocols. Seller types are uniformly distributed on $[0, 1]$, and $(L, H) = (1, 3)$.

Formally, we consider the following protocol: $t \equiv v$, $A_S = \{a_0\}$, $A_B(v) = \{\{v\}, \emptyset\}$ for each $v \in \{H, L\}$, and $s \equiv a_B$, so that the buyer can choose whether to disclose his value. We argue that the set of equilibrium outcomes coincides with all implementable outcomes spanned by the low-value flagging signals. Indeed, if an L -buyer sends message \emptyset with a positive probability, then the seller with a sufficiently low type θ sets price L after observing message \emptyset . But then, an H -buyer strictly prefers sending \emptyset to $\{H\}$. Thus, there could be only two kinds of equilibrium: (i) an L -buyer mixes between messages $\{L\}$ and \emptyset and an H -buyer sends \emptyset with probability 1, and (ii) an L -buyer sends message $\{L\}$ for sure, and an H -buyer sends \emptyset or $\{H\}$ revealing his value H . In fact, any such buyer strategy can arise in equilibrium, because an L -buyer never obtains a positive surplus and thus is indifferent between any messages, and any possible deviation by an H -buyer can be deterred by the seller's skeptical belief that places probability 1 on H .

The low-value flagging characterization, combined with [Proposition 5](#), implies that if the seller's type is uniformly distributed, the equilibrium outcomes of the voluntary disclosure game span the left boundary of the surplus set (see [Figure 2](#)), which highlights the potential inefficiency of this protocol. This inefficiency remains even if the seller can communicate before the buyer's disclosure—i.e., if A_S is general. In that case, in equilibrium, each message sent by the seller induces some low-value flagging signal, but any two low-value flagging

signals are ranked in Blackwell informativeness. Hence, all seller types send a message that induces the same low-value flagging signal.

Request-Consent Protocol. The seller’s action space A_S is given by the set of all signals with $S = [0, 1]$, and the buyer’s action space is given by $A_B(H) = A_B(L) = \{accept, reject\}$. The seller first chooses a signal, and then the buyer observes the requested signal (but not its realization) and decides whether to accept it. The seller obtains the requested signal if the buyer chooses *accept* and does not observe any additional information otherwise.

In this protocol, if the buyer is perfectly informed (i.e., $t \equiv v$), then any implementable outcome can arise in some equilibrium: Take any implementable outcome and the public signal \mathcal{I} that implements it. The following equilibrium has all seller types obtain signal \mathcal{I} . On the equilibrium path, all seller types request signal \mathcal{I} and the buyer accepts it regardless of his value. The seller’s deviation to an off-path signal can be deterred if the buyer, following the deviation, believes that the seller’s type is $\theta = 0$. The buyer with this belief thinks that the seller will set price L , so the buyer never strictly benefits from providing additional information. In turn, if the seller believes that any buyer who deviates and rejects the signal request has value H , then the buyer finds it optimal to accept signal \mathcal{I} .

In contrast, if the buyer is not informed when deciding on the request acceptance (i.e., $t \equiv t_0$), then the set of equilibrium outcomes is smaller and coincides with the set of implementable outcomes such that the buyer is at least as well off as under no data provision. Indeed, if all seller types request the same signal \mathcal{I} , the uninformed buyer will accept it if and only if signal \mathcal{I} weakly increases his ex ante payoff; the seller cannot use the off-path belief punishment because the buyer is uninformed. The result then follows from [Proposition 2](#). Importantly, in some cases there may be no public signal that strictly increases buyer surplus (see [Figure 2 \(b\)](#)). In those cases, the unique equilibrium outcome of the request-consent protocol with an uninformed buyer would be no data collection.

4.4 Many Values

There are multiple ways in which our setting can be extended. One is to allow many possible buyer values, $V = \{v_1, \dots, v_n\}$ with $v_1 < \dots < v_n$, $n > 2$. The seller's type $\theta \in \Delta(V)$ is now multidimensional, and the complete characterization of implementable outcomes is not attainable because it involves multidimensional screening. Nevertheless, we show that our main qualitative findings extend to this setting and provide necessary clarifications.

