Data Collection by an Informed Seller*

Shota Ichihashi   Alex Smolin

April 16, 2022

Abstract

A seller faces a consumer with an uncertain value for the product. The seller has imperfect private information about the value and requests additional data to set the price. The consumer can decline any request. The consumer's willingness to provide data depends on his belief about the seller's type which in turn depends on the request. We show that the type uncertainty limits the scope of data collection: All equilibrium payoffs are spanned by fully pooling equilibria in which the seller collects the same data regardless of the type. The seller's private information lowers efficiency and profits, but benefits the consumer by fueling his skepticism and preventing excessive data collection. Having less private information may enable the seller to collect more data directly from the consumer and may lower the overall consumer welfare.

*Ichihashi: Bank of Canada, shotaichihashi@gmail.com. Smolin: Toulouse School of Economics, University of Toulouse Capitole and CEPR, alexey.v.smolin@gmail.com. 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). The opinions expressed in this article are authors' own and do not reflect the views of their respective institutions.

1 Introduction

The data collection has become an important and ubiquitous transaction between companies and consumers in the digital economy. The companies collect personal information and track consumers as the consumers browse entertainment portals, post content on social media, and buy products on e-commerce websites. This data is used to adapt the online services directly at the point of collection or later on if transferred elsewhere. When being online, each consumer makes an implicit or explicit decision of whether and what data to provide. For example, a website may ask the consumer whether to accept web cookies or to use their social media account to login to use its service. These collection requests and consent decisions shape information flows in the digital economy and the ensued allocation of goods and services.

Despite the prevalence of data collection, consumers are often uncertain about the benefits and costs of providing data to companies. A survey from 2019 reports that 59% of U.S. adults say that they understand little about what companies do with the collected data (Auxier et al. (2019)). In contrast, companies usually know well the use of data collection. Thus, when a firm requests consumers to provide data, the situation is likely to involve information asymmetry: A company knows more about the impact of data collection than a consumer.

In this paper, we develop a parsimonious Bayesian model of data collection in a buyer-seller setting that incorporates the stylized features outlined above. The seller has a product for sale. The consumer has an uncertain value for the product that can be high or low with some prior probability. The seller has some private information about the value captured by her "type"—an interim probability that the value is high. To improve her pricing decision, the seller can request from the consumer additional data, which we model as a statistical signal that can be arbitrarily informative about the value. The consumer decides whether to provide the signal, without knowing the value or the signal realization, but knowing their joint distribution. If the consumer accepts the request, the seller obtains the signal and tailors the price to the value assessment conditional on the signal realization; if the consumer rejects the request, the seller prices the product based solely on her private information. Afterwards, the consumer observes the price, learns the value, and decides whether to purchase the product.

In this setting, the consumer's benefit from providing data depends on the seller's type. If the seller's type is sufficiently high, then in the absence of additional data she would set a high price, and the consumer may benefit from providing some data to persuade the seller to lower the price. In contrast, if the seller's type is low, then she would set a low price in the absence of data, and the consumer cannot benefit from data provision since it could only increase the price. The consumer does not observe the seller's type directly, but he may infer the type from the requested signal, if different types request different signals. However, for such signalling to occur, it must be compatible with the seller's own incentives—different types must be willing to request different data.

We characterize the set of all equilibrium outcomes, i.e., the generated surplus and its redistribution across players, in perfect Bayesian equilibria of this game. In doing so, we consider separately the case of binary seller types (Section 3.2) and the case of continuously distributed types with full support (Section 3.3). Two general observations emerge. First, the right for consent empowers the consumer and enables him to guarantee a certain payoff level. Intuitively, since the consumer can decline all data requests, he cannot be made worse off than under providing no information. Second, the consumer's uncertainty about the seller's type limits the scope of data collection and request differentiation: All equilibrium outcomes can be spanned by fully pooling equilibria in which all seller types request the same signal. In other words, the differentiated data collection, while being able to promote efficiency, is hindered by adverse selection across seller types into different data requests. In some cases, e.g., with the uniformly distributed types, these two properties prevent the seller from collecting any data.

At the same time, any pooling data collection can arise in equilibrium as long as it benefits the consumer on average. In such an equilibrium, whenever the seller deviates and requests a different signal, the consumer rejects the request, believing that the request comes from a seller who would use the signal to extract more surplus. We use our results to provide an alternative, geometric characterization of the set of equilibrium outcomes and place bounds on the complexity of equilibrium signals.

Our equilibrium characterization reveals how the seller's private information changes the data collection patterns and welfare of the consumer and the seller. The consumer is weakly better off when the seller type is private as opposed to being observed by the consumer. If the seller's type were directly observed, then the seller would request a type-dependent signal that attains the efficient allocation while giving the consumer the same payoff as under no data collection. However, if the consumer does not know the seller's type, then there may be equilibria in which data collection strictly benefits the consumer. In such an equilibrium, the seller's attempt to collect more profitable data is deterred by the consumer's skepticism. At the same time, the seller's profits and total surplus in equilibrium are lower compared to the case in which the consumer observes the seller's type. Facing uncertainty regarding the seller's type, the consumer prefers not to provide detailed information that is necessary for some seller types to price efficiently.

Our results imply that the seller may benefit from committing to not collect extraneous information on consumers. Collecting information without a consumer's consent may kill the future opportunity for the seller to collect data directly from the consumer, because the seller's private information makes the consumer more uncertain and skeptical about the data use. This observation highlights a potential benefit for firms from avoiding tracking consumers without their explicit consent, e.g., via the access to third-party cookies. Our analysis also reveals that such a strategic commitment to not track consumers, while seemingly privacy friendly, could in fact harm consumers by enabling collection of more data.

Related literature.— First and foremost, our paper relates to the strand of theoretical literature that studies possible welfare impacts of price discrimination and market segmentation (Bergemann, Brooks, and Morris, 2015; Roesler and Szentes, 2017; Haghpanah and Siegel, 2019; Ichihashi, 2020; Shi and Zhang, 2020; Haghpanah and Siegel, 2021; Rhodes and Zhou, 2021). These papers typically take market segmentation, or information structure, as exogenously given and either identify the set of all outcomes that could possibly arise or an outcome that maximizes some welfare criterion. We contribute to this literature by endogenizing the information structure via a protocol of data requests widespread in practice. The set of equilibrium outcomes under such timing is smaller than the set of all feasible outcomes, enabling us to make more precise predictions.

Second, our paper contributes to the recent and rapidly growing literature on online privacy and privacy regulation (Acquisti, Taylor, and Wagman, 2016; Choi, Jeon, and Kim, 2019; Fainmesser, Galeotti, and Momot, 2021; Argenziano and Bonatti, 2021; Bergemann, Bonatti, and Gan, forthcoming). We contribute to this literature by highlighting information asymmetry whereby a firm knows more about the impacts of data collection than consumers—a situation that is prevalent in practice but not well understood in the literature.

Finally, we build on several existing methodologies. Our mechanism-design approach to study equilibrium outcomes follows Myerson (1983), and our seller can be viewed as his informed principal.[1] The mechanism-design machinery has recently been used in the literature on information provision by Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017), Bergemann, Bonatti, and Smolin (2018), Smolin (2020), and Yang (2022). The flexible way of designing information collection follows the Bayesian persuasion literature (e.g., Rayo and Segal (2010) and Kamenica and Gentzkow (2011)). The equilibrium reduction to pooling outcomes builds on the extreme-point analysis of Kleiner, Moldovanu, and Strack (2021). The convex-hull characterization of the equilibrium-outcome set follows Doval and Smolin (2021).

2 Model

A seller (she) has a product to sell to a consumer (he). The ex post value of the product to the consumer is $v \in V = \{L, H\}$ with $H > L > 0$, distributed according to $\mu_0 \in \Delta(V)$.[2] Initially, neither the seller nor the consumer knows v. However, the seller observes a signal that induces a spread of interim beliefs θ distributed over $\Theta \subseteq \Delta(V)$ according to measure F, so that $\mathbb{E}_F[\theta] = \mu_0$. We call the seller's interim belief $\theta \in \Theta$ the seller *type*. The seller observes her type, but the consumer knows only type distribution F.

The seller can request additional data from the consumer. We model data as a statistical signal 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 distributions $\{\pi(\cdot|v)\}_{v \in V}$ over S. The seller can request any signal.

[1] Nonetheless, our equilibrium characterization—that all equilibrium outcomes are spanned by pooling equilibria—is not a direct implication of and stronger than the Inscrutability Principle.

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

The data-collection game proceeds as follows. First, the nature draws θ according to F and v according to θ.[3] Second, the seller privately observes her type θ and chooses signal $\mathcal{I}$ to request from the consumer. Third, the consumer observes the signal requested and, without knowing v, θ, or s, decides whether to accept the request and provide the data. If the consumer accepts, the seller observes the signal realization s drawn from $\pi(\cdot|v)$ and posts a price $p \in \mathbb{R}$ for the product; otherwise, the seller posts a price based solely on her interim belief θ. Finally, the consumer observes value v and price p, and decides whether to buy the product at that price. If the consumer buys, he obtains payoff $v - p$ and the seller obtains payoff p. Otherwise, both players obtain zero payoffs.

The solution concept is a perfect Bayesian equilibrium as defined by Fudenberg and Tirole (1991). According to this concept, the players' actions cannot reveal information that they do not have, be it on or off the equilibrium path. In particular, the seller's data request can only provide information about her type, whereas the consumer's acceptance decision cannot reveal any information about the value.

For any given strategy profiles adopted by the players, we denote the consumer's expected payoff when the seller's type is θ by $U(\theta)$ and the seller's expected payoff by $\Pi(\theta)$. These interim payoffs give rise to the ex ante payoffs by aggregating across all types.

Definition 1. An *outcome* of any given players' strategies is the pair of the corresponding ex ante payoffs of the consumer and the seller, $\left(\int_\Theta U(\theta)\mathrm{d}F(\theta),\ \int_\Theta \Pi(\theta)\mathrm{d}F(\theta)\right)$.

We call an outcome of any pair of equilibrium strategy profile an *equilibrium outcome*. We call an outcome of strategies that maximize the sum of players' ex ante payoffs an *efficient outcome*. Our primary goal is to characterize the set of all equilibrium outcomes, $\mathcal{E}$, and to assess welfare impacts of data collection.

In addition, we are interested in the patterns of equilibrium data collection. We call the behavior induced by some strategy profile *full data collection* if all types collect completely informative signal $\overline{\mathcal{I}}$, and we call it *no data collection* if all types collect completely uninformative signal $\underline{\mathcal{I}}$. We call a data collection strategy *differentiated* if at least two seller types request different signals; otherwise, we call the data collection strategy *non-differentiated*.

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

Our secondary goal is to understand how much data is collected in equilibrium and whether differentiated data collection expands the set of equilibrium outcomes of the game.

2.1 Discussion of Modeling Assumptions

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

Data as a Signal. As common in the literature on information design, we model data as a statistical signal informative about consumer's value. In doing so, we abstract away from the algorithmic implementation of data analysis and focus directly on the assessments the seller makes about the value. Different signals correspond to different kinds of data the seller may request from the consumer, e.g., access to web cookies, geolocation, or permission to track consumer behavior on the seller's website. In turn, a signal realization corresponds to the outcome of the data collection registered as an entry in a database, e.g., the content of consumer cookies, his geolocation, or the web history of behavior. We allow the seller to request any kind of signal. This assumption captures the idea of flexible data collection and reflects the plurality of data the seller may request from the consumer.

Uninformed Consumer. The consumer is uninformed at the time of data request but learns the true value before purchasing the product. In other words, when deciding whether to purchase the product, the consumer does not make mistakes, but when asked for data, the consumer assesses the request from the prior perspective. This informational assumption is driven by the typical features of online interactions: First, a decision to provide data often comes much earlier in time than purchasing decisions on the platform, typically at a registration; second, to assess the value, the consumer needs to see the product and to face a concrete decision problem of whether to purchase the product at a given price. Also, we note that the alternative case of an informed consumer would give raise to equilibria with information unraveling, in which the consumer provides complete information driven by the threat to be perceived of having a high value. We believe such unraveling behavior is already well understood in the literature and may be less relevant in practice since it relies on the excessive awareness and sophistication of consumers.

Informed Seller. We assume that the seller is more informed than the consumer at the time of a data request. There are at least two reasons for that to happen in practice. First, the seller may have access to some information irrespectively of the consumer's will. Such information may come from technological constraints, such as the consumer's IP-address, or be an uncontrolled consequence of online interactions, such as some information on past purchasing decisions. Second, the seller may be more informed about the general quality of her product. A higher-quality seller would then naturally believe that the consumer's value is more likely to be high. We analyze how this seller heterogeneity affects equilibria and whether seller's interim information is revealed in equilibrium by her data request.

