Data Provision to an Informed Seller^{*}

Shota Ichihashi Alex Smolin

June 4, 2025

Abstract

A monopoly seller is privately and imperfectly informed about the buyer's value of the product. A designer can provide the seller with additional information, which the seller uses to price discriminate the buyer. We demonstrate the difficulty of screening the seller's information: When the buyer's value is binary, no combination of buyer surplus and seller profit can be implemented other than those achieved by providing the same information to all seller types. We use the result to characterize the set of implementable welfare outcomes 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 the Editor, Nageeb Ali; three anonymous referees; Ricardo Alonso, James Best, Teck Yong Tan, Jidong Zhou, and our discussant, Daniele Condorelli; and 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 is an important aspect of the digital economy. One prominent use of consumer data is price discrimination: A seller uses data to learn about consumer preferences and tailor prices accordingly. Given its potential harm, various parties—such as policymakers, platforms, and consumers themselves—attempt to control the flow of consumer data. However, they likely have limited ability to control the flow of data across sellers because of information asymmetry—i.e., sellers may be privately informed about their market demand and customers. This raises questions about the scope of data provision to privately informed sellers and its impact on consumer surplus, seller profits, and efficiency.

We study information provision to a privately informed seller. A monopoly seller has a product for which a buyer has an uncertain binary value: high or low. The seller's interim belief that the value is high—called the seller's type—is her private information. The seller may obtain additional information, which we model as a statistical signal about the value. We take a mechanism design approach and consider a designer who offers the seller a menu of signals. Given the menu, the seller selects a signal, learns about the buyer's value, then sets a price. Finally, the buyer decides whether to buy the product.

We study the set of all welfare outcomes the designer can implement with an arbitrary menu of signals. If the seller's type is observable, the set becomes the "surplus triangle" in Bergemann, Brooks, and Morris (2015)—i.e., the designer can attain any outcome such that the total surplus is no higher than the efficient level, the buyer's surplus is nonnegative, and the seller's profit is no lower than what she would earn without additional information.

The designer can no longer attain the surplus triangle when the seller's type is private. For example, the first-best buyer surplus, which attains the efficient allocation and gives the seller the same profit as under uniform pricing, 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 seller types would choose signals intended for other types in order to extract surplus from the buyer.

We show that the designer cannot effectively screen the seller: An outcome is implementable through some menu of signals if and only if it is implementable by providing the same signal to all seller types. This result highlights the inherent difficulty in eliciting the seller's private information through personalized information provision.

To establish this result, we first identify structural properties that govern how a menu of signals affects prices across various seller types and buyer values. These properties parallel those of one-dimensional incentive compatibility with transfers. In our setup of data screening without transfers, the properties imply a monotone likelihood structure of the prices set by different seller types. We then use this lemma to explicitly construct, for any given menu of signals, a single signal that replicates the seller's equilibrium pricing behavior induced by the original menu of signals. The signal we construct is Blackwell more informative than each signal in the original menu. However, the signal is designed so that no seller type benefits from the additional information.

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 providing the same signal to all seller types, but the signal does 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, such as one that maximizes buyer surplus, requires at most two signal realizations.

We apply the above results to solve the case in which 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. Providing any signal reduces buyer surplus, because a large mass of seller types will use the signal to extract surplus from the buyer. Also, 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 analysis reveals two insights that are more broadly relevant. First, it highlights the trade-off between increasing consumer surplus and total surplus regarding the provision of consumer data to sellers. As we show in Section 4, the trade-off holds beyond binary values. Second, in Appendix B, we show that the surplus set for binary values subsumes the set of equilibrium outcomes of various games in which the buyer and the seller exchange information



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

with each other, including voluntary disclosure of the buyer's value and a "request–consent" protocol.

Related literature Our paper relates to the recent theoretical literature on third-degree price discrimination and market segmentation. Bergemann, Brooks, and Morris (2015) show that all individually rational outcomes can arise in a single-product monopoly setting. The result is extended to multiproduct markets (Haghpanah and Siegel (2022, 2023)); competitive markets (Shi and Zhang (2020); Rhodes and Zhou (2024); Elliott, Galeotti, Koh, and Li (2022)); and two-sided markets (Condorelli and Szentes (2022)).¹ We contribute to this literature by clarifying how the seller's private information limits possible welfare outcomes and creates the trade-off between consumer welfare and efficiency.²

¹Relatedly, Roesler and Szentes (2017) and Deb and Roesler (2023) analyze the informational impacts in single-product monopoly pricing and second-degree price discrimination, respectively. Our paper is also related to recent work that refines Bergemann et al. (2015) based on how a seller endogenously obtains information, such as Ali, Lewis, and Vasserman (2023); Fainmesser, Galeotti, and Momot (2023a); and Fallah, Jordan, Makhdoumi, and Malekian (2024).

²To focus on the role of the seller's private information, we abstract away from other important forces relevant to consumer data and data regulation. See, e.g., Acquisti, Taylor, and Wagman (2016); Choi, Jeon, and Kim (2019); Argenziano and Bonatti (2021); Bergemann, Bonatti, and Gan (2022); and Fainmesser, Galeotti, and Momot (2023b).