First of all, private information held by the seller generally limits the set of implementable outcomes, as in the case of binary values. However, in the case of many values, the limits are less stark, since public signals do not necessarily span all implementable outcomes. Differentiated data provision can be welfare enhancing in this setting, and it can even attain the buyer's first-best outcome, which attains efficiency without increasing the seller's profits. The following example illustrates such a possibility.

Example 1. (Benefits of Heterogeneous Data Provision) Let the possible values be $V = \{L, M, H\} = \{1, 3, 4\}$. Let there be three seller types, $\theta_1 = (1/2, 1/4, 1/4)$, $\theta_2 = (1/2, 1/2, 0)$, and $\theta_3 = (1/2, 0, 1/2)$, where the three types are equally likely. In the absence of additional data, the types set prices $p_1 = M$, $p_2 = M$, and $p_3 = H$, and buyer surplus is $U_0 = 1/12$.

We derive the buyer-optimal outcome and show that it is efficient and cannot be attained by a public signal. Namely, consider a menu in which the types are provided with the following signals:

$$\begin{array}{c|cc} \mathcal{I}(\theta_1) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 1/2 & 1/2 \\ H & 1/2 & 1/2 \end{array}, \quad \begin{array}{c|cc} \mathcal{I}(\theta_2) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 1/2 & 1/2 \\ H & 1/2 & 1/2 \end{array}, \quad \begin{array}{c|cc} \mathcal{I}(\theta_3) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 2/3 & 1/3 \\ H & 1/3 & 2/3 \end{array}, \quad (12)$$

so that each signal has two possible realizations with the likelihood function π as presented in the tabular form. It is straightforward to verify that no type wants to deviate. Moreover, the outcome is efficient and each type earns the same profits as in the absence of additional information: Given these signals, types θ_1 and θ_2 set prices $p = L$ and $p = M$ after signal realizations s_1 and s_2 , respectively, and type θ_3 sets prices $p = L$ and $p = H$, respectively. The

buyer surplus is then maximal among all feasible allocations, given the individual rationality of the seller.

However, we cannot implement this outcome by providing the same signal to all seller types: Any efficient public signal must give strictly positive rents to type θ_1 . Indeed, for signal $\mathcal{I} = (S, \pi)$ to attain efficiency, it must be that whenever $\pi(s_0|L) > 0$ for some $s_0 \in S$ all types charge $p = L$ following that signal, which means that $\pi(s_0|L) \geq 2\pi(s_0|M)$ to persuade type θ_2 and $\pi(s_0|L) \geq 3\pi(s_0|H)$ to persuade type θ_3 ; in turn, these incentive constraints together guarantee that θ_1 also charges a low price. Since type θ_1 never charges a price below $p = L$, her profit decreases as we increase the frequency of signal s_0 —i.e., as we increase $\pi(s_0|L)$, $\pi(s_0|M)$, and $\pi(s_0|H)$ subject to the incentive constraints. Therefore, of all efficient public signals, the profit of type θ_1 is minimized at the signal

$$\hat{\mathcal{I}} \left| \begin{array}{cc} s_1 & s_2 \\ \hline L & 1 \quad 0 \\ M & 1/2 \quad 1/2 \\ H & 1/3 \quad 2/3 \end{array} \right. , \tag{13}$$

which is efficient and maximizes the probability of type θ_1 setting price $p = L$ while ensuring that type θ_1 offers the remaining buyers the second-lowest price $p = M$. The resulting θ_1 's minimal profit equals $19/12$, which is strictly greater than her outside option of $3/2$ and, as such, necessarily results in lower buyer surplus than under the multi-item mechanism presented above. \square

Intuitively, as in most mechanism-design problems, the multidimensionality of the seller's type gives rise to a richer structure of incentive constraints and implementable outcomes. In particular, the buyer-optimal efficient signals are not necessarily Blackwell-comparable and can appeal to the “right” buyer type. As a result, complete characterization of the set of implementable outcomes becomes more complicated. Nevertheless, we show that the welfare implications of [Section 4.1](#) generalize: Seller heterogeneity introduces a trade-off between efficiency and consumer welfare.¹⁴

¹⁴All conditions of [Proposition 9](#) hold, for example, whenever F has full support over $\Delta(V)$.