Consumer's Options. We assume that the consumer's consent is required for data collection: The consumer has the right to reject the data request and still have the possibility of a trade. If the consumer did not have that right, the seller could extort complete data provision by threatening to decline her services to the consumer. While present at some extent in practice, we believe that such data extortion is not prevalent because it could trigger a consumer backlash and, moreover, is explicitly prohibited by regulatory acts such as the General Data Protection Regulation (GDPR) in the European Union.[4] Our results shed light on how much welfare this right for consent brings to consumers.

Price Discrimination. We assume that the seller can use the collected data to engage in third-degree price discrimination. In practice, while sellers may be reluctant to display different base prices to different consumers, there are at least two ways, both widespread in the digital economy, by which the sellers may price discriminate indirectly. First, the sellers may offer to the consumers personalized discounts. By setting the base price high and changing the size or frequency of discounts, the seller can effectively implement any discriminatory pricing. Second, the sellers may offer the consumers personalized recommen-

[4]For example, GDPR stipulates that a company cannot require consent to data processing as a condition of using the service, unless some information, such as mailing address, is necessary for providing the service. See, e.g., https://gdpr.eu/gdpr-consent-requirements/.

dations, leading them towards the products that vary in price but are similar in quality. By having a sufficiently rich inventory of such products, the sellers could again engage in price discrimination.

3 Analysis

With a small abuse of notation, in what follows we redefine the seller's type to be the interim probability that she assigns to $v = H$, so that $\theta \in [0, 1]$. Omitted proofs are in the Appendix.

3.1 A Single Seller Type

We begin our analysis by studying a benchmark in which the seller's type space is a singleton, $\Theta = \{\theta_0\}$. In this case, the seller type is commonly known between the seller and the consumer.[5] This case allows to present the basic trade-offs the players face and highlight the restrictions put on data collection by strategic incentives.

The equilibria that emerge in this case depend on the optimal pricing in the absence of additional information. This pricing is determined by the choice of the seller between the low price $p = L$, which guarantees a trade and secures a profit of L, and the high price $p = H$, which is accepted only by the high-value consumer and results in the expected profit $\theta_0 H$. The optimal prior pricing is thus determined by a threshold belief

$$\hat{\theta} \triangleq \frac{L}{H}. \tag{1}$$

If $\theta_0 < \hat{\theta}$, then in the absence of additional information the seller would charge a low price. This is the best outcome for the consumer: The trade is efficient and the seller charges the lowest price in the support of her best responses. The consumer understands it and would never provide any data that would sacrifice some of his surplus, i.e., the signal that could possibly induce a high price. As a result, the equilibrium outcome is unique and is the same

[5]Clearly, this setting also addresses any subgame in the benchmark in which the seller type is directly observed by the consumer.

as under no data collection.[6]

In contrast, if $\theta_0 > \hat{\theta}$, then in the absence of additional information the seller would charge a high price. This is the worst outcome for the consumer, which makes him willing to give away data easily: The consumer strictly prefers to provide any data that could possibly induce the low price. This possibility allows the seller to extract full surplus. In a unique equilibrium the seller requests full information; otherwise, she could always request a signal that would be arbitrarily close to full information and be strictly preferred by both herself and the consumer, which is a contradiction.[7] The following proposition summarizes this discussion.

Proposition 1. (Single-Type Equilibrium) *The equilibrium outcome in the data-collection game with a single type θ_0 is as follows. If $\theta_0 < \hat{\theta}$, then the equilibrium outcome is unique and obtained by no data collection. If $\theta_0 > \hat{\theta}$, then the equilibrium outcome is unique and obtained by full data collection. In either case, the outcome is efficient and the consumer does not benefit from data collection.*

Proposition 1 provides sharp predictions on outcomes when the seller has no private information. Either the consumer expects a high price and provides full information to the seller or the consumer expects a low price and does not provide any valuable information. This finding contrasts with the outcomes achievable under *exogenous* information structure. As Bergemann, Brooks, and Morris (2015) show, if the information available to the seller is fixed exogenously, then *any* feasible outcome that gives the seller at least her no-information profit can be achieved. However, in practice the information is collected by the sellers under the consumer consent, and the information structure arises endogenously in equilibrium. As we showed, this may lead to extreme information outcomes and may deprive the consumer of all rents. However, as we show in the next section, the scope of information collection is more limited if the consumer is uncertain about the information the seller already possesses prior to the data request.

[6] This argument uses the "don't signal what you don't know" property of a PBE: Since the consumer is uninformed, his decision does not convey information about the valuation.

[7] Bergemann, Brooks, and Morris (2015) establish the existence of such a signal.

3.2 Binary Seller Types

We proceed by studying the case of minimal seller heterogeneity, when the type takes one of two possible values $\Theta = \{\theta_1, \theta_2\}$, $0 < \theta_1 < \theta_2 < 1$ with the respective probabilities $f_1 \triangleq \Pr(\theta_1) > 0$ and $f_2 \triangleq \Pr(\theta_2) > 0$. This case allows to clarify the role of the seller's heterogeneity and outline some general features of the equilibrium behavior.

The players' choices and incentives are similar to the case of a single type, but have two important differences. First, the seller's data request and the pricing strategy may depend on her type. Second, the consumer is uncertain about the seller's type and makes the inference about it from the data request, by Bayes' rule if the request is expected in equilibrium and arbitrarily if not. As with a single type, the equilibria that emerge depend on the seller's optimal price in the absence of additional data, which depends on whether the type is above or below the threshold $\hat{\theta}$. If θ_1 and θ_2 set the same interim price, be it low or high, we call them *congruent*.

The equilibrium analysis of the cases of congruent types resembles the analysis of a single type. If $\theta_1 < \theta_2 < \hat{\theta}$, then in the absence of additional data both seller types charge low prices. This is the best outcome for the consumer. The consumer understands it and would never provide any data that would sacrifice some of his surplus. As a result, the equilibrium outcome is unique and is the same as under no data collection. If $\hat{\theta} < \theta_1 < \theta_2$, then in the absence of additional information both seller types would charge high prices. This is the worst outcome for the consumer, which makes him willing to give away data easily. This possibility allows the seller to extract full surplus. In equilibrium, both types request full information; if one of them did not, then she could always request a more informative and profitable signal that would be preferred by both herself and the consumer, which is a contradiction. The equilibrium outcome is again unique and corresponds to full data collection.

The case of non-congruent types, $\theta_1 < \hat{\theta} < \theta_2$, is richer. The consumer would like to provide data to type θ_2 to persuade her to set a lower price, but not to type θ_2. At the same time, the seller is more disciplined in her requests because the consumer may suspect that the request comes from type θ_1. We solve this case by making the following observations.

First, any equilibrium should give the consumer at least the utility that he would obtain by providing no information at all; this property should hold after every equilibrium request and thus holds on average. Second, any pooling strategy profile in which both types request the same signal that on average benefits the consumer can be supported in equilibrium: On the equilibrium path, the consumer does not know which type requests the data and is willing to accept the request; off the equilibrium path, the consumer holds cautious beliefs concentrated on type θ_1 and reject all requests, deterring the seller's deviations.

It remains to show whether pooling equilibria span all equilibrium outcomes. It is easy to see that any outcome of a fully separating equilibrium can be replicated in a pooling equilibrium. Indeed, in such equilibrium type θ_1 is perfectly identified by her requests and thus cannot obtain any surplus above L, since it would have to come at the expense of the consumer. Hence, type θ_1 cannot benefit from requesting signal $\mathcal{I}(\theta_2)$ and if obtaining it she would optimally keep setting the price $p = L$. As a result, this separating equilibrium can be replicated by a pooling equilibrium in which both types request signal $\mathcal{I}(\theta_2)$.

The case of semi-separating equilibria, in which the types randomize over several information structures, is less straightforward, because the data request does not have to fully reveal the type and thus even the low type may benefit from data collection. We study this case by using the mechanism-design tools of Myerson (1983), anticipating the continuum-type analysis of the next section. By the revelation principle, any equilibrium outcome can be implemented in a *direct* equilibrium in which the consumer accepts the on-path data requests, each signal realization recommends a price, and each type follows the recommendation. Thus, in a direct equilibrium types θ_1 and θ_2 request respectively the signals

$$\begin{array}{c|cc} \mathcal{I}(\theta_1) & s_{1L} & s_{1H} \\ \hline L & 1-\alpha_1 & \alpha_1 \\ H & 1-\beta_1 & \beta_1 \end{array} \qquad \text{and} \qquad \begin{array}{c|cc} \mathcal{I}(\theta_2) & s_{2L} & s_{2H} \\ \hline L & 1-\alpha_2 & \alpha_2 \\ H & 1-\beta_1 & \beta_2 \end{array}, \tag{2}$$

the consumer is willing to provide the data, and type θ_i sets prices $p = L$ and $p = H$ after observing s_{iL} and s_{iH}, respectively.

For the equilibrium to be compatible with the seller's incentives, no type should benefit from mimicking the data request and the pricing responses of another type. These necessary incentive constraints can be written as:

$$(1-\alpha_i)L+\theta_i(\alpha_i L+\beta_i(H-L)) \geq (1-\alpha_j)L+\theta_i(\alpha_j L+\beta_j(H-L)), \tag{3}$$

for $i=1,2, j\neq i$. We show in the Appendix that the inequalities (3) imply that $\alpha_2 \geq \alpha_1$ and $\beta_2 \geq \beta_1$. That is, in any equilibrium the higher seller type charges the price $p=H$ more often than the lower seller type, for both values. This monotonicity is intuitive and reminiscent of the allocation monotonicity in a mechanism design problem, but it is stronger in that it applies to a likelihood function in each state and not only to their linear combination $\alpha_i L+\beta_i(H-L)$.

Importantly, this monotonicity enables a construction of a pooling equilibrium in which both types request the signal

$$\begin{array}{c|ccc} \hat{\mathcal{I}} & s_{LL} & s_{LH} & s_{HH} \\ \hline L & 1-\alpha_2 & \alpha_2-\alpha_1 & \alpha_1 \\ H & 1-\beta_2 & \beta_2-\beta_1 & \beta_1 \end{array} . \tag{4}$$

The defining feature of this signal is that by merging signals s_{LL} and s_{LH} one obtains the signal $\mathcal{I}(\theta_1)$ whereas by merging signals s_{LH} and s_{HH} one obtains the signal $\mathcal{I}(\theta_2)$; that is, $\hat{\mathcal{I}}$ is Blackwell more informative than either of the direct signals. We claim that a pooling equilibrium in which both types collect signal $\hat{\mathcal{I}}$ exists. Moreover, the seller types optimally respond to this data differently and in a way that replicates the direct allocation for both types, and thus leads to the same equilibrium outcome. Indeed, the informational content of realization s_{LL} of $\hat{\mathcal{I}}$, its likelihood description, equals the informational content of realization s_{2L} of $\mathcal{I}(\theta_2)$. Hence, both types optimally respond by low prices to it as justified by the θ_2's response in the direct equilibrium. By the analogous argument both types optimally respond by high prices to s_{HH}. However, the types' responses differ to realization s_{LH}. The choice between the low and the high price after s_{LH} is precisely the choice between the two allocations of the direct equilibrium. By incentive compatibility, type θ_1 prefers to set low price, whereas θ_2 prefers to set high price. The resulting allocation replicates the direct allocation and the consumer is willing to provide the data requested. The pooling equilibrium is completed by supporting the deviating requests by skeptical beliefs.

For any seller's type $\theta \in \Theta$ and signal $\mathcal{I}$, let $U(\theta, \mathcal{I})$ denote the consumer's expected payoff when a type-θ seller prices optimally based on signal $\mathcal{I}$, breaking ties in favor of a lower price.[8] We can then summarize the discussion in the following proposition.