We use a mechanism-design framework to study provision of consumer information to a privately informed monopolist, which connects our work with that of Baron and Myerson (1982). Similar mechanism-design machinery for information provision is used by Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017); Bergemann, Bonatti, and Smolin (2018); Smolin (2023); and Yang (2022). Specifically, our screening impossibility result complements the findings of Kolotilin et al. (2017), as we discuss in Section 3.2. These papers allow for a fully flexible way of designing information, following the Bayesian persuasion literature (Rayo and Segal (2010); Kamenica and Gentzkow (2011)).

2 Model

The setting involves a seller, a buyer, and a designer. 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 has a private type, which is captured by $\theta \in \Theta \subseteq [0, 1]$ and represents her belief that v = H. The type is distributed according to measure $F \in \Delta([0, 1])$.³

The seller seeks to obtain additional information, which we model as a statistical signal about the buyer's value, v. Formally, a signal $\mathcal{I} = (S, \pi)$ consists of a set S of signal realizations 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 $\pi(v)$.

At the outset, the designer posts a menu \mathcal{M} of signals. Then, the game proceeds as follows. First, the seller's type θ is drawn according to F and the buyer's value v according to θ .⁴ 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 $s \in S$ drawn according to $\pi(v)$ and posts a price $p \in \mathbb{R}$. 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* $a: V \times \Theta \to [0,1] \times \mathbb{R}$, which specifies, for each

³Given set X, we write $\Delta(X)$ for the set of all probability distributions on X.

⁴By Bayes' rule, this timing is equivalent to one in which the value is drawn before the type.

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

3 Surplus Set Characterization

We begin our analysis with the benchmark case in which the designer observes the seller's type. We then proceed to the primary scenario of unobservable types.

3.1 Observable Seller Type

Suppose that the seller's type is known to be θ_0 . Then, the efficient total surplus is $\overline{W}(\theta_0) \triangleq \theta_0 H + (1 - \theta_0)L$ and the seller's profit from optimal uniform pricing is $\underline{\Pi}(\theta_0) \triangleq \max\{\theta_0 H, L\}$. Bergemann et al. (2015) show that the set of all implementable outcomes coincides with the "surplus triangle."

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

The result generalizes to the case in which the seller's type is drawn according to distribution F and the designer observes the realized type. In this case, 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

$$\underline{\Pi} \triangleq \int_0^1 \underline{\Pi}(\theta) \mathrm{d}F(\theta) \quad \text{and} \quad \overline{\mathbf{W}} \triangleq \int_0^1 \overline{\mathbf{W}}(\theta) \mathrm{d}F(\theta). \tag{1}$$

From the ex ante perspective, $\underline{\Pi}$ is the profit the seller can guarantee and \overline{W} is the maximum feasible total surplus. We show that the set of implementable outcomes equals the surplus triangle that uses these aggregate values (omitted proofs are in Appendix A).

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 $U \ge 0$, $\Pi \ge \underline{\Pi}$, and $U + \Pi \le \overline{W}$.

The result implies that the buyer-optimal outcome $(\overline{W} - \underline{\Pi}, \underline{\Pi})$, which is efficient and gives the seller the same profit as under no additional information, is implementable. Thus, if the seller's type is observable, there is no trade-off between consumer surplus and efficiency.

3.2 Unobservable Seller Type

We now turn to the main setup, in which the seller's type is private. Our starting observation is that some points in the surplus triangle are no longer implementable, because they violate the incentive compatibility constraint in the screening problem.

For example, suppose that the seller's type is either $\theta_1 \in (0, L/H)$ or $\theta_2 \in (L/H, 1)$. The optimal uniform price is L for type θ_1 and H for type θ_2 . At the buyer-optimal outcome, type θ_1 should continue setting price L, and for efficiency, type θ_2 must obtain a signal that reveals the buyer's value to be H with a positive probability.⁵ However, type θ_1 could then strictly benefit from the signal for θ_2 , which violates the incentive compatibility constraint. Thus, the seller's private information shrinks the set of implementable outcomes.

The example is not specific to the buyer-optimal outcome: We show that the designer can never meaningfully allocate different signals to different types when types are private.

Proposition 1. (Screening Impossibility) Any implementable allocation rule can be implemented by a menu that consists of a single signal.

To prove this result, we first analyze the structure of the seller's incentive compatibility constraints. 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 $v \in \{H, L\}$,

⁵By 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.

we can view a direct mechanism as sending a value-dependent price recommendation based 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 $\beta(\theta) \geq \alpha(\theta)$.⁶ We can then express signal $\mathcal{I}(\theta)$ in matrix form as

$$\frac{\mathcal{I}(\theta)}{v = L} \frac{s_L}{1 - \alpha(\theta)} \frac{s_H}{\alpha(\theta)} \cdot (2)$$

$$v = H \frac{1 - \beta(\theta)}{1 - \beta(\theta)} \frac{\beta(\theta)}{\beta(\theta)}$$

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

$$\Pi(\theta, \theta') \triangleq (1 - \theta)(1 - \alpha(\theta'))L + \theta((1 - \beta(\theta'))L + \beta(\theta')H).$$
(3)

Incentive compatibility requires that $\Pi(\theta, \theta) \ge \Pi(\theta, \theta')$ 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

⁶Also, without loss of generality, in what follows we set $\alpha(1) = \beta(1) = 1$.

implementable allocation rule.