Proposition 9. (Efficiency and Buyer Surplus) *Assume that the type distribution F places probability 1 on the interior of $\Delta(V)$. Then the following hold:*

1. *If F admits a positive density over an open set of types who, in the absence of data, charge prices $p > v_1$, then any implementable efficient outcome gives a strictly positive rent to the seller.*
2. *If there is a strictly positive measure of types in any neighborhood of the type that places probability 1 on v_n , then the only implementable efficient outcome is that under full information.*

Proposition 9 is a counterpart of Proposition 6 for the case of many values. The first part of Proposition 9 is based on the fact that signals that yield no rents to the seller must be carefully tailored to the seller's type. Specifically, consider the types over which F has a density. Efficiency requires that these types occasionally set price $p = v_1$ to serve the lowest-value buyer, and therefore must be given additional information. However, if some type weakly prefers to set the lowest price after receiving some signal realization, nearby types would strictly prefer to set the lowest price following that realization and would therefore strictly prefer to change their prior action. This means that they can earn strictly positive rents from that signal and must earn strictly positive rents when faced with any efficient menu.

For the second part of Proposition 9, we can show that under the stated condition, there exists a sequence of types with the following properties. First, types converge to the extreme type that places probability 1 on the highest possible value. Second, along the sequence, the corresponding signals, which induce efficient outcomes, converge to the fully informative signal. In particular, we show that the types along this sequence need to be persuaded to charge all possible prices and thus must be provided with progressively more detailed information. The limit argument, coupled with incentive compatibility, then implies that all types must be offered a fully informative signal.

5 Conclusion

We studied the provision of consumer data to a monopoly seller who has private belief about the value of its product. Without private information, the designer can flexibly provide the seller with data to distribute efficient total surplus between the buyer and the seller. We show that the seller’s private information prevents the designer from providing different signals to different seller types, which leads to the impossibility of effective screening. Our results provide insights into data policies for policymakers and businesses. For policymakers or platforms that aim to control the flow of consumer data to sellers, the results highlight the trade-off between consumer protection and efficiency driven by adverse selection regarding data use. From the seller’s perspective, our results clarify why the collection of third-party data without consumer consent might harm profits. To clarify the economic intuition, we have focused on a simple setting in which the buyer has binary values and the seller uses data for price discrimination. However, the issues that arise when a firm’s private information hinders the effective allocation of consumer data—and the resulting challenges faced by regulators and firms with regard to data policies—would be relevant in broader contexts.

References

- ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): “The Economics of Privacy,” *Journal of Economic Literature*, 54, 442–492.
- ALI, S. N., G. LEWIS, AND S. VASSERMAN (forthcoming): “Voluntary Disclosure and Personalized Pricing,” *Review of Economic Studies*.
- ALONSO, R. AND O. CAMARA (2016): “Bayesian Persuasion with Heterogeneous Priors,” *Journal of Economic Theory*, 165, 672–706.
- ALONSO, R. AND K. ZACHARIADIS (2021): “Persuading Large Investors,” *Working paper*.
- ARGENZIANO, R. AND A. BONATTI (2021): “Information Revelation and Privacy Protection,” *Working paper*.

- AUMANN, R. AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*, MIT Press.
- BARON, D. P. AND R. B. MYERSON (1982): “Regulating a Monopolist with Unknown Costs,” *Econometrica*, 911–930.
- BERGEMANN, D., A. BONATTI, AND T. GAN (2022): “The Economics of Social Data,” *The RAND Journal of Economics*, 53, 263–296.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2018): “The Design and Price of Information,” *American Economic Review*, 108, 1–48.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- CHOI, J. P., D.-S. JEON, AND B.-C. KIM (2019): “Privacy and Personal Data Collection with Information Externalities,” *Journal of Public Economics*, 173, 113–124.
- CONDORELLI, D. AND B. SZENTES (2022): “Buyer-Optimal Platform Design,” *Working paper*.
- DEB, R. AND A.-K. ROESLER (2022): “Multi-dimensional Screening: Buyer-Optimal Learning and Informational Robustness,” *Working paper*.
- DOVAL, L. AND A. SMOLIN (2023): “Persuasion and Welfare,” *Working paper*.
- ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2022): “Market Segmentation through Information,” *Working paper*.
- FAINMESSER, I. P., A. GALEOTTI, AND R. MOMOT (2022): “Digital privacy,” *Management Science*.
- HAGHPANAH, N. AND R. SIEGEL (2022): “The Limits of Multiproduct Price Discrimination,” *American Economic Review: Insights*, 4, 443–58.
- (forthcoming): “Pareto Improving Segmentation of Multi-product Markets,” *Journal of Political Economy*.

- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KOLOTILIN, A., T. MYLOVANOV, A. ZAPECHELNYUK, AND M. LI (2017): “Persuasion of a Privately Informed Receiver,” *Econometrica*, 85, 1949–1964.
- MYERSON, R. B. (1982): “Optimal Coordination Mechanisms in Generalized Principal–Agent Problems,” *Journal of Mathematical Economics*, 10, 67–81.
- (1983): “Mechanism Design by an Informed Principal,” *Econometrica*, 1767–1797.
- RAYO, L. AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949–987.
- RHODES, A. AND J. ZHOU (2022): “Personalized Pricing and Competition,” *Working paper*.
- ROESLER, A.-K. AND B. SZENTES (2017): “Buyer-Optimal Learning and Monopoly Pricing,” *American Economic Review*, 107, 2072–80.
- SHAKED, M. AND J. G. SHANTHIKUMAR (2007): *Stochastic Orders*, Springer.
- SHI, X. AND J. ZHANG (2020): “Welfare of Price Discrimination and Market Segmentation in Duopoly,” *Working paper*.
- SMOLIN, A. (forthcoming): “Disclosure and Pricing of Attributes,” *The RAND Journal of Economics*.
- YANG, K. H. (2022): “Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences,” *American Economic Review*, 112, 1364–93.

Appendix: Proofs Omitted from the Main Text

Proof of Claim 2. To show the “only if” direction, take any feasible welfare outcome

$$(U, \Pi) = \left(\int_0^1 U(\theta) dF(\theta), \int_0^1 \Pi(\theta) dF(\theta) \right).$$

[Claim 1](#) implies that for each θ , we have $U(\theta) \geq 0$, $\Pi(\theta) \geq \underline{\Pi}(\theta)$, and $U(\theta) + \Pi(\theta) \leq \overline{W}(\theta)$. Integrating both sides of each inequality with F , we obtain $U \geq 0$, $\Pi \geq \underline{\Pi}$, and $U + \Pi \leq \overline{W}$.

To show the “if” direction, take any $(U, \Pi) \in \mathbb{R}^2$ such that $U \geq 0$, $\Pi \geq \underline{\Pi}$, and $U + \Pi \leq \overline{W}$. The point (U, Π) belongs to the triangle whose vertices are $(0, \overline{W})$, $(0, \underline{\Pi})$, and $(\overline{W} - \underline{\Pi}, \underline{\Pi})$. Let $\alpha, \beta \in [0, 1]$ with $\alpha + \beta \leq 1$ satisfy

$$\begin{aligned} (U, \Pi) &= \alpha(0, \overline{W}) + \beta(0, \underline{\Pi}) + (1 - \alpha - \beta)(\overline{W} - \underline{\Pi}, \underline{\Pi}) \\ &= \left(\int_0^1 (1 - \alpha - \beta)[\overline{W}(\theta) - \underline{\Pi}(\theta)]dF(\theta), \int_0^1 \alpha\overline{W}(\theta) + \beta\underline{\Pi}(\theta) + (1 - \alpha - \beta)\underline{\Pi}(\theta)dF(\theta) \right). \end{aligned}$$

For each θ , the point

$$(U(\theta), \Pi(\theta)) \triangleq \left((1 - \alpha - \beta)[\overline{W}(\theta) - \underline{\Pi}(\theta)], \alpha\overline{W}(\theta) + \beta\underline{\Pi}(\theta) + (1 - \alpha - \beta)\underline{\Pi}(\theta) \right)$$

is in the surplus triangle and thus feasible. Aggregating the welfare outcome $(U(\theta), \Pi(\theta))$ across possible θ , we conclude that $(U, \Pi) = \left(\int_0^1 U(\theta)dF(\theta), \int_0^1 \Pi(\theta)dF(\theta) \right)$ is also feasible. \square