Proposition 2. (Two-Type Equilibrium) *The equilibrium outcomes in the data-collection game with binary types are as follows. If $\theta_1 < \theta_2 < \hat{\theta}$, then the equilibrium outcome is unique and obtained by no data collection. If $\hat{\theta} < \theta_1 < \theta_2$, then the equilibrium outcome is unique and obtained by full data collection. If $\theta_1 < \hat{\theta} < \theta_2$, then all equilibrium outcomes are spanned by pooling equilibria in which both seller types request signals $\mathcal{I}$ that satisfy*

$$f_1 U(\theta_1, \mathcal{I}) + f_2 U(\theta_2, \mathcal{I}) \geq f_1 U(\theta_1, \underline{\mathcal{I}}) + f_2 U(\theta_2, \underline{\mathcal{I}}). \tag{5}$$

Proposition 2 highlights the power and the limits of data collection by the seller. On the one hand, the seller can collect data even if the consumer knows that it might be used against her. This is particularly evident in the case of high and congruent types, that is $\theta_2 > \theta_1 > \hat{\theta}$, in which the seller extorts full information about the consumer's valuation. On the other hand, the data collection is limited by the consumer's option to reject the incoming requests: In any equilibrium, the consumer is weakly better off than under no data collection. Moreover, the consumer's caution may deter the seller's off-path deviations and support the beneficial data collection: Any pooling data collection that benefits the consumer can be supported in equilibrium. However, the seller's adverse selection across the data requests limits the use of differentiated data collection: Only the outcomes achieved by pooling equilibria can arise in equilibrium. In the next section we show that these results are not an artifact of the seller having only two types and hold for *any* type distribution over $[0, 1]$.

Propositions 1 and 2 illustrate how the seller's private information changes the welfare implication of data collection (Section 4 presents formal results for a continuum of types). First, the seller's private information reduces efficiency and profits. If the seller's type were

[8]Even though we specify the tie-breaking rule in the definition of $U(\theta, \mathcal{I})$, we recall that our solution concept does not impose any tie-breaking assumption.

known, she would price efficiently and leave the consumer the same payoff as under no data collection. The seller's private information discourages the consumer from providing precise data, which lowers total surplus and the seller profits. At the same time, the consumer's ex ante payoffs can be higher when the seller has private information: The consumer's skepticism disciplines the seller and prevents the collection of signals that are more profitable but less beneficial to the consumer than the equilibrium signal. Such an outcome can arise only when the consumer is uncertain about the seller's type.

The results also have an implication on the seller's data policy. If the seller has a known type $\mu_0 = \mathbb{E}[\theta] > \hat{\theta}$ she could obtain a fully informative signal. However, once the seller acquires private information, which moves her belief into $\theta_1 < \hat{\theta}$ or $\theta_2 > \hat{\theta}$, she can never obtain full information. Thus the seller can be worse off when she has more information about the consumer prior to the data request. To avoid such a situation, the seller may commit to not collect data without a consumer's consent. Such a data policy, however, may decrease consumer surplus.

3.3 Continuum of Seller Types

We proceed with studying the general case of a continuum of types. This case features more possibilities for differentiated data collection, because different types may request different signals with complex pooling patterns across types. At the same time, the more signals are requested in equilibrium, the more on-path deviations each type has. Despite the daunting multiplicity of incentives, we show that the equilibria of this general case preserve the qualitative features of the benchmarks studied in the previous sections.

If the types are congruent, i.e., if all seller types set the same price in the absence of data, then the data collection takes extreme forms by the same arguments as in the binary-type case. If all types satisfy $\theta > \hat{\theta}$, then they all collect full information; if all types satisfy $\theta < \hat{\theta}$, then they all effectively collect no information. Any equilibrium is efficient, and the consumer welfare is the same as under no data collection.

For the rest of this section, we study the case of non-congruent types. In particular, we consider a continuum of seller types distributed over the unit interval $[0, 1]$ according to an arbitrary cumulative distribution F with density f which is strictly positive everywhere.

Proposition 3. (Pooling Outcomes) *The set of equilibria has the following properties:*

1. *An outcome arises in some equilibrium if and only if it arises in a pooling equilibrium.*

2. *A pooling equilibrium in which all seller types request and obtain signal $\mathcal{I}$ exists iff $\mathcal{I}$ benefits consumer relative to no data collection, $\int_\Theta U(\theta,\mathcal{I})\mathrm{d}F(\theta) \geq \int_\Theta U(\theta,\underline{\mathcal{I}})\mathrm{d}F(\theta)$.*

3. *In any equilibrium, no seller type obtains a fully informative signal. In the seller-preferred equilibrium, the consumer's payoff is the same as under no data collection.*

The proof of Part 2 follows the same logic as the proof of Proposition 2: Any signal $\mathcal{I}$ collected in a pooling equilibrium must increase the consumer's ex ante payoff as otherwise he would refuse to provide data. Conversely, for any such signal $\mathcal{I}$, we can construct a pooling equilibrium in which the consumer agrees to provide $\mathcal{I}$ and rejects any deviant request, believing that it comes from seller type $\theta < \hat{\theta}$.

To prove Part 1, we begin with analyzing an auxiliary mechanism-design problem in which a fictitious principal commits to a mechanism, without transfers, that provides the seller with a signal as a function of her reported type. After obtaining a signal, the seller sets a price and the consumer makes an optimal purchase decision. The incentive compatibility of an auxiliary mechanism captures the incentive of each seller type in the game regarding whether to deviate and request an "on-path" signal, which another seller type would request in equilibrium. The set of the implementable outcomes in this mechanism-design problem contains all equilibrium outcomes in the data collection game.

The analysis of the mechanism-design problem consists of two parts. First, we characterize the implementable outcomes in terms of a profile $\Pi = (\Pi(\theta))_{\theta\in\Theta}$ of the seller's interim payoffs. For any θ, $\Pi(\theta)$ must be between the revenue $\underline{\Pi}(\theta)$ from no information and the revenue $\overline{\Pi}(\theta)$ from perfect information. In addition, any implementable Π must be convex. We show that these conditions fully characterize all implementable outcomes: Π is implementable if and only if it is convex and between $\underline{\Pi}$ and $\overline{\Pi}$.

We then show that any implementable Π can be implemented by a mechanism that allocates to all seller types the same signal. To prove this result, we first characterize the "extreme points" of the set of implementable outcomes, using results by Kleiner, Moldovanu,

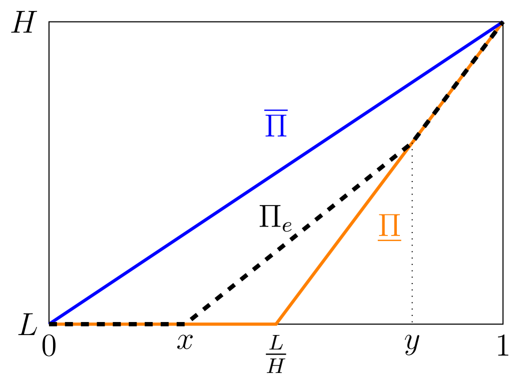

Figure 1: Equilibrium interim payoffs $\Pi(\theta)$ of the seller. $\overline{\Pi}(\theta) = \theta H$, $\underline{\Pi}(\theta) = \max\{L, \theta H\}$.

and Strack (2021) on extreme points and majorization. Figure 1 depicts $\overline{\Pi}$, $\underline{\Pi}$, and an extreme point Π_e. For any extreme point Π_e, we can find two numbers, $x \in [0, \frac{L}{H}]$ and $y \in [\frac{L}{H}, 1]$, such that $\Pi_e = \underline{\Pi}$ if $\theta \in [0, x] \cup [y, 1]$, and Π_e is a straight line on $[x, y]$. We then prove, by construction, that any extreme point Π_e is implementable by a mechanism that gives all seller types the same signal. Because any implementable Π is a convex combination of extreme points, the principal can implement any such Π by randomizing over signals according to the weights of the convex combination.

Finally, observe that any equilibrium profile Π is implementable in the auxiliary mechanism-design problem, because we can use the seller's equilibrium strategy as an incentive compatible mechanism. By Lemma 2 there is a signal $\mathcal{I}$ such that if all seller types obtain $\mathcal{I}$ then they receive payoffs Π. We then construct a pooling equilibrium of the data collection game, in which all seller types collect $\mathcal{I}$. Both for the seller and the consumer, this pooling equilibrium attains the same outcome as the original equilibrium.

Remark 1. Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) study the design of persuasion mechanisms for a privately informed receiver, and show that any implementable outcome can arise in a mechanism that provides the same information to all receiver types. Our result on the auxiliary mechanism-design problem is similar to theirs; however, the two approaches are different and complementary. In their model, the receiver privately knows the cost of taking an action, but shares the same belief as the sender about the state. This assumption enables them to view the receiver's interim utility function (i.e., our Π) as the distribution of posterior beliefs that applies to all receiver types. We cannot use this approach because the same signal induces different posterior beliefs depending on the seller's type. Conversely, our approach is not applicable in their model, because the set of uncertain states is an interval and thus the set of implementable payoff profiles of the receiver cannot be described by a majorization constraint.

Equilibrium Outcome Characterization. According to Proposition 3, to characterize the set of all equilibrium outcomes it suffices to characterize the set of outcomes of pooling equilibria, in which all seller types obtain the same signal. To characterize the latter, it suffices to first understand the players' payoffs after any given signal realization and then to aggregate these payoffs across realizations of a signal.

Towards the former, for any given signal $\mathcal{I}$ and a realization s, define by $\mu \in [0, 1]$ the posterior belief that s induces in a hypothetical uninformed agent with a prior belief μ_0, $\mu \triangleq \Pr(v = H \mid s, \mu_0)$. We call this belief μ a *basic* posterior belief. By Bayes' rule, the same signal realization observed by the informed seller of type θ results in the posterior belief $t \triangleq \Pr(v = H \mid s, \theta)$ equal to[9]

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

The posterior belief increases both in μ and in θ, equals zero if either of the arguments equals zero, and equals one if either of the arguments equals one. The formula (6) applies to all seller types. Therefore, the basic posterior belief μ fully determines the distribution of the seller's posterior beliefs. The higher is μ, the higher is the corresponding distribution of the seller posterior beliefs, in the sense of a first-order stochastic dominance—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 the high price if $t > \hat{\theta}$ and sets the low price if $t < \hat{\theta}$. Equivalently, if 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 sets a low price if $\theta < \tilde{\theta}(\mu)$ where the threshold type $\tilde{\theta}$ is uniquely defined by $t(\mu, \tilde{\theta}(\mu))) = \hat{\theta}$.

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

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' equilibrium payoffs after a signal realization s that induces a basic posterior belief μ can be written as:

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

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

We can use this indirect payoff functions to characterize the set of equilibrium outcomes. Define by $\mathrm{graph}(U,\Pi)$ a graph of the vector function that keeps track of players' payoffs $(U(\mu),\Pi(\mu))$ at different basic beliefs $\mu \in \Delta(V)$. Any signal $\mathcal{I}$ may allow for multiple signal realizations and hence induce multiple beliefs μ. By Bayes' rule, these beliefs must in expectation equal the prior belief μ_0. Therefore, any distribution of beliefs μ induces an outcome in a convex hull of $\mathrm{graph}(U,\Pi)$ 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 the graph such that the first component is μ_0 can be induced by some distribution over the points of the graph with the average belief equal to μ_0, and hence can be achieved by some signal (cf. Doval and Smolin (2021)). Combining this argument with Proposition 3, we obtain the following characterization.

Proposition 4. (Equilibrium Outcomes) *The set of equilibrium outcomes in the data collection game with seller types distributed according to distribution F is*

$$\mathcal{E} = \{(x,y) : (\mu_0,x,y) \in \mathrm{co}(\mathrm{graph}(U,\Pi)), x \geq \underline{U}\}. \tag{9}$$

For any given type distribution, Proposition 4 enables a geometric characterization of the set of equilibrium outcomes. Moreover, it highlights the effect that the seller heterogeneity captured by the type distribution has on the equilibrium outcomes. First, the type distribution determines consumer's outside option $\underline{U}$, his expected payoff after providing no additional data to the seller. The more likely the seller is to charge the high price, the lower is $\underline{U}$. Second, the type distribution affects the players' indirect payoffs according to (7) and

(8). This impact is more subtle because these payoffs enter the characterization via a convex hull of their graph.

In addition, Proposition 4 places bounds on the necessary complexity of equilibrium signals. Namely, equations (7) and (8) reveal that the indirect payoffs are continuous functions of basic belief μ, because $\tilde{\theta}(\mu)$ is continuous and the types are continuously distributed over [0,1]. Hence, graph(U, Π) features a single connected component and by Fenchel-Bunt's theorem every point in its convex hull can be generated by a randomization over at most $|V| - 1 + 2 = 3$ points. Such randomization corresponds to a signal with at most 3 signal realizations. We obtain the following.

Corollary 1. (Signal Complexity) *Any equilibrium outcome can be obtained by a pooling equilibrium in which the seller requests a signal with at most 3 signal realizations.*

Corollary 1 shows that even when there are many seller types, the ex ante consumer and seller payoffs can be obtained in equilibria in which the seller requests a coarse data that features at most 3 labels. We stress however that these results are concerned with the *ex ante* players' payoffs, averaged across all seller types. Richer data, with many signal realizations, may induce richer payoff profiles across types because it could enable more diverse responses from different types of the seller.