Lemma 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) $(\beta(\theta_3) - \beta(\theta_2))(\alpha(\theta_2) - \alpha(\theta_1)) \le (\beta(\theta_2) - \beta(\theta_1))(\alpha(\theta_3) - \alpha(\theta_2)).$

Lemma 1 follows from the seller's incentive constraints.⁷ Point 1 means that if the seller faces a larger fraction of higher-value buyers, then regardless of the mechanism, the seller sets a high price with higher probability for both low-value and high-value buyers. Thus, high-value buyers impose negative externalities on low-value buyers by increasing the likelihood that low-value buyers face high prices.⁸ Point 2 disciplines the relative price impact of an increase in the seller's type on low-value and high-value buyers. The condition is written as ratio monotonicity—i.e., $\frac{\alpha(\theta_3)-\alpha(\theta_2)}{\beta(\theta_3)-\beta(\theta_2)} \geq \frac{\alpha(\theta_2)-\alpha(\theta_1)}{\beta(\theta_2)-\beta(\theta_1)}$ —which means that increases in the seller's type disproportionately affect low-value buyers at higher types compared with lower types.

To provide additional intuition, we can interpret Lemma 1 in the context of the standard monopoly screening, such as that of Mussa and Rosen (1978), by viewing $\alpha(\theta)$ and $\beta(\theta)$ as "price" and "quality," respectively. Note that the seller's incentive compatibility constraints remain the same after we apply any θ -dependent positive affine transformation to the seller's profit. In particular, consider the following transformed profit, $\hat{\Pi}(\theta, \theta')$, defined as follows:

$$\hat{\Pi}(\theta, \theta') \triangleq \frac{1}{L(1-\theta)} \Pi(\theta, \theta') - \frac{1}{1-\theta} = -\alpha(\theta') + \left(\frac{H}{L} - 1\right) \frac{\theta}{1-\theta} \beta(\theta').$$

The transformed profit $\hat{\Pi}(\theta, \theta')$ is equivalent to consumer utility in Mussa and Rosen (1978), where $\alpha(\theta')$ and $\beta(\theta')$ are the price and quality, respectively.

In the screening problem that results from the transformed payoff $\hat{\Pi}(\theta, \theta')$, the seller now acts as a buyer. Point 1 of Lemma 1 is a standard property of such a problem: A "buyer"

⁷The conditions in Lemma 1 are necessary but not sufficient for implementation. For example, the conditions hold for $\alpha(\theta) = \beta(\theta) = 0$ for all $\theta \in \Theta$, but such an outcome may not be implementable because any type $\theta > L/H$ will set price H with positive probability under any signal.

⁸This point resonates with Galperti, Levkun, and Perego (2024a) and Galperti, Liu, and Perego (2024b), who study the impacts of negative externalities from high-value buyers to low-value buyers on the value of data records and the functioning of a data market, respectively.

of a higher type, who cares more about quality, purchases higher quality and pays more. In our model, this means that the seller who faces a larger fraction of high-value buyers is willing to accept a higher "price" (i.e., a higher probability of charging price H to low-value buyers and missing out on trade with them) in order to enjoy better "quality" (i.e., a higher probability of correctly charging price H to high-value buyers). Point 2—again written as $\frac{\alpha(\theta_3)-\alpha(\theta_2)}{\beta(\theta_3)-\beta(\theta_2)} \geq \frac{\alpha(\theta_2)-\alpha(\theta_1)}{\beta(\theta_2)-\beta(\theta_1)}$ —means that the marginal price per unit of quality is increasing in the current quality level (cf. Johnson and Myatt 2003). Indeed, if this inequality were violated, then type θ_2 , who is willing to pay an extra $\alpha(\theta_2) - \alpha(\theta_1)$ to purchase a quality increment of $\beta(\theta_2) - \beta(\theta_1)$, would find it even cheaper and thus strictly profitable to pay another $\alpha(\theta_3) - \alpha(\theta_2)$ to further increase quality by $\beta(\theta_3) - \beta(\theta_2)$.

Proceeding with the analysis, observe that Claim 3 implies that Lemma 1 applies to any menu of signals, once we view $\alpha(\theta)$ and $\beta(\theta)$ as the equilibrium pricing probabilities of each type θ under a given menu. In particular, if the designer provides the same signal to all types, sellers of different types optimally respond in such a way that the induced pricing behavior conforms with Lemma 1. Proposition 1 states that the opposite is also true—i.e., the allocation rule of *any* mechanism can be replicated by providing the same signal to all seller types. We are now ready to prove this result.

Proof of Proposition 1. The proof is constructive. We prove the statement for $\Theta = [0, 1]$; a fortiori, this proves the statement for any $\Theta \subset [0, 1]$. Take any implementable allocation rule. Claim 3 ensures that the allocation rule can be implemented by a direct mechanism, $(\alpha(\theta), \beta(\theta))_{\theta \in \Theta}$. We construct a signal $\hat{\mathcal{I}}$ such that if the designer provides $\hat{\mathcal{I}}$ to all seller types, the same outcome $(\alpha(\theta), \beta(\theta))_{\theta \in \Theta}$ arises in equilibrium. As we show, signal $\hat{\mathcal{I}}$ is Blackwell more informative than any of the signals in the original mechanism; however, no type can benefit from the extra informativeness.