Proof of [Proposition 1](#). Part 1. Fix any $\theta_2 \geq \theta_1$. The system of mutual incentive constraints is

$$(1 - \alpha(\theta_1))L + \theta_1(\alpha(\theta_1)L + \beta(\theta_1)(H - L)) \geq (1 - \alpha(\theta_2))L + \theta_1(\alpha(\theta_2)L + \beta(\theta_2)(H - L)), \quad (14)$$

$$(1 - \alpha(\theta_2))L + \theta_2(\alpha(\theta_2)L + \beta(\theta_2)(H - L)) \geq (1 - \alpha(\theta_1))L + \theta_2(\alpha(\theta_1)L + \beta(\theta_1)(H - L)). \quad (15)$$

Summing over the inequalities (14) and (15) and using the fact that $\theta_2 \geq \theta_1$ we obtain

$$\alpha(\theta_2)L + \beta(\theta_2)(H - L) \geq \alpha(\theta_1)L + \beta(\theta_1)(H - L). \quad (16)$$

In turn, (14) and (16) together imply $\alpha(\theta_2) \geq \alpha(\theta_1)$ because

$$(\alpha(\theta_2) - \alpha(\theta_1))L \geq \theta_1(\alpha(\theta_2)L + \beta(\theta_2)(H - L) - \alpha(\theta_1)L - \beta(\theta_1)(H - L)) \geq 0. \quad (17)$$

Finally, (15) and (17) imply $\beta(\theta_2) \geq \beta(\theta_1)$ because

$$(\beta(\theta_2) - \beta(\theta_1))\theta_2(H - L) \geq (\alpha(\theta_2) - \alpha(\theta_1))(1 - \theta_2)L \geq 0. \quad (18)$$

Part 2. Fix any $\theta_1 \leq \theta_2 \leq \theta_3$. The system of incentive constraints of type θ_2 toward types θ_1 and θ_3 can be written as

$$(\beta(\theta_2) - \beta(\theta_1))\theta_2(H - L) \geq (\alpha(\theta_2) - \alpha(\theta_1))(1 - \theta_2)L, \quad (19)$$

$$(\beta(\theta_3) - \beta(\theta_2))\theta_2(H - L) \leq (\alpha(\theta_3) - \alpha(\theta_2))(1 - \theta_2)L. \quad (20)$$

By property in Part 1, all sides of (19) and (20) are positive. Multiplying the respective smaller and larger parts and dividing the resulting inequality by $\theta_2(1 - \theta_2)(H - L)L$, we obtain the desired inequality of Part 2. \square

Proof of Proposition 5. Step 1. We use the “boundary” to mean the boundary of \mathcal{E} . By Proposition 4 and Corollary 1, any extreme point on the boundary can arise with a signal that has two signal realizations. We parameterize such signals by (α, β) , $\beta \geq \alpha$ as

$$\begin{array}{c|cc} \mathcal{I}(\theta) & s_L & s_H \\ \hline v = L & 1 - \alpha & \alpha \\ v = H & 1 - \beta & \beta \end{array} \quad (21)$$

Given any such signal, we write the objective $\lambda_U U + \lambda_\Pi \Pi$ in terms of the highest and lowest types that respond to the signal realizations. In particular, for any given (α, β) with $\beta \geq \alpha$, we can find two cutoffs x, y with $x \leq y$ such that types below x set price L after both realizations; types above y set price H after both realizations; and types between x and y will set prices H and L after s_H and s_L , respectively. Cutoffs x and y solve

$$\begin{aligned} L &= (1 - \alpha)L + x(\alpha L + \beta(H - L)), \\ yH &= (1 - \alpha)L + y(\alpha L + \beta(H - L)), \end{aligned}$$

which implies

$$x = \frac{\alpha L}{\alpha L + \beta(H - L)}, \quad y = \frac{(1 - \alpha)L}{H - \alpha L - \beta(H - L)}.$$

Alternatively, we can write (α, β) as functions of (x, y) :

$$\alpha = \frac{x(Hy - L)}{L(y - x)}, \quad \beta = \frac{(Hy - L)(1 - x)}{(y - x)(H - L)}.$$

Cutoffs (x, y) can arise under some signal if and only if $0 \leq x \leq \frac{L}{H} \leq y \leq 1$. The interim profit of type $\theta \in [x, y]$ is

$$\Pi(x, y | \theta) = (1 - \theta)(1 - \alpha)L + \theta((1 - \beta)L + \beta H) = L + \frac{(Hy - L)(\theta - x)}{y - x}.$$