Uniform Type Distribution. The analysis so far has established the general properties of equilibria. However, it is silent about a few key questions, such as when the seller's private information enables the consumer to have a positive rent, or how the set of equilibria depends on possible values, L and H. To tackle these questions and obtain further intuitions, 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 the consumer believes that all seller types are equally likely.

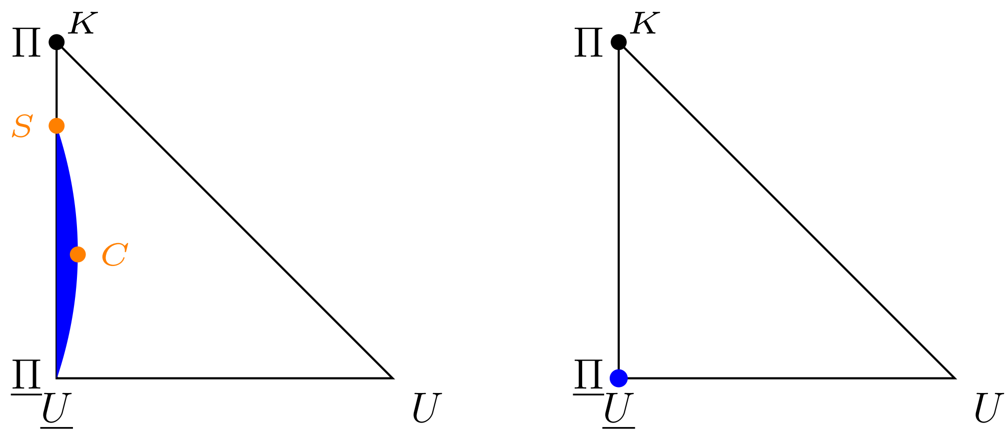

(a) $L/H = 1/3$ (b) $L/H = 2/3$

Figure 2: Equilibrium outcomes with seller types uniformly distributed on $[0, 1]$. Black lines bound feasible, individually rational outcomes. Blue color marks equilibrium outcomes. Points S and C are the seller- and the consumer-preferred equilibria, respectively. Point K is the equilibrium outcome when the seller's type is observed by the consumer.

Proposition 5. (Uniform Types) *Let the seller type be uniformly distributed on* $[0, 1]$.

1. *If* $\frac{L}{H} \geq \frac{1}{2}$, *then the equilibrium outcome is unique and achieved by no data collection.*

2. *If* $\frac{L}{H} < \frac{1}{2}$, *then the set of equilibrium outcomes has a non-empty interior. Its Pareto frontier is spanned by pooling equilibria with high-value targeting: For every* $\alpha \in [0, 1]$, *there is an equilibrium that maximizes* $\alpha U + (1 - \alpha)\Pi$ *in which all seller types collect signal* $\mathcal{I} = (S, \pi)$ *such that* $S = \{\ell, h\}$, $\pi(\ell|L) = 1$, *and* $\pi(h|H) = \rho_\alpha \in (0, 1)$, *where* ρ_α *is increasing in* α. *In the consumer-optimal equilibrium the seller obtains a signal with* $\rho_1 = \frac{H-2L}{H-L}$.

The proof of Proposition 5 uses the technique developed for Proposition 3. The ex ante payoffs of the seller and the consumer are linear in the profile Π of the seller's interim payoffs. Thus, there is an extreme point of implementable outcomes that maximizes a social welfare, i.e., a convex combination of the two ex ante payoffs. Since any extreme point is characterized by $(x, y) \in \mathbb{R}^2$ (see Figure 1) we can reduce the social welfare maximization problem to a two-dimensional maximization problem, which we can solve in a closed form.

Intuitively, data collection affects the consumer's interim payoffs in two ways. On the one hand, data collection can benefit the consumer by persuading some seller types $\theta > \hat{\theta}$ to set price L even when the consumer has value H. On the other hand, data collection may harm the consumer if $\theta < \hat{\theta}$ because such seller types would set price L without data. Since all equilibria are essentially pooling, any equilibrium data collection increases the consumer's interim payoffs for some seller types and decrease them for other types. If $\hat{\theta}$ is high, i.e., a large fraction of seller types set price L without data, then the negative effect dominates and no data collection arises as a unique outcome (Part 1). If $\hat{\theta}$ is low, there is a continuum of equilibria with a non-degenerate Pareto frontier (Part 2).

Figure 2 depicts the set of equilibrium outcomes. Both panels show that the seller's private information keeps equilibrium outcomes far away from the efficient outcomes. In the left panel, the set of equilibrium outcomes has a non-empty interior. The seller-preferred equilibrium (Point S) maximizes both the seller's profits and total surplus, but gives the consumer the same payoff as under no data. As the social welfare places a larger weight α on the consumer's payoff, the welfare-maximizing outcome moves along the frontier of the equilibrium outcomes toward the southeast direction, eventually hitting Point C, the consumer-optimal outcome. Correspondingly, as α increases, the equilibrium signal becomes more informative, i.e., ρ_α in the proposition increases in α. In the right panel, the seller's private information prevents any data collection and decreases profits without changing consumer surplus.

4 Implications of Seller's Private Information

We have shown that any equilibrium is essentially pooling, i.e., all seller types obtain the same, partially informative signal. Because the equilibrium signal is never fully revealing, a type-θ seller with a large θ sometimes sets price H when the consumer has value L, which leads to efficiency loss. The seller's private information thus decreases efficiency. It also decreases profits, because if the consumer observed θ, the seller profits would be the efficient total surplus minus the consumer's payoff under no data. Thus in some equilibrium—such as the seller-preferred equilibrium or the equilibrium with no data collection—the seller's private information decreases profits without changing the consumer's payoffs, leading to a Pareto inferior outcome. However, in some equilibria the consumer enjoys a strictly higher payoff than when the seller's type is observed (see Figure 2). The following result summarizes this discussion (the following results assume that the seller has a continuum of types).

Corollary 2. (Outcome Limits) *In any equilibrium, total surplus and the seller's ex ante payoffs are strictly lower and the consumer's ex ante payoff is weakly higher when the seller's type is private than when it is observed by the consumer.*

Another implication of the above analysis is that the seller can be worse off by having more information prior to requesting data from the consumer:

Corollary 3. (Cookie Curse) *The seller can be worse off from being more informed about the consumer prior to the data request. Formally, regardless of which equilibrium is played, the seller's equilibrium payoffs are strictly lower when she has private type $\theta \sim F$ than when she has a known type $\mu_0 = \mathbb{E}_F[\theta]$, provided $\mu_0 > \hat{\theta}$. Under the same condition, the consumer's payoff is weakly greater.*

To see why Corollary 3 holds, observe that if the seller is not informed and has the prior belief μ_0 above $\hat{\theta}$, then she can request and collect full information from the consumer, extracting full surplus. However, if this seller becomes imperfectly informed and spreads her prior belief into $\theta \sim F$ then full information can never be collected afterwards and full surplus cannot be extracted. More generally, having more interim information may hurt the seller because it may impede collection of the additional data voluntarily supplied by the consumer. This means, for example, that an online platform may benefit from adopting a policy that restricts the platform's access to third-party web cookies. Such a restriction, while seemingly privacy friendly, could in fact harm users by enabling the platform to collect excessive amounts of data by direct requests.

5 Discussion

5.1 Policy Implication: Enhancing Consumer's Control over Data

The seller's private information limits the potential of data to improve efficiency and consumer welfare. We now ask whether a certain privacy regulation can improve such a situation.

In the spirit of the GDPR, we examine a regulation that renders a consumer full control of their personal data. Such a regulation changes the game as follows: First, the seller observes her type θ (we assume a continuum type space). Second, the consumer chooses a menu of signals $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$ without knowing v, θ, or s. The seller chooses a signal from the menu, observes a signal realization, and then sets a price. Finally, the consumer observes v and makes a purchase decision. We may interpret a menu $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$ as a choice of a regulator that specifies data available to the seller on behalf of the consumer.

Proposition 6. *Suppose that the consumer has full control over his data. If the seller's type is observed by the consumer, the equilibrium is efficient and gives the seller the same payoff as under no data collection. If the seller's type is private, the equilibrium obtains the same outcome as the consumer-optimal equilibrium of the original data collection game.*

The result implies that the seller's private information limits the effectiveness of a regulation that gives consumers control over their data. If the seller's type is observed, the consumer could provide different signals to different seller types in order to extract full value of data without sacrificing efficiency.[10] In contrast, when the seller's type is unobservable, the optimal menu for the consumer is to provide all seller types with a signal that the seller would collect in the consumer-optimal equilibrium of the original game. Thus even though the regulation eliminates other consumer-suboptimal outcomes, it does not improve consumer welfare beyond what could arise when the seller requests data. The regulation may also reduce total surplus. For example, if the seller's type is uniformly distributed, the seller-optimal equilibrium maximizes total surplus across all equilibria (see Figure 2(a)).[11] Thus the regulation, which obtains the consumer-optimum, may strictly decrease total surplus.

5.2 Equilibrium Refinement

The potential multiplicity of equilibria poses a question of whether stronger solution concepts or some criteria refine the set of equilibria. Restricting attention to a continuum of types,

[10]So that each type faces the consumer-optimal segmentation of Bergemann, Brooks, and Morris (2015).

[11]The proof of Proposition 5 shows that under the uniform type distribution the seller's profit and total surplus are maximized by the same equilibrium.

we show that the pooling equilibria of our model survive the intuitive criterion of Cho and Kreps (1987), and also consist of perfect conditional equilibrium distributions of Myerson and Reny (2020), which extend the notion of sequential equilibrium to infinite games (see Appendix F for omitted proofs).

Regarding the intuitive criterion, suppose that the seller sets the lowest optimal price at any posterior belief. Our game then becomes the following signaling game: First, the sender privately learns her type θ and chooses a signal to request. Second, the receiver decides whether to accept the request. Given the requested signal and acceptance decision, the payoffs of the sender and the receiver are those of the seller and the consumer, respectively, given the optimal pricing and purchase decision. The set of perfect Bayesian equilibria of this signaling game coincides with the set of equilibria (of our model) in which the seller sets the lowest optimal price after any history. This interpretation enables us to apply the intuitive criterion of Cho and Kreps (1987).

When the type space of the seller contains $\theta = 0$, all equilibria pass the intuitive criterion for a trivial reason: The seller type $\theta = 0$ is indifferent across all signals, and if the consumer believes that $\theta = 0$ then he is indifferent between providing and not proving any signal, and in particular, it is optimal for him to not provide data. These observations imply that no seller type has a "profitable deviation" described in the intuitive criterion. A similar argument holds even if we exclude degenerate types $\theta = 0$ and $\theta = 1$ so that the type space is $\Theta = (0, 1)$. In such a case, any pooling equilibrium of the above signaling game can be approximated by equilibria that pass the intuitive criterion.

Regarding the relationship to sequential equilibria, all equilibrium strategy profiles of our game are also perfect conditional equilibrium distributions of Myerson and Reny (2020). To see the intuition, take a pooling equilibrium in which the seller collects signal $\mathcal{I}$. We can then construct a sequence (a net, to be precise) of strategy profiles converging to the equilibrium strategy profile such that along the sequence, perturbations that request $\mathcal{I}' \neq \mathcal{I}$ are sufficiently likely to be from low seller types. As a result it is approximately optimal for the consumer to reject any deviant request, rendering the equilibrium strategy profile a perfect conditional equilibrium distribution.

5.3 General Multiple Values

There are many ways in which our baseline model can be extended. One natural extension is the setting with many possible consumer values, $V = \{v_1, ..., v_n\}$, $n > 2$. In that case, the seller's type is multidimensional and the comprehensive equilibrium analysis is not attainable because it requires solving a multidimensional screening problem. However, we show that our main qualitative findings extend to this setup.

In particular, the welfare results of Corollary 2 extend to this case verbatim. Assume that the seller type θ is drawn from a distribution that has a strictly positive density everywhere on $\Delta(V)$. We then have the following result.