The signal realization space of $\hat{\mathcal{I}}$ is S = [0, 1]. The likelihood function π is such that for all $x \in [0, 1]$, $\Pr(s \leq x \mid v = L) = \alpha(x)$ and $\Pr(s \leq x \mid v = H) = \beta(x)$. Doing so is possible, because the first property of Lemma 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 second property of Lemma 1 implies that signal $\hat{\mathcal{I}}$ satisfies the monotone likelihood

ratio property.⁹ 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 $s > \tilde{s}$. However, the choice between different thresholds is equivalent to the choice between different signals in the original mechanism, because pooling signal realizations s below and above θ results in signal $\mathcal{I}(\theta) = (\alpha(\theta), \beta(\theta))$. Hence, by 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.

Proposition 1 is reminiscent of and complementary to the equivalence between experiments and persuasion mechanisms established by Kolotilin et al. (2017). However, our analysis differs from theirs in three respects. First, in our model, the seller's type is correlated with the value, which affects the seller's assessment of and response to a signal. Second, we obtain a stronger equivalence result, which applies to allocation rules and not just to the seller's interim utility. In our setting, the stronger result is necessary for welfare analysis because the seller's interim utilities alone do not determine buyer surplus. Third, our result requires a different proof technique. Kolotilin et al. (2017) establish their equivalence by studying the seller's interim utilities and appealing to an existence argument. In contrast, we operate in the more detailed space of allocation rules and, for any given implementable allocation rule, explicitly construct a single signal that implements it.

Like the result of Kolotilin et al. (2017) relies on binary actions, Proposition 1 relies on binary values. If there are more than two values, the belief space is multidimensional and the structural properties of incentive compatibility lose their one-dimensional structure. As a result, for any given signal that is Blackwell-more-informative than all signals in the menu, we cannot ensure that this joint signal would not spill relevant information across types.

In fact, as we demonstrate in the example below, differentiated information provision can be welfare-enhancing even with three values, because buyer-optimal signals can appeal to the "right" seller type.

⁹The condition $(\alpha(\theta_3) - \alpha(\theta_2))(\beta(\theta_2) - \beta(\theta_1)) \ge (\alpha(\theta_2) - \alpha(\theta_1))(\beta(\theta_3) - \beta(\theta_2))$ is a general definition for CDF α dominating CDF β in the monotone likelihood ratio order (see Theorem 1.C.5 of Shaked and Shanthikumar (2007)).

Example 1. (Three Values) Let the possible values be $V = \{L, M, H\} = \{1, 3, 4\}$. There are three seller types, $\theta_1 = (1/2, 1/4, 1/4)$, $\theta_2 = (1/2, 1/2, 0)$, and $\theta_3 = (1/2, 0, 1/2)$, which are equally likely. In the absence of additional information, 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 single signal. Consider a menu in which types are provided with the following signals:

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 single 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 $\pi(s_0|L) \ge 2\pi(s_0|M)$ to persuade type θ_2 and $\pi(s_0|L) \ge 3\pi(s_0|H)$ 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 single signals, the profit of type θ_1 is minimized at the signal

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.

3.3 Implementable Outcome Characterization

In this section, we leverage Proposition 1 to provide results that help us characterize the set of implementable welfare outcomes. Specifically, the results enable us to (i) cast the problem of finding extreme implementable outcomes (such as the buyer-optimal outcome) as Bayesian persuasion (Lemma 2) and (ii) bound the number of signal realizations necessary to attain implementable outcomes (Corollary 1).

To begin, we follow Alonso and Camara (2016) and show that the posterior belief of one seller type pins down that of all seller types. Consider a hypothetical seller who has a prior belief $\mu_0 \triangleq \mathbb{E}[\theta]$ that v = H. We call the posterior belief of the hypothetical seller a *basic* posterior belief. Suppose that the hypothetical seller and the seller with type θ face the same signal and observe its realization. Then, by Bayes' rule, the basic posterior belief μ of the hypothetical seller and the posterior belief t of type θ must satisfy the following equation:¹⁰

$$t = \frac{\theta \mu (1 - \mu_0)}{\theta \mu (1 - \mu_0) + (1 - \theta)(1 - \mu)\mu_0}.$$
(6)

Let $t(\mu, \theta)$ denote the right-hand side of Equation 6. Thus, the basic posterior belief μ

 $^{^{10}}$ Equation 6 must hold regardless of what signal and signal realization induce these posterior beliefs.

determines the posterior beliefs of all seller types.

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

We now define the welfare functions that, for any basic posterior belief, aggregate consumer surplus and seller profit across all seller types. By Bayes' consistency, if the seller's posterior belief is t, the probability of v = H is indeed t. Thus, the players' expected payoffs after any signal realization that induces a basic posterior belief μ are written as

$$U(\mu) = \int_0^{\tilde{\theta}(\mu)} t(\mu, \theta) (H - L) \,\mathrm{d}F(\theta),\tag{7}$$

$$\Pi(\mu) = \int_0^{\theta(\mu)} L \,\mathrm{d}F(\theta) + \int_{\tilde{\theta}(\mu)}^1 t(\mu,\theta) H \,\mathrm{d}F(\theta).$$
(8)

We use these functions to characterize the set of implementable outcomes. Define by graph(U, II) a graph of the vector function that keeps track of players' payoffs (U(μ), II(μ)) at different basic beliefs μ . Any signal is characterized by a distribution of the basic beliefs, which by Bayes' rule must average to the prior belief μ_0 . Therefore, any signal implements an outcome in a convex hull of graph(U, II) 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 graph(U, II) such that the first component is μ_0 can be implemented by some signal (cf. Doval and Smolin (2024)).