The interim profits of types $\theta \leq x$ and $\theta \geq y$ are L and θH , respectively. The ex ante seller profit is

$$\begin{aligned} \Pi(x, y) &= \int_0^x L d\theta + \int_x^y L + \frac{Hy - L}{y - x}(\theta - x) d\theta + \int_y^1 \theta H d\theta \\ &= Ly + \frac{1}{2}(Hy - L)(y + x) - (Hy - L)x + \frac{H}{2}(1 - y^2). \end{aligned} \quad (22)$$

The buyer surplus is

$$\begin{aligned} U(x, y) &\triangleq (H - L) \int_0^x \theta d\theta + (1 - \beta)(H - L) \int_x^y \theta d\theta \\ &= (H - L) \int_0^x \theta d\theta + \left(H - L - \frac{(Hy - L)(1 - x)}{y - x} \right) \int_x^y \theta d\theta \\ &= (H - L) \int_0^y \theta d\theta - \frac{(Hy - L)(1 - x)}{y - x} \int_x^y \theta d\theta \\ &= \frac{1}{2}(H - L)y^2 - \frac{1}{2}(Hy - L)(1 - x)(y + x). \end{aligned} \quad (23)$$

Step 2. We characterize signals that span the “right” boundary that corresponds to $\lambda_U \geq 0$. Take any $(\lambda_U, \lambda_\Pi) \in \mathbb{R}^2$ with $\lambda_U \geq 0$. Because $\Pi(x, y)$ is linear in x and $U(x, y)$ is convex in x , $W(x, y)$ is convex in x and maximized at $x = 0$ or $x = L/H$. Note that $(x, y) = (L/H, y)$ implies $(\alpha, \beta) = (1, 1)$ and $(x, y) = (0, L/H)$ leads to $(0, 0)$, but both are the same

uninformative signal. Thus we can without loss of generality assume $x = 0$ and focus on the problem $\max_{\frac{L}{H} \leq y \leq 1} W(0, y)$, where

$$\begin{aligned} W(0, y) &= \lambda_U \left[\frac{1}{2}(H - L)y^2 - \frac{1}{2}(Hy - L)y \right] + \lambda_\Pi \left[Ly + \frac{1}{2}(Hy - L)y + \frac{H}{2}(1 - y^2) \right] \\ &= \frac{\lambda_U}{2} Ly(1 - y) + \frac{\lambda_\Pi}{2} (Ly + H). \end{aligned}$$

If $\lambda_U = 0$, then the function is maximized at $y = 1$ for $\lambda_\Pi > 0$ and at $y = 0$ for $\lambda_\Pi < 0$. Thus, the point that maximizes the seller profit is attained by the fully informative signal, and the point that minimizes the seller profit is attained by the uninformative signal.

If $\lambda_U > 0$, the function $W(0, y)$ is strictly concave in y so we can use the first-order condition to determine an interior solution:

$$y = \frac{\lambda_U + \lambda_\Pi}{2\lambda_U} = \frac{1}{2} \left(1 + \frac{\lambda_\Pi}{\lambda_U} \right).$$

If $\frac{L}{H} \leq \frac{1}{2} \left(1 + \frac{\lambda_\Pi}{\lambda_U} \right) \leq 1$, then the optimal y is $\frac{1}{2} \left(1 + \frac{\lambda_\Pi}{\lambda_U} \right)$. Otherwise, there is a corner solution $y = \frac{L}{H}$ or $y = 1$ for low or large values of λ_Π , respectively.

Plugging $x = 0$ into (α, β) above, we obtain $\alpha = 0$ —i.e., the signal sends realization s_H only if the value is H , so it is a high-value flagging signal. Thus each point on the right boundary arises under some high-value flagging signal. As we move the right boundary from the seller-worst point to the seller-optimal point, cutoff y increases (or equivalently, β increases) and the corresponding signal changes from the uninformative signal to the fully informative signal. If $\frac{L}{H} \geq \frac{1}{2}$, the buyer optimal point is $\lambda_\Pi = 0$, so we have $y^* = L/H$ —i.e., the uninformative signal maximizes the buyer surplus. If $\frac{L}{H} < \frac{1}{2}$, the buyer optimal point is $y^* = 1/2$ —i.e., a partially informative high-value flagging signal maximizes buyer surplus. In this case, plugging the optimal (x, y) into (α, β) , we obtain $\alpha = 1$ and $\beta = \frac{H-2L}{H-L}$.