Proposition 7. (Many Values) *The total surplus and the seller profits are strictly lower and the consumer's ex ante payoffs are weakly higher when the seller's type is private than when observed by the consumer. There exists a pooling equilibrium in which all seller types request and obtain signal $\mathcal{I}$ if and only if the signal $\mathcal{I}$ increases the consumer's payoff relative to no data collection, $\int_\Theta U(\theta, \mathcal{I})\mathrm{d}F \geq \int_\Theta U(\theta, \underline{\mathcal{I}})\mathrm{d}F$.*

The first part of the proposition states that the welfare implication of the seller's private information, established in the previous section, does not depend on the binary value assumption. The inefficiency of equilibrium relies on the observation that some seller types require signals that are arbitrarily close to the fully informative signal in order to price efficiently. Such an outcome leaves the consumer zero surplus, which cannot happen in equilibrium.

The second part of the proposition characterizes pooling equilibria but does not state that any equilibrium outcome can be obtained in a pooling equilibrium. Clearly, as in the case of binary valuations, the types' selection into different data requests limits the scope of differentiated data collection. However, in the general case of many values the limits are less stark and *some* differentiated data collection can be welfare enhancing. The following example illustrates.

Example 1. (Limits of Pooling Equilibria) Let the possible values be $V = \{L, M, H\} = \{1, 3, 4\}$. Let there be four seller types, $\theta_1 = (1/2, 1/4, 1/4)$, $\theta_2 = (1/2, 1/2, 0)$, $\theta_3 = (1/2, 0, 1/2)$, and $\theta_4 = (1, 0, 0)$, with the uniform type distribution $f = (1/4, 1/4, 1/4, 1/4)$. (Type θ_4 is a dummy to simplify the deterrence of off-path requests.) In the absence of additional data, the types set prices $p_1 = 3$, $p_2 = 3$, $p_3 = 4$, and $p_4 = 1$, which leads to the type payoffs $\Pi_1 = 3/2, \Pi_2 = 3/2, \Pi_3 = 2$, and $\Pi_4 = 1$ and consumer surplus $U_0 = 1/16$.

We now construct a semi-separating equilibrium that maximizes the consumer surplus across all equilibrium outcomes. Namely, let types θ_1, θ_2, and θ_3 request the following signals:

$$\begin{array}{c|cc} \mathcal{I}(\theta_1) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 1/2 & 1/2 \\ H & 1/2 & 1/2 \end{array}, \qquad \begin{array}{c|cc} \mathcal{I}(\theta_2) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 1/2 & 1/2 \\ H & 1/2 & 1/2 \end{array}, \qquad \begin{array}{c|cc} \mathcal{I}(\theta_3) & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 2/3 & 1/3 \\ H & 1/3 & 2/3 \end{array}, \tag{10}$$

so that each signal has two possible realizations with the likelihood function π as presented in the tabular form. Let type θ_4 request an arbitrary signal $\mathcal{I}(\theta_4)$ that is different from these signals. The consumer provides $\mathcal{I}(\theta_i)$ when requested and rejects all other requests by believing that they come from type θ_4.

It is straightforward to verify that this strategy profile forms an equilibrium. No type wants to deviate to any other data requested on the equilibrium path. Off-equilibrium requests are deterred by the consumer's rejection. The consumer benefits from each data request.

Moreover, the equilibrium allocation is efficient, the good is traded with probability one, and each seller type earns the same profits as in the absence of data collection. Indeed, in equilibrium, types θ_1 and θ_2 set prices $p = L$ and $p = M$ after signal realizations s_1 and s_2 respectively, type θ_3 sets prices $p = L$ and $p = H$ respectively, and type θ_4 always sets price $p = L$.[12] It follows that the consumer's surplus is maximal among all feasible allocations given the individual rationality of the seller, and thus maximal among all equilibrium outcomes.

However, no such outcome can be obtained in a pooling equilibrium because, as we show, any efficient pooling equilibrium must give strictly positive rents to type θ_1. Indeed, for signal $\mathcal{I} = (S, \pi)$ to induce efficient allocation, 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 that $\pi(s_0|L) \geq 3\pi(s_0|H)$ to persuade type θ_3; these incentive

[12] Under type θ_4, there are no consumers with valuation $v = M$, so her setting the price $p = H$ after s_2 doesn't hurt efficiency.

constraints together guarantee that θ_1 also charges low price. Since type θ_1 never charges a price below $p = L$, her profit decreases as one increases the frequency of signal s_0, i.e., as one increases $\pi(s_0|L)$, $\pi(s_0|M)$, and $\pi(s_0|H)$ subject to the incentive constraints. Therefore, among all efficient pooling equilibria the profit of type θ_1 is minimized at the signal

$$\begin{array}{c|cc} \hat{\mathcal{I}} & s_1 & s_2 \\ \hline L & 1 & 0 \\ M & 1/2 & 1/2 \\ H & 1/3 & 2/3 \end{array}, \tag{11}$$

which happens to be efficient and maximizes the probability of type θ_1 setting the price $p = L$ while ensuring that she serves the rest of the consumers at 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 the lower consumer surplus than under the semi-separating equilibrium presented above. □

It should not be surprising that the differentiated data collection can be beneficial in the general case with many values. After all, this case features multi-dimensional seller types and gives rise to a richer structure of incentive-compatible allocations. However, our main economic messages continue to hold: The uncertainty about the seller's type limits the scope of data collection and can improve consumer welfare.

6 Conclusion

Despite the prevalence of personal data collection in the digital economy, consumers typically know less than companies about the benefits and costs of data collection. We study the implication of such information asymmetry in a stylized model of data collection by an informed seller—a seller requests data from a consumer to tailor pricing decision, with the seller privately informed about how data affect product prices. We show that all equilibrium outcomes are spanned by pooling equilibria in which all seller types collect the same data. The seller's private information and the resulting lack of differentiated data collection hurt efficiency and profits. At the same time, the consumer's skepticism about the seller's type prevents her from collecting profitable data that harm the consumer. As a result, the consumer's payoff is higher when the seller has private information. The seller can increase profits by committing to not collect data without the consumer's consent. Such a commitment, however, can harm consumers.

References

Acquisti, A., C. Taylor, and L. Wagman (2016): "The Economics of Privacy," *Journal of Economic Literature*, 54, 442–92.

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*, MIT Press.

Auxier, B., L. Rainie, M. Anderson, A. Perrin, M. Kumar, and E. Turner (2019): "Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information," *Pew Research Center*, accessed at https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/.

Bergemann, D., A. Bonatti, and T. Gan (forthcoming): "The Economics of Social Data," *RAND Journal of Economics*.

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.

Cho, I.-K. and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102, 179–221.

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.

Doval, L. and A. Smolin (2021): "Information Payoffs: An Interim Perspective," *Working Paper*.

Fainmesser, I. P., A. Galeotti, and R. Momot (2021): "Digital Privacy," *Working paper*.

Fudenberg, D. and J. Tirole (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236–260.

Haghpanah, N. and R. Siegel (2019): "Pareto Improving Segmentation of Multi-product Markets," *Working paper*.

——— (2021): "The Limits of Multi-Product Price Discrimination," *American Economic Review: Insights*.

Ichihashi, S. (2020): "Online Privacy and Information Disclosure by Consumers," *American Economic Review*, 110, 569–95.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kleiner, A., B. Moldovanu, and P. Strack (2021): "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89, 1557–1593.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85, 1949–1964.

Myerson, R. B. (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 1767–1797.

Myerson, R. B. and P. J. Reny (2020): "Perfect Conditional ε-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions," *Econometrica*, 88, 495–531.

Phelps, R. R. (2001): *Lectures on Choquet's Theorem*, Springer.

Rayo, L. and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

Rhodes, A. and J. Zhou (2021): "Personalized Pricing and Privacy Choice," *Working paper.*

Roesler, A.-K. and B. Szentes (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107, 2072–80.

Shi, X. and J. Zhang (2020): "Welfare of Price Discrimination and Market Segmentation in Duopoly," *Working paper.*

Smolin, A. (2020): "Disclosure and Pricing of Attributes," *Working paper.*

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

Appendix

A Proof of Proposition 2: Equilibrium Outcomes with Binary Types

Most of the proof is presented in the main text. It is left to show that the incentive compatibility implies the likelihood monotonicity in each state: $\alpha_2 \geq \alpha_1$ and $\beta_2 \geq \beta_1$. The system of necessary incentive constraints (3) can be written as:

$$(1-\alpha_1)L + \theta_1(\alpha_1 L + \beta_1(H-L)) \geq (1-\alpha_2)L + \theta_1(\alpha_2 L + \beta_2(H-L)), \tag{12}$$

$$(1-\alpha_2)L + \theta_2(\alpha_2 L + \beta_2(H-L)) \geq (1-\alpha_1)L + \theta_2(\alpha_1 L + \beta_1(H-L)). \tag{13}$$

Summing over the inequalities (12) and (13) and using the fact that $\theta_2 > \theta_1$ we obtain:

$$\alpha_2 L + \beta_2(H - L) \geq \alpha_1 L + \beta_1(H - L). \tag{14}$$

In turn, (12) and (14) together imply that $\alpha_2 \geq \alpha_1$:

$$(\alpha_2 - \alpha_1)L \geq \theta_1(\alpha_2 L + \beta_2(H - L) - \alpha_1 L - \beta_1(H - L)) \geq 0. \tag{15}$$

Finally, (13) and (15) imply that $\beta_2 \geq \beta_1$:

$$(\beta_2 - \beta_1)\theta_2(H - L) \geq (\alpha_2 - \alpha_1)(1 - \theta_2)L \geq 0. \tag{16}$$

It follows that the pooling signal (4) is feasible, which completes the proof.

B Proof of Proposition 3: Equilibrium Outcomes with a Continuum of Types

Auxiliary Mechanism-Design Problem

In the data collection stage, the seller has two kinds of deviations: The seller can request a deviant signal that other seller types request in equilibrium, or a deviant signal that no seller type requests. To understand the seller's incentive regarding the former deviation, we study an auxiliary mechanism-design problem in which a fictitious principal commits to a mechanism, without transfers, that provides the seller with a signal as a function of her reported type. After obtaining a signal, the seller sets a price and the consumer makes an optimal purchase decision. The set of all implementable outcomes in this mechanism-design problem provides an upper bound on the set of equilibrium outcomes in the data collection game. In what follows, we fully characterize this set and show that it, in fact, can be spanned by pooling equilibria.

To characterize implementable outcomes, we can without loss of generality focus on mechanisms that take the form $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$. If the seller reports θ, then she obtains signal characterized by $(r_\theta(h|H), r_\theta(h|L))$. This signal sends realization h or ℓ, which is a recommendation to set price equal to H or L, respectively. For $s \in \{h, \ell\}$ and $v \in \{H, L\}$, $r_\theta(s|v) \in [0,1]$ is the probability with which the signal draws realization s given the realized value v.

A mechanism is *incentive-compatible* if the seller finds it optimal to report her true type and follow the price recommendation. We define the *obedient* seller as the seller who commits to follow the price recommendation but strategically reports her type. (We borrow terminologies from Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017).)

A profile of the seller's interim payoffs $(\Pi(\theta))_{\theta\in\Theta}$ is *implementable* if there exists an incentive-compatible mechanism in which a type-θ seller obtains payoff $\Pi(\theta)$. The following lemma characterizes all implementable payoff profiles. To state the result, let $\underline{\Pi}(\theta) := \max\{\theta H, L\}$ denote the seller's optimal revenue from no information, and let $\overline{\Pi}(\theta) := \theta H + (1-\theta)L$ denote the optimal revenue from perfect information about the value.

Lemma 1. *A profile of seller's interim payoffs* $(\Pi(\theta))_{\theta\in\Theta}$ *is implementable if and only if* $\Pi(\cdot)$ *is a convex function such that*

$$\forall \theta \in [0,1], \quad \underline{\Pi}(\theta) \le \Pi(\theta) \le \overline{\Pi}(\theta). \tag{17}$$

A mechanism of the form $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ *that implements* $\Pi(\cdot)$ *is unique up to a measure-zero subset of* $\Theta = [0,1]$.

Proof. First, we show the "only if" direction of the first part. Inequalities (17) hold because the seller's interim payoff for any type is between the revenue under no information and the revenue under perfect information. Moreover, in any incentive-compatible mechanism,

$$\begin{aligned}
\Pi(\theta) &= \theta\left[Hr_\theta(h|H) + Lr_\theta(\ell|H)\right] + (1-\theta)Lr_\theta(\ell|L) \\
&= \theta\left[Hr_\theta(h|H) + L(1-r_\theta(h|H))\right] + (1-\theta)L(1-r_\theta(h|L)) \\
&= \theta\left[(H-L)r_\theta(h|H) + L\right] + (1-\theta)L(1-r_\theta(h|L)) \\
&= \theta\left[(H-L)r_\theta(h|H) + Lr_\theta(h|L)\right] + L(1-r_\theta(h|L)).
\end{aligned}$$

The incentive-compatibility constraint for the obedient seller is

$$\Pi(\theta) \geq \theta \left[H r_{\theta'}(h|H) + L r_{\theta'}(\ell|H) \right] + (1-\theta) L r_{\theta'}(\ell|L), \forall \theta, \theta' \in \Theta.$$

The standard argument of mechanism design implies that

$$\Pi(\theta) = L + \int_0^\theta \left[(H-L) r_x(h|H) + L r_x(h|L) \right] dx, \forall \theta \in [0,1] \tag{18}$$

and that $(H-L)r_\theta(h|H) + Lr_\theta(h|L)$ is non-decreasing in θ. Thus, $\Pi(\cdot)$ is convex.