Proposition 2. (Implementable Outcomes) The set of all implementable outcomes is

$$\mathcal{E} = \{ (x, y) : (\mu_0, x, y) \in \operatorname{co}(\operatorname{graph}(\mathbf{U}, \Pi)) \},$$
(9)

where functions U and Π are given by (7) and (8).

Proposition 2 enables us to cast the problem of finding extreme implementable outcomes as Bayesian persuasion. Assume that the type distribution F is continuous, which implies that the indirect payoffs (i.e., (7) and (8)) are continuous in basic belief μ . Then, by Proposition 2, the set \mathcal{E} is compact and convex. As a result, 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 split of basic posterior beliefs to maximize

$$W^{\lambda}(\mu) \triangleq \lambda_U \mathcal{U}(\mu) + \lambda_{\Pi} \Pi(\mu) \tag{10}$$

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

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

$$\max_{\tau \in \Delta(\Delta(V))} \mathbb{E}_{\mu \sim \tau}[W^{\lambda}(\mu)] \quad \text{subject to} \int_{\Delta(V)} \mu \tau(d\mu) = \mu_0.$$
(11)

The above two results bound the number of signal realizations that are necessary for implementation. If F is continuous, the indirect payoffs are continuous functions of basic belief μ and graph(U, \Pi) features a single connected component. Thus, by Fenchel-Bunt's theorem, every point in the convex hull of 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. Also, any extreme point of \mathcal{E} is implemented by a solution of a Bayesian persuasion problem, which has at most |V| = 2 signal realizations.

Corollary 1. Suppose that the seller type is continuously distributed. Any implementable outcome is implementable by providing all seller types with a signal with at most 3 signal realizations. Any extreme point of the set of implementable outcomes is implementable by a signal with at most 2 signal realizations.

3.4 Uniform Type Distribution

We now leverage the results in the previous section to analytically solve the case in which the seller's type is uniformly distributed on [0, 1]. First, we introduce two classes of signals. Consider a signal whose likelihood function in tabular form is

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 3. (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 result builds on Lemma 2 and Corollary 1. Any extreme point of \mathcal{E} solves a Bayesian persuasion problem (11) and is generated by a signal with two signal 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, information provision affects 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, information 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 providing the same signal to all seller types, information 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: 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} .

Figure 2 depicts the sets of implementable outcomes for two concrete cases of value distribution. When the seller's type is unobserved, the sets of implementable outcomes are smaller than the entire surplus triangle. 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.¹¹ The two boundaries meet at the "north" end of \mathcal{E} , which is the seller-optimal outcome and is 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 Proposition 3: 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 information. In this case, adverse selection is so severe that any additional information would on average hurt the buyer.

¹¹This observation holds for any (L, H); see the proof of Proposition 3 in Appendix A, where we characterize those flagging signals in closed form.

4 Beyond Binary Values

Even though Proposition 1 fails for more than two values, some of the welfare implications demonstrated in Section 3.4 hold more generally. Formally, suppose that the buyer's value can take more than two values, i.e., $V = \{v_1, \ldots, v_n\}$ with $v_1 < \cdots < v_n$ and n > 2. The following result shows that to attain the efficient outcome, the designer must typically give a positive rent to the seller—and, in some cases, the full surplus. Thus, the trade-off between consumer welfare and efficiency, observed in Figure 2, is a general phenomenon.

Proposition 4. (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 has 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 positive rent to the seller.
- 2. If there is a 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.

For Part 1, consider the types over which F has a positive density. Efficiency requires that these types occasionally set the lowest price $p = v_1$ to serve the lowest-value buyer, and thus obtain additional information. However, if some type weakly prefers to set $p = v_1$ after some signal realization, nearby types would strictly prefer to set the same 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 Part 2, under the stated condition, there is a sequence of types with the following properties. First, the types converge to the extreme type that places probability 1 on the highest value, v_n . Second, along the sequence of types, the corresponding signals converge to the fully informative signal. In particular, types along this sequence need to be persuaded to charge all possible prices depending on the buyer's value, 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.¹²

¹²Note that the condition of Part 2 is violated in Example 1, which is why it was possible in that case for the buyer-optimal mechanism to lead to efficiency.

5 Conclusion

We studied the provision of consumer information to a monopoly seller who has a private belief about the binary value of its product. We established the impossibility of effectively screening the seller via personalized information provision and presented a method to derive the set of possible welfare outcomes. Our results highlight the trade-off between consumer protection and efficiency, driven by adverse selection in the use of data. While our model is stylized, it helps clarify how a firm's private information may hinder the effective allocation of consumer data—an issue that should be relevant across a broad range of contexts.

References

- ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): "The Economics of Privacy," *Journal* of Economic Literature, 54, 442–492.
- ALI, S. N., A. KLEINER, AND K. ZHANG (2024): "From Design to Disclosure," arXiv preprint arXiv:2411.03608.
- ALI, S. N., G. LEWIS, AND S. VASSERMAN (2023): "Voluntary Disclosure and Personalized Pricing," *Review of Economic Studies*, 90, 538–571.
- ALONSO, R. AND O. CAMARA (2016): "Bayesian Persuasion with Heterogeneous Priors," Journal of Economic Theory, 165, 672–706.
- ARGENZIANO, R. AND A. BONATTI (2021): "Information Revelation and Privacy Protection," *Working paper*.
- AUMANN, R. AND M. MASCHLER (1995): Repeated Games with Incomplete Information, Cambridge, MA: 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," 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 (2023): "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Review of Economic Studies*, rdad100.
- DOVAL, L. AND A. SMOLIN (2024): "Persuasion and Welfare," Journal of Political Economy, 132, 2451–2487.
- ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2022): "Market Segmentation Through Information," *Working paper*.
- FAINMESSER, I. P., A. GALEOTTI, AND R. MOMOT (2023a): "Consumer Profiling via Information Design," Available at SSRN 4655468.