Step 3. We characterize the “left” boundary that corresponds to $\lambda_U < 0$. The seller profit $\Pi(x, y)$ is linear in y and the buyer surplus is strictly concave in y . Because $\lambda_U < 0$, the function $W(x, y)$ is strictly convex in y , which means that the optimal y will be $\frac{L}{H}$ or 1. Because $y = \frac{L}{H}$ is equivalent to $(x, y) = (\frac{L}{H}, 1)$, we can without loss of generality assume

that $y = 1$. We have

$$\begin{aligned} W(x, 1) &= \lambda_U \left[\frac{1}{2}(H - L) - \frac{1}{2}(H - L)(1 - x)(1 + x) \right] + \lambda_\Pi \left[L + \frac{1}{2}(H - L)(1 + x) - (H - L)x \right] \\ &= \frac{\lambda_U}{2}(H - L)x^2 + \frac{\lambda_\Pi}{2} [2L + (H - L)(1 - x)]. \end{aligned}$$

The first-order condition with respect to x yields

$$\lambda_U x - \frac{\lambda_\Pi}{2} = 0 \iff x = \frac{\lambda_\Pi}{2\lambda_U}. \quad (24)$$

Thus any point on the left boundary can arise under a low-value flagging signal. As we move the left boundary from the seller-worst point to the seller-optimal point by raising λ_Π , the corresponding signal changes from no disclosure $(x, y) = (L/H, 1)$ to full disclosure $(x, y) = (0, 1)$. \square

Proof of Proposition 6. To prove part 1, take any efficient outcome. By Proposition 2, all types can be assumed to obtain the same signal. If the signal is fully revealing, then the seller obtains a strictly positive rent because almost all seller types belong to the interior $(0, 1)$ of the type space. If the signal is not fully revealing, then there must exist signal realization s such that the posterior is non-degenerate and induces all types to set price L . By the martingale property of belief, there must exist another signal realization s' under which types in $(L/H, 1)$ optimally set price H . Because the outcome is efficient, s' perfectly reveals that the value is H . Then, there are two possibilities. First, if all types in $(L/H, 1)$ are concentrated on a single point, then distribution F , which is non-degenerate, must place a positive probability on types in $[0, L/H]$. Those types would earn a strictly higher profit than without additional information, because they would strictly benefit from observing s' . Second, if types strictly above L/H are not concentrated on a single point, then below the highest type in the support of F there is a positive measure of types who strictly prefer to set price L after observing signal realization s . These types earn a strictly higher profit than without additional information.

To prove part 2, consider an implementable outcome in which a positive mass of types do not obtain the fully informative signal. Profit $\Pi(\theta)$ of any such type satisfies $\Pi(\theta) <$

$\theta H + (1 - \theta)L$. [Proposition 2](#) implies that there exists a public signal \mathcal{I} that implements the same outcome. Signal \mathcal{I} is not fully informative, because otherwise we would have $\Pi(\theta) = \theta H + (1 - \theta)L$ for all types. Thus there is a set of signal realizations that can arise under both $v = H$ and $v = L$ with positive probabilities. After observing such realizations, types that are sufficiently close to 1 set price H , which leads to no trade when $v = L$. Because the set of any such types has a strictly positive measure, the efficiency loss is strictly positive. \square

Proof of [Proposition 7](#). Part 1. If the seller's type is known to be $\theta_0 > L/H$, then in the absence of additional data the seller would set price H , leading to zero buyer surplus. As a result, at the constrained seller-optimal outcome the seller extracts full surplus. In contrast, if the seller is privately informed, then the buyer surplus under no data provision is positive, because the seller sets price L whenever her type is in $(0, L/H)$. As a result, the buyer surplus under the constrained seller-optimal outcome is strictly positive, and $0 = U < U'$, $\overline{W}(\theta_0) = \Pi > \Pi'$.