To show the "if" direction, take any Π that is convex and between $\underline{\Pi}$ and $\overline{\Pi}$. Fix any $\theta \in [0,1]$. We show that there is a pair $(r_\theta(h|H), r_\theta(h|L)) \in [0,1]^2$ such that

$$\Pi(\theta) = \theta \left[(H-L) r_\theta(h|H) + L r_\theta(h|L) \right] + L(1 - r_\theta(h|L)), \tag{19}$$

$$\Pi'(\theta) = (H-L) r_\theta(h|H) + L r_\theta(h|L). \tag{20}$$

and that the pair is unique up to a measure-zero subset of $[0,1]$. In equation (20), whenever Π is not differentiable at θ (which occurs only at countably many θ's) we take Π' as the right derivative. Plugging (20) into (19) and rearranging the resulting equation, we obtain

$$\Pi(\theta) = \theta \Pi'(\theta) + L(1 - r_\theta(h|L)) \tag{21}$$

$$\iff r_\theta(h|L) = \frac{\theta \Pi'(\theta) - \Pi(\theta) + L}{L}. \tag{22}$$

Plugging it further into (20) we obtain

$$\Pi'(\theta) = (H-L) r_\theta(h|H) + \theta \Pi'(\theta) - \Pi(\theta) + L \tag{23}$$

$$\iff r_\theta(h|H) = \frac{\Pi(\theta) + (1-\theta)\Pi'(\theta) - L}{H-L}. \tag{24}$$

We now show that $(r_\theta(h|H), r_\theta(h|L))$ is a well-defined recommendation rule. First, we show $r_\theta(h|L) \geq 0$, which is written as $\theta\Pi'(\theta) - \Pi(\theta) + L \geq 0$. Because Π is convex and $\Pi(0) = L$, we have $\Pi(\theta) = L + \int_0^\theta \Pi'(x)dx \leq L + \theta\Pi'(\theta)$, so we obtain the desired inequality. Second, we show $r_\theta(h|L) \leq 1$, which is written as $\theta\Pi'(\theta) - \Pi(\theta) \leq 0$. If $\theta \geq \frac{L}{H}$, $\Pi'(\theta) = H$ and

$\Pi(\theta) = \theta H$ maximize $\Pi'(\cdot)$ and minimize $\Pi(\cdot)$. Even in such a case, we have $\theta H - \theta H = 0$. Thus $\theta\Pi'(\theta) - \Pi(\theta) \leq 0$ for all $\theta \geq \frac{L}{H}$. If $\theta < \frac{L}{H}$, $\Pi'(\theta) = \frac{H-L}{1-\theta}$ and $\Pi(\theta) = L$ maximize $\Pi'(\cdot)$ and minimize $\Pi(\cdot)$. Then $\theta\Pi'(\theta) - \Pi(\theta) \leq 0$ becomes $\theta \leq \frac{L}{H}$, which is true. Thus $r_\theta(h|L) \leq 1$. Third, we show $r_\theta(h|H) \geq 0$, which is written as $\Pi(\theta) + (1-\theta)\Pi'(\theta) - L \geq 0$, which is true because $\Pi(\theta) \geq L$ and $\Pi'(\theta) \geq 0$. Finally, we show $r_\theta(h|H) \leq 1$, which is written as $H \geq \Pi(\theta) + (1-\theta)\Pi'(\theta)$. The inequality holds because $H \geq \Pi(\theta) + \int_\theta^1 \Pi'(x)dx \geq \Pi(\theta) + (1-\theta)\Pi'(\theta)$.

Next, we show $r_\theta(h|H) \geq r_\theta(\ell|L)$, which is written as

$$
\begin{aligned}
&\frac{\Pi(\theta) + (1-\theta)\Pi'(\theta) - L}{H-L} \geq \frac{\theta\Pi'(\theta) - \Pi(\theta) + L}{L} \\
\iff & L\Pi(\theta) + L(1-\theta)\Pi'(\theta) - L^2 \geq (H-L)L - H\Pi(\theta) + H\theta\Pi'(\theta) + L\Pi(\theta) - L\theta\Pi'(\theta) \\
\iff & L\Pi'(\theta) \geq HL - H\Pi(\theta) + H\theta\Pi'(\theta) \\
\iff & \Pi(\theta) - L \geq \Pi'(\theta)\left(\theta - \frac{L}{H}\right).
\end{aligned}
$$

For $\theta \geq \frac{L}{H}$, the inequality holds because it holds when Π is minimized and Π' is maximized, i.e., $\Pi(\theta) = \theta H$ and $\Pi'(\theta) = H$. For $\theta < \frac{L}{H}$, the inequality holds because the left-hand side is non-negative but the right-hand side becomes negative. Thus, the signal defined by equations (19) and (20) are well defined.

Finally, we show that $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ is incentive-compatible. First, suppose that the seller reports her true type. Because the seller has binary choices (price H or L), it is optimal to follow the price recommendation if and only if doing so is better than ignoring information. Because $\Pi(\theta) \geq \underline{\Pi}(\theta)$, the seller prefers to follow the price recommendation, having reported her true type. Suppose that the seller misreports her type. Because $\Pi(\theta) \geq \underline{\Pi}(\theta)$, misreporting and ignoring information reduce her payoff. Misreporting and following the price recommendation also reduce her payoff, because the envelope formula (18) and the monotonicity of $(H-L)r_\theta(h|H) + Lr_\theta(h|L)$ (by (20)) ensure the incentive compatibility for the obedient seller. Finally, misreporting and doing the opposite of the price recommendation (e.g., setting price H after realization ℓ) reduce the seller's payoff, because $r_\theta(h|H) \geq r_\theta(h|L)$. Thus mechanism $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ is incentive-compatible, and the seller's interim payoff at θ is $\Pi(\theta)$ because of (19).

To conclude the proof, observe that if mechanism $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ is incentive-compatible and implements $(\Pi(\theta))_{\theta\in\Theta}$, then equation (19) must hold for all θ and equation (20) for almost all θ. Thus, a mechanism of the form $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ that implements $(\Pi(\theta))_{\theta\in\Theta}$ is unique except possibly on a measure-zero subset of Θ. $\square$

The next lemma shows that any implementable outcome can be implemented by a mechanism that allocates to the seller the same signal, irrespectively of the announced type.

Lemma 2. *Take any profile $(\Pi(\theta))_{\theta\in\Theta}$ that is implementable by some mechanism. Then it is implementable by a mechanism in which all types obtain the same signal. Moreover, the consumer's ex ante payoffs are necessarily the same in both mechanisms.*

Proof. The proof uses results on majorization of Kleiner, Moldovanu, and Strack (2021). Recall that $\underline{\Pi}(\theta) = \max\{L, \theta H\}$ is the seller's optimal revenue under no information, and its right derivative is given by $\underline{\Pi}'(\theta) = 0$ for $\theta < \frac{L}{H}$ and $\underline{\Pi}'(\theta) = H$ if $\theta \geq \frac{L}{H}$.

We show that the set of implementable outcomes coincides with the set of functions majorized by $\underline{\Pi}'$, $MPS(\underline{\Pi}')$, where in the notation of Kleiner, Moldovanu, and Strack (2021),

$$MPS(\underline{\Pi}') := \{\Pi' : [0,1] \to \mathbb{R} : \Pi' \text{ is right continuous, non-decreasing, and majorized by } \underline{\Pi}'\}.$$

Indeed, take any implementable profile $\Pi = (\Pi(\theta))_{\theta\in\Theta}$. We have $\underline{\Pi}(x) \leq \Pi(x)$ for all $x \in [0,1]$ and $\underline{\Pi}(1) - \underline{\Pi}(0) = H - L$. These conditions are equivalent to

$$\int_x^1 \Pi'(s)ds \leq \int_x^1 \underline{\Pi}'(s)ds, \forall x \in [0,1] \quad \text{and} \quad \int_0^1 \Pi'(s)ds = \int_0^1 \underline{\Pi}'(s)ds, \tag{25}$$

respectively, and thus $\Pi' \in MPS(\underline{\Pi}')$.

Conversely, take any $\Pi' \in MPS(\underline{\Pi}')$ and define $\hat{\Pi}(x) = L + \int_0^x \Pi'(s)ds$ for all x. First, Π' is non-decreasing, so $\hat{\Pi}$ is convex. Second, $\int_x^1 \Pi'(s)ds \leq \int_x^1 \underline{\Pi}'(s)ds$ and $\int_0^1 \Pi'(s)ds = \int_0^1 \underline{\Pi}'(s)ds$ together imply $\int_0^x \underline{\Pi}'(s)ds \leq \int_0^x \Pi'(s)ds$, which is equivalent to $\underline{\Pi}(x) - L \leq \int_0^x \Pi'(s)ds$ resulting in $\underline{\Pi}(x) \leq L + \int_0^x \Pi'(s)ds = \hat{\Pi}(x)$. Finally, suppose to the contrary that $\hat{\Pi}(x) > \overline{\Pi}(x)$ for some $x \in (0,1)$. Because $\hat{\Pi}(0) = \overline{\Pi}(0) = L$, the convexity of $\hat{\Pi}$ and the linearity of $\overline{\Pi}$ on $[0,1]$ imply $\hat{\Pi}(y) > \overline{\Pi}(y)$ for all $y \in [x,1]$. Then, $\int_0^1 \Pi'(s)ds = \hat{\Pi}(1) - \hat{\Pi}(0) > \overline{\Pi}(1) - \overline{\Pi}(0) = H - L = \underline{\Pi}(1) - \underline{\Pi}(0)$, which is a contradiction. Thus, $\hat{\Pi}$ is implementable.

We now show that the principal can implement every extreme point of $MPS(\underline{\Pi}')$ by providing the same signal to all seller types. Theorem 1 of Kleiner, Moldovanu, and Strack (2021) implies that for any extreme point Π_e' of $MPS(\underline{\Pi}')$, there is an interval $[\underline{\theta}, \overline{\theta}) \subset [0,1]$ such that

$$\Pi_e'(\theta) = \begin{cases} 0 & \text{if} \quad \theta \leq \underline{\theta}, \\ \frac{\int_{\underline{\theta}}^{\overline{\theta}} \underline{\Pi}'(\theta)d\theta}{\overline{\theta}-\underline{\theta}} & \text{if} \quad \theta \in [\underline{\theta}, \overline{\theta}), \\ H & \text{if} \quad \theta \geq \overline{\theta}. \end{cases}$$

The implementable profile Π_e that corresponds to Π_e' is defined by $\Pi_e(x) = L + \int_0^x \Pi_e'(s)ds$. In terms of Π_e, the above condition implies that $\Pi_e(\theta)$ equals $\underline{\Pi}(\theta)$ if $\theta < \underline{\theta}$ or $\theta \geq \overline{\theta}$, and $\Pi_e(\theta)$ is a line that connects $(\underline{\theta}, \underline{\Pi}(\theta))$ and $(\overline{\theta}, \underline{\Pi}(\theta))$ on $[\underline{\theta}, \overline{\theta})$. Plugging Π_e into (22) and (24), we obtain the following:

$$r_\theta(h|L) = \begin{cases} 0 & \text{if} \quad \theta \leq \underline{\theta}, \\ \frac{\underline{\theta}\int_{\underline{\theta}}^{\overline{\theta}} \underline{\Pi}'(\theta)d\theta}{(H-L)(\overline{\theta}-\underline{\theta})} & \text{if} \quad \theta \in [\underline{\theta}, \overline{\theta}), \\ 1 & \text{if} \quad \theta \geq \overline{\theta}, \end{cases} \quad r_\theta(h|H) = \begin{cases} 0 & \text{if} \quad \theta \leq \underline{\theta}, \\ \frac{(1-\underline{\theta})\int_{\underline{\theta}}^{\overline{\theta}} \underline{\Pi}'(\theta)d\theta}{(H-L)(\overline{\theta}-\underline{\theta})} & \text{if} \quad \theta \in [\underline{\theta}, \overline{\theta}), \\ 1 & \text{if} \quad \theta \geq \overline{\theta}. \end{cases}$$

Note that signal $(r_\theta(h|H), r_\theta(h|L))$ is uninformative if $\theta \leq \underline{\theta}$ or $\theta > \overline{\theta}$. Thus, we can implement Π_e' by publicly disclosing information about v according to

$$\pi(h|L) = \frac{\underline{\theta}\int_{\underline{\theta}}^{\overline{\theta}} \underline{\Pi}'(\theta)d\theta}{(H-L)(\overline{\theta}-\underline{\theta})} \tag{26}$$

$$\pi(h|H) = \frac{(1-\underline{\theta})\int_{\underline{\theta}}^{\overline{\theta}} \underline{\Pi}'(\theta)d\theta}{(H-L)(\overline{\theta}-\underline{\theta})}. \tag{27}$$

The incentive compatibility implies that a seller type $\theta < \underline{\theta}$ or $\theta \geq \overline{\theta}$ ignores this information and behaves according to $(r_\theta(h|H), r_\theta(h|L))$.