——— (2023b): "Digital privacy," *Management Science*, 69, 3157–3173.

- FALLAH, A., M. I. JORDAN, A. MAKHDOUMI, AND A. MALEKIAN (2024): "The Limits of Price Discrimination Under Privacy Constraints," *arXiv preprint arXiv:2402.08223*.
- GALPERTI, S., A. LEVKUN, AND J. PEREGO (2024a): "The Value of Data Records," *Review of Economic Studies*, 91, 1007–1038.
- GALPERTI, S., T. LIU, AND J. PEREGO (2024b): "Competitive Markets for Personal Data," Available at SSRN 4865950.
- HAGHPANAH, N. AND R. SIEGEL (2022): "The Limits of Multiproduct Price Discrimination," American Economic Review: Insights, 4, 443–58.

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

- JOHNSON, J. P. AND D. P. MYATT (2003): "Multiproduct Quality Competition: Fighting Brands and Product Line Pruning," *American Economic Review*, 93, 748–774.
- 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.
- MUSSA, M. AND S. ROSEN (1978): "Monopoly and Product Quality," Journal of Economic Theory, 18, 301–317.
- 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 (2024): "Personalized Pricing and Competition," American Economic Review, 114, 2141–2170.
- 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*, New York, NY: Springer.
- SHI, X. AND J. ZHANG (2020): "Welfare of Price Discrimination and Market Segmentation in Duopoly," *Working paper*.
- SMOLIN, A. (2023): "Disclosure and Pricing of Attributes," Rand Journal of Economics, 54, 570–597.

YANG, K. H. (2022): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," American Economic Review, 112, 1364–93.

Appendix A: Proofs Omitted from the Main Text

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

$$(\mathbf{U}, \mathbf{\Pi}) = \left(\int_0^1 \mathbf{U}(\theta) \mathrm{d}F(\theta), \int_0^1 \mathbf{\Pi}(\theta) \mathrm{d}F(\theta)\right).$$

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

To show the "if" direction, take any $(U, \Pi) \in \mathbb{R}^2$ such that $U \ge 0$, $\Pi \ge \underline{\Pi}$, and $U + \Pi \le \overline{W}$. 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 $\alpha + \beta \le 1$ satisfy

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

For each θ , the point

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

is in the surplus triangle and thus feasible. Aggregating the welfare outcome $(U(\theta), \Pi(\theta))$ across possible θ , we conclude that $(U, \Pi) = \left(\int_0^1 U(\theta) dF(\theta), \Pi(\theta) dF(\theta)\right)$ is also feasible. \Box *Proof of Lemma 1.* Part 1. Fix any $\theta_2 \ge \theta_1$. The system of mutual incentive constraints is

$$(1 - \alpha(\theta_1))L + \theta_1(\alpha(\theta_1)L + \beta(\theta_1)(H - L)) \ge (1 - \alpha(\theta_2))L + \theta_1(\alpha(\theta_2)L + \beta(\theta_2)(H - L)),$$
(13)

$$(1 - \alpha(\theta_2))L + \theta_2(\alpha(\theta_2)L + \beta(\theta_2)(H - L)) \ge (1 - \alpha(\theta_1))L + \theta_2(\alpha(\theta_1)L + \beta(\theta_1)(H - L)).$$
(14)

Summing over the inequalities (13) and (14) and using the fact that $\theta_2 \ge \theta_1$ we obtain

$$\alpha(\theta_2)L + \beta(\theta_2)(H - L) \ge \alpha(\theta_1)L + \beta(\theta_1)(H - L).$$
(15)

In turn, (13) and (15) together imply $\alpha(\theta_2) \ge \alpha(\theta_1)$ because

$$(\alpha(\theta_2) - \alpha(\theta_1))L \ge \theta_1(\alpha(\theta_2)L + \beta(\theta_2)(H - L) - \alpha(\theta_1)L - \beta(\theta_1)(H - L)) \ge 0.$$
(16)

Finally, (14) and (16) imply $\beta(\theta_2) \ge \beta(\theta_1)$ because

$$(\beta(\theta_2) - \beta(\theta_1))\theta_2(H - L) \ge (\alpha(\theta_2) - \alpha(\theta_1))(1 - \theta_2)L \ge 0.$$
(17)

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

$$(\beta(\theta_2) - \beta(\theta_1))\theta_2(H - L) \ge (\alpha(\theta_2) - \alpha(\theta_1))(1 - \theta_2)L, \tag{18}$$

$$(\beta(\theta_3) - \beta(\theta_2))\theta_2(H - L) \le (\alpha(\theta_3) - \alpha(\theta_2))(1 - \theta_2)L.$$
(19)

By the property in Part 1, all sides of (18) and (19) are nonnegative. Multiplying the respective smaller and larger sides and dividing the resulting inequality by $\theta_2(1-\theta_2)(H-L)L$, we obtain the desired inequality of Part 2.