Part 2. If the seller's type is known to be $\theta_0 < L/H$, then in the absence of additional data the seller sets price L . Because this is the uniquely best possible outcome for the buyer, the constrained seller-optimal outcome is the same as under no data provision. In contrast, when the seller is privately informed, then she has a strictly positive measure of types above L/H at which she sets price H and earns strictly higher profits than under no data provision. Thus $\underline{\Pi}(\theta_0) = \Pi < \Pi'$, $\overline{W}(\theta_0) - \underline{\Pi}(\theta_0) = U > U'$. \square

Proof of [Proposition 9](#). Part 1. Assume the stated condition holds. Consider a direct menu that leads to an efficient outcome. Let $\tilde{\Theta} \subseteq \Delta(V)$ be the open set over which F has a positive density. By the condition described in Part 1, we can take $\tilde{\Theta}$ so that any type in $\tilde{\Theta}$ chooses a price strictly above the lowest possible value v_1 in the absence of additional information. In this set, a positive measure of types assign positive probability to $v = v_1$, so there exists a type $\tilde{\theta} \in \tilde{\Theta}$ such that the direct signal $\mathcal{I}(\tilde{\theta})$ recommends $p = v_1$ with a strictly positive probability. Let $\mu(\tilde{\theta})$ be the posterior belief of $\tilde{\theta}$ after that recommendation. If types in an open neighborhood of $\tilde{\theta}$ observe recommendation $p = v_1$, they would have posterior beliefs over an open neighborhood of $\mu(\tilde{\theta})$, in accordance with equation (5). Because pricing

indifference curves in $\Delta(V)$ have measure zero, a strictly positive measure of types in $\tilde{\Theta}$ would strictly prefer to follow that recommendation, and would thus strictly benefit from $\mathcal{I}(\tilde{\theta})$. It follows by incentive compatibility that the seller's rents are strictly positive.

Part 2. Assume the stated condition holds. Consider any direct menu that leads to an efficient outcome. There exists $\bar{\varepsilon} > 0$ such that for all $0 < \varepsilon < \bar{\varepsilon}$, type θ_ε , defined as

$$\theta_\varepsilon \simeq (\varepsilon^{n-1}, \varepsilon^{n-2}, \dots, \varepsilon^2, \varepsilon, 1 - \sum_{k=1}^{n-1} \varepsilon^k), \quad (25)$$

belongs to Θ and prices efficiently after observing direct signal $\mathcal{I}(\theta_\varepsilon)$, where the approximation means being in an ε^n -neighborhood.

Start with value $v = v_1$. Since θ_ε attaches strictly positive probability to $v = v_1$ and prices efficiently, $\mathcal{I}(\theta_\varepsilon)$ recommends $p = v_1$ at value v_1 with probability 1. For the recommendation to be incentive compatible, this recommendation must be sent with probability $O(\varepsilon^{k-1})$ at all values v_k , $1 < k \leq n$.¹⁵ Proceed to value $v = v_2$. Because θ_ε attaches strictly positive probability to $v = v_2$ and prices efficiently, and $\mathcal{I}(\theta_\varepsilon)$ recommends price $p = v_1$ with probability $O(\varepsilon)$ at $v = v_2$, it must be that $\mathcal{I}(\theta_\varepsilon)$ recommends price $p = v_2$ with probability $1 - O(\varepsilon) = O(1)$ at $v = v_2$. For the recommendation to be incentive compatible, it must be sent with probability $O(\varepsilon^{k-2})$ at all values v_k , $2 < k \leq n$. Proceeding analogously for all higher values, we obtain that signal $\mathcal{I}(\theta_\varepsilon)$ must take the following form:

θ_ε	$\mathcal{I}(\theta_\varepsilon)$	$p = v_1$	$p = v_2$	\dots	$p = v_n$
ε^{n-1}	$v = v_1$	1	0	\dots	0
ε^{n-2}	$v = v_2$	$O(\varepsilon)$	$O(1)$	\dots	0
\dots	\dots	\dots	\dots	\ddots	0
$1 - \sum_{k=1}^{n-1} \varepsilon^k$	$v = v_n$	$O(\varepsilon^{n-1})$	$O(\varepsilon^{n-2})$	\dots	$O(1)$

(26)

When $\varepsilon \rightarrow 0$, $\mathcal{I}(\theta_\varepsilon)$ converges to a fully informative signal. Since ε can be set arbitrarily small, incentive compatibility implies that each type $\theta \in \Theta$ earns a maximal possible rent and thus is offered a fully informative signal. □

¹⁵That is, the recommendation probability is bounded by $L \cdot \varepsilon^{k-1}$ for some fixed L .