Now, take any $\Pi' \in MPS(\underline{\Pi}')$. The set $MPS(\underline{\Pi}')$ is compact and convex in the norm topology. Choquet's theorem implies that we can find a probability measure w on the set of extreme points of $MPS(\underline{\Pi}')$ such that

$$\Pi' = \int \Pi'_e dw(e). \tag{28}$$

Here, the integral is a Bochner integral and the equality means that $V(\Pi') = \int V(\Pi'_e)dw(e)$ for any continuous, linear functional V on $MPS(\underline{\Pi}')$ (see Phelps (2001) and Footnote 14 of Kleiner et al. (2021)). We take $V : f \mapsto \int_0^\theta f(x)dx$ for any $\theta \in [0,1]$ and obtain

$$\Pi(\theta) = \int \Pi_e(\theta)dw(e), \forall \theta \in [0,1]. \tag{29}$$

Equation (29) means that the designer can implement the seller's interim payoffs Π by providing a signal that randomizes over different signals that implement extreme points Π_e according to probability measure w regardless of the reported type.

Finally, the last part of Lemma 1 implies that if two mechanisms implement the same profile $\Pi(\cdot)$, then the seller's pricing behaviors are the same for almost every $\theta \in [0,1]$. As a result, the consumer's ex ante payoffs are the same in the two mechanisms. □

Equilibrium Characterization

We are now ready to prove Proposition 3.

Proof. The "if" direction of Part 1 holds by definition. To show the "only if" direction, take any equilibrium E and let Π denote the seller's interim payoff function. Let $\mathcal{I}_\theta$ be the signal that type θ obtains in the equilibrium; U^* be the consumer's ex ante payoff in the equilibrium; and $U_\emptyset$ be his ex ante payoff when all seller types set prices optimally given no information. If the consumer rejects the request by a type-θ seller in the equilibrium, we define $\mathcal{I}_\theta = \underline{\mathcal{I}}$, as an uninformative signal.

The profile Π is implementable in a mechanism that allocates signal $\mathcal{I}_\theta$ given report θ, because the seller's sequential rationality in equilibrium E implies incentive compatibility of the mechanism. Lemma 2 then implies that there exists a signal $\mathcal{I}^*$ such that if all seller types obtain $\mathcal{I}^*$ and set the price optimally, the resulting seller payoffs equal Π. We use signal $\mathcal{I}^*$ to construct the following pooling equilibrium E_P: All seller types send the request for signal $\mathcal{I}^*$, which the consumer accepts. Whenever the seller deviates and requests a different signal, the consumer believes that the seller's type lies in $[0, \frac{L}{H})$ and rejects the request. This strategy profile and a belief system form an equilibrium. Indeed, it is optimal for the consumer to reject any deviant request because the seller would set the lowest price, equal to L, absent additional information. It is also optimal for the consumer to accept the on-path request $\mathcal{I}^*$. If the consumer does so, his payoff is U^*. If the consumer rejects the request, his payoff is $U_\emptyset$. However, the last part of Lemma 2 implies that the consumer's ex ante payoffs are the same in equilibria E and E^*. In the original equilibrium E, the consumer must prefer his equilibrium strategy to rejecting all requests so it must be that $U^* \geq U_\emptyset$. Therefore, any equilibrium outcome can arise in a pooling equilibrium.

The "if" direction of Part 2 holds because in equilibrium the consumer weakly prefers accepting the request and providing signal $\mathcal{I}$ to rejecting the request. The "only if" direction holds by the same argument as in the previous paragraph. □

C Proof of Proposition 5: Uniform Type Distribution

Proof. Recall from the proof of Proposition 3 that an implementable profile Π of the seller's interim payoffs determines the seller's pricing behavior $(r_\theta(h|H), r_\theta(h|L))_{\theta\in\Theta}$ for almost every θ through equations (22) and (24). The consumer obtains a positive payoff iff his value is H but the seller sets price L, which occurs with an interim probability of

$$\theta(1 - r_\theta(h|H)) = \theta\left(1 - \frac{\Pi(\theta) + (1-\theta)\Pi'(\theta) - L}{H - L}\right) = \theta\frac{H - \Pi(\theta) - (1-\theta)\Pi'(\theta)}{H - L}.$$

The consumer's ex ante payoff is then

$$\int_{[0,1]} \theta\left[H - \Pi(\theta) - (1-\theta)\Pi'(\theta)\right] \mathrm{d}F(\theta).$$

Given Π, the seller's ex ante payoff is, by definition, $\int_{[0,1]} \Pi(\theta)\mathrm{d}F(\theta)$. The ex ante payoffs of the consumer and the seller are linear in Π. Thus the social surplus,

$$\alpha \left[\int_{[0,1]} \theta \left[H - \Pi(\theta) - (1-\theta)\Pi'(\theta)\right] \mathrm{d}F(\theta)\right] + (1-\alpha)\int_{[0,1]} \Pi(\theta)\mathrm{d}F(\theta),$$

is maximized by an extreme point of the set of implementable Π's.

We now find an extreme point that maximizes the social surplus. By the arguments presented in the proof of Lemma 2, for any extreme point Π there exist $x \leq \frac{L}{H}$ and $y \geq \frac{L}{H}$ such that $\Pi(\theta) = \underline{\Pi}(\theta)$ if $\theta \leq x$ or $\theta \geq y$ (Figure 1). We can then identify the set of extreme points of $MPS(\underline{\Pi}')$ with

$$X := \left\{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq \frac{L}{H} \leq y \leq 1\right\}. \tag{30}$$

For each $(x, y) \in X$, the corresponding extreme point of $MPS(\underline{\Pi}')$ (in terms of profits, not their derivatives) is

$$\Pi(\theta) = \begin{cases} L & \text{if} \quad \theta \leq x \\ L + \frac{Hy-L}{y-x}(\theta - x) & \text{if} \quad x < \theta < y \\ \theta H & \text{if} \quad y \leq \theta. \end{cases}$$

Plugging this expression into the consumer's payoffs for the uniform type distribution, we obtain

$$\begin{aligned}
U(x,y) \triangleq& \int_0^1 \theta \left[H - \Pi(\theta) - (1-\theta)\Pi'(\theta)\right] \mathrm{d}F(\theta) \\
=& \int_0^x \theta \left[H - L\right] \mathrm{d}F(\theta) + \int_x^y \theta \left[H - L - \frac{Hy-L}{y-x}(\theta - x) - (1-\theta)\frac{Hy-L}{y-x}\right] \mathrm{d}F(\theta) \\
=&(H-L)\int_0^x \theta \mathrm{d}\theta + \left(H - L - \frac{(Hy-L)(1-x)}{y-x}\right)\int_x^y \theta \mathrm{d}\theta \\
=&(H-L)\int_0^y \theta \mathrm{d}\theta - \frac{(Hy-L)(1-x)}{y-x}\int_x^y \theta \mathrm{d}\theta \\
=&\frac{1}{2}(H-L)y^2 - \frac{1}{2}(Hy-L)(1-x)(y+x).
\end{aligned} \tag{31}$$

Similarly, the seller's ex ante payoff is

$$
\begin{aligned}
\Pi(x,y) \triangleq & \int_0^1 \Pi(\theta)\mathrm{d}\theta \\
=& \int_0^x L\mathrm{d}\theta + \int_x^y L + \frac{Hy-L}{y-x}(\theta - x)\mathrm{d}\theta + \int_y^1 \theta H \mathrm{d}\theta \\
=& Ly + \frac{1}{2}(Hy-L)(y+x) - (Hy-L)x + \frac{H}{2}(1-y^2).
\end{aligned} \tag{32}
$$

Define the social surplus $W(x,y) \triangleq \alpha U(x,y) + (1-\alpha)\Pi(x,y)$. Because $U(x,y)$ is strictly convex in x and $\Pi(x,y)$ is convex in x, $W(x,y)$ is maximized at the extreme values $x=0$ or $x=\frac{L}{H}$. The case $x=\frac{L}{H}$ is redundant since that outcome corresponds to no data collection and can equivalently be captured by $(x,y)=(0,\frac{L}{H})$. Hence, we can set $x=0$, and our problem reduces to $\max_{y\in[\frac{L}{H},1]} W(0,y)$, which we solve next.

First, consider the case $\frac{L}{H} \geq \frac{1}{2}$. We have $\frac{\partial U}{\partial y}(0,y) = (H-L)y - Hy + \frac{1}{2}L = L(\frac{1}{2}-y)$, which is negative for all $y > \frac{L}{H}$. As $y \in [\frac{L}{H},1]$, the consumer's payoff $U(0,y)$ is uniquely maximized at $y=\frac{L}{H}$ which corresponds to no data collection $(x,y)=(0,\frac{L}{H})$. It follows that no data collection is a unique equilibrium outcome. To the contrary, suppose that there exists a pooling equilibrium in which the consumer obtains payoff $U(0,\frac{L}{H})$ and the seller obtains a payoff strictly higher than $\Pi(0,\frac{L}{H})$. Let Π denote the seller's interim payoff profile in the equilibrium. Function Π cannot be an extreme point of $MPS(\underline{\Pi}')$, so it is a non-trivial convex combination of the extreme points. Some of the extreme points in the convex combination must differ from no data collection and thus give the consumer a strictly lower payoff than no data collection, which is a contradiction. Thus if $\frac{L}{H} \geq \frac{1}{2}$ the equilibrium outcome is unique and corresponds to no data collection.

Second, consider the case $\frac{L}{H} < \frac{1}{2}$. In that case $U(0,y)$ is strictly concave in y and is maximized at $y=\frac{1}{2}$, Thus, there is a unique $\overline{y} > \frac{1}{2}$ such that $U(0,\frac{L}{H}) = U(0,\overline{y})$. The function $W(0,y)$ is strictly concave in y, and the first-order condition is

$$\begin{aligned}\frac{\partial W(0,y)}{\partial y} &= \alpha\left[(H-L)y - \frac{1}{2}(2Hy-L)\right] + (1-\alpha)\left[L + \frac{1}{2}(2Hy-L) - Hy\right]\\ &= -\alpha L y + \frac{1}{2}L\\ &= 0,\end{aligned}$$

which implies that a unique maximizer of $W(0,y)$ is

$$y = \frac{1}{2\alpha}.$$

If $\alpha < \frac{1}{2\overline{y}}$, then the welfare-maximizing equilibrium is a pooling equilibrium in which all types collect a signal that corresponds to $(x,y) = (0,\overline{y})$.