Proof of Proposition 3. Step 1. We use the "boundary" to mean the boundary of \mathcal{E} . By Lemma 2 and Corollary 1, any extreme point on the boundary can arise with a signal that has two signal realizations. We parameterize such signals by $(\alpha, \beta), \beta \geq \alpha$ as

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 ywill set prices H and L after s_H and s_L , respectively. Cutoffs x and y solve

$$L = (1 - \alpha)L + x (\alpha L + \beta (H - L)),$$

$$yH = (1 - \alpha)L + y (\alpha L + \beta (H - L)),$$

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 $0 \le x \le \frac{L}{H} \le y \le 1$. The interim profit of type $\theta \in [x, y]$ is

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

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

$$\Pi(x,y) = \int_0^x Ld\theta + \int_x^y L + \frac{Hy - L}{y - x}(\theta - x)d\theta + \int_y^1 \theta Hd\theta$$
$$= Ly + \frac{1}{2}(Hy - L)(y + x) - (Hy - L)x + \frac{H}{2}(1 - y^2).$$
(21)

The buyer surplus is

$$U(x,y) \triangleq (H-L) \int_0^x \theta d\theta + (1-\beta)(H-L) \int_x^y \theta d\theta$$

= $(H-L) \int_0^x \theta d\theta + \left(H-L - \frac{(Hy-L)(1-x)}{y-x}\right) \int_x^y \theta d\theta$
= $(H-L) \int_0^y \theta d\theta - \frac{(Hy-L)(1-x)}{y-x} \int_x^y \theta d\theta$
= $\frac{1}{2}(H-L)y^2 - \frac{1}{2}(Hy-L)(1-x)(y+x).$ (22)

Step 2. We characterize signals that span the "right" boundary that corresponds to $\lambda_U \ge 0$. Take any $(\lambda_U, \lambda_\Pi) \in \mathbb{R}^2$ with $\lambda_U \ge 0$. 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 $\max_{\frac{L}{H} \le y \le 1} W(0, y)$, where

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

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

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

The first-order condition with respect to x yields

$$\lambda_U x - \frac{\lambda_\Pi}{2} = 0 \iff x = \frac{\lambda_\Pi}{2\lambda_U}.$$
(23)

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).

Proof of Proposition 4. 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 (6). 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 $\overline{\varepsilon} > 0$ such that for all $0 < \varepsilon < \overline{\varepsilon}$, type θ_{ε} , defined as

$$\theta_{\varepsilon} \simeq (\varepsilon^{n-1}, \varepsilon^{n-2}, \dots, \varepsilon^2, \varepsilon, 1 - \sum_{k=1}^{n-1} \varepsilon^k),$$
(24)

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 , $1 < k \leq n$.¹³ 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:

¹³That is, the recommendation probability is bounded by $L \cdot \varepsilon^{k-1}$ for some fixed L.

When $\varepsilon \to 0$, $\mathcal{I}(\theta_{\varepsilon})$ 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.

Appendix B: Communication Protocols

The goal of this appendix is to demonstrate that our characterization of implementable outcomes helps us understand the equilibrium outcomes of various communication games in which the seller and the buyer exchange information. Specifically, we first establish the equivalence between our mechanism-design approach and general communication games. We then consider particular communication games—such as cheap-talk communication, voluntary disclosure, and a request-and-consent protocol—and depict how the equilibrium outcomes of these games are mapped into the set of implementable outcomes in our model.

We begin by describing a general communication game, in which both the buyer and the seller may be initially imperfectly informed about the buyer's value, v, and both players may take an action that affects what information the seller will learn about v.

Formally, the game consists of the buyer and the seller, who have private types $t \in T$ and $\theta \in \Theta$, respectively. The joint distribution of (v, t, θ) is such that t and θ are independent conditional on v.

The game consists of a communication stage and a trade stage. The communication stage is described by a protocol \mathcal{P} , which is a triple, $(A_S, \{A_B(t)\}_{t\in T}, \Psi)$, where A_S is the seller's action space and $A_B(t)$ is the action space of the buyer with type t.¹⁴ Depending on an application, the buyer's action could be a message he sends to the seller or some action the buyer takes, such as whether to accept the seller's data request. The last element of the protocol \mathcal{P} is the signal scheme, $\Psi = (S, \pi)$. Here, S is the set of signal realizations the seller can possibly observe in the protocol, and $\pi : V \times A_S \times A_B \to \Delta(S)$ is the likelihood function of different signal realizations, which can depend on the value and the players' actions.

In the communication stage, the seller and the buyer take the protocol as exogenous and move in turn. First, the seller takes an action. Second, the buyer observes the seller's action and chooses his own action. A seller strategy is $\sigma_S : \Theta \to \Delta(A_S)$ and a buyer strategy is $\sigma_B : T \times A_S \to \Delta(A_B(t))$. At the end of the communication stage, the seller observes signal realization s according to Ψ .

The communication stage is followed by the trade stage, in which the seller sets a price,

¹⁴This 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.

p. The buyer learns the value v, observes the price, and decides whether to purchase the good. If trade occurs, the buyer's payoff is v - p and the seller's payoff is p; otherwise, both players obtain a payoff of 0. The solution concept is perfect Bayesian equilibrium.

Any equilibrium in a protocol results in an allocation rule $a: V \to [0, 1] \times \mathbb{R}$, which, as in the main text, 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.

Given any equilibrium of any protocol, the designer in our baseline model can replicate the same outcome by offering the seller a menu of signals, and the converse is also true.

Proposition 5. An allocation rule can arise in an equilibrium with some communication protocol if and only if it can arise in an equilibrium under some menu of signals.

Proof. The proof follows the revelation principle argument of Myerson (1982, 1983), but is adapted to our environment. For the "if" direction, note that any menu mechanism is a communication protocol such that $A_B(t) \equiv \{a_0\}$. For the "only if" direction, consider any protocol and an equilibrium in it. In this equilibrium, each action of the seller induces a signal $\mathcal{I}(a) = (S, \pi(a))$ with $\pi(a) : V \to \Delta(S)$, where the likelihood function averages over the buyer's equilibrium strategy and the signal scheme. 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. Also, 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.

Below, we describe several protocols that select different subsets of implementable outcomes, as depicted in Figure 3. We assume that the seller's type distribution F has a full support on [0, 1], which implies that the seller's type distribution *conditional on the buyer's* value, $v \in \{H, L\}$, has a full support on (0, 1) for each v.



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

Cheap Talk. Suppose that the buyer knows the value at the outset and sends a cheap-talk message—i.e., $t \equiv v$, $A_S = \{a_0\}$, $A_B(H) = A_B(L) = A_B$ for some nonempty set 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 weakly lower prices for all seller types and strictly lower prices for types above but close to L/H. 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. Thus, cheap-talk communication leads to no information.

Voluntary Disclosure. We turn to the buyer's voluntary disclosure of verifiable information, which corresponds to the following protocol: $t \equiv v$, $A_S = \{a_0\}$, $A_B(v) = \{\{v\}, \emptyset\}$ for each $v \in \{H, L\}$, and $s \equiv a_B$. That is, the buyer knows v at the outset and chooses whether to disclose his value.¹⁵

In this game, the set of equilibrium outcomes coincides with all implementable outcomes spanned by low-value flagging signals. As a starting observation, take any equilibrium and suppose that an *L*-buyer sends message \emptyset with a positive probability. After observing message \emptyset , the seller with a sufficiently low type θ sets price *L*. But if so, an *H*-buyer sends \emptyset with

¹⁵If the seller is uninformed, this setting is a special case of Ali, Lewis, and Vasserman (2023).

probability 1, because sending $\{H\}$ leads to price H for sure, whereas sending \emptyset secures a positive probability of price L.¹⁶ Thus, if an L-buyer sends message \emptyset with a positive probability, an H-buyer sends message \emptyset with probability 1.

Any equilibrium must then take one of the two forms: (i) an *L*-buyer sends message \emptyset with a positive probability and an *H*-buyer sends \emptyset with probability 1, or (ii) an *L*-buyer sends message $\{L\}$ with probability 1. In equilibrium (ii), we can without loss of generality assume that an *H*-buyer sends \emptyset with probability 1, because whether an *H*-buyer sends \emptyset or $\{H\}$ does not affect the seller's posterior belief; either way, the seller learns that v = H.

The buyer's equilibrium strategy in equilibrium (i) or (ii) can be written in a tabular form as follows:

$$\begin{array}{c|ccccc}
\mathcal{I} & L & \emptyset \\
\hline
v = L & 1 - \alpha & \alpha \\
v = H & 0 & 1
\end{array}$$
(26)

where $\alpha > 0$ in equilibrium (i) and $\alpha = 0$ in equilibrium (ii). Thus, the buyer's equilibrium strategy serves as a low-value flagging signal.

Conversely, for any low-value flagging signal (i.e., for any $\alpha \in [0, 1]$), there is an equilibrium in which the buyer's equilibrium strategy coincides with (26). Indeed, an *L*-buyer is indifferent between any messages because she will always earn a payoff of 0. The deviation of an *H*-buyer to the off-path message $\{H\}$ can be deterred by the seller's belief that places probability 1 on *H*. In particular, the two ends of low-value flagging signals—no disclosure $(\alpha = 1)$ and full revelation $(\alpha = 0)$ —can arise in some equilibria.

This observation implies that voluntary disclosure decreases the buyer's ex ante payoff, because any low-value flagging signal increases the probability that an H-buyer faces price H.¹⁷ Also, Proposition 3 implies that if the seller's type is uniformly distributed, the equi-

¹⁶An *H*-buyer and an *L*-buyer have different beliefs on the seller's type, because the distribution of the seller's type θ conditional on v differs between v = H and v = L. However, all the arguments in this section hold because we only use the fact that the belief of each-type buyer on the seller's type has a full support on (0, 1).

¹⁷This observation is consistent with Example 5 of Ali, Kleiner, and Zhang (2024), which shows that the buyer surplus is always 0 under voluntary disclosure when the buyer's value is binary and the seller's type is known to be $\theta_0 > L/H$.

librium 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.

Request-Consent Protocol. In this game, the seller chooses a signal to request and the buyer decides whether to accept it. The seller's action space A_S equals the set of all signals with S = [0, 1], and the buyer's action space equals $A_B(H) = A_B(L) = \{accept, reject\}$. The seller first chooses a signal, 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 initially perfectly informed (i.e., $t \equiv v$), then any implementable outcome can arise in some equilibrium: Take any implementable outcome and the 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 uninformed 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 1.