If $\alpha \in [\frac{1}{2\overline{y}}, 1]$, then $\frac{1}{2\alpha} \leq \overline{y}$ so the welfare-maximizing equilibrium is a pooling equilibrium in which all seller types collect a signal that corresponds to $(x,y) = (0, \frac{1}{2\alpha})$, specifically:

$$\pi(h|L) = \frac{\theta\Pi'(\theta) - \Pi(\theta) + L}{L} = \frac{\theta\frac{Hy-L}{y-x} - L - \frac{Hy-L}{y-x}\theta + L}{L} = 0 \tag{33}$$

and

$$\pi(h|H) = \frac{L + \frac{Hy-L}{y-x}\theta + (1-\theta)\frac{Hy-L}{y-x} - L}{H-L} = \frac{Hy-L}{y(H-L)} = \frac{H\min(\overline{y}, \frac{1}{2\alpha}) - L}{\min(\overline{y}, \frac{1}{2\alpha})(H-L)} \in (0,1). \tag{34}$$

Equivalently, the resulting signal takes the form of $\pi(\ell|L) = 1$ and $\pi(\ell|H) = 1 - \rho_\alpha \in (0,1)$. Different $\alpha \in (\frac{1}{2\overline{y}}, 1)$ correspond to distinct equilibrium outcomes that deliver to the seller and the consumer the payoffs strictly higher than under no data collection. As no data collection remains a possible equilibrium outcome, the set of equilibrium outcomes has a non-empty interior whenever $\frac{L}{H} < \frac{1}{2}$. $\square$

D Proof of Corollary 2: Welfare Implications of Private Information

Proof. Without loss of generality (by Proposition 3), take any pooling equilibrium in which the seller obtains signal $\mathcal{I} = (S, \pi)$. Because $\mathcal{I}$ is not fully informative, there is a realization $s \in S$ that can arise under both $v = H$ and $v = L$ with positive probabilities. After observing s, the seller with a sufficiently large θ sets price H, which leads to no trade when $v = L$, i.e., any equilibrium is inefficient. The second part of the result holds because if the seller's type is known, the outcome is efficient and the consumer's payoff is the payoff under no data collection. Finally, take any pooling equilibrium. If signal $\mathcal{I}$ gives the consumer a strictly higher payoff than under no data collection, then we can find another pooling equilibrium in which the seller collects signal $\mathcal{I}$ and $\overline{\mathcal{I}}$ with probabilities close to 1 and 0, respectively. This equilibrium gives the seller a strictly higher payoff; therefore the seller-preferred equilibrium gives the consumer the same payoff as under no data collection. $\square$

E Proof of Proposition 7: General Multiple Values

Proof. The proof of the second part follows from the case of binary values: Take any signal $\mathcal{I}$ that weakly increases the consumer's payoff. There is a pooling equilibrium in which all seller types obtains $\mathcal{I}$, and the consumer rejects any deviant request, believing that it comes from seller types that would set price $\min V$ under no data. Conversely, if the seller collects a signal in a pooling equilibrium then it must weakly increase the consumer's ex ante payoff because of the consumer's sequential rationality.

For the first part, we show that if the consumer observes the seller's type, the equilibrium is efficient and the consumer's payoff is the same under no data. Indeed, if the consumer knows that the seller's type is θ then the seller can request a signal that obtains the efficient outcome and gives the consumer the same payoff as under no data. Such a signal exists for any θ because of Bergemann, Brooks, and Morris (2015). Because the seller can also choose a signal that attains profits arbitrarily close to $\mathbb{E}_{v\sim\theta}[v] - U(\theta, \underline{\mathcal{I}})$ but gives the consumer a strictly higher payoff than $U(\theta, \underline{\mathcal{I}})$, the unique equilibrium outcome is such that the allocation is efficient and the consumer's payoff is the same as under no data.

In contrast, when the seller's type is unobservable, any equilibrium is inefficient. Without loss of generality assume $v_1 > v_2 > \cdots > v_n$. For each $\varepsilon \in (0,1)$, define a seller type θ_ε as $\theta_\varepsilon(v_k) = \frac{\varepsilon^{k-1}}{\sum_{k=1}^n \varepsilon^{k-1}}$. Suppose to the contrary that there is an efficient equilibrium. Let $\mathcal{I}_\varepsilon = (V, \pi_\varepsilon)$ denote the signal such that $\pi_\varepsilon(v|v')$ (i.e., the probability of signal realization v when the true value is v') is the probability with which a type-θ_ε seller sets price v when the true value is v' in the efficient equilibrium. For every $k > 1$, a type-θ_ε seller prefers price v_k to price v_1 given signal realization v_k:

$$\begin{aligned} & v_1 \cdot \theta_\varepsilon(v_1)\pi_\varepsilon(v_k|v_1) \le v_k \cdot \theta_\varepsilon(v_1)\pi_\varepsilon(v_k|v_1) + v_k(1-\theta_\varepsilon(v_1)) \\ \Longleftrightarrow\ & \pi_\varepsilon(v_k|v_1) \le \frac{v_k(1-\theta_\varepsilon(v_1))}{(v_1 - v_k)\theta_\varepsilon(v_1)}, \end{aligned}$$

which approaches 0 as $\varepsilon \to 0$ because $\frac{1-\theta_\varepsilon(v_1)}{\theta_\varepsilon(v_1)} \to 0$. Thus as $\varepsilon \to 0$, $\pi_\varepsilon(v_k|v_1) \to 0$ for all $k > 1$. Next, for any $k > 2$, a type-θ_ε seller prefers price v_k to price v_2 given signal realization v_k:

$$\begin{aligned} & v_2 \cdot [\theta_\varepsilon(v_2)\pi_\varepsilon(v_k|v_2) + \theta_\varepsilon(v_1)\pi_\varepsilon(v_k|v_1)] \\ \le & v_k \cdot [\theta_\varepsilon(v_2)\pi_\varepsilon(v_k|v_2) + \theta_\varepsilon(v_1)\pi_\varepsilon(v_k|v_1)] + v_k(1-\theta_\varepsilon(v_1) - \theta_\varepsilon(v_2)), \end{aligned}$$

which is equivalent to

$$(v_2 - v_k)\pi_\varepsilon(v_k|v_2) + (v_2 - v_k)\frac{\theta_\varepsilon(v_1)}{\theta_\varepsilon(v_2)}\pi_\varepsilon(v_k|v_1) \le v_k \frac{1-\theta_\varepsilon(v_1)-\theta_\varepsilon(v_2)}{\theta_\varepsilon(v_2)},$$

which implies

$$\pi_\varepsilon(v_k|v_2) \le \frac{v_k}{v_2 - v_k} \cdot \frac{1-\theta_\varepsilon(v_1)-\theta_\varepsilon(v_2)}{\theta_\varepsilon(v_2)}.$$

As $\varepsilon \to 0$, the right-hand side goes to 0, and thus $\pi_\varepsilon(v_k|v_2) \to 0$. Repeating the same procedure, we can show that $\lim_{\varepsilon\to 0} \pi_\varepsilon(v_k|v_\ell) = 0$ for all $k > \ell$, i.e., the probability with which the seller sets a strictly lower price than the true value goes to 0 as $\varepsilon \to 0$. Because the equilibrium is efficient, the probability with which the seller sets a strictly higher price than the true value is always 0. As a result, as $\varepsilon \to 0$, $\pi_\varepsilon(v_k|v_k) \to 1$ for all $k = 1, \ldots, n$. A type-θ_ε seller obtains a signal that is more weakly informative than $\mathcal{I}_\varepsilon$; otherwise, she would not be able to price according to $\mathcal{I}_\varepsilon$. Thus in any efficient equilibrium, we can find seller types that obtain signals arbitrarily close to the fully informative one. This in turn implies that all seller types extract the efficient total surplus; otherwise, the seller can deviate to a signal that is sufficiently close to the fully informative signal. Such an equilibrium gives the consumer a strictly lower payoff than under no data, so we obtain a contradiction.

To conclude the proof, observe that when the seller has private information, the equilibrium total surplus is strictly lower than the benchmark of known types, and the consumer's ex ante payoff must be at least the payoff under no data. Thus the seller's profit is strictly lower when her type is private. □

F Appendix for Section 5.2: Equilibrium Refinement

Intuitive Criterion

We consider a signaling game described in Section 5.2 and show that all equilibria pass the intuitive criterion of Cho and Kreps (1987). First, assume $\Theta = [0, 1]$. Take any perfect Bayesian equilibrium and any off-path signal $\mathcal{I}$. According to the intuitive criterion, we first consider the set Θ' of types that are strictly worse off by requesting $\mathcal{I}$ than following the equilibrium choice, when the consumer holds most favorable beliefs after observing $\mathcal{I}$. Observe that $0 \notin \Theta'$ because seller type 0 is indifferent across all signals. The intuitive criterion then considers the seller-worst best responses of the consumer that are consistent with his beliefs supported on $\Theta \setminus \Theta'$. This set contains $\theta = 0$, so according to the criterion, the consumer is allowed to believe that $\theta = 0$ and optimally reject data request $\mathcal{I}$. If so, the seller cannot earn strictly higher payoffs by requesting $\mathcal{I}$ than following the equilibrium choice. As a result, all equilibria pass the intuitive criterion.

The above argument is not specific to the existence of a degenerate type $\theta = 0$. To see this, suppose that the type space is an open interval, $\Theta = (0, 1)$. We consider two cases: First, suppose that the only equilibrium outcome is no data collection. Then no types are excluded in the first step of the intuitive criterion. In the second step, if the consumer believes that the seller's type is $\theta < \hat{\theta}$ then he prefers not to provide any data. Given such a response, the seller never obtains a higher payoff by requesting a non-equilibrium signal. Thus the equilibrium passes the intuitive criterion.

Second, suppose that some pooling equilibrium delivers to the consumer a strictly higher payoff than under no data collection. Let $\mathcal{I}^*$ denote the equilibrium signal in such an equilibrium. Take any pooling equilibrium and let $\mathcal{I}$ denote the equilibrium signal. For any $n, m \in \mathbb{N}$, we construct the following pooling equilibrium, E_n. Let Π_n denote the seller's interim profits when she obtains signals $\mathcal{I}^*$ and $\mathcal{I}$ with probabilities $\frac{1}{n}$ and $1 - \frac{1}{n}$, respectively. The outcome corresponding Π_n gives the consumer a strictly higher payoff than under no data collection, so there is a pooling equilibrium consistent with Π_n. Then define $\hat{\Pi}_n(\theta) \triangleq \max(\Pi_n(\theta) - \frac{1}{m(n)}, \underline{\Pi}(\theta))$, where $m(n) \geq n$ is sufficiently large so that the pooling outcome that corresponds to $\hat{\Pi}$ gives the consumer a higher payoff than under no data collection. For any n, the pooling equilibrium that corresponds to $\hat{\Pi}_n$ passes the intuitive criterion. Note that in this equilibrium, we have $\hat{\Pi}_n(\theta) = \underline{\Pi}(\theta)$ for a sufficiently small $\theta > 0$, so such types are not excluded in the first step of the intuitive criterion. By the same argument as above, the consumer can believe that the seller's type is sufficiently close 0 (and thus $\theta < \hat{\theta}$) and optimally reject any "deviant" request in the second step of the criterion. By taking $n \to \infty$, the outcomes in $\{E_n\}_{n=1}^{\infty}$ converge to the outcome of the original equilibrium with $\mathcal{I}$.

Perfect Conditional Equilibrium Distribution (Myerson and Reny, 2020)

Any pooling equilibrium of the data collection game is a perfect conditional equilibrium distribution of Myerson and Reny (2020), who extend the idea of sequential equilibrium to infinite games (see their paper for the definition and terminologies). To see this, take any pooling equilibrium E^* and let $\mathcal{I}^*$ denote the equilibrium signal. Take any $\varepsilon > 0$. Take any finite set of signals H that does not contain $\mathcal{I}^*$. For each $n \in \mathbb{N}$ we define the seller's data request strategy $s^{n,H}$ as follows: Pick any $\delta > 0$. Any seller types $\theta < \delta n^{-1}$ request signal $\mathcal{I}^*$ with probability $1 - \frac{1}{n}$ and each signal $\mathcal{I} \in H$ with probability $\frac{1}{n|H|}$. Any seller types $\theta \geq \delta n^{-1}$ request signal $\mathcal{I}^*$ with probability $1 - \delta n^{-1}$ and each signal $\mathcal{I} \in H$ with probability $\frac{\delta}{n|H|}$. Given H and n, choose δ sufficiently small so that the consumer's payoff difference between providing and not providing signal $\mathcal{I} \in H$ conditional on that the seller adopts $s^{n,H}$ and requests $\mathcal{I}$ is less than ε. Such a construction is possible because as $\delta \to 0$, the consumer's belief after observing $\mathcal{I}$ converges to a degenerate distribution at $\theta = 0$. The net $(s^{n,H})_{n \in \mathbb{N}, H}$ is an admissible net of the seller's data request strategies in the sense of Myerson and Reny (2020), and the tails of the net give every signal positive probability. Here, the partial order on the set of indices is such that (n, H) is greater than (n', H') if $n \geq n'$ and $H \supset H'$. The construction of the perturbation for the consumer's request approval, the seller's pricing, as well as nature's draw of the seller's type, signal realizations, and the consumer's values are straightforward. By keeping the probabilities of these perturbations sufficiently small, the resulting net of the strategy profiles and nature perturbations are admissible and consist of a conditional ε-equilibrium. Because the limit of this net is always the strategy profile under E^*, the strategy profile is a perfect conditional equilibrium distribution of Myerson and Reny (2